

Voices of *Her*: Analyzing Gender Differences in the AI Publication World

Yiwen Ding^{1,*} Jiarui Liu^{1,*} Zhiheng Lyu^{2,*} Kun Zhang^{5,6} Bernhard Schölkopf^{3,4}
Zhijing Jin^{3,4,†} Rada Mihalcea^{1,†}

¹University of Michigan, ²University of Hong Kong, ³Max Planck Institute, ⁴ETH Zürich
{dyiwen, ljrjerry, mihalcea}@umich.edu
zhiheng.lyu.cs@gmail.com bs@tue.mpg.de jinzhi@ethz.ch

Abstract

While several previous studies have analyzed gender bias in research, we are still missing a comprehensive analysis of gender differences in the AI community, covering diverse topics and different development trends. Using the AI SCHOLAR dataset of 78K researchers in the field of AI, we identify several gender differences: (1) Although female researchers tend to have fewer overall citations than males, this citation difference does not hold for all academic-age groups; (2) There exist large gender homophily in co-authorship on AI papers; (3) Female first-authored papers show distinct linguistic styles, such as longer text, more positive emotion words, and more catchy titles than male first-authored papers. Our analysis provides a window into the current demographic trends in our AI community, and encourages more gender equality and diversity in the future.¹

1 Introduction

Motivated by the spirit of the ACL Year-Round Mentorship Program² to support junior researchers to understand how a career path in NLP is, we want to answer this question technically, namely, what are the causal factors for academic success.

Although nearly half of the world population is female (Ritchie and Roser, 2019), the proportion of female researchers in science fields is often disproportionately smaller (Robnett, 2016; Hand et al., 2017). Specifically, in the research community of AI, we find that female researchers constitute only 17.99% of all the scholars in the field of AI with more than 100 citations, as collected in the AI SCHOLAR dataset (Jin et al., 2022). This fraction is even smaller in some subdomains of AI such as




Figure 1: Cumulative number of AI scholars and female scholars ratio vs. years to publish the first paper. Female scholars are taking an increasing percentage of all AI scholars, but the ratio is still small (around 18%).

computer vision (CV), where only 15.64% of the researchers are females.

Despite this strong gender gap among AI researchers, comprehensive research on the state of the field is yet to be conducted. Most existing research work has been done on certain subdomains of AI, such as the NLP community (Vogel and Jurafsky, 2012; Schluter, 2018; Mohammad, 2020), or has addressed research aspects such as the values listed in top-cited papers (Birhane et al., 2022). To the best of our knowledge, this study is the first to conduct an up-to-date AI community-wise comprehensive analysis since the study of Stathoulopoulos and Mateos-Garcia (2019).

In this paper, we look into distinct features of the female subgroup in the AI community, and conduct comprehensive statistical analyses from a diverse range of perspectives: basic scholar profile statistics, citation trends, coauthorship, and linguistic styles of papers. The main findings from our study are as follows:

1. Although female AI researchers tend to have fewer overall citations than males, this citation difference does not always hold for all

* Equal contribution.

† Equal supervision.

¹Our code and data have been uploaded to the submission system, and will be open-sourced upon acceptance.

²<https://mentorship.aclweb.org/>

academic-age groups or all time stages in one’s career.

2. There exists large homophily in genders of the first author, last author, and the majority of the authors on AI papers, such as a high correlation between male last authors and the majority of authors being male. This gender homophily pattern in AI echoes the observation of [Schluter \(2018\)](#) in NLP.
3. Female first-authored papers tend to have distinct linguistic patterns such as more words about positive emotion, longer text, and more catchy titles.

Our findings contribute suggestions and supporting evidence to future AI community organizers or individuals who want to push for informed community changes.

2 Data Collection and Cleaning

AI SCHOLAR. We use all the scholar information from the most recent collection of researchers in the field of AI, i.e., the AI SCHOLAR dataset ([Jin et al., 2022](#)),³ which contains all the scholars in the field of AI with at least 100 citations according to Google Scholar. The data consists of 78K scholars with tags related to AI such as artificial intelligence (AI) and machine learning (ML), or subdomains of AI such as computer vision (CV) and natural language processing (NLP). It only includes scholars with at least 100 citations, an approximate cut-off for the long-tail since it is not feasible to include all scholar profiles. We discuss the limitations of using this dataset in Section 7. Throughout the paper, we use the term “AI researchers” to denote the set of scholars in the AI SCHOLAR dataset.

For each AI researcher, the AI SCHOLAR dataset collects information such as the name, affiliation, up to five domain tags, total citations, citations by year, and all their papers with title, year, and the number of citations.

Since the total number of papers is massive (2.8M papers for the 78K AI researchers), we use the random subset of papers provided by [Jin et al. \(2022\)](#). They collect 100K papers with detailed information such as abstracts and full names of all the coauthors. Among the 100K papers with detailed information, we further filter out papers with

empty abstracts and keep 91K papers, which we denote as “AI papers” in our analysis.

Identifying Female Researchers. Since the focus of this paper is to analyze the female subgroup in the AI community, we have to find a way to identify AI researchers that are female. Admittedly, this is a daunting task due to two main concerns. First, gender is a continuum that goes beyond the male/female binary distinction. Second, there are no computational methods to identify the gender of a researcher that are perfectly correct and perfectly ethical. A possible way is to collect as many self-reports of gender as possible, but this method will be largely time-consuming on the scale of 78K, and also might lead to a large selection bias in the data, since the collected responses might be of a small number and not an *i.i.d.* subset of the entire data. After balancing all the ethical and practical concerns, we decided to follow the practice from [Mohammad \(2020\)](#), who classified gender by collecting first names that correspond to male and female genders more than 95% of the time in the merged records of the US Social Security Administration’s published database of names and genders along with the PubMed authors with genders, as well as using the hand-labeled author genders by [Vogel and Jurafsky \(2012\)](#) to correct for wrongly classified names. Using this conservative but ethical approach, we obtain 7,036 female authors and 32,074 male authors from the 78,066 AI researchers, and leave the author names that cannot be classified as “unclassified.”

We acknowledge that the name-gender records that we use have limited representations of names from all cultural backgrounds. In our paper, we make an effort to keep the errors modular, such that future work can use our analysis framework on a more accurate set of female AI researchers to produce more accurate insights.

Population with Unclassified Gender. Since we choose to stick with our ethical standards, including not using any name- or photo-based classifier, we have to leave out a large set of AI researchers whose gender cannot be identified. To address this, we would like to frame the scope our analysis to this subset of AI researchers whose gender can be identified. And also we believe that despite this limitation, this study is still more meaningful to the community than not drawing any conclusions. In

³Dataset is available at <https://github.com/causalNLP/AI-Scholar>

Appendix B, we analyze the coverage and properties of this subset.

3 Analysis of Scholar Profiles

We first analyze the basic scholar profiles to compare the general statistics with those of the female subgroup. Our analyses answer the following questions: (1) What percentage of female scholars are there in AI and in each subdomain? (2) What are the scholarly indices of the female researchers doing in AI? And (3) How do the analyses differ if we consider the scholarly trends before and after 2012 – a year that corresponds approximately to the time when deep learning started to become widely used (among others, it is the publication year of AlexNet (Krizhevsky et al., 2012)) – w.r.t. different academic age groups, and in academia vs. industry? Note that for the scope of this paper, we focus on overall trends, and we encourage future work to dive into causal analysis.

3.1 Female Percentage

We first check the size of the female subgroup in AI and various subdomains of AI.




Figure 2: Female scholar percentage (F%) by subdomains of AI collected from Google Scholar profiles.

In Figure 2, we can see that there are 17.99% female scholars among all AI researchers with classified gender, and this percentage varies across the different subdomains of AI that scholars self-label on their Google Scholar profiles. The representation of females is relatively more pronounced in areas such as natural language processing (27%) and human-computer interaction (22%), and less seen in areas such as physics and robotics, both with only 13%. We discuss the experimental details in Appendix A, including how we manually clean and cluster these tags as well as count normalization. Note that to get an informative percentage of female researchers (denoted as “F%”), when we

calculate the percentages throughout this paper, we consider female scholars among all scholars whose gender is classified, because the non-trivial size of the unclassified group may make the percentage of female scholars look disproportionately small, thus not very informative for understanding the statistics.

3.2 Profile Statistics

Next, we look into the profiles of AI researchers and calculate overall scholarly statistics, reported in Table 1.

