

It's a Man's Wikipedia? Assessing Gender Inequality in an Online Encyclopedia

Wagner, Claudia; Garcia, David; Jadidi, Mohsen; Strohmaier, Markus

Veröffentlichungsversion / Published Version
Konferenzbeitrag / conference paper

Zur Verfügung gestellt in Kooperation mit / provided in cooperation with:
GESIS - Leibniz-Institut für Sozialwissenschaften

Empfohlene Zitierung / Suggested Citation:

Wagner, C., Garcia, D., Jadidi, M., & Strohmaier, M. (2015). It's a Man's Wikipedia? Assessing Gender Inequality in an Online Encyclopedia. In *Proceedings of the 9th International AAAI Conference on Weblogs and Social Media* (pp. 454-463). Palo Alto, CA: Association for the Advancement of Artificial Intelligence (AAAI). <https://nbn-resolving.org/urn:nbn:de:0168-ssoar-66086-2>

Nutzungsbedingungen:

Dieser Text wird unter einer Deposit-Lizenz (Keine Weiterverbreitung - keine Bearbeitung) zur Verfügung gestellt. Gewährt wird ein nicht exklusives, nicht übertragbares, persönliches und beschränktes Recht auf Nutzung dieses Dokuments. Dieses Dokument ist ausschließlich für den persönlichen, nicht-kommerziellen Gebrauch bestimmt. Auf sämtlichen Kopien dieses Dokuments müssen alle Urheberrechtshinweise und sonstigen Hinweise auf gesetzlichen Schutz beibehalten werden. Sie dürfen dieses Dokument nicht in irgendeiner Weise abändern, noch dürfen Sie dieses Dokument für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, aufführen, vertreiben oder anderweitig nutzen.

Mit der Verwendung dieses Dokuments erkennen Sie die Nutzungsbedingungen an.

Terms of use:

This document is made available under Deposit Licence (No Redistribution - no modifications). We grant a non-exclusive, non-transferable, individual and limited right to using this document. This document is solely intended for your personal, non-commercial use. All of the copies of this documents must retain all copyright information and other information regarding legal protection. You are not allowed to alter this document in any way, to copy it for public or commercial purposes, to exhibit the document in public, to perform, distribute or otherwise use the document in public.

By using this particular document, you accept the above-stated conditions of use.

It's a Man's Wikipedia?

Assessing Gender Inequality in an Online Encyclopedia

Claudia Wagner
GESIS & Uni. Koblenz
claudia.wagner@gesis.org

David Garcia
ETH Zürich
dgarcia@ethz.ch

Mohsen Jadidi
GESIS
mohsen.jadidi@gesis.org

Markus Strohmaier
GESIS & Uni. Koblenz
markus.strohmaier@gesis.org

Abstract

Wikipedia is a community-created encyclopedia that contains information about notable people from different countries, epochs and disciplines and aims to document the world's knowledge from a neutral point of view. However, the narrow diversity of the Wikipedia editor community has the potential to introduce systemic biases such as gender biases into the content of Wikipedia. In this paper we aim to tackle a sub problem of this larger challenge by presenting and applying a computational method for *assessing gender bias on Wikipedia along multiple dimensions*. We find that while women on Wikipedia are covered and featured well in many Wikipedia language editions, *the way women are portrayed* starkly differs from the way men are portrayed. We hope our work contributes to increasing awareness about gender biases online, and in particular to raising attention to the different levels in which gender biases can manifest themselves on the web.

Introduction

Wikipedia aims to provide a platform to freely share the sum of all human knowledge. It represents an influential source of information on the web, containing encyclopedic information about notable people from different countries, epochs and disciplines that is used for learning and educational purposes worldwide. Wikipedia is also a community-created effort driven by a self-selected set of editors. The demographic characteristics of this set of editors is known: it is predominately white and male (Lam et al. 2011; Collier and Bear 2012; Hill and Shaw 2013).

This known gender bias in the population of editors has the potential to introduce gender biases into the contents of Wikipedia as well. For example, the population bias might lead to differences in the ways women and men are portrayed on Wikipedia. It might also mimic or even exaggerate inequalities that are already existing in the real world. At the same time, assessing the manifold and subtle ways in which gender biases can manifest themselves has been challenging, and we know little about the different dimensions

of gender biases on Wikipedia. Yet, due to the influential nature of Wikipedia, it is important to reveal, assess and correct such biases, if they exist. This paper tackles a sub-part of this larger challenge.

