

Representation of Women in HPC Conferences

Eitan Frachtenberg
eitan@reed.edu
Reed College
Portland, Oregon, USA

Rhody D. Kaner
rkaner25@gmail.com
Reed College
Portland, Oregon, USA

ABSTRACT

Women are acutely underrepresented in the HPC workforce. Addressing this gap requires accurate metrics on the representation of women and its associated factors. The goal of this paper is to provide current, broad, and reproducible data on this gender gap. Specifically, this study provides in-depth statistics on women’s representation in HPC conferences, especially for authors of peer-reviewed papers, who serve as the keystone for future advances in the field.

To this end, we analyzed participant data from nine HPC and HPC-related peer-reviewed conferences. In addition to gender distributions, we looked at post-publication citation statistics of the papers and authors’ research experience, country, and work sector.

Our main finding is that women represent only 10% of all HPC authors, with large geographical variations and small variations by sector. Representation is particularly low at higher experience levels. This 10% ratio is lower than even the 20–30% ratio in all computer science.

CCS CONCEPTS

- **Social and professional topics** → **Industry statistics; Gender; Geographic characteristics**; Computing occupations;
- **General and reference** → Metrics.

KEYWORDS

high-performance computing (HPC), women in science, underrepresentation, bibliometrics

ACM Reference Format:

Eitan Frachtenberg and Rhody D. Kaner. 2021. Representation of Women in HPC Conferences. In *The International Conference for High Performance Computing, Networking, Storage and Analysis (SC '21)*, November 14–19, 2021, St. Louis, MO, USA. ACM, New York, NY, USA, 12 pages. <https://doi.org/10.1145/3458817.3476164>

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.

SC '21, November 14–19, 2021, St. Louis, MO

© 2021 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-8442-1/21/11... \$15.00

<https://doi.org/10.1145/3458817.3476164>

1 INTRODUCTION

The gender gap in numerous scientific fields is an active area of research [8, 9, 19]. The field of computer science (CS) is well-known for its gender disparity and consequent significant societal effects, such as inequality in economic opportunities for women and an undersupply of researchers and engineers in the rapidly growing discipline [34, 36].

The gender gap among CS researchers is particularly noteworthy: the people who participate in research, publish about it, and have their research acknowledged for its value are predominantly men [6]. Numerous studies estimate that only about 20%–30% of the CS research community are women [9, 24, 35, 51–53]. Although some recent indications show these numbers could be growing, they remain low, and the rate of growth remains slow [51].

In some subfields of CS, the gender gap has not yet been thoroughly studied [16]. One example is High-Performance Computing (HPC), a large and vibrant research field with significant impact on the economy and the sciences, including CS, Physics, Biology, Epidemiology, Climate, and others.

Despite the importance of the field, less than 17% of its workforce is estimated to be women, although the exact number is difficult to ascertain [10, 19, 47]. Consequently, it is the stated goal of the organizers of leading HPC conferences to increase the diversity of its attendees and researchers. For example, SC’s website from 2017 (our year of study) states the following on the benefits of diversity:

As the leading international HPC conference SC has an ongoing commitment to diversify the HPC community. Our commitment to diversity aims to broaden participation from all aspects of the community, and provide a more inclusive experience for everyone.

In 2015, 14% of attendees were women – our goal is [to] increase this to 20% by 2020. Supercomputing is publishing our demographics. In addition, other leading HPC organizations around the world are also doing workforce demographics. By working together to publicly document today’s workforce, we can begin to understand what is needed to properly diversify the workforce and create a more inclusive working environment.

These contributions are a positive step in creating a more diverse group of attendees, presenters, authors, and exhibitors in the years to come. There are other contributions that can be made to increase our efforts, and they start with you. Start a conversation within your organization

about how you can become part of the effort, publish your data, and share your results. We cannot get to where we are going if we do not know from where we are starting.¹

This paper heeds the call to publish such data. Its primary goal is to reduce some of the uncertainty around gender representation by computing reproducible metrics within the HPC community, which can serve as a benchmark against which future progress can be measured. Although our current study cannot fully address the complex question of why women remain underrepresented in HPC, it attempts to quantify this underrepresentation with two goals in mind: to facilitate subsequent study of its causes and to establish a baseline against which to measure the effects of future efforts and policies to increase gender diversity.

The technical objectives and contributions of this paper are to quantify more precisely the ratio of women in various HPC conference roles, and to examine the relationships to various demographic and conference factors. We estimate the rate of women’s participation in HPC research by using the proxy metric of the female author ratio (FAR) in a set of peer-reviewed systems conferences. This popular approach has been previously tested with various researcher populations [13, 29, 33]. In addition to computing gender ratios, we also collected and analyzed demographic and bibliometric data in an attempt to expose and examine how these factors correlate with female researcher ratios.

The next section explains our data and methodology in detail, and in particular, how we produced accurate estimates of FAR. We then dive into the data to look at the ratio of women researchers from three perspectives: conferences, papers, and researchers (demographics). Finally, the last section of this paper concludes and suggests future work.

2 METHODOLOGY AND DATA

As in most subfields of CS, the primary channel for publishing research results in HPC is peer-reviewed conferences [18, 20, 38, 48, 49]. The delineation of HPC within CS is not always clear, so different conferences will have varying numbers of HPC-specific papers, sometimes depending on the subjective views of the reader. For this study, we collected broad data from nine HPC-focused and HPC-related conferences (Table 1). We chose to focus on a cross-section of HPC conferences from a single publication year (2017) to reduce time-related variance, but we also look at other years for the two flagship conferences, SC and ISC.

The SC conference takes place in the US every November and is the largest HPC conference in the world, attracting thousands of attendees. SC is also notable for its diversity initiatives, including an inclusivity chair, on-site childcare, a code of conduct, and diversity-related content, which have all been suggested as ways to attract women to the conference [10]. The other flagship HPC conference, ISC, takes place in Germany every summer and also enlists an inclusivity chair.

Table 1: HPC-related conferences, including start date, number of published papers, total number of published authors, acceptance rate, and country code.