		Avg	Min	25th	50th	75th	Max
Citations:	All	2,129.54	100	214	475	1,345	533,757
	F.	1,762.11	100	197	414	1,165	209,549
h-Index:	All	14.03	1	7	10	16	266
	F.	13.25	1	6	9	15	211
# Papers:	All	67.44	1	17	32	68	3,000
	F.	60.20	1	16	29	64	2,125
AcadAge:	All	16.89	2	10	14	20	73
	F.	16.33	2	10	14	20	73
Active Yrs.:	All	15.47	1	8	13	19	73
	F.	14.87	1	8	12	19	72
F. Coauthor (%)	All	9.58	0	5.21	8.70	12.50	48.72
	F.	14.53	1.64	9.09	13.04	18.58	48.72
Academia (%)	All	60.25	0	0	100	100	100
	F.	61.79	0	0	100	100	100
Big 10 (%)	All	6.61	0	0	0	0	100
	F.	15.97	0	0	0	0	100

Table 1: Statistics of Google Scholar profiles. We compare the statistics for the total population of all researchers (“All”) with those of the female subgroup (“F.”), w.r.t. the average; minimum; 25th, 50th, and 75th percentile; and maximum values. The reported statistics include citations, h-indices, number of papers (# Papers), academic age (calculated by subtracting the year of the first paper from the current year), the number of active years (calculated by subtracting the year of the first paper from the year of the last paper), percentage of females among their coauthors (F. Coauthor), proportion of researchers who are affiliated with academia, and the proportion of researchers who are affiliated with the most frequently appearing ten organizations among AI researchers (Big 10). See Appendix C for a more comprehensive table including the standard deviation.

As we can see in Table 1, the average citation number for female researchers is 1.7K, which is 367 less than that of all AI researchers. If we look closely at the detailed information on the citation distribution, we can see that this gap may be attributed to the difference in highly cited scholars. Here, the citation difference is moderate until the 50th percentile, with only –61 difference, but the gap drastically increases in higher percentiles, such as –180 in the 75th percentile, and finally –324K

in the maximum. Across other statistics, we can see a similar trend in the h-index and the number of papers.

The most significant differences can be seen on the percentage of female coauthors. In Table 1, average scholars have 9.58% female coauthors among all coauthors, but female scholars have 14.53% female coauthors. It is also noteworthy that the percentage of female coauthors for average scholars is less than 10% until the 50th percentile, which can demonstrate that some fields or coauthorship sub-networks have a very low representation of females.

As some additional notes, we can see that, on average, females’ academic age is slightly younger by 0.56 years. And, if we account for scholars who have stopped publishing by reporting the number of active years (the number of years between the year that a scholar first published a paper and the year that they published their last paper), then we can see a slightly larger gap of 0.6 years. This shows a slight trend that females stop publishing a bit earlier.

In the last row of Table 1, we can see that female researchers are more concentrated in the most frequently-appearing ten organizations: Google, Stanford, CMU, MIT, Amazon, UCB, Microsoft, Facebook, IBM and Apple. See the implementation details to extract these organizations in Appendix D.1. There could be many potential explanations, such as that the big organizations have a stronger diversity requirement to bridge the gender gap, or it could be that females who persist in the research field are very talented, among many other possibilities.

3.3 Varying Views

After analyzing the overall statistics, we also perform analyses on different subsets of the data: (1) scholarly statistics before and after 2012 when the wide use of deep learning has started; (2) citations by different academic age groups and at different career stages; and (3) academic vs. industry affiliations.

The 2012 AI Wave. Since a domain such as AI can go through many ups and downs, we also want to check how the statistics differ before and after the time when deep learning achieved the first round of large empirical success. As a rough estimate, we take the year 2012, when the highly

influential paper, AlexNet (Krizhevsky et al., 2012) on ImageNet (Deng et al., 2009), was published.

In Table 2, we compare statistics of AI researchers who published their first paper after 2012 (the post-2012 generation) and before 2012 (the pre-2012 generation). We can see that it takes clearly fewer years for a researcher to reach 100 citations in the post-2012 generation, due to the surge of research on AI after 2012. Some differences between female scholars among all scholars are slightly larger in the post-2012 generation than in the pre-2012 generation.

		Post-2012	Pre-2012
# Papers/Yr.:	All	2.86±2.78	3.96±4.32
	F.	2.62±3.12	3.77±4.05
Citations/Paper:	All	47.93±146.58	42.06±89.33
	F.	45.02±129.13	39.73±91.68
Yrs. to Reach 100 Cit. (↓):	All	4.82±1.89	6.77±3.69
	F.	5.03±1.94	6.93±3.67
<i>Most Cited Paper</i>			
Avg. Citations:	All	369.41±2208.56	692.35±2835.00
	F.	316.49±1661.53	553.63±1829.00
Most Common Yr.:	All	2018	2011
	F.	2018	2011

Table 2: Statistics of AI researchers who published their first paper after 2012 (the Post-2012 Generation) and before 2012 (the Pre-2012 Generation). For the most cited paper of each scholar, we list the average citations (Avg. Citations), and the most common year for the most cited paper (Most Common Yr.).

AcadAge	Cit. by 5th Yr		By 10th Yr		By 15th Yr		By 20th Yr		By 25th Yr		By 2022		Total Ratio
	All/F	All/F	All/F	All/F	All/F	All/F	All/F	All/F	All/F	All/F	All/F	All/F	
0-5	200/188												1.06
6-10	127/114		275/248										1.11
11-15	79/72		300/253		418/349								1.20
16-20	81/77		324/299		599/544		724/698						1.04
21-25	97/95		371/366		793/789		1,029/988		1,209/1,215				1.00
>25	109/104		391/383		891/886		1,519/1,473		1,983/1,872		2,090/1,961		1.07

Table 3: Median citations from different academic ages from scholars of different academic age groups (average scholars statistics / female scholars statistics).

Academic Age. We also explore the citation differences across different academic age groups, inspired by the analysis of NLP scholars Mohammad (2020). We separate the citations of all scholars and female scholars across two dimensions: each age group (e.g., 0 – 5, 6 – 10, ...), and every 5-year window for each group (e.g., citations by 5th year, 10th year, ...).

With this more time-specific view, in Table 3, we can see almost equality in several academic age groups with certain time spans, such as the 21 – 25 academic age group, where the female overall citation by the 25th year (1,215) is even higher than

		Affiliation Is Industry Only	
		No	Yes
Stops Publishing:	No	41,217	25,484
	Yes	2,749	3,521

Table 4: Contingency table of being in the industry only and stopping publishing since 2018. We get a p-value of $1e-169$ by χ^2 test, confirming a strong correlation between being in industry only and stopping publishing.

the average (1,209). The citation difference that we see in previous sections can be attributed to more specific age groups and time, such as the 11 – 15 age group, and the 16 – 20 age group.

Dropout and Industry. Some possible alternative reasons why a scholar has fewer citations could just be a matter of career choice. We want to account for the affiliation difference of scholars (i.e., whether a scholar is in the industry or not) and its correlation with some results that can affect citations (e.g., whether the scholar stops publishing). Therefore, we calculate the correlation between whether a scholar is only affiliated with industry and whether they stopped publishing recently. Note that we take the year 2018 as an empirical threshold for the recent stop in publication, because some domains may take longer to publish and 2018 is a relatively reasonable year that avoids the effects of the COVID-19 outbreak.

In Table 4, we can see that a χ^2 test confirms a strong correlation between a stop in publication and being exclusively affiliated with the industry. When conditioning on all people that keep publishing, the number in academia is almost twice that in the industry. Additionally, when conditioning on all people that stop publishing, there are 28% more people in the industry than that in academia. We include a fine-grained analysis by academic age in Appendix E.3.

4 Analysis of Citation Time Series

To take the analysis one step further, we perform a more fine-grained analysis of the scholar statistics.

Time Series Clustering. We are interested in patterns in the scholar citation time series. Inspired by the time series construction by Tanveer et al. (2018), we take the citations-by-year data of all the 78K scholars, normalize them by the average citation number, and linearly interpolate the citation time span to the largest number of active years, so

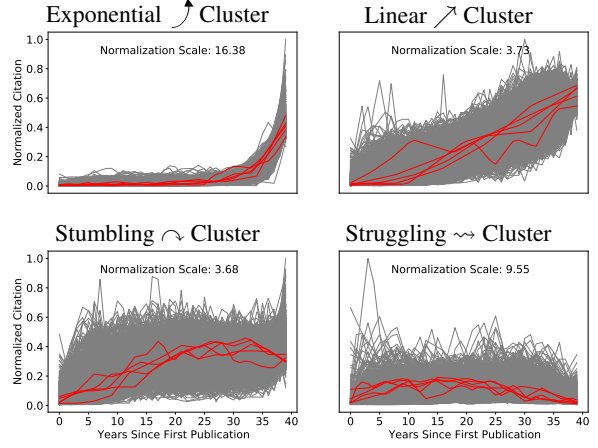


Figure 3: Four main types of time series clusters of AI researchers' citations.