Objectives: In particular, the overall goal of this work is to *assess potential gender inequalities in Wikipedia articles* along different dimensions.

Approach: To assess the extent to which Wikipedia suffers from potential gender bias, we analyze articles about notable people in six language editions along four different gender bias dimensions: coverage bias, structural bias, lexical bias and visibility bias. *Coverage bias* determines differences between the number of notable women and men portrayed on Wikipedia. For example, one might hypothesize that notable men are more likely to be covered by Wikipedia. *Structural bias* quantifies gender homophily/disassortativity, i.e. gender-specific tendencies to preferably link articles of notable people with the same or different gender. For example, one might hypothesize that articles about women have more links to men than vice versa. *Lexical bias* reveals inequalities in the words used to describe notable men and women on Wikipedia. For example, articles about women are potentially more likely to mention their family (husband or kids) than articles about men. *Visibility bias* reflects how many articles about men or women make it to the front page of Wikipedia. Again, one can hypothesize that articles about men might have better chances to be selected.

Contributions & Findings: We present and apply a computational method for *assessing gender bias on Wikipedia along multiple dimensions*. We find that most Wikipedia language editions exhibit a slight over-representation of women, but the proportional differences in the coverage of men and women are not significant. That means, men and women are covered equally well in all six Wikipedia language editions. Also on the visibility level, we do not find any evidence for male-bias in the selection procedure of articles that are featured on the startpage of the English Wikipedia. These are encouraging findings suggesting that the Wikipedia editor community is sensible to gender in-

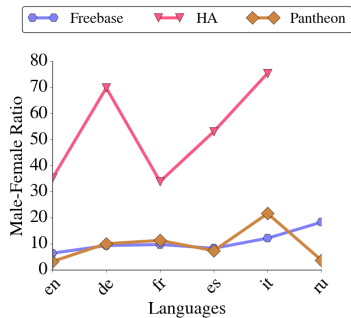


Figure 1: **Male-Female Ratio:** The ratio of men and women in our reference datasets that are born in a country where one of the six languages is predominantly spoken. Across all language editions the local heroes of a country tend to be predominantly male. For example, if we look at notable people in freebase we find between 7 and 12 times more men than women depending on which countries we consider.

equalities¹ and covers notable women and men equally well. However, we also find that *the way women are portrayed* on Wikipedia starkly differs from the way men are portrayed. We find evidence for both structural and lexical gender biases. On a structural level, we observe an asymmetry: Women on Wikipedia tend to be more linked to men than vice versa. On a lexical level we find that especially romantic relationships and family-related issues are much more frequently discussed on Wikipedia articles about women than men.

Materials & Methods

In the following we discuss our data collection and our methodology that allows to systematically explore gender inequalities on Wikipedia on multiple dimensions.

Datasets

To estimate the bias on Wikipedia that goes beyond the bias in the offline world, ideally one would have a complete list of notable people available that is (a) not biased and (b) independent from Wikipedia. Since it is impossible to obtain such a list, we use the following three collections of notable

¹also cf. http://meta.wikimedia.org/wiki/Gender_gap

Table 1: **Statistics of the datasets:** The number of articles and median article length of all Wikipedia articles that belong to one of the notable people from our three reference datasets.

	Freebase	HA	Pantheon
Total Num Articles	109,481	4,002	11,341
Female Articles	12,685	88	1,496
Male Articles	96,796	3,914	9,845
Median Num Words	458	1,121	1,106
Female			
Median Num Words	412	820	1,017
Male			

people as *reference datasets*, each having different strength and weaknesses:

Freebase: We use a collection of around 120k notable people that has been used in previous research for studying the mobility of notable people (Schich et al. 2014) and was obtained from freebase. Freebase contains data harvested from sources such as Wikipedia, NNDB, FMD and MusicBrainz, as well as individually contributed data from users. We only take individuals into account for which gender and basic bibliographic information (i.e., full birth and death date and birth and death location) is available. Freebase directly links to Wikipedia articles in different language editions, if articles about the entity are available.

Pantheon: Pantheon is a project developed by the Macro Connections group at the MIT Media Lab that is collecting, analyzing, and visualizing data on historical cultural popularity and production. The Pantheon dataset (Amy Yu and Hidalgo 2013) contains information on 11,340 biographies that have presence in more than 25 languages in the Wikipedia (as of May 2013) and provides links to Wikipedia articles about these people.