Conference	Date	Papers	Authors	Acceptance	Country
CCGrid	2017-05-14	72	296	0.252	ES
IPDPS	2017-05-29	116	447	0.228	US
ISC	2017-06-18	22	99	0.333	DE
HPDC	2017-06-28	19	76	0.190	US
ICPP	2017-08-14	60	234	0.286	UK
EuroPar	2017-08-30	50	179	0.284	ES
SC	2017-11-13	61	325	0.187	US
HiPC	2017-12-18	41	168	0.223	IN
HPCC	2017-12-18	77	287	0.438	TH

The seven other conferences we investigate are not exclusive to HPC, but contain numerous HPC papers. None of them displayed diversity initiatives for 2017.

For each conference, we downloaded all papers and gathered information about all authors, program committee (PC) members, and other roles. Because our main metric for the gender gap counts the percentage of women among authors, we collected the names and author positions of all 1885 unique coauthors. Systems papers typically list the primary contributor in the leading (first) position and senior authors last, so we examined the gender of both first and last authors. Additionally, we collected information on researchers in the following conference roles:

- program committee (PC) chairs, who coordinate the review activities (36 total).
- PC members, who conduct most of the paper reviews and therefore have a direct influence on which papers get accepted (908 total).
- Keynote speakers (30 total), panelists (106 total), and session chairs (158 total), who have no direct influence on the population of authors, but who represent the “face” of the conference to attendees. The visibility of women in such role models may have an indirect impact or appeal for women practitioners [10–12].

For this study, the most critical piece of information on these researchers is their *perceived gender*. Gender is a complex, multifaceted identity [27], and understanding its nuances can improve science and engineering [26, 44]. However, most bibliometric studies still rely on binary genders—either collected by the journal, or inferred from first name—because that is the only designator available to them [3, 9, 24, 35, 51–53]. In the absence of self-identified gender information for our population, we also necessarily compromised on using binary gender designations. We therefore use the gender terms “women” and “men” interchangeably with the sex terms “female” and “male”.

The conferences in our dataset did not collect or share gender information of participants, so we had to collect this information ourselves. Similar studies have typically used automated gender-inference services based on forename and

¹<https://sc17.supercomputing.org/inclusivity/research-and-tools/index.html>

sometimes country of origin [25, 28]. These statistical approaches can be reasonably accurate for names of Western origin, and especially for male names, but less accurate for women and names of Asian origin [9, 33, 39].

For greater accuracy, we opted instead to rely primarily on a manual approach that can overcome the limitations of name-based inference. We manually assigned the gender of 95.18% of the researchers, for whom we could identify an unambiguous web page with a recognizable gendered pronoun or absent that, a photo.² For 1.79% others, we used genderize.io’s automated gender designations if it was at least 70% confident about them [39]. The remaining 144 persons were not assigned a gender and were excluded from most analyses. This method provided more gender data and higher accuracy than automated approaches based on forename and country, especially for women [28, 29, 33, 43, 51]. This labor-intensive approach does introduce the prospect of human bias and error. Nevertheless, based on a separate author survey we conducted where we found no discrepancies between assigned gender and self-selected gender, we believe such errors to be limited.

Conferences also do not generally offer information on authors’ demographics, but we were able to unambiguously link approximately two thirds (68.3%) of researchers in our dataset to a Google Scholar (GS) profile [50]. For each author and PC member, we collected all metrics in their GS profile, such as total previous publications (ca. 2017), h-index, etc. Note that we found no GS profile for about a third of the researchers (31.71%), and these researchers appear to be less experienced than researchers with a profile. We therefore collected a comparable proxy metric for author experience, total number of past publications, from a different source, the Semantic Scholar database [21].

We also looked up each author’s affiliation institute on GS to find their country of residence and work sector whenever they could be unambiguously inferred, using hand-coded regular expressions. We combined this data with regional information from <https://github.com/mlledoze/countries>.

Many authors also included their email address in the full text of the paper, from which we inferred more timely affiliation and country information than from GS. From authors’ affiliations, we broadly categorized their work sector as either “COM” for industry (8.6% of all unique authors and PC members), “EDU” for academia, (72.8%), and “GOV” for government and national labs (18.6%).

In addition to researcher information, we gathered various statistics on each conference, either from its web page, proceedings, or directly from its chairs. We collected data about review policies, important dates, the composition of its technical PC, and the number of submitted papers, among others. Finally, from each conference’s web site and proceedings we collected information on any explicit policies the conference made to increase attendance diversity, such as a diversity chair or code of conduct.

²For example, many LinkedIn profiles may lack a photo, but include a gendered pronoun in the recommendations section.

Statistics

For statistical testing, group means were compared pairwise using Welch’s two-sample t-test; differences between distributions of two categorical variables were tested with the χ^2 test; and comparisons between two numerical variables were evaluated with Pearson’s product moment correlation coefficient. All statistical tests are reported with their p-values and degrees of freedom where applicable.

Code and data availability

The complete dataset and source code necessary to reproduce this analysis can be found in the artifact appendix, as well as at <https://github.com/eitanf/sysconf>. The specific analyses of this article are in the file `pubs/whpc/women-hpc.Rmd`.

Limitations

Our study uses the FAR proxy metric to estimate women’s participation in systems research, as do comparable studies estimating the gender gap in other fields [13, 29, 33]. FAR has been found to correlate tightly with the gender ratios across scientific disciplines [24]. Nevertheless, it is important to keep in mind that FAR may undercount women if men are more likely to submit papers or have them accepted. We found no evidence to suggest that such undercounting would be worse for HPC than for other fields in comparable studies. We also considered the opposite hypothesis, that men may be undercounted in our dataset. Again, our data does not provide any strong evidence for this possibility, which runs counter to the findings of women’s underrepresentation in other scientific fields [3, 5].

In the literature, we found few controlled experiments that evaluate the peer-review process on both accepted and rejected papers, and they are typically limited in scope to a single conference or journal [2, 37, 40, 46]. We chose an observational approach of accepted papers that allowed us to sample the entire HPC field of study and produce metrics that are comparable with those in other fields. The main limitation of this approach is that it may miscount women if there is significant gender bias in the publication or review processes. That said, employing the same approach as other studies implies that the resulting statistics are directly comparable to theirs. Moreover, our survey results indicate that such peer-review bias may be limited [17].

Our methodology is also constrained by the manual collection of data. The effort involved in compiling all the necessary data limits the scalability of our approach to additional conferences or years. Furthermore, the manual assignment of genders is a laborious process, prone to human error. Nevertheless, such errors appear to be smaller in quantity and bias than those of automated approaches, as discussed previously.