		Exponential	Linear	Stumbling	Struggling
# Scholars:	All	8,008	41,698	16,857	1,565
	F	1,535	7,831	2,818	239
	Ratio (All:F)	5.22:1	5.32:1	5.98:1	6.56:1
Citation:	All	2,472.55	2,292.87	2,114.19	787.80
	F	1,569.49	1,897.35	1,895.13	438.15
	Ratio	1.58:1	1.21:1	1.12:1	1.80:1
h-Index:	All	16.21	14.34	13.34	9.86
	F	14.68	13.69	12.56	8.50
	Ratio	1.10:1	1.05:1	1.06:1	1.16:1
AcadAge:	All	17.40	14.64	19.24	21.95
	F	16.89	14.53	18.25	19.92
	Ratio	1.03:1	1.01:1	1.05:1	1.10:1
Stop Pub.:	All	0.25%	2.58%	22.98%	44.53%
	F	0.14%	2.35%	24.45%	44.27%
	Ratio	1.79:1	1.10:1	0.94:1	1.01:1

Table 5: Scholar statistics in each cluster.

that we can focus on the shape of the citation times series and stay agnostic with respect to different academic ages. We apply K-Means clustering on time series (Tavenard et al., 2020) and introduce our implementation details in Appendix D.2.

We further manually group the multiple clusters generated by the algorithm into four main types according to human-interpretable shape patterns. For notation convenience, we manually assign some easy-to-remember names to the four cluster types: the exponential cluster (\uparrow), linear cluster (\nearrow), stumbling cluster (\curvearrowright), and struggling cluster (\rightsquigarrow). For each type, we visualize a representative cluster in Figure 3, and plot all machine-identified clusters in Appendix F.1.

Cluster Statistics. In Table 5, we can see that the majority of the scholars are in the linear cluster, which is the most common time-series pattern. The exponential pattern is substantially rarer than the linear pattern, only 1/5 by the number of scholars,

but with the largest h-indices across all clusters. In the exponential cluster, although female scholars have a higher representation than in other clusters, the average citations for females is lower than those of the linear and stumbling cluster, which may be explained by the previous observations in Section 3.2 that top percentile citations in all AI scholars are higher than those in female scholars, and this phenomenon might be more pronounced in the exponential cluster.

The smallest cluster is the struggling cluster, where the scholars experience fluctuations in citations but no clear pattern of overall increase. This cluster is also the one that correlates with the highest ratio of scholars that stop publishing, with a percentage of 44+% for both average scholars and female scholars, followed by the stumbling cluster with 22+% percentage of people who stop publishing. All these percentages are substantially higher than the stop-publishing percentage in the exponential and linear clusters.

Female Subgroups in the Clusters. For each cluster that we identify, we also show the female ratio of each feature for the clusters in Table 5. Across all the clusters, female researchers are always less than 1/5 of the population, and experience fewer citations (e.g., with an All:F ratio being 1.58:1 in the exponential cluster, and 1.80:1 in the struggling cluster), and lower h-indices, which are correlated with citations. In the stumbling cluster, we see that female researchers are 1.47% more likely to stop publishing.

As an additional note, we also conduct additional analysis to focus on female subgroups in NLP, and find that female researchers in NLP have higher citations than average female scholars, which is a trend across all clusters. We include detailed results in the Appendix Table 18.

5 Analysis of Co-Authorship Patterns

We also address co-authorship patterns, and conduct analyses to answer the following questions: (1) Do female scholars tend to have more diverse collaborators? And (2) Are there certain gender patterns in different author roles, and what does that indicate for mentor-mentee relationship?

Aggregated Coauthor Statistics. We are interested in the question “Does diversity attract diversity?” A potential angle to understand this is to compare the characteristics of all AI scholars’

coauthors and female scholars’ coauthors.

	All	F.
F. Coauthors % (\uparrow)	9.58 \pm 14.27	14.53 \pm 26.89
Coauthors’ Domain Diversity (\uparrow)	2.48 \pm 0.99	2.44 \pm 0.96
% Coauthors in Freq. Ten Orgs. (\downarrow)	6.76 \pm 14.15	6.98 \pm 14.43
% Coauthors in AI Scholars (\downarrow)	18.28 \pm 14.39	19.85 \pm 14.80

Table 6: Diversity indices among the coauthors of general AI scholars and female scholars. We use \uparrow and \downarrow to indicate a higher or lower number in this indicator might represent more diversity. Implementation details are in Appendix D.3.

As we can see in Table 6, female scholars have a much larger percentage of female coauthors (14.53%), which is +4.95% by absolute value higher than that of the average AI scholars. Female researchers’ other diversity indices are slightly lower, which might correlate with the previous finding that female scholars are more concentrated as the ten most frequent organizations. The dynamics of collaboration could be worth exploration in future studies.

Author Lists of Papers. Furthermore, we calculate the statistics based on the author lists of the AI papers in the AI SCHOLAR dataset Jin et al. (2022). In Table 7, we investigate that, given the last author’s gender, what are some gender patterns in the first author role, and the majority of the authors.

There are some noteworthy conditional probabilities showing large gender disparity. For example, among all papers with male last authors, there are some astoundingly strong gender disparities – 1:4.61 female-to-male ratio in the first authors, and 1:48.47 female-to-male ratio for the gender of the majority authors. Among female last authors, the first author role reaches more gender balance, and the gender ratio of majority authors are reversed, with almost two times more female-majority papers than male-majority papers. This echoes with the gender homophily observation in Schluter (2018), although the previous study only focuses on the NLP domain.

Moreover, previous papers have suggested using the relationship between the first author and the last author as a proxy for mentee-mentor relation (Schluter, 2018). From Table 7, it seems that male mentors (using the last author as a proxy) tend to take more male mentees (using the first author as a proxy), while female mentors are more balanced, although the ratio is still not equal, perhaps limited by the disparity in the sheer amount of female

	1st Author F:M	Majority Authors F:M
Last F.	1:1.75	1:0.54
Last M.	1:4.61	1:48.47

Table 7: Given the last author’s gender, we show the female-to-male gender ratio (F:M) in the first author role (1st Author) and majority authors (i.e., >50% of the authors). We calculate the statistics based on the author lists of the AI papers in the AI SCHOLAR dataset Jin et al. (2022).

researchers in the AI community.

6 Analysis of Female-Authored Papers

	Citations		# Coauthors	
	Avg	>95th (4,599 Papers)	Avg	>95th (5,615 Papers)
1st F.	33	305 (6.63%)	5	468 (8.33%)
1st M.	42	1,337 (29.07%)	5	1,147 (20.43%)
>50% F.	29	285 (6.20%)	8	601 (10.70%)
>50% M.	42	2,309 (50.20%)	6	2,852 (50.79%)
Last F.	32	233 (5.07%)	5	304 (5.41%)
Last M.	42	1,425 (30.98%)	4	1,093 (19.47%)

Table 8: For each group of papers with a certain author gender information, we calculate the average citations and the average number of coauthors. In addition, we also check each coauthor’s gender group’s presence in the 95th percentile of paper citations (i.e., >113 citations) and the number of coauthors (i.e., >10 coauthors).

Since our study features a comprehensive bottom-up analysis of the female subgroup, we have covered statistics related to individual scholars and coauthorship, and, finally, in this section, we analyze statistics of female-authored papers.

6.1 General Paper Statistics

We first calculate some general statistics of papers with different author gender information in Table 8. We can see that papers by female authors tend to have more coauthors, while male authors tend to have on average higher citations than those by female authors at the same authorship position or majority representation, for example, 13 more citations on average for male-majority papers than female-majority papers. Moreover, this disparity is very large if we zoom into the top papers. Specifically, we take papers over the 95th percentile (with over 113 citations), for example, 2.3K of these papers have a male-majority author list, in contrast to the 285 papers with a female-majority author list (8:1).

LIWC Category & Top 5 Freq Words	All	1st F.
Positive Emotion {well, important, energy, better, support}	1.98	2.05 (↑3.71%)
Female References {female, her, women, females, she}	0.01	0.02 (↑44.48%)
Achievement {first, work, efficient, obtained, better}	2.15	2.19 (↑1.92%)
Certainty {all, accuracy, specific, accurate, total}	0.88	0.91 (↑3.73%)
Interrogatives {which, when, where, how, whether}	0.88	0.91 (↑3.44%)
Past Focus {used, was, been, were, obtained}	1.96	2.14 (↑9.34%)
Present Focus {is, are, be, can, have}	6.29	6.26 (↓0.56%)

Table 9: Linguistic features extracted by LIWC have the most difference between female scholars and male scholars. Each number means occurrence per string (which is abstract). The number in the parentheses shows the relative difference. We also show the top 5 words from score-All. We compare features of general abstracts (using the 83K random sample), and features of abstracts of female-authored papers. See std, full word category, etc. in the appendix.