Human Accomplishment: The third dataset which we use is compiled from a book called “Human Accomplishment” (Murray 2003) (short HA) and contains information on 4,002 eminent individuals from arts and sciences who made a significant contribution prior to 1950. The inventories were constructed by Charles Murray using linguistic records, such as encyclopedia entries from a number of different languages and sources. Also this dataset has biases since e.g. Murray relied mainly on materials in Roman-alphabet languages. To find Wikipedia articles about those individuals, we use the Wikipedia search API and search for the full name. To select the right search result from the list we compare the birth date, birth location, death date and death location of the candidates in the search results with the person we are looking for.

Data Collection Procedure: We crawled the content of articles about people in our reference datasets using Wikipedia’s API in November 2014. For the English Wikipedia, the articles that have been featured at the front page in the last few years were extracted from the “Today’s Featured Article” archive². Table 1 provides the basic statistics for each dataset and Figure 1 shows the ratio between men and women that are born in a country where one of the six languages we studied is predominantly spoken. The overlap between the three reference datasets is very low. For example, for those people from our reference datasets which we could map to the English Wikipedia the Jaccard coefficient is 0.016 for freebase and HA, 0.035 for freebase and pantheon and 0.097 for pantheon and HA. The six language editions that we explore in this study are those which had the highest coverage of notable men and women from our largest reference dataset, freebase.

²http://en.wikipedia.org/wiki/Wikipedia:Today%27s_featured_article

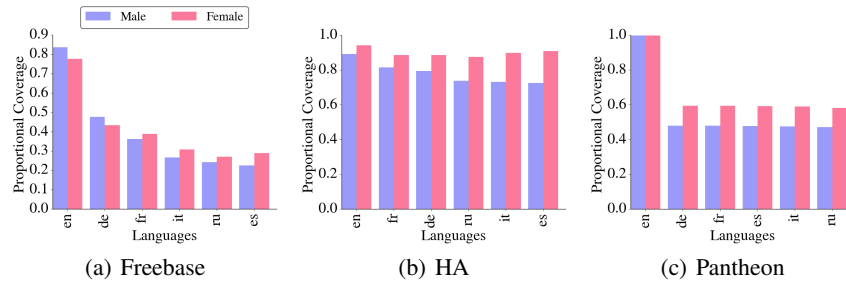


Figure 2: **Coverage Bias:** Proportional coverage of notable women and men. Surprisingly, in most language editions the proportion of notable women covered is slightly higher than the proportion of notable men.

Measuring Gender Inequality

We propose to analyze gender inequality on Wikipedia on the following four dimensions: which notable men or women are presented on Wikipedia (coverage bias)? How are they presented (lexical bias)? What structure emerges from the hyperlink network of articles (structural bias)? And which articles get featured on the startpage of Wikipedia (visibility bias)?

Coverage Bias: To estimate coverage bias we compare the proportions of notable men and women of different reference datasets that are covered by Wikipedia. Ideally, a reference dataset consists of an unbiased list of people who should be presented on Wikipedia. It is important to understand that a biased reference dataset will obviously impact our results. If, for example, our reference dataset is already biased towards men (i.e., it covers only extremely famous women but also less famous men) than the proportion of women who are represented on Wikipedia would probably be higher than the proportion of men. To address this issue we analyze the coverage using several independent reference datasets (Jaccard coefficient between the three datasets ranges from 0.0 to 0.12 for different language editions), assuming that each of them will have a different bias and seeking patterns that exist across all three datasets.

Further, gender-differences in the extent to which men and women are covered on Wikipedia may exist. Therefore, we also analyse the article length distribution of men and women.

Structural Bias: We analyze the patterns of gender assortativity based on the probability that an article about a person of one gender links to an article about a person of the other gender. We compare the probability that a link ends in an article of gender g_2 given that it comes from an article of gender g_1 with the probability that a link ends in an article of gender g_2 regardless of the gender of its origin:

$$L(g_1, g_2) = \log \left(\frac{P(to = g_2 | from = g_1)}{P(to = g_2)} \right) \quad (1)$$

where $P(to = g_2 | from = g_1)$ is the conditional distribution that an edge links to an article of gender g_2 given that it comes from an article of gender g_1 , and $P(to = g_2)$ is the probability that any link ends in an article of gender g_2 regardless of the gender of its origin. L measures the log likelihood ratio between edge probabilities, comparing the