Even with manual gender assignment, 3.03% of researchers still have unassigned gender. Although this ratio is small, and smaller than that of most other studies we reviewed, we nevertheless performed a sensitivity analysis to examine its effect. We first artificially set the gender of all 144 unassigned researchers to women, and then to men, and recomputed

all statistical analyses. None of our observations were subsequently changed in either direction or statistical significance, which justified our decision to omit these missing data points from the analysis.

3 CONFERENCE METRICS

We analyze the conferences in our set through the lens of the ratio of women in different roles (Fig. 1). In this section, we first look at the ratio of women across all conferences and papers in our set, and then focus on the flagship HPC conferences only.




Figure 1: Representation of women across conference roles.

3.1 Authors

We start by looking at the ratio of female authors (FAR). Since authors are peer-reviewed and ideally selected from the overall researcher population based on the merit of their papers alone, this sample population may be the most representative of the overall female population across publishing HPC researchers. The result, if indeed indicative of the overall population, points to a particularly low representation of women. About 9.9% of all 2236 authors were women, significantly lower than the 20–30% rates for the rest of CS [9, 51, 53]. It is particularly low for the two HPC flagship conferences, with only 8.12% female authors in SC and 5.77% in ISC (excluding the few authors for whom we have no gender information).

It is possible these results are too pessimistic because we are only looking at accepted papers, rather than all submitted papers. For example, survivorship bias [41] or peer-review bias [31] would lead to a higher percentage of women among submitted papers than the percentage we observe across accepted papers. However, variations in the visibility of the author’s gender do not support this hypothesis. For example, SC and ISC are the only double-blind conferences in our dataset, where the identity of authors are hidden from reviewers, and yet these conferences show the lowest female representation (7.57% in SC and ISC combined vs. 10.52% in the other conferences, $\chi^2 = 3.133, p = 0.0767$). In contrast, the ratio of female lead authors in the seven single-blind conferences, where their gender could potentially be more prominent, is generally no lower than the female author ratio, and often higher.

As for gender bias in the peer-review process, we might expect it to lead to fewer women in lead author positions when their identity is known to the reviewer. Instead, the overall ratio of female lead authors in single-blind conferences is almost double that of double-blind conferences (although the difference is not statistically significant, so we cannot completely rule out review bias without additional information on rejected papers: 11.79% vs. 6.17%, $\chi^2 = 1.662, p = 0.197$).

The last author position in systems papers is often reserved for the head of the research lab or the most senior researcher in the team. The representation of women in last author position is slightly worse than that of the overall author population (8.4% vs. 9.9%, $\chi^2 = 0.724, p = 0.395$). Again, the difference is statistically nonsignificant, but it is nevertheless consistent with other studies that found that women continue to senior academic ranks at a lower rate than men [10, 14, 16, 19, 34, 42].

3.2 Program Committee

Next, we examine gender representation among two of the most important elected roles of conference participants, the program committee (PC) chairs and members. These roles have direct bearing on the technical content and author selection of a conference, and therefore have strong influence over the author population. Most conferences only have a handful of PC and area chairs, so it is difficult to draw statistical

conclusions from the smaller sample size. Nevertheless, it is instructive to note that four of the nine conferences appointed no women at all to serve as chairs.

Among the 1220 total PC members, 18.46% are women (with repeats, meaning that the same person is counted multiple times if they serve on more than one PC). This rate is about twice the rate among authors. The SC conference invited the most women to its PC, both in absolute numbers (225) and in relative terms (29.6%). But even excluding the data from SC, the ratio of women among PCs is still 16.1%, significantly higher than that of authors.

Conceivably, women are not as poorly represented in this population because conferences often aspire to intentionally increase diversity, which can directly lead to selecting more women for this role. SC's example supports this explanation, with its explicit push for diversity. Then again, the other conference with a diversity chair, ISC, had a more average ratio of women in the PC ($\approx 16\%$).

Even if we assume that the PC ratios are more representative of the overall female HPC population than the authors' statistics, they are still far from equal. They are also likely insufficient on their own as catalysts to increase the ratios among authors, as the two metrics appear to be unrelated.

3.3 Visible Roles

The last three conference roles we look at are keynote speakers, session chairs, and panelists. These roles are collateral to the main technical content of a conference, but they can be very visible to conference attendees, and as such, can represent the "face" of the conference [11, 12, 15, 30]. A conference with low female visibility in these roles could conceivably further deter women from joining or staying in the conference or field. Like PC chairs, keynote speakers are few, and once more we see four conferences with no women at all in this role. Even more striking is the finding that three conferences had zero female session chairs out of a total of 45 session chairs: HPDC, HPCC, and HiPC. Only SC shows a ratio that is approaching gender parity, again perhaps because of its explicit push for diversity and inclusivity. As for panelists, the samples sizes are too small for statistical analysis.

3.4 Case Studies: SC and ISC

Both SC and ISC are the largest HPC conferences (by attendance). Both also have the distinction, among our set, of appointing a diversity/inclusivity chair, collecting and reporting demographic data, requiring attendees to adhere to a code of conduct, and in SC's example, also providing childcare and special meetings on diversity. All of these diversity-focused policies likely contribute to a more inclusive conference experience for attendees [1, 4, 10]. The SC conference also exhibited a higher ratio of women in invited roles than other conferences.

Regrettably, neither conference has consistently shared author demographics to evaluate changes over time. We therefore expanded our data to compute FAR in both conferences

for the five year period starting in 2016. Throughout this period, women's attendance rate at SC remained near constant at around 13%-14% (SC only shared FAR for 2018 at 12%). For ISC, FAR values were in the range of 5%-9%, lower than the average conference in our dataset.

How are these low rates reconciled with these conferences' focus on inclusivity policies? We can think of a few potential explanations:

- It is possible that these policies were in fact more reactive than proactive, in an attempt to improve previous statistics.
- It is also possible that their effects can only be measured over longer periods of times.
- In a specialized research field like HPC, it is also plausible that inclusivity initiatives are only one of the selection criteria when choosing a conference to publish in, and that other criteria such as conference date, location, prestige, and subfield take precedence.