6.2 Linguistic Features of Titles and Abstracts

Next, we take into consideration the titles and abstracts of all the papers, and calculate their linguistic features.

Frequencies of Different Word Categories. We first look at the word categories and their frequency by the Linguistic Inquiry and Word Count (LIWC) 2015 (Pennebaker et al., 2001). We show in Table 9 a selection of features on which female first-authored papers show a clear difference from average papers, and the comprehensive list of all features on all scholars, female first-author papers, female-majority papers, and female last-authored papers are in Appendix H.3.

In Table 9, for example, female first authors tend to use more words about positive emotion, such as “better” and “support.” Moreover, we can see that female first authors usually use more female references in their papers, which might be due to more female researchers publishing gender-related papers.

We also find it very interesting that the interrogative words are more dominant in female-authored papers, which is probably explained by the writing style difference that female first authors tend to use longer sentences and more commas in their writing (a more detailed analysis of which can be seen in the next paragraph), which might indicate the use of more clauses. Another interesting fact is that

past focus words are more used by female first authors, whereas male first authors’ time orientation is a more present focus, which we believe explains another aspect of the writing style difference.

Comprehensive List of Features. Apart from the word categories, we also calculate general writing features in Table 10. Some distinct features include that female first-author papers tend to have more words in the titles, less use of acronyms, but still more catchy titles. And in the abstract, female first-author papers have more sentences, a larger vocabulary, and more words, while male first-author papers are simpler according to the Flesch readability score (Talbur, 1986). Also, female first-author papers tend to include numbers more frequently.

Feature	F.	M.
Title Features		
# Words	10.08 \pm 4.18	9.44 \pm 4.12
Has Acronym	3.78%	3.86%
Catchy titles	14.96%	14.17%
Abstract Features		
# Sentences	7.32 \pm 4.24	6.97 \pm 4.18
# Vocabulary	105.27 \pm 43.22	101.41 \pm 45.75
# Words	160.46 \pm 82.70	153.24 \pm 79.89
TTR	0.64 \pm 0.07	0.63 \pm 0.08
MATTR density	91.12 \pm 5.37	90.97 \pm 6.44
Comma count	22.57 \pm 15.82	21.15 \pm 14.66
Flesch Readability (\uparrow)	10.25 \pm 22.72	10.94 \pm 30.02
# Syllables/Word	2.02 \pm 0.16	2.01 \pm 0.21
Difficult Word Ratio	0.30 \pm 0.06	0.30 \pm 0.06
Passive Speech (\downarrow)	0.64 \pm 0.47	0.64 \pm 0.47
Uncertainty Tone	4.78 \pm 0.21	4.79 \pm 0.20
Abstract Content		
Available on GitHub	0.49 \pm 0.10	0.49 \pm 0.10
Proposed a Dataset	0.49 \pm 0.12	0.49 \pm 0.12
Proposed a Task	0.53 \pm 0.11	0.53 \pm 0.11
SOTA Results	0.60 \pm 0.09	0.60 \pm 0.09
Has Numbers	50.34%	47.98%
Has Questions	1.74%	1.77%

Table 10: Linguistic features of papers with female first authors, male first authors, and all. See implementation details in Appendix H.1.

A Case Study of Title Styles. We introduce in detail our identification of catchy titles and findings. We consider a standard, straightforward paper title as mostly a declarative expression that contains the name of the task and the name of the methodology, while a catchy title is more riveting or humorous, which may involve more diverse forms including questions, quotations, exclamations, and others. According to these motivations, we build a set of linguistic rules to identify titles that carry catchy styles. The detailed algorithm is in Appendix H.2.

Example Titles from Male First-Authored Papers

- Information Power Grid: The new frontier in parallel computing?
- A systematic review of solid-pseudopapillary neoplasms: Are these rare lesions?
- Dengue fever again in Pakistan: Are we going in the right direction

Example Titles from Female First-Authored Papers

- “I want to slay that Dragon!” – Influencing choice in interactive storytelling
- Biting off more than we could chew – A surprising find on biopsy!
- ‘Spam, Spam, Spam, Spam... Lovely Spam!’ Why Is Bluespam Different?

Table 11: Stylish titles selected from paper titles of the top 5 scholars that have the largest number of stylish titles among male scholars and female scholars.

On our self-annotated test set of 1,000 paper titles randomly sampled from the AI SCHOLAR dataset, our binary classification algorithm achieves 86.3% F1, with 81.1% precision and 92.3% recall, which is significantly higher than the direct application of general catchy website title detection (Mathur, 2020) with only 13.2% F1 scores on our test set.

Apart from the overall observation in Table 10 that female first-author papers have more catchy titles in general, we can also see from the example titles in Table 11 that even among catchy titles, male and female authors tend to have different nature of attractiveness in titles, perhaps more creativity, vividness, and humor, at least from a rough glance in our data. A fine-grained analysis could be interesting for future work.

7 Conclusion

In this work, we investigated the gender differences in the AI publication world from a comprehensive range of perspectives: basic scholar profile statistics, citation trends, coauthorship, and linguistic styles of papers. We identified that the female subgroup overall still shows underrepresentation and disadvantages in the AI community. However, there are also distinct characteristics of the female subgroup that makes it unique from the general population. Our analysis provides a window to look at the current trends in our AI community, and encourages more gender equality and diversity in the future.

Limitations

It is very challenging to conduct such a large-scale and diverse-view study on gender differences in the AI publication world. Our limitations are mainly

from three perspectives: the inherent difficulties of identifying each term, inevitable noises in the data, and our method which is mainly correlational analysis.

The inherent difficulties of identifying each term are the largest limitation and constraint for this type of study, starting from the difficult process to decide some ethically-sensitive terms such as gender based on balancing concerns over both feasibility and ethics, to deciding some human-interpretable but slightly subjective categories such as catchy titles, which is a balance over reader-friendliness of the results and objectivity of the feature identification.

Another challenge is the inevitable noises in the data. Also, the most important noises come from the identification of gender, where we have to stay relatively conservative and leave a large portion of the author genders undecided, not to mention the errors for researchers whose names can be matched with names in the database but the self-identification of gender could still vary case by case. Apart from this, there are also various other noises such as selection biases. For example, not all AI researchers establish a Google Scholar profile or tag themselves in the AI domain on their profile, among many other sources of noise.

The third limitation is that our study is mainly based on analysis over correlations. It is not suggestive to directly use the study to guide interventions or decision-making, since our conclusions have not nailed down to causal factors of the disparities related to gender. In future work, it is very welcome to investigate more and use causal inference to identify, for example, mediators of academic success that provide equal opportunities for all genders.

Ethical Considerations

The ethical considerations of this study mainly overlap with our limitations. There is no perfect way when it comes to conclusions related to gender. We deeply understand that gender is highly personal and diverse in nature. In this study, we have to take a difficult step to balance the practical and ethical concerns, since the large-scale statistical analysis needs to be based on the identification of gender in a relatively scalable way. We do not wish to harm anyone, while in the meantime we try to bring as informative analysis that could be helpful for the community to understand the under-representation of the female subgroup on various

axes as possible. We are very welcome for follow-up discussions on the ethics of this study, and we are open to improvements accordingly.

Acknowledgment

We appreciate constructive discussions with Saif M. Mohammad (Senior Research Scientist at National Research Council Canada) and his generous sharing of previous experience in his NLP Scholar project.

This material is based in part upon works supported by the German Federal Ministry of Education and Research (BMBF): Tübingen AI Center, FKZ: 01IS18039B; by the Machine Learning Cluster of Excellence, EXC number 2064/1 – Project number 390727645; by the Precision Health Initiative at the University of Michigan; by the John Templeton Foundation (grant #61156); by a Responsible AI grant by the Haslerstiftung; and an ETH Grant (ETH-19 21-1).

Zhijing Jin is supported by Ph.D. fellowships from the Future of Life Institute and Open Philanthropy. We also thank OpenAI for granting Zhijing free access to their API of the GPT series through the Researcher Access Program.