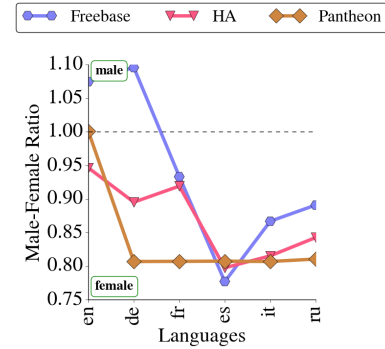


Figure 3: **Coverage Gap:** Ratio between the number of notable men and women from three different reference lists that are covered on different language editions of Wikipedia.

posterior probability of finding a gender at the edge of a link given that we know the gender of its origin, and comparing it with the base rate of linking to an article of gender g_2 . This way, positive values of L indicate increased connectivity from g_1 to g_2 , and negative values the opposite, and define a assortativity matrix of the four combinations of genders that measures the tendencies to connect within and across genders.

For the case of same gender connections we use the standard definition of assortativity (Newman 2003):

$$\frac{\sum_g P(from = g, to = g) - P(from = g) * P(to = g)}{1 - \sum_g P(from = g) * P(to = g)} \quad (2)$$

For the case of asymmetry across genders, we compare the entries of L from one gender to the other, as $A = L(F, M) - L(M, F)$. Positive values of A will indicate a stronger tendency of articles about women to connect to articles about men than the opposite, controlling for the difference in in-degrees and sizes of both genders.

The finding of gender assortativity and asymmetry between genders requires a test that allows us to compare our empirical estimates against null models of the network. For that reason, we set up numerical simulations of three different null models: a *randomized gender* model in which we shuffle the genders of nodes; a *randomized link end* model in which we rewire links to random articles, maintaining out

degrees but fully randomizing in-degree; and a *randomized link origin* model, in which we maintain link ends but rewire their origin to an article sampled at random, which maintains in-degrees but randomizes out degrees. We run each simulation 10,000 times, recording values of assortativity and asymmetry to measure the mean and 95% confidence intervals of these two statistics under each null model.

Structural biases can also manifest in the centrality measures, as suggested by the *Smurfette principle* (Pollitt 1991). That means, women can be positioned in the periphery of a network with a core composed of men. In that case the centrality of women would be lower. We operationalize centrality on Wikipedia as a quantification of importance, measuring the in-degree and k-core-ness of an article. The in-degree of article p is trivially calculated as the amount of articles that link to article p , and the in k-core-ness is computed through a pruning mechanism based on in-degree (Giatsidis, Thilikos, and Vazirgiannis 2013).

Lexical Bias: To explore gender-specific lexical inequalities on Wikipedia we use an open vocabulary approach, inspired by (Schwartz et al. 2013). An open-vocabulary approach is not limited to predefined word lists, but linguistics are automatically determined from the text. We compute the tfidf scores of the word stems obtained from a Snowball Stemmer and use them as features to train a Naive Bayes classifier. The classifier determines which words are most effective in distinguishing the gender of the person an article is about. Log likelihood ratios $L(word, g)$ are used for comparing different feature-outcome relationships.

$$L(word, g) = \log \left(\frac{P(word|g)}{P(word)} \right) \quad (3)$$

where $P(word|g)$ is the conditional distribution that a word shows up in an article about a person given that the person's gender is g , and $P(word)$ is the probability that a word shows up in any article regardless of the gender of the person the article is about.

The *Finkbeiner test* (Finkbeiner 2013) suggests that articles about women often emphasize the fact that she is a woman, mention her husband and his job, her kids and child care arrangements, how she nurtures her underlings, how she was taken aback by the competitiveness in her field and how she is such a role model for other women. Also the historian Gillian Thomas who investigated the role of women in Britannica states in her book (Thomas 1992) that as contributors, women were relegated to matters of “social and purely feminine affairs” and as subjects, women were often little more than addenda to male biographies (e.g., Marie Curie as the wife of Pierre Curie).