Another possible motivation for these policies is that these initiatives do indeed help with boosting inclusivity in the short term, by improving the subjective experience of women, even if not their overall numbers [4]. For example, ISC reported in 2019 [1] that "89.6% of all surveyed attendees this year agreed that ISC is a conference that makes all attendees feel welcome".³

4 PAPERS

4.1 Topic

Some of the conferences in our set advertise a broader scope that includes topics that are incidental to HPC. We can therefore evaluate whether the representation of women changes when we limit our observations to strictly HPC papers. To this end, we skimmed each paper and tagged it as "HPC" if its topic related directly to high performance hardware or software.⁴ The resulting list of 178 papers is approximately a third of the complete list of 518 papers in our 9 conferences, and can help us determine what the representation of women looks like among paper authors when we restrict ourselves to only HPC research.

Of the 812 authors of HPC-only papers with known gender, 10.1% were women, slightly higher than the 9.9% in the overall conference author population ($\chi^2 = 4.656$, $p = 0.031$).

When we look specifically at lead authors, the picture is similar: of the 172 HPC papers for which we know the first author's gender, 19 are women (11.05%), a statistically nonsignificant difference from the overall ratio of 10.86% ($\chi^2 = 0.0547$, $p = 0.8151$).

These results suggest that gender representation among these HPC authors is in line with that of the overall representation across conferences in our set. Subsequently, we will

³Gender breakdown for this question was unavailable, so men could comprise the vast majority of respondents to this survey, as they did among authors. The report also does not specify whether the remaining respondents found the conference unwelcoming, or simply abstained from giving an opinion.

⁴The determination of this tag was based on our experience in HPC research and is therefore subjective and imprecise.

return to analyzing the complete set of papers throughout the rest of this study, for greater coverage.

4.2 Reception

One way to measure a paper’s influence is through its citations over time. We have allowed our dataset to age to the point where all papers could be discovered and cited by other researchers. The density plot in Fig. 2 shows the distributions of numbers of citations of the papers 36 months after publication, broken up by the gender of the lead author.

Many of the 53 papers with a woman as lead author are concentrated on the left side of the chart, with fewer citations than the 435 papers written by male first authors.

Women in lead author position average more citations than men overall (13.04 vs. 10.55). However, there is one outlier at the long tail of the distribution. This paper, led by a female author, garnered more than 450 citations as of this writing, but does not appear to be strongly related to HPC [32]. When we exclude this paper, the mean citations for women’s papers drops to 7.63, lower than men’s ($t = -2.18$, $df = 86$, $p = 0.032$).

In addition to citation count, fewer papers written by a female lead author accomplish wide recognition. Google Scholar tracks the number of papers that were cited 10 or more times as the i-10 index. By this metric, only 23% of female-led papers exceeded this threshold, compared to 38% for men ($\chi^2 = 3.69$, $p = 0.055$).

In both cases, the results may not be statistically significant and the sample size may be too small to establish a clear difference in citations by gender. It may be interesting to follow up on this analysis in regular intervals in the future and observe how the difference in reception behaves over time.

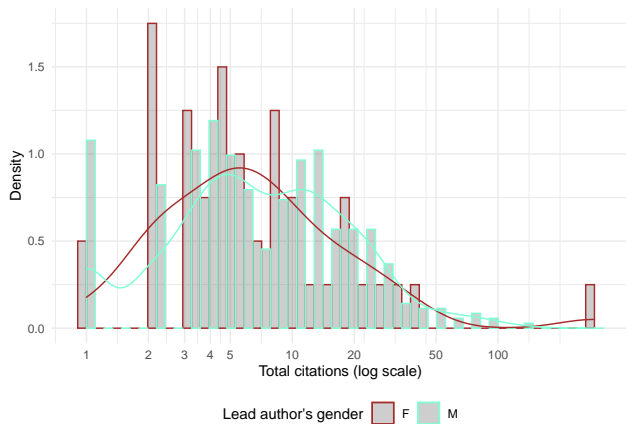


Figure 2: Distribution of paper citations three years after publication.

5 RESEARCHER DEMOGRAPHICS

There are a total of 3456 authors and PC members in our dataset, spanning numerous countries and institutions. We

can estimate the distribution of these demographics by looking at email addresses or GS affiliations. Using regular expressions, we mapped most domains to an institution’s country and sector, and marked the rest as unknown. We also collected the various GS bibliometrics for each researcher to estimate their research experience, which is our next analysis.

5.1 Research Experience

We can approximate the experience of researchers with proxy metrics such as total past publications or h-index [23]. We started by manually identifying the unique GS profile of researchers whenever possible (in our data, 69.65% of researchers with known gender had a uniquely identifiable GS profile). The profile contains various bibliometric measures that can approximate the experience and research influence of a person. For example, Fig. 3 depicts the distribution of the total number of previous publications (circa their conference’s date), while Fig. 4 shows the distribution of the h-index metric. We also collected data on the number of past publications of authors only from a different source, Semantic Scholar, which has 100% author coverage (Fig. 5). Google Scholar and Semantic Scholar use different data and algorithms for questions such as name disambiguation, resulting in low correlation ($r = 0.334$, $p < 0.0001$). Nevertheless, the overall similarity in distributions supports the following observations.

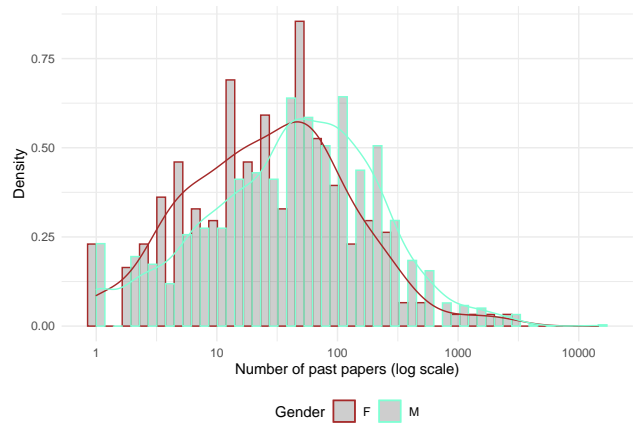


Figure 5: Distribution of authors’ past publications from Semantic Scholar by gender.