References

- Abeba Birhane, Pratyusha Kalluri, Dallas Card, William Agnew, Ravit Dotan, and Michelle Bao. 2022. [The values encoded in machine learning research](#).
- Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. [Imagenet: A large-scale hierarchical image database](#). In *2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009)*, 20-25 June 2009, Miami, Florida, USA, pages 248–255. IEEE Computer Society.
- Sarah Hand, Lindsay Rice, and Eric Greenlee. 2017. Exploring teachers’ and students’ gender role bias and students’ confidence in stem fields. *Social Psychology of Education*, 20(4):929–945.
- Matthew Honnibal and Ines Montani. 2017. spaCy 2: Natural language understanding with Bloom embeddings, convolutional neural networks and incremental parsing. To appear.
- Zhijing Jin, Zhiheng Lyu, Yiwen Ding, Mrinmaya Sachan, Kun Zhang, Rada Mihalcea, and Bernhard Schoelkopf. 2022. [AI Scholars: A dataset for NLP-involved causal inference](#).
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. [ImageNet classification with deep convolutional neural networks](#). In *Advances in Neural*

- Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, pages 1106–1114.
- Saurabh Mathur. 2020. [clickbait-detector](#).
- Saif M. Mohammad. 2020. [Gender gap in natural language processing research: Disparities in authorship and citations](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7860–7870, Online. Association for Computational Linguistics.
- James W Pennebaker, Martha E Francis, and Roger J Booth. 2001. Linguistic inquiry and word count: Liwc 2001. *Mahway: Lawrence Erlbaum Associates*, 71(2001):2001.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-bert: Sentence embeddings using siamese bert-networks](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics.
- Hannah Ritchie and Max Roser. 2019. Gender ratio. *Our World in Data*. <https://ourworldindata.org/gender-ratio>.
- Rachael D Robnett. 2016. Gender bias in stem fields: Variation in prevalence and links to stem self-concept. *Psychology of women quarterly*, 40(1):65–79.
- 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. Association for Computational Linguistics.
- Konstantinos Stathoulopoulos and Juan C Mateos-Garcia. 2019. Gender diversity in ai research. *Available at SSRN 3428240*.
- John Talburt. 1986. [The flesch index: An easily programmable readability analysis algorithm](#). In *Proceedings of the 4th Annual International Conference on Systems Documentation, SIGDOC '85*, page 114–122, New York, NY, USA. Association for Computing Machinery.
- M. Iftekhar Tanveer, Samiha Samrose, Raiyan Abdul Baten, and M. Ehsan Hoque. 2018. [Awe the audience: How the narrative trajectories affect audience perception in public speaking](#). In *Proceedings of the 2018 CHI Conference on Human Factors in Computing Systems, CHI '18*, page 1–12, New York, NY, USA. Association for Computing Machinery.
- Romain Tavenard, Johann Faouzi, Gilles Vandewiele, Felix Divo, Guillaume Androz, Chester Holtz, Marie Payne, Roman Yurchak, Marc Rußwurm, Kushal Kolar, and Eli Woods. 2020. [Tslearn, a machine learning toolkit for time series data](#). *Journal of Machine Learning Research*, 21(118):1–6.
- Adam Vogel and Dan Jurafsky. 2012. [He said, she said: Gender in the ACL Anthology](#). In *Proceedings of the ACL-2012 Special Workshop on Rediscovering 50 Years of Discoveries*, pages 33–41, Jeju Island, Korea. Association for Computational Linguistics.

A Domain Tag Cleaning

Among 30,596 unique domains all scholars have, we manually extract 26 general domains and merge them with their sub-domains. To account for the fact that some scholars might label themselves with more than one domain, we normalize the count by $1 / \text{the number of domains they identify themselves with}$.

Limitations of Self-Labeling: We acknowledge that some domains have fewer samples, which may lead to a deviation in the female percentage. However, it should be noted that the female percentage in NLP, AI, and CV and their ranking in Figure 2 are matched with the result in Table 15.

B Analysis of the Population with Unclassified Gender

Our data is inclusive for various ethnicities, as shown in the left subfigure of Figure 4. However, our dataset subsamples certain groups such as east Asians, Indians and so on, as in the right subfigure of Figure 4. Some cases might be intractable. For example, Chinese names have gender markers only in their own writing system, so the gender markers are lost if we use the Romanized spelling of Chinese names on Google Scholar, leaving it only possible to classify the gender using researchers’ photos, which is unethical.




Figure 4: Ethnicity distributions of classified (left) and unclassified (right) researchers.

Further, the left-out population carry overall similar characteristics with our gender-identified population. From the analysis of profile statistics of unclassified researchers in Figure 5, we can see that the citation distribution and starting years of the two population are roughly similar, with the unclassified population is slightly younger and thus less cited.




Figure 5: Left: Citation distributions of the researcher population with classified and unclassified gender. Right: Histogram of the publication starting year (i.e., the year of the first paper according to Google Scholar) of the researcher population with classified and unclassified gender.

C General Profile Statistics

We calculate more statistics of AI researchers’ profiles. Table 12 shows standard deviations of the features in Table 1. Table 13 includes the statistics of citation within different year spans, from which we can see that female scholars’ citation is generally less than all scholars’. Table 14 confirms that female scholars take a higher percentage in younger academic age groups than in senior groups.

		Avg (std)	Min	25th	50th	75th	Max
Citations:	All	2,129.54 \pm 8,639.88	100	214	475	1,345	533,757
	F.	1,762.11 \pm 6,246.01	100	197	414	1,165	209,549
h-Index:	All	14.03 \pm 13.24	1	7	10	16	266
	F.	13.25 \pm 12.67	1	6	9	15	211
# Papers:	All	67.44 \pm 127.00	1	17	32	68	3,000
	F.	60.20 \pm 103.96	1	16	29	64	2,125
AcadAge:	All	16.89 \pm 10.65	2	10	14	20	72
	F.	16.33 \pm 9.74	2	10	14	20	73
Active Yrs.:	All	15.47 \pm 10.75	1	8	13	19	72
	F.	14.87 \pm 9.88	1	8	12	19	72
F. Coauthor (%)	All	9.58 \pm 6.18	0	5.21	8.70	12.50	48.72
	F.	14.53 \pm 7.52	1.64	9.09	13.04	18.58	48.72

Table 12: A more comprehensive version of Table 1 including the standard deviation.

		Avg (std)	Min	25th	50th	75th	Max
Citations	< Yr 2012	715.97 \pm 2841.15	1	18	83	371	122,289
	< Yr 2012 (F)	555.41 \pm 1832.53	1	16	75	343	31,130
>= Yr 2012	>= Yr 2012	2484.98 \pm 8908.21	0	280	698	1,885	467,586
	>= Yr 2012 (F)	2120.10 \pm 6901.53	0	261	639	1,670	203,008
>= Yr 2012 (newbies)	>= Yr 2012 (newbies)	905.14 \pm 4504.45	2	165	285	621	287,603
	>= Yr 2012 (newbies, F)	734.65 \pm 3063.95	15	156	250	521	79,245

Table 13: Descriptive statistics of citation breakdown. We compare the citation for the total population with the citation only for female scholars, where female citations are generally fewer.

D Implementation Details

D.1 Academia and Orgs

Identification of Academia Status: We define whether a scholar belongs to academia by their description in the GS profile. We use keyword matches such as “university”, “professor” etc., to determine their academic status. If there is no evidence that a scholar is in academia, we will label the scholar in industry.

Extraction Method for Top10 Organizations:

We use the description of each scholar (e.g., Professor of Computer Science, University of Michigan) in our AI SCHOLAR to classify their organizations. Google Scholar itself has a unique code for a wide range of organizations, and the discrepancy in position description will not affect the organization code. Thus we first cluster the organizations with unique codes and get 3568 organizations in total. For those without unique code from GS, we use Named Entity Recognition by [Honribal and Montani \(2017\)](#) to filter out the plain organization (ORG) in their description. Then we employ sentence embedding followed by a fast clustering

Academic Age	# F	# M	# All	F Rate
0-5	454	1,848	5,271	24.56
5-10	2,242	9,136	24,458	24.54
10-15	1,877	8,270	20,965	22.69
15-20	1,126	5,391	11,951	9.42
20-25	496	2,573	5,338	9.29
25-30	235	1,371	2,598	9.05
30-35	86	665	1,152	12.9
35-40	44	394	704	6.25

Table 14: The number of female scholars, male scholars, and total scholars in different groups of academic age. Female scholars take up a much higher proportion in younger academic age groups than in senior groups, while for male scholars the opposite is true.

algorithm ([Reimers and Gurevych, 2019](#)) with a cosine similarity threshold of 0.75 to cluster the organization, which results in 220 clusters. With the organization results combined from the above two methods, we obtain the top 10 most frequent organizations as Google, Stanford, CMU, MIT, Amazon, UCB, Microsoft, Facebook, IBM, and Apple.

D.2 Time Series Implementation

We simplify the method from [Tanveer et al. \(2018\)](#), where they first smooth the trajectories by a 5-point average kernel and standardize the trajectories by subtracting the time average and dividing by the time-wise standard deviation. Instead, we interpolate the citation time span to the longest active academic age and normalize the trajectories by their average citation number to focus on the relative rises and falls. In addition, we use TimeSeriesKMeans with DTW metric to cluster the trajectories, instead of density-based clustering (DBSCAN) employed by [Tanveer et al. \(2018\)](#), as DTW metric is used specifically to collect time series of similar shapes.