We create the following three categories of words that capture some aspects that could be over-represented in articles about women according to what Thomas observed in the Britannica and what the Finkbeiner test suggest:

- *Gender* category contains words that emphasize that someone is a man or woman (i.e., man, women, mrs, mrs, lady, gentleman)
- *Relationship* category consists of words about romantic relationships (e.g., married, divorced, couple, husband, wife)

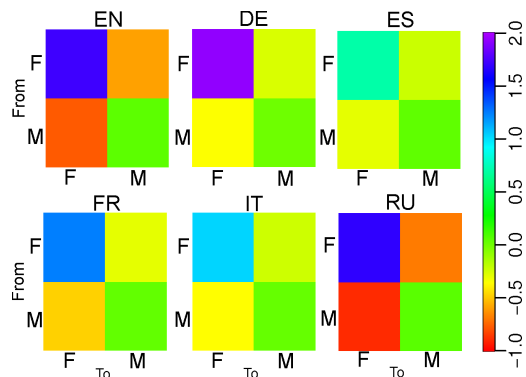


Figure 4: **Structural Assortativity and Asymmetry Bias:** Logarithmic assortativity matrices for the hyperlink networks of articles about notable men and women in six language editions of Wikipedia. Assortativity of connections within genders becomes apparent for the minority class, women. All language editions show an asymmetry of connectivity across genders. The strongest assortativity and asymmetry is visible in the English and Russian Wikipedia.

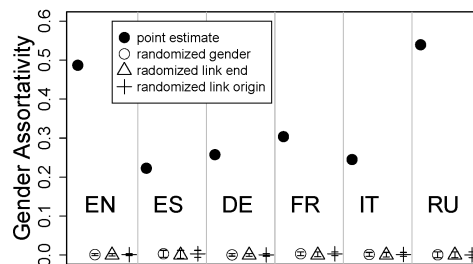


Figure 5: **Significance of Structural Assortativity Bias:** Point estimates of gender assortativity in six language editions and comparison with the three null reference models. Error bars (smaller than symbol size) show 95% confidence intervals over 10,000 simulations of each model. The empirical estimates are significant in comparison to the narrow confidence interval of the null models.

- *Family* category aggregates words about family relations (e.g., kids, children, mother, grandmother).

All other words that cannot be assigned to the above mentioned categories fall into the category *Others*. To gain further insights into the types of words that have the highest log likelihood ratio for articles about men or women, native speakers of each language manually code the 150 words which are most useful for differentiating articles about men and women in each language edition.

Visibility Bias: To estimate visibility bias we simply compare the proportions of notable men and women of different reference datasets that got featured on the startpage of the English Wikipedia. We test the significance of the difference in proportions between men and women that got featured using a Chi-Square test.

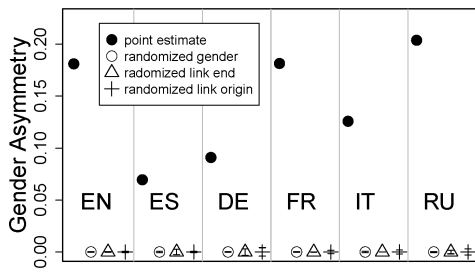


Figure 6: **Significance of Structural Asymmetry Bias:** Arithmetic mean of point estimates of gender asymmetry for men and women in six language editions and comparison with the three null reference models. Error bars (smaller than symbol size) show 95% confidence intervals over 10,000 simulations of each model. The empirical estimates are significant in comparison to the narrow confidence interval of the null models.

Results

In the following, we present our empirical results on gender inequality on Wikipedia.

Coverage Bias

Figure 2 shows that the best coverage across languages is achieved for people that made significant contributions to science and arts before 1950 and are therefore listed in the HA reference dataset. Across all three reference datasets we consistently observe that *women are not* - as initially hypothesized - underrepresented on Wikipedia, but are even *slightly overrepresented* (cf. Figure 3). Also when looking at article notable distributions of men and women, we see that *articles about women tend to be longer* than articles about men (cf. Table 1) in all three datasets. This could potentially be the result of the effort of Wikipedians to improve the coverage of minorities such as women or it can be a side product of a bias in our reference datasets which may only include very notable women, but may also cover less notable men. We addressed the later issue by selecting several reference datasets which we hope are not all subject to the same bias.

Structural Bias

Figure 4 shows the logarithmic assortativity matrices of articles about men and women in six different language editions of Wikipedia based on our largest reference dataset, Freebase. The assortativity of connections within genders becomes apparent for the minority class, women, in all cases (cf. high values of $L(F, F)$). The matrices also provide a comparison across genders: $L(F, M)$ and $L(M, F)$ are both slightly negative in all language edition, which means that women connect less to men and men less to women than we would expect. All language editions show an asymmetry of connectivity across genders, even when we correct for overall incidence in Equation 1. The value of $L(F, M)$ tends to be higher than $L(M, F)$, which means that men link even less to women than women to men.