First, we note that all distributions are right-skewed, with a few researchers with thousands of publications while most have fewer than 100—not surprising, given that many authors are still students. PC members generally have more experience (publications) than authors, especially among women. This gap suggests that perhaps more of the female authors are novices, in relative terms. Another interesting observation is that the male authors’ distributions “pull to the right”. In other words, there appear to be relatively more male authors in experienced or senior positions. This disparity may be related to the observations that women do not

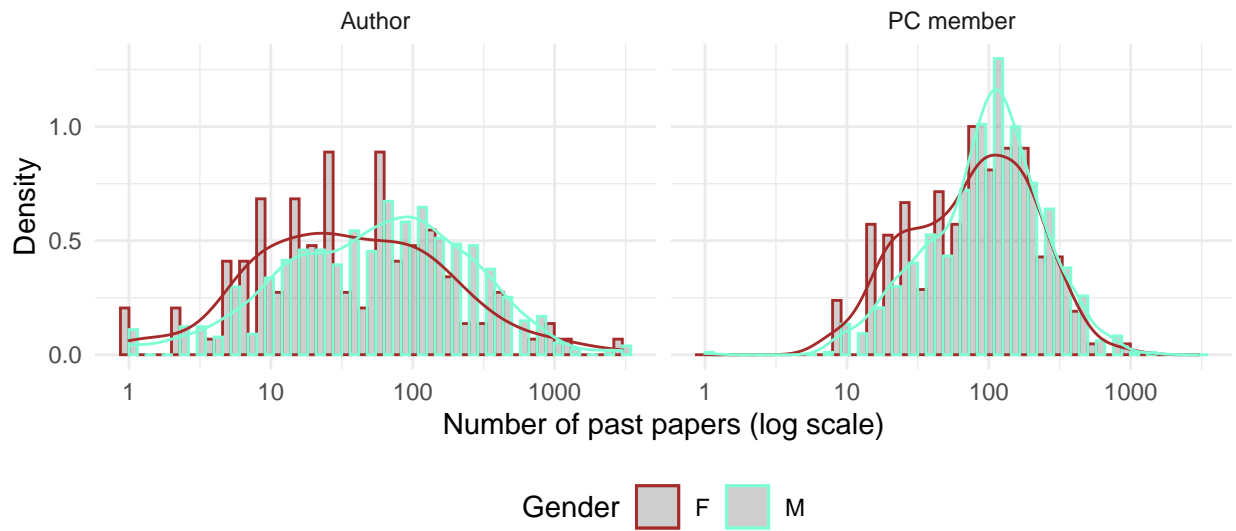


Figure 3: Distribution of past publications (ca. 2017) from Google Scholar by gender and role.

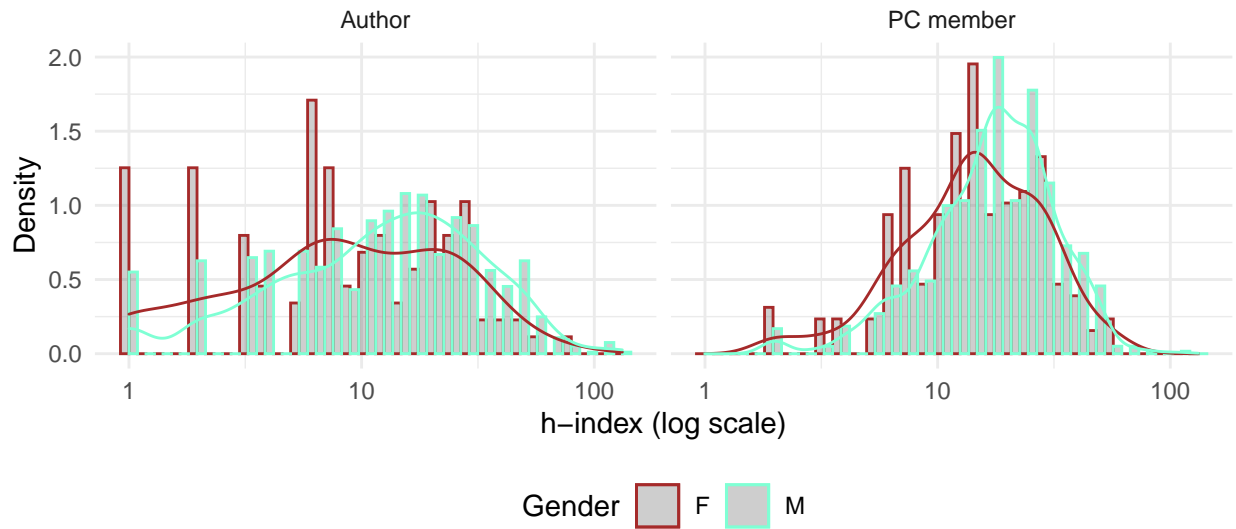


Figure 4: Distribution of H-index (ca. 2017) by gender and role.

continue to senior research positions at the same rate as men [16, 19, 34, 42].

To confirm these observations, we stratified all researchers with a known h-index in our dataset into the following three groups: those with a metric value of less than 13, those with a value between 13 and 18, and the rest. Following Hirsch’s categorization, we arbitrarily named these groups novices, mid-career, and experienced [23]. As Fig. 6 shows, a smaller proportion of women than men reach the senior ranks of research. The experience gap is particularly noticeable among

authors (44.8% novice female authors compared to 36.4% men, $\chi^2 = 7.419$, $p = 0.00645$), and has also been observed in other engineering fields [16, 53].

Table 2: Top ten countries by number of researchers.

Country	% Women	Total
United States	15.38	1408
China	10.43	200
France	13.61	147
Germany	8.63	139
Spain	8.94	123
India	5.63	72
Switzerland	14.06	64
Japan	1.59	63
United Kingdom	7.69	52
Canada	6.82	44




Figure 6: Experience bands by gender, all researchers.

5.2 Geography

When we look at the geographical distribution of researchers, it may not be surprising that most hail from the West, although China, India, and Japan are also prominent (Table 2 and Fig. 7). The United States has the highest percentage of women of any country with more than 15 total researchers in our dataset, but it is still far from gender parity. On the opposite end, Japan stands out among economically developed countries with particularly low female representation.

Another way to break down these numbers is by geographic region and conference role (Table 3). This perspective exposes the stark lack of geographic diversity in the field, with a full half of the 1787 identified authors associated with a US email domain, and another 14.33% from Western Europe.

It is notable, however, that Western reviewers are not overrepresented compared to authors, as has been observed in journals in other fields [7]. For example, the percentage of PC members from the US (52.57%) or Western Europe (16.36%) are similar to those of authors in these regions (50.2% and (14.33%, respectively).