Algorithm 1 Generating the Time Series Clusters

Input: Citations vs. years for 78k scholars

Output: 9 Clusters of relative citations time trends

```

max_num ← max(len(years_list))
cites_upd_list ← empty list
for years, cites in zip(years_list, cites_list) do
    itp_cites ← interp(cites, max_num)
    cites_upd_list.append(itp_cites/μ(itp_cites))
end for
model ← TimeSeriesKMeans(cites_upd_list,
    n_clusters=9, metric='dtw')
```





Figure 6: Female, non-female and male scholars coauthorship.

D.3 Coauthor Diversity

Without specific clarifications, we use “all coauthors” in the main text by default, as in Table 6. We measure this by joining the features of coauthors of all papers that a scholar has. Figure 6 plots the percentage of scholars with the different number of coauthors. Female scholars tend to write a paper with slightly more coauthors.

In addition, the group of people, **scholars’ female coauthors** in our dataset, is different from what we mention as **female scholars**. Although many of our experiments mainly work on female coauthors that are also in our dataset (who we have a detailed analysis of their features), they may or may not be findable in our dataset, depending on whether they listed themselves as in AI fields and whether their citations are over 100.

To get the coauthors’ domain diversity of a scholar, we union sets of domain tags for “all coauthors” of the scholar, and divide the set size by the number of coauthors.

E Additional Basic Stats

E.1 General AI Subdomains

We calculate the female scholar percentage in some main AI domains in Table 15. We check the percentage of papers that have a female first author and female last author. The table shows that the computer vision domain has the lowest female percentage whereas the natural language processing domain has the highest female percentage. This trend also extends to the female first author paper and female last author paper in every domain.

	% F. Scholars	Paper 1st F. (%)	Paper Last F. (%)
All	17.99	21.66	16.94
AI	17.27	21.12	17.43
CV	15.57	19.66	14.30
ML	17.08	21.18	16.26
NLP	24.89	27.14	23.31

Table 15: Female author rate in 4 different fields. We calculate the rate by $\# \text{ females} / (\# \text{ females} + \# \text{ males})$ for papers. Note that the female percentage in NLP is the highest.

E.2 Age Groups

From Figure 7, we can see that the citation difference is not very large at the beginning of the career, but as we proceed to academic age groups of 15, the difference gradually shows up, and becomes larger in more senior academic age groups.




Figure 7: Citations by academic age of female scholars and total scholars. The plot shows that the gap in total citations between female scholars and total scholars keeps widening, as the academic age increases from 15. In Table 16, the column “Male Citation / Female Citation” shows a similar trend in another angle.

A. Age	Avg Citations	F. Citations	M. Citations	M:C:F:C
Overall	2122.80 ± 8626.45	1757.09 ± 6236.91	2694.15 ± 10921.73	1.5333
0 - 5	361.57 ± 706.41	317.38 ± 413.36	385.57 ± 984.26	1.2148
6 - 10	675.11 ± 2443.80	590.68 ± 2425.48	736.15 ± 2651.44	1.2463
11 - 15	1182.22 ± 5045.14	1142.19 ± 7064.93	1398.30 ± 6042.53	1.2242
16 - 20	1997.78 ± 6098.71	1787.87 ± 4258.76	2316.87 ± 7897.54	1.2959
21 - 25	3336.25 ± 7973.33	3483.79 ± 8469.01	3801.13 ± 9051.56	1.0911
≥26	7053.99 ± 18813.80	5354.27 ± 9918.48	8218.15 ± 22185.58	1.5350

Table 16: Average citations of all scholars and female scholars in different academic age groups. Female scholars’ average citations are less than average citations in nearly all academic ages, and vice versa for male scholars.

E.3 Age-Specific Dropout

In Table 17, we show relations between the dropout rate and academic status given different academic age spans. In general, scholars that are in the industry have a much higher dropout rate. In addition,

the dropout rate first does up and then down as the academic ages grow, and in academic age 6 - 10 (the time right after Ph.D.), female scholars are less likely to dropout in the industry while male scholars are not.

F Additional Analysis of Clusters

F.1 Machine-Identified 9 Clusters

Figure 12 plots all 9 clusters generated by the Time-SeriesKMeans method. As we can see from these 9 clusters, cluster 3, 5, and 8 shows the linear growth pattern of citation; cluster 1, 6, and 7 shows a common trend from rising to decline; cluster 2 and 6 shows the exponential growth; and cluster 9 shows the struggling trend of citation. Thus we manually group them into 4 general patterns and select representative cluster in Figure 3.

F.2 NLP Scholars and Time Series Cluster

The data of average academic ages per cluster is in Table 18. With the cluster label and academic age for each scholar, the Pearson coefficient is -0.039 and the p-value is 3.72e-24 using Pearson’s χ^2 test. Therefore, it shows a strong correlation between the cluster a scholar belongs to and their academic age.

	Exponential	Stumbling	Linear	Struggling
F. NLP / F. Total	137 / 718	678 / 3830	311 / 1534	32 / 131
M. NLP / M. Total	332 / 3028	2039 / 16565	1041 / 7640	79 / 728
F.% ratio in NLP	29.21	24.95	23.00	28.83
F. NLP academic age	17.50±9.58	16.04±10.56	19.95±9.83	21.75±8.71
M. NLP academic age	18.76±10.88	16.45±11.20	21.43±10.87	24.18±12.13
F. academic age	16.98±9.12	14.74±9.80	18.40±9.24	19.92±8.18
M. academic age	18.75±11.50	15.72±10.78	20.66±10.99	23.54±10.92
F. NLP citation	1884±4771	2002±3867	2244±9023	392±375
M. NLP citation	2413±4051	2841±7565	1996±4175	860±2242
F. citation	1569±3826	1897±5545	1895±9233	438±698
M. citation	3398±15660	2939±11550	2660±10068	827±2332

Table 18: NLP scholars count and # female_nlp / (# female_nlp + # male_nlp) for each cluster. Exponential growth clusters have a larger Female ratio in NLP. Female NLP scholars also have higher average citations than nearly all female scholars in all fields, while male scholars are not.

G Additional Analysis of Coauthorship

We plot a heatmap of statistics in Table 7 for better visualization.




Figure 13: Heatmap corresponding to table 7. The dataframe is constructed using boolean value. There are 91790 papers which includes unclassified gender paper.

H Linguist Style

H.1 General linguistic statistics

We follow the same features set of Jin et al. (2022) but extend their linguistic analysis in terms of stylish titles and separation of female- or male-first author papers.

H.2 Stylish Titles

The algorithm for detecting stylish paper titles is in algorithm 2. We use Part of Speech tags to capture certain syntactic characteristics. Despite its simplicity, the detection result fairly conforms with human labels of stylish titles in our understanding. Table 19 and Table 20 get the top 10 male scholars and top 5 female scholars that have the most number of stylish paper titles, and show their stylish title examples.

Algorithm 2 Algorithms of the stylish title detector

Input: The title of a paper

Output: A boolean whether the title is stylish

```

title remove noises and convert to lowercase
if special punctuation in title then
    return True
else if 1st or 2nd personal pronouns in title then
    return True
else if meaningful numeric values in title then
    return True
end if
return False

```

H.3 Full List of LIWC Features

Table 21 and Table 22 show a full list of word categories along with their frequency by using LIWC

Academic Age	F. Dropout (%) Among All → Among Industry	M. Dropout (%) All → Industry	F. Industry (%) All → Dropout	M. Industry (%) All → Dropout
0 – 5			28.99 → 0	33.17 → 0
6 – 10	3.10 → 2.71 (↓0.39)	3.59 → 3.94 (↑0.35)	38.17 → 33.33 (↓4.84)	42.30 → 46.34 (↑4.04)
11 – 15	10.89 → 13.07 (↑2.18)	11.41 → 13.36 (↑1.95)	40.05 → 48.08 (↑8.03)	46.70 → 54.68 (↑7.98)
16 – 20	14.95 → 21.51 (↑6.56)	14.07 → 17.71 (↑3.64)	43.58 → 62.71 (↑19.13)	48.64 → 61.25 (↑12.61)
21 – 25	8.49 → 13.85 (↑5.36)	13.00 → 19.83 (↑6.83)	35.65 → 58.18 (↑22.53)	45.46 → 69.37 (↑23.91)
26 – 30	10.80 → 18.03 (↑7.23)	10.53 → 16.99 (↑6.46)	33.80 → 56.41 (↑22.61)	42.24 → 68.12 (↑25.88)
31 – 35	6.53 → 9.76 (↑3.23)	9.28 → 14.95 (↑5.67)	33.47 → 50.00 (↑16.53)	37.64 → 60.66 (↑23.02)
35 – 40	8.70 → 5.88 (↓2.82)	10.36 → 16.13 (↑5.77)	29.57 → 20.00 (↓9.57)	34.83 → 54.22 (↑19.39)

Table 17: Female scholars' and male scholars' academic dropout rate (no paper published since 2018 Jan), given the total number of people of that gender and total number of people of that gender in the industry. Female scholars and male scholars industry rates given the total number of people of that gender and the total number of people of that gender who have dropped out.