Figures 5 and 6 show the arithmetic mean of the empirical point estimates of assortativity and asymmetry for both gender, in comparison with the values in the three null models. It is evident that the three randomization methods destroy any kind of assortativity or asymmetry pattern, and that the empirical estimates are significant in comparison to the narrow confidence interval of the null models. Assortativity is positive in all cases, indicating that *articles about people with the same gender tend to link to each other*. For the case of asymmetry, there is a positive value of A (which we defined as $A = L(F, M) - L(M, F)$) in all six language editions, validating our observation that *articles about women tend to link more to articles about men than the opposite*.

The above results show the existence of assortativity and asymmetry across genders controlling for degree. However, structural biases can also manifest in the centrality measures, as suggested by the *Smurfette principle* (Pollitt 1991). To test the existence of this principle, we compare in-degree and k-core-ness of articles about men and women on Wikipedia. Figure 7 shows the complementary cumulative density functions $P(d_i > D)$ for in-degree and $P(k_i > K)$ for in k-core-ness in the six networks. An initial observation reveals that, in general, the tail of in-degree and in k-core-ness of male articles is longer than for women articles, which is specially pronounced in the case of k-core-ness of German and Russian. We validate the above observations by measuring the distance between the two distributions and test the significance of the distance through a two-tailed Wilcoxon tests and Kolmogorov-Smirnov test (cf. Table 2). Our results highlight that, according to their in-degree distribution, men are indeed significantly more central in all language editions with $p < 0.05$ except in the Spanish one where men and women are equally central. The k-core-ness distributions suggest that in all language editions except the Spanish, the Italian and the French one, men are more central than women. This indicates, in some language editions like the English, the Russian and the German one, men are always significantly more central than women, no matter how we measure centrality.

Lexical Bias

Our lexical analysis reveals that articles about women tend to emphasize the fact that they are about a women (i.e., they contain words like “woman”, “female” or “lady”), while articles about men don’t contain words like “man”, “masculine” or “gentleman”. The lower salience of male-related words in articles about men can be related to the concept of male as the *null gender* (Fox, Johnson, and Rosser 2006), which suggests that there is a social bias to assume male as the standard gender in certain social situations. This would imply that male-defining words are not necessary because the context already defines the gender of the person the article talks about. This seems to be a plausible assumption due to the imbalance between the number of articles about men and women (cf. Table 1).

We also noticed that the relationship status and family related issues seem to be more extensively discussed in articles about woman since words like “married”, “divorced”, “children” or “family” are much more frequently used in ar-

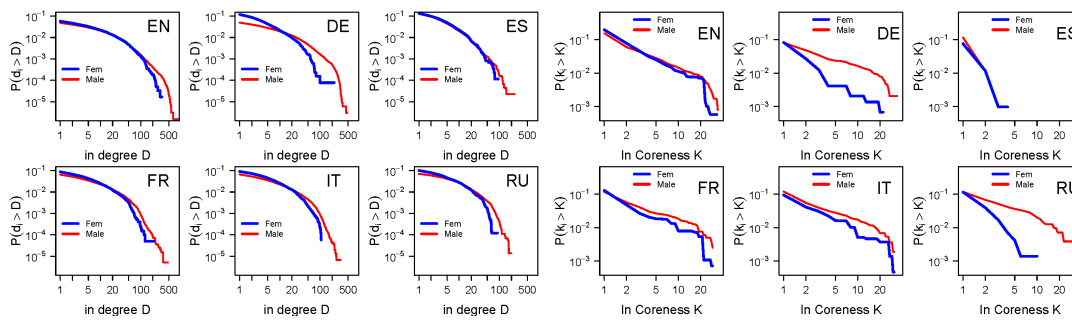


Figure 7: **Structural Centrality Bias:** Complementary cumulative density function of the in-degree distributions (left) and in k-core decompositions (right) of articles about men and women in six language editions. In some language editions like the English (EN), the Russian (RU) and the German (DE) one, men are always significantly more central than women, no matter how we measure centrality, while in others like the Spanish (ES) one, women and men are either equally central or women are more central.