In terms of representation of women, none of the larger regions deviate much from the $\approx 10\%$ overall percentage. The regions with fewer than approximately 25 authors exhibit more variance, due to the small denominators (for example, a single female author from Eastern Europe would have bumped the percentage from zero to $\approx 8\%$). It is therefore difficult to generalize conclusions about representation of women to these regions because of their small sample size.

Table 3: Representation of women by region and role, sorted by total authors.

Region	Authors		PC members	
	% Women	Total	% Women	Total
Northern America	9.78	930	24.47	523
Western Europe	8.98	256	16.35	159
Eastern Asia	11.94	201	2.90	69
Southern Europe	6.60	106	12.50	80
Northern Europe	7.69	65	8.00	50
Southern Asia	6.35	63	5.00	20
South America	8.33	36	27.27	11
Australia and New Zealand	8.33	24	0.00	14
Western Asia	27.27	22	12.50	24
South-Eastern Asia	5.00	20	0.00	4
Eastern Europe	0.00	12	11.76	17
Western Africa	50.00	2	0.00	0
Central America	100.00	1	0.00	0
Central Asia	0.00	1	0.00	0
Northern Africa	0.00	1	0.00	0

5.3 Work Sector

From researchers’ affiliations we can broadly categorize their sector as either “COM” for industry (8.6% of unique researchers), “EDU” for academia (72.8%), or “GOV” for government and national labs (18.6%). The relative abundance of government researchers is not surprising, given that many of the largest HPC systems in the world are operated by government labs.

Fig. 8 depicts the percentage of women in each sector, by role. It shows that among PC members, there is higher representation of women in government and academia than in industry ($\chi^2 = 0.522, p = 0.77$). Although not statistically significant, this finding is consistent with past studies that found relatively fewer female engineers in industry research positions overall [16, 22, 45].

The differences across sectors for authors are even smaller, and also nonsignificant ($\chi^2 = 1.629, p = 0.443$). The smaller difference may stem from the larger sample size and from the fact that authors are all selected by peer review, presumably a less subjective process than that of appointment to a committee [31]. Either way, our data suggests that there are no gender differences in HPC based on work sector alone.




Figure 8: Representation of women by sector and role.

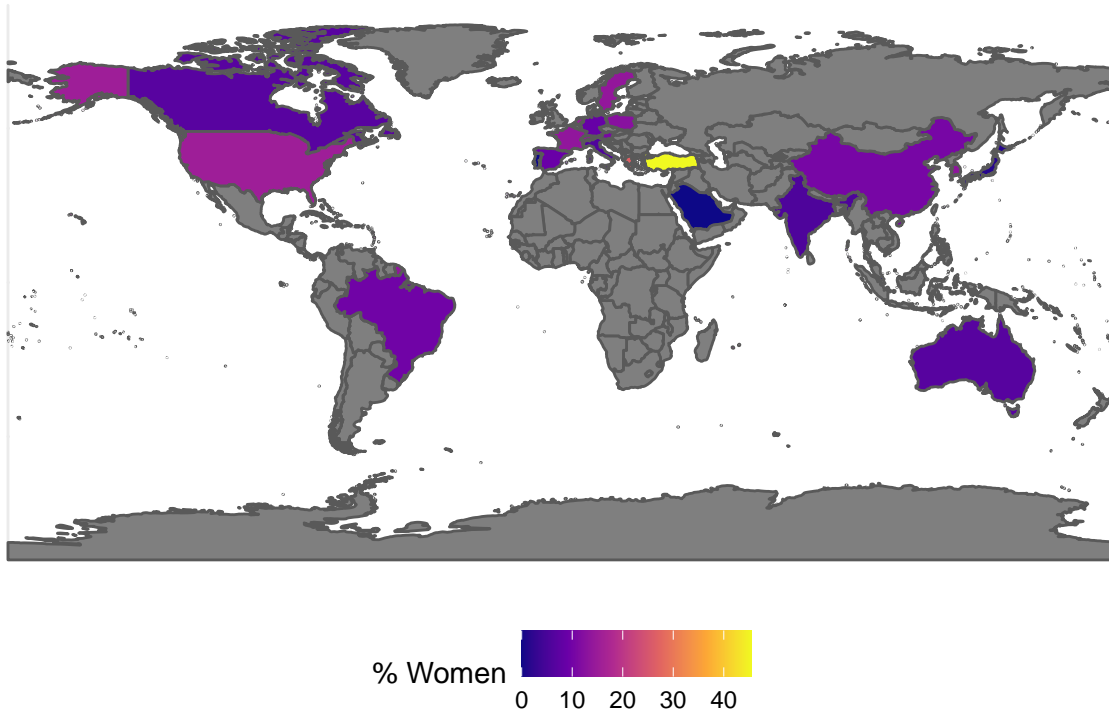


Figure 7: Representation of women for the 25 countries with at least 10 authors in the dataset.

6 CONCLUSION

The main goal of this study was to compute accurate and recent statistics about the representation of women in a cross section of leading HPC conferences. The picture that emerges shows a field that is still far from achieving gender equality. In most countries and geographical regions, fewer than one in ten HPC authors is a woman. Men also occupy most other visible conference roles and exhibit more research experience and higher citation counts than women as lead authors.

Nevertheless, there are some indications that representation and diversity are gaining importance in the field. Both SC and ISC are actively measuring inclusivity metrics, and all ACM conferences now require a code of conduct. As the two flagship HPC conferences, and the only two in our set with a dedicated diversity chair and inclusivity policies, they may serve as a benchmark to estimate the effect of diversity efforts. At the very least, these efforts are designed to make conferences more hospitable to all attendees, which may result in long-term improvements in gender representation. We plan to follow up and collect additional statistics at regular intervals to evaluate this hypothesis.

Other future work includes expanding this analysis to the larger set of 56 conferences we have collected from all subfields

of computer systems. We also plan to address deeper gender questions that emerge from the data, such as the differences in collaboration patterns between women and men and the causal effects that lead to lower representation of women.

ACKNOWLEDGMENTS

We would like to thank Kelly Shaw for spearheading this research and providing detailed feedback along the way. We would also like to thank Josh Reiss for his part in the collection of gender data. This work was supported by a grant from the Social Justice Research and Education Fund at Reed College.