2015 (Pennebaker et al., 2001).

Scholar name	# Stylish titles	Portion (%)	Sample titles
T. Y. W. (M)	294	28.25%	Does the Photographic Angle of Incidence Alter the Measured Fractal Dimension of the Retinal Vasculature? We can save not only lives, but also quality of life: submandibular gland-sparing neck dissection Erratum to: Is Sensory Loss an Understudied Risk Factor for Frailty? A Systematic Review and Meta-analysis
F. M. (M)	243	22.62%	A systematic review of solid-pseudopapillary neoplasms: Are these rare lesions? CT during arterial portography for the preoperative evaluation of hepatic tumors: how, when, and why? Bikeshare: Barriers, facilitators and impacts on car use
M. P. (M)	236	43.89%	Breaking the spell: Religion as a natural phenomenon Are we explaining consciousness yet? Speaking for our selves: An assessment of multiple personality disorder
S. G. (M)	236	31.81%	Is baseline autonomic tone associated with new onset atrial fibrillation?: Insights from the framingham heart study Biventricular pacing: more is better! 6 Field evaluation of insecticides and neem formulations for management of brinjal shoot and fruit borer, <i>Leucinodes orbonalis</i> Guenee in brinjal
J. B. (M)	234	29.83%	Energy, EROI and quality of life Integrated child development services (ICDS) scheme: a journey of 37 years Two methods for load balanced distributed adaptive integration
D. D. R. (M)	222	28.41%	Information Power Grid: The new frontier in parallel computing? Depth-first vs. best-first search Top 10 algorithms in data mining. Survey paper
J. C. (M)	219	27.63%	Do LGBT workplace diversity policies create value for firms? Peering vs. transit: Performance comparison of peering and transit inter-connections 4 Strong Association Between the-308 TNF Promoter Polymorphism and Allergic Rhinitis in Pakistani Patients
G. K. (M)	218	47.63%	Working knowledge: How organizations manage what they know Thinking for a living: How to get better performances and results from knowledge workers Saving IT's soul: Human-centered information management.
J. H. (M)	212	29.67%	Pharmacotherapy-based problems in the management of diabetes mellitus: Needs much more to be done! Long run relationship between gold prices, oil prices and Karachi stock market Dengue fever again in Pakistan: Are we going in the right direction
T. D. (M)	210	31.27%	Cloning, characterization and localization of a novel basic peroxidase gene from <i>Catharanthus roseus</i> Technology Packages: Solar, biomass and hybrid dryers Spinal tuberculosis with concomitant spondylolisthesis: coexisting entities or 'cause and effect'?

Table 19: Top 10 male scholars sorted by the number of stylish titles.




Figure 8: Linear




Figure 9: Stumbling



Figure 10: Exponential




Figure 11: Struggling

Figure 12: Citation trend of 78k scholars over their active years in 9 clusters. We further manually group them into 4 types. Clusters in Figure 8 show linear growth, Figure 9 shows a common trend from rise to decline, Figure 10 shows exponential growth, and clusters in Figure 11 is struggling. Each grey line represents the trend of an AI scholar, and five of them are randomly sampled and labeled red for easy reading. The title of each plot contains the number of scholars in that cluster. The plot also labels the scale for each cluster.

Scholar name	# Stylish titles	Portion (%)	Sample titles
P. V. (F)	155	44.87%	'Spam, spam, spam, spam... Lovely spam!' Why is Bluespam different? One world one dream? Sports blogging at the Beijing Olympic Games Forget me (in Europe), forget me not (outside Europe)
A. P. (F)	130	27.86%	Is the wolf angry or... just hungry? Tell me that bit again... bringing interactivity to a virtual storyteller "I want to slay that Dragon!" – Influencing choice in interactive storytelling
K. D. (F)	123	33.55%	How may I serve you?: A robot companion approaching a seated person in a helping context I could be you: The phenomenological dimension of social understanding Robots we like to live with! A developmental perspective on a personalized, life-long robot companion
M. H. (F)	99	46.73%	Defining profiling: A new type of knowledge? Location Data, Purpose Binding and Contextual Integrity: What's the Message? Dualism is dead. Long live plurality (instead of duality)
P. S. (F)	95	24.81%	Does Your Food Affect Your Intelligence? Taking leads out of nature, can nano deliver us from COVID-like pandemics? Biting off more than we could chew – A surprising find on biopsy!

Table 20: Top 5 female scholars sorted by the number of stylish titles.

Category		Score (All)	Score (Female Abstracts)		
			1st=F	≥50% F	Last=F
Word Count					
Summary Language Variables					
Words/Sentence					
Words > 6 Letters					
Linguistic Dimensions					
Total Function Words	{the, of, and, a, to}	43.91 ± 16.39	45.04 ± 16.30	44.43 ± 16.12	44.79 ± 16.73
Total Pronouns	{that, this, we, which, it}	4.21 ± 3.06	4.39 ± 3.08	4.28 ± 3.05	4.48 ± 3.17
Personal Pronouns	{we, our, they, them, us}	1.06 ± 1.38	1.17 ± 1.47	1.11 ± 1.41	1.19 ± 1.51
1st Person Singular	{i, mine, my, im, me}	0.01 ± 0.16	0.01 ± 0.21	0.01 ± 0.19	0.01 ± 0.23
1st Person Plural	{we, our, us, lets, ourselves}	0.90 ± 1.25	0.97 ± 1.28	0.93 ± 1.26	0.98 ± 1.31
2nd Person	{you, your, u, ya, ye}	0.01 ± 0.17	0.01 ± 0.30	0.01 ± 0.25	0.01 ± 0.32
3rd Person Singular	{his, her, he, she, him}	0.01 ± 0.17	0.01 ± 0.17	0.01 ± 0.15	0.01 ± 0.15
3rd Person Plural	{they, them, themselves, their, theirs}	0.14 ± 0.42	0.16 ± 0.46	0.15 ± 0.45	0.17 ± 0.48
Impersonal Pronouns	{that, this, which, it, these}	3.15 ± 2.31	3.22 ± 2.30	3.17 ± 2.28	3.29 ± 2.36
Articles	{the, a an}	10.27 ± 5.21	10.18 ± 5.21	10.00 ± 5.16	10.09 ± 5.20
Prepositions	{of, to, in, for, with}	17.34 ± 6.86	18.02 ± 6.91	17.76 ± 6.87	17.81 ± 7.01
Auxiliary Verbs	{is, are, be can, have}	5.58 ± 3.35	5.59 ± 3.31	5.62 ± 3.32	5.57 ± 3.40
Common Adverbs	{such, also, when, only, where}	1.88 ± 1.81	1.94 ± 1.82	1.86 ± 1.80	1.97 ± 1.87
Conjunctions	{and, as, or, also, but}	5.64 ± 3.19	5.99 ± 3.26	5.92 ± 3.24	5.92 ± 3.26
Negations	{not, without, no, cannot, negative}	0.36 ± 0.68	0.38 ± 0.71	0.38 ± 0.69	0.39 ± 0.70
Other Grammar					
Common Verbs	{is, are, be, using, based}	8.85 ± 4.65	9.01 ± 4.61	8.98 ± 4.63	8.99 ± 4.70
Common Adjectives	{as, different, new, more, than}	4.67 ± 3.13	4.87 ± 3.19	4.83 ± 3.17	4.88 ± 3.21
Comparisons	{as, different, more, than, most}	2.43 ± 2.15	2.56 ± 2.21	2.55 ± 2.24	2.56 ± 2.22
Interrogatives	{which, when, where, how, whether}	0.88 ± 1.08	0.91 ± 1.12	0.90 ± 1.11	0.95 ± 1.15
Numbers	{two, one, first, three, single}	0.75 ± 1.20	0.76 ± 1.19	0.75 ± 1.14	0.74 ± 1.12
Quantifiers	{more, each, both, most, all}	1.94 ± 1.87	1.95 ± 1.86	1.93 ± 1.84	1.93 ± 1.83
Psychological Processes					
Affective Processes	{well, important, problems, energy, problem}	2.89 ± 2.46	2.98 ± 2.47	3.00 ± 2.53	3.02 ± 2.56
Positive Emotion	{well, important, energy, better, support}	1.98 ± 1.88	2.05 ± 1.93	2.06 ± 1.96	2.11 ± 1.99
Negative Emotion	{problems, problem, low, critical, difficult}	0.88 ± 1.39	0.88 ± 1.36	0.89 ± 1.40	0.86 ± 1.37
Anxiety	{uncertainty, pressure, uncertainties, risk, risks}	0.15 ± 0.57	0.15 ± 0.55	0.16 ± 0.58	0.14 ± 0.54
Anger	{critical, attacks, argue, dominant, arguments}	0.13 ± 0.54	0.15 ± 0.57	0.14 ± 0.53	0.15 ± 0.58
Sadness	{low, lower, failure, missing, suffer}	0.20 ± 0.59	0.21 ± 0.62	0.21 ± 0.60	0.18 ± 0.54
Social Processes	{we, our, provide, they, provides}	2.81 ± 2.75	3.25 ± 3.06	3.13 ± 2.99	3.30 ± 3.04
Family	{family, families, parents, pregnancy, son}	0.03 ± 0.25	0.03 ± 0.23	0.04 ± 0.32	0.03 ± 0.33
Friends	{contact, neighborhood, neighboring, neighbors, date}	0.04 ± 0.29	0.04 ± 0.27	0.05 ± 0.31	0.05 ± 0.35
Female References	{female, her, women, females, she}	0.01 ± 0.17	0.02 ± 0.21	0.02 ± 0.23	0.02 ± 0.21
Male References	{his, male, he, men, son}	0.02 ± 0.22	0.02 ± 0.23	0.03 ± 0.24	0.02 ± 0.24

Table 21: Linguistic features extracted by LIWC. Each number means occurrence per string (which is abstract). We also show the top 5 words from score-All. We compare features of general abstracts (using the 83K random sample), and features of abstracts of female-authored papers. Among female-authored papers, we analyze papers whose first author is female (1st=F), the last author is female (last=F), and over 50% female authors.

Category		Score (All)	Score (Female Abstracts)		
			1st=F	≥50% F	Last=F
Cognitive Processes	{using, based, or, used, results}	10.94 ± 5.81	11.37 ± 5.85	11.26 ± 5.85	11.59 ± 6.01
Insight	{information, learning, analysis, knowledge, recognition}	3.77 ± 2.95	4.06 ± 3.05	4.04 ± 3.04	4.17 ± 3.14
Causation	{using, based, used, results, use}	3.44 ± 2.43	3.48 ± 2.44	3.45 ± 2.44	3.50 ± 2.45
Discrepancy	{problems, problem, need, could, if}	0.56 ± 0.96	0.57 ± 0.94	0.57 ± 0.95	0.60 ± 0.99
Tentative	{or, most, may, some, any}	1.72 ± 1.89	1.70 ± 1.84	1.68 ± 1.82	1.75 ± 1.87
Certainty	{all, accuracy, specific, accurate, total}	0.88 ± 1.14	0.91 ± 1.14	0.88 ± 1.13	0.94 ± 1.21
Differentiation	{or, different, not, than, other}	1.57 ± 1.69	1.65 ± 1.74	1.63 ± 1.73	1.69 ± 1.77
Perceptual Processes	{show, images, search, fuzzy, image}	1.42 ± 1.86	1.47 ± 1.88	1.42 ± 1.87	1.40 ± 1.86
See	{show, images, search, image, shows}	0.85 ± 1.35	0.86 ± 1.32	0.85 ± 1.35	0.82 ± 1.30
Hear	{noise, noisy, music, voice, speech}	0.17 ± 0.74	0.21 ± 0.81	0.19 ± 0.75	0.20 ± 0.80
Feel	{fuzzy, flexible, weight, weighted, hand}	0.25 ± 0.80	0.23 ± 0.78	0.23 ± 0.75	0.23 ± 0.77
Biological Processes	{clinical, expression, face, medical, physical}	1.16 ± 2.20	1.37 ± 2.41	1.44 ± 2.54	1.18 ± 2.30
Body	{face, blood, hand, heart, neurons}	0.28 ± 0.98	0.31 ± 1.04	0.31 ± 1.03	0.25 ± 0.93
Health	{clinical, medical, physical, health, diagnosis}	0.71 ± 1.67	0.85 ± 1.83	0.92 ± 1.95	0.76 ± 1.80
Sexual	{prostate, pregnancy, sex, ovarian, arousal}	0.02 ± 0.25	0.02 ± 0.32	0.02 ± 0.28	0.03 ± 0.33
Ingestion	{expression, water, weight, expressions, expressed}	0.16 ± 0.68	0.19 ± 0.75	0.20 ± 0.81	0.16 ± 0.67
Drives	{we, approach, our, first, over}	6.65 ± 4.26	6.92 ± 4.21	6.82 ± 4.27	7.07 ± 4.40
Affiliation	{we, our, social, communication, interaction}	1.62 ± 1.92	1.81 ± 2.03	1.76 ± 2.00	1.84 ± 2.03
Achievement	{first, work, efficient, obtained, better}	2.15 ± 1.99	2.19 ± 1.97	2.15 ± 1.98	2.23 ± 2.05
Power	{over, high, order, large, important}	2.11 ± 2.06	2.15 ± 2.07	2.15 ± 2.08	2.21 ± 2.12
Reward	{approach, obtained, approaches, better, best}	1.10 ± 1.30	1.11 ± 1.29	1.11 ± 1.30	1.16 ± 1.36
Risk	{problems, problem, security, difficult, lack}	0.52 ± 1.02	0.50 ± 0.97	0.51 ± 1.01	0.50 ± 1.00
Time Orientations					
Past Focus	{used, was, been, were, obtained}	1.96 ± 2.23	2.14 ± 2.39	2.19 ± 2.45	2.01 ± 2.25
Present Focus	{is, are, be, can, have}	6.29 ± 3.65	6.26 ± 3.63	6.16 ± 3.65	6.37 ± 3.71
Future Focus	{may, then, will, prediction, future}	0.61 ± 1.08	0.60 ± 1.05	0.63 ± 1.09	0.65 ± 1.11
Relativity	{in, on, at, approach, new}	10.84 ± 5.64	11.09 ± 5.66	11.00 ± 5.61	10.95 ± 5.70
Motion	{approach, approaches, behavior, changes, increase}	1.44 ± 1.62	1.47 ± 1.61	1.44 ± 1.61	1.47 ± 1.63
Space	{in, on, at, into, both}	6.96 ± 4.05	7.10 ± 4.07	7.05 ± 4.03	6.98 ± 4.03
Time	{new, present, first, when, then}	2.40 ± 2.20	2.44 ± 2.25	2.44 ± 2.23	2.43 ± 2.23
Personal Concerns					
Work	{performance, learning, analysis, paper, applications}	4.53 ± 3.58	4.78 ± 3.70	4.74 ± 3.69	4.95 ± 3.91
Leisure	{novel, expression, channels, videos, play}	0.48 ± 1.01	0.51 ± 1.08	0.51 ± 1.06	0.48 ± 1.00
Home	{address, family, home, neighborhood, neighboring}	0.12 ± 0.45	0.11 ± 0.44	0.12 ± 0.44	0.12 ± 0.47
Money	{investigate, cost, investigated, free, economic}	0.42 ± 1.02	0.43 ± 0.99	0.42 ± 1.00	0.44 ± 1.03
Religion	{beliefs, moral, sacrificing, monkeys, agnostic}	0.01 ± 0.17	0.02 ± 0.18	0.02 ± 0.19	0.02 ± 0.21
Death	{mortality, die, mortality, deaths, death}	0.04 ± 0.31	0.04 ± 0.32	0.04 ± 0.36	0.04 ± 0.33
Informal Language	{well, o, da, en, um}	0.16 ± 0.66	0.16 ± 0.65	0.16 ± 0.68	0.15 ± 0.64
Swear Words	{retardation, dummy, screws, screw, retarded}	0.00 ± 0.05	0.00 ± 0.04	0.00 ± 0.03	0.00 ± 0.03
Netspeak	{o, da, em, k, mm}	0.04 ± 0.50	0.04 ± 0.53	0.04 ± 0.53	0.04 ± 0.51
Assent	{k, indeed, agree, absolutely, cool}	0.01 ± 0.13	0.01 ± 0.15	0.01 ± 0.13	0.01 ± 0.12
Nonfluencies	{well, um, mm, er, ah}	0.11 ± 0.37	0.12 ± 0.37	0.11 ± 0.36	0.11 ± 0.35
Fillers	{rrani, rranr}	0.00 ± 0.01	0.00 ± 0.00	0.00 ± 0.00	0.00 ± 0.00

Table 22: Following Table 21.