REFERENCES

- [1] 2019. ISC 2019 Post-Conference Summary. https://www.isc-hpc.com/files/isc_events/documents/ISC2019_Summary.pdf
- [2] Robert Beverly and Mark Allman. 2013. Findings and implications from data mining the IMC review process. *ACM SIGCOMM Computer Communication Review* 43, 1 (2013), 22–29. <https://dl.acm.org/doi/pdf/10.1145/2427036.2427040>
- [3] Vijay Bhagat. 2018. Data and Techniques Used for Analysis of Women Authorship in STEMM: A Review. *Feminist Research* 2, 2 (2018), 77–86.
- [4] Rosie Campbell. 2018. In Defence of Diversity Measures. <https://medium.com/@RosieCampbell/in-defence-of-diversity-measures-48e4702b1dbd>
- [5] Stephen J Ceci and Wendy M Williams. 2011. Understanding current causes of women's underrepresentation in science. *Proceedings of the National Academy of Sciences* 108, 8 (2011), 3157–3162. <https://www.pnas.org/content/pnas/108/8/3157.full.pdf>
- [6] Suw Charman-Anderson, Lauren Kane, and Al Meadows. 2017. Championing the success of women in science, technology, engineering, maths, and medicine: A collection of thought pieces from members of the academic community. <https://doi.org/10.6084/m9.figshare.5463502.v1>
- [7] Clarivate Analytics. 2018. Global State of Peer Review. <https://publons.com/static/Publons-Global-State-Of-Peer-Review-2018.pdf>
- [8] J. McGrath Cohoon and William Aspray. 2006. A Critical Review of the Research on Women's Participation in Postsecondary Computing Education. Women and Information Technology. In *Women and Information Technology: Research on Underrepresentation*, J. McGrath Cohoon and William Aspray (Eds.). Cambridge Massachusetts, the MIT Press, Chapter 5, 137–180.
- [9] J. McGrath Cohoon, Sergey Nigai, and Joseph Kaye. 2011. Gender and Computing Conference Papers. *Commun. ACM* 54, 8 (Aug. 2011), 72–80. <https://doi.org/10.1145/1978542.1978561>
- [10] Toni Collins. 2016. Improving Diversity at HPC Conferences and Events. http://www.hpc-diversity.ac.uk/sites/default/files/images/Improving_Diversity_conferences.pdf
- [11] James RA Davenport, Morgan Foesneau, Erin Grand, Alex Hagen, Katja Poppenhaeger, and Laura L Watkins. 2014. Studying Gender in Conference Talks—data from the 223rd meeting of the American Astronomical Society. *arXiv preprint arXiv:1403.3091* (2014). <https://arxiv.org/pdf/1403.3091>
- [12] Lizanne DeStefano. 2018. Analysis of MICRO Conference Diversity Survey Results. <https://www.microarch.org/docs/diversity-survey-2018.pdf>
- [13] Elsevier. 2017. Gender in the Global Research Landscape. <https://www.elsevier.com/research-intelligence/campaigns/gender-17>
- [14] Elsevier. 2020. The Researcher Journey Through a Gender Lens. <https://www.elsevier.com/research-intelligence/resource-library/gender-report-2020>
- [15] Heather L Ford, Cameron Brick, Karine Blaufuss, and Petra S Dekens. 2018. Gender inequity in speaking opportunities at the American Geophysical Union Fall Meeting. *Nature communications* 9, 1 (2018), 1–6. <https://doi.org/10.1038/s41467-018-03809-5>
- [16] Mary Frank Fox. 2006. Women, men, and engineering. In *Women, Gender, and Technology*, Mary Ann Fox, Deborah G. Johnson, and Sue V. Rosser (Eds.). University of Chicago Press Urbana, IL, Chapter 2, 47–59.
- [17] Eitan Frachtenberg and Noah Koster. 2020. A Survey of Accepted Authors in Computer Systems Conferences. *PeerJ Computer Science* 6 (2020), e299. <https://doi.org/10.7717/peerj-cs.299>
- [18] Massimo Franceschet. 2010. The role of conference publications in CS. *Commun. ACM* 53, 12 (Dec. 2010), 129–132.
- [19] Athina Frantzana. 2019. *Women's representation and experiences in the high performance computing community*. Ph.D. Dissertation. The University of Edinburgh.
- [20] Jill Freyne, Lorcan Coyle, Barry Smyth, and Pdraig Cunningham. 2010. Relative status of journal and conference publications in computer science. *Commun. ACM* 53, 11 (2010), 124–132.
- [21] Suzanne Fricke. 2018. Semantic scholar. *Journal of the Medical Library Association* 106, 1 (2018), 145.
- [22] Gita Ghiasi, Vincent Larivière, and Cassidy R Sugimoto. 2015. On the compliance of women engineers with a gendered scientific system. *PLoS one* 10, 12 (2015), e0145931. <https://journals.plos.org/plosone/article?id=10.1371/journal.pone.0145931>
- [23] Jorge E Hirsch. 2005. An index to quantify an individual's scientific research output. *Proceedings of the National Academy of Sciences* 102, 46 (2005), 16569–16572.
- [24] Luke Holman, Devi Stuart-Fox, and Cindy E. Hauser. 2018. The gender gap in science: How long until women are equally represented? *PLoS biology* 16, 4 (2018), e2004956. <https://journals.plos.org/plosbiology/article?id=10.1371/journal.pbio.2004956>
- [25] Junming Huang, Alexander J Gates, Roberta Sinatra, and Albert-Laszlo Barabasi. 2019. Historical comparison of gender inequality in scientific careers across countries and disciplines. *arXiv preprint arXiv:1907.04103* (2019).
- [26] Joy L. Johnson, Lorraine Greaves, and Robin Repta. 2009. Better science with sex and gender: facilitating the use of a sex and gender-based analysis in health research. *International journal for equity in health* 8, 1 (2009), 1–11.
- [27] Joy L. Johnson and Robin Repta. 2012. Sex and gender. *Designing and conducting gender, sex, and health research* 1737 (2012).
- [28] Fariba Karimi, Claudia Wagner, Florian Lemmerich, Mohsen Jaididi, and Markus Strohmaier. 2016. Inferring Gender from Names on the Web: A Comparative Evaluation of Gender Detection Methods. In *Proceedings of the 25th International Conference Companion on World Wide Web (Montréal, Québec, Canada) (WWW'16 Companion)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, Switzerland, 53–54. <https://doi.org/10.1145/2872518.2889385>
- [29] Vincent Larivière, Chaoqun Ni, Yves Gingras, Blaise Cronin, and Cassidy R Sugimoto. 2013. Bibliometrics: Global gender disparities in science. *Nature News* 504, 7479 (2013), 211.
- [30] Trang T. Le, Daniel S. Himmelstein, Ariel A. Hippen Anderson, Matthew R. Gazzara, and Casey S. Greene. 2020. Analysis of ISCB honorees and keynotes reveals disparities. *bioRxiv* (2020). <https://doi.org/10.1101/2020.04.14.927251> arXiv:https://www.biorxiv.org/content/early/2020/04/14/2020.04.14.927251.full
- [31] Carole J Lee, Cassidy R Sugimoto, Guo Zhang, and Blaise Cronin. 2013. Bias in peer review. *Journal of the American Society for Information Science and Technology* 64, 1 (2013), 2–17. http://faculty.washington.edu/c3/Leet_et_al.2013.pdf
- [32] Xueping Liang, Sachin Shetty, Deepak Tosh, Charles Kamhoua, Kevin Kwiat, and Laurent Njilla. 2017. Prochain: A blockchain-based data provenance architecture in cloud environment with enhanced privacy and availability. In *Proceedings of the 17th IEEE/ACM international symposium on cluster, cloud and grid computing*. IEEE Press, 468–477.
- [33] Sandra Mattauch, Katja Lohmann, Frank Hannig, Daniel Lohmann, and Jürgen Teich. 2020. A Bibliometric Approach for Detecting the Gender Gap in Computer Science. *Commun. ACM* 63 (2020), 74–80. <https://doi.org/10.1145/3376901>
- [34] Mary C Mattis. 2007. Upstream and downstream in the engineering pipeline: What's blocking US women from pursuing engineering careers. In *Women and minorities in science, technology, engineering and mathematics: Upping the numbers*. Edward Elgar Cheltenham, UK, Chapter 14, 334–362.
- [35] National Science Board (US). 2020. *Science & engineering indicators*. National Science Board. <https://ncses.nsf.gov/pubs/nsb20201/u-s-s-e-workforce>
- [36] Nature Editorials. 2018. Science benefits from diversity. *Nature* 558 (June 2018). <https://doi.org/10.1038/d41586-018-05326-3>
- [37] Bryan Parno, Ulfar Erlingsson, and Will Enck. 2017. Report on the IEEE S&P 2017 submission and review process and its experiments. <http://www.ieee-security.org/TC/Reports/2017/SP2017-PCChairReport.pdf>
- [38] David Patterson, Lawrence Snyder, and Jeffrey Ullman. 1999. Evaluating Computer Scientists and Engineers For Promotion and Tenure. *Computing Research News* (September 1999). <http://www.cra.org/resources/bp-view/evaluating-computer-scientists-and-engineers-for-promotion-and-tenure/>
- [39] Lucia Santamaria and Helena Mihaljevic. 2018. Comparison and benchmark of name-to-gender inference services. *PeerJ Computer Science* 4 (2018), e156.
- [40] Nihar B Shah, Behzad Tabibian, Krikamol Muandet, Isabelle Guyon, and Ulrike Von Luxburg. 2018. Design and analysis of the NIPS 2016 review process. *The Journal of Machine Learning Research* 19, 1 (2018), 1913–1946.
- [41] Michael Shermer. 2014. How the survivor bias distorts reality. *Scientific American* 1 (2014), 106.
- [42] Gerhard Sonnert, Mary Frank Fox, and Kristen Adkins. 2007. Undergraduate Women in Science and Engineering: Effects of Faculty, Fields, and Institutions Over Time. *Social Science Quarterly* 88, 5 (2007), 1333–1356. <https://doi.org/10.1111/j.1540-6237.2007.00505.x>

- arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1540-6237.2007.00505.x>
- [43] Flaminio Squazzoni, Giangiaco Bravo, Pierpaolo Dondio, Mike Farjam, Ana Marusic, Bahar Mehmami, Michael Willis, Aliaksandr Birukou, and Francisco Grimaldo. 2020. No evidence of any systematic bias against manuscripts by women in the peer review process of 145 scholarly journals. *SocXiv preprint* (2020). <https://doi.org/10.31235/osf.io/gh4rv>
- [44] Cara Tannenbaum, Robert P Ellis, Friederike Eyssel, James Zou, and Londa Schiebinger. 2019. Sex and gender analysis improves science and engineering. *Nature* 575, 7781 (2019), 137–146.
- [45] Yu Tao. 2016. Where Do They Do Engineering? Gender Differences in Employment at the Position Level among Engineering Doctorate Recipients. *Journal of Women and Minorities in Science and Engineering* 22, 1 (2016).
- [46] Andrew Tomkins, Min Zhang, and William D Heavlin. 2017. Reviewer bias in single-versus double-blind peer review. *Proceedings of the National Academy of Sciences* 114, 48 (2017), 12708–12713. <https://www.pnas.org/content/114/48/12708>
- [47] Tiffany Trader. 2016. A Celebration of Women in HPC. <https://www.hpcwire.com/2016/07/07/celebration-women-hpc/>. Retrieved Nov. 18, 2019.
- [48] Moshe Y Vardi. 2009. Conferences vs. journals in computing research. *Commun. ACM* 52, 5 (2009), 5–5. <https://cacm.acm.org/magazines/2009/5/24632-conferences-vs-journals-in-computing-research/fulltext>
- [49] Moshe Y. Vardi. 2021. Reboot the Computing-Research Publication Systems. *Commun. ACM* 64, 1 (2021), 7. <https://cacm.acm.org/magazines/2021/1/249441>
- [50] Rita Vine. 2006. Google scholar. *Journal of the Medical Library Association* 94, 1 (2006), 97.
- [51] Lucy Lu Wang, Gabriel Stanovsky, Luca Weihs, and Oren Etzioni. 2021. Gender Trends in Computer Science Authorship. *Commun. ACM* 64, 3 (Feb. 2021), 78–84. <https://doi.org/10.1145/3430803>
- [52] Samuel F Way, Daniel B Larremore, and Aaron Clauset. 2016. Gender, productivity, and prestige in computer science faculty hiring networks. In *Proceedings of the 25th International Conference on World Wide Web*. 1169–1179.
- [53] Stuard Zweben and Betsy Bizot. 2018. 2017 CRA Taulbee Survey. *Computing Research News* 30, 5 (May 2018). <https://cra.org/crn/category/2018/vol-30-no-5/>