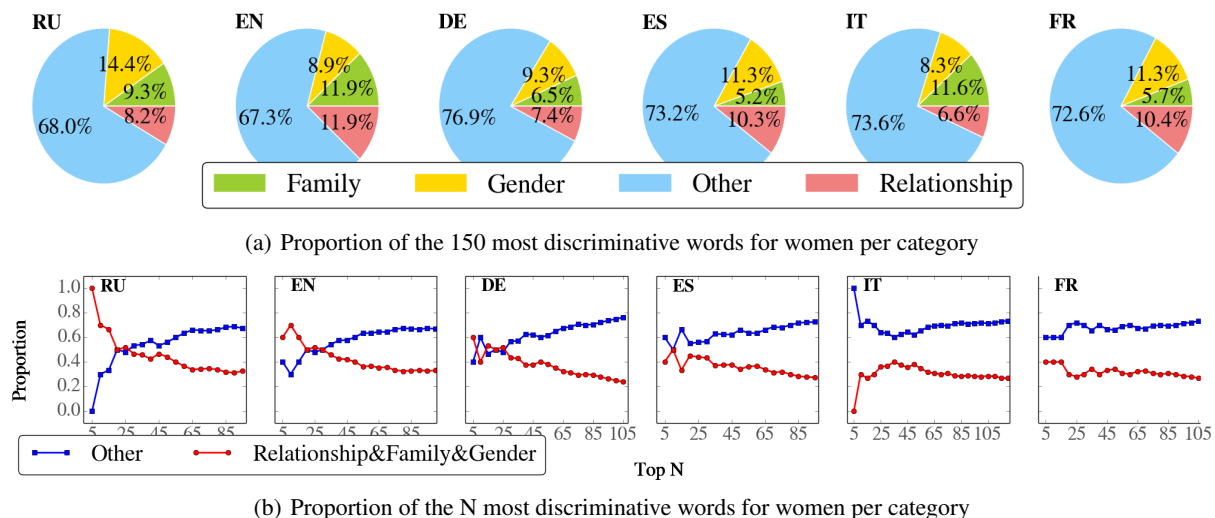


Figure 8: **Lexical Bias:** The proportion of the 150 most discriminative words of articles about women that belong to different categories. In all language editions between 32% and 23% of the 150 most indicative words for women belong to one of the three categories, while only between 0% and 4% of the most discriminative words for men belong to one of these categories. In some language edition, like the Russian (RU), the English (EN) and the German (DE) one, the proportion of the most discriminative words that belong to one of these three categories is especially high among the top words.

ticles about women. This confirms that *men and women are indeed presented differently* on Wikipedia and that those differences go beyond what we would expect due to the history of gender inequalities - i.e., the fact that it was more difficult for women to become famous in the past, amongst others because of unequal access to resources and the fact that the history was mainly documented through the eyes of men. We leave the question of investigating if the lexical bias on Wikipedia reflects the lexical bias from the general media or if the Wikipedia editor community introduces an additional bias because of their narrow demographics for future work.

We use log likelihood ratios for comparing different word-gender relationships. Not surprisingly, the most indicative words for men are often related to certain domains or fields (e.g., certain sports or professions). For example,

the most discriminative word stems for men in the English Wikipedia are “basebal”, “footbal” and “infantri” and an article that contains a word with the stem “basebal” is 11.5 times more likely to be about a man than a woman.

For women the picture is different since among the most discriminative words for women, words like “husband”, “female” and “woman” can be found. To gain more insights into those difference, we use the previously introduced categories of words and manually code the words with the highest likelihood ratio for men or women. Our results clearly show that across all language editions almost all words that fall into the category Family, Relationship or Gender, reveal a high likelihood ratio for women. Figure 8(a) shows that between 32% and 23% of the 150 most indicative words for women belong to one of the three categories. Note that for

men only 0% and 4% of the most discriminative words belong to one of these categories. That means, words that fall into one of those categories indeed indicate that an article is about a woman which suggests that lexical gender inequalities are present on Wikipedia. Especially, in the Russian and English Wikipedia, we can see that the majority of the 25 most discriminative words of females fall into one of those three categories (cf. Figure 8(b)).

What are these words that fall into the categories Family, Relationship or Gender and discriminate men and women? Table 3 and 4 show the word stems with the highest gender-specific log-likelihood ratio that belong to one of the three categories. Almost all of them are indicative for women which means that words which are indicative for men tend not to fall into these categories. One can further see that, for instance, in the English Wikipedia an article about a notable person that mentions that the *person is divorced is 4.4 times more likely to be about a woman* rather than a man. We observe similar results in all six language editions. For example, in the German Wikipedia an article that mentions that a person is divorced is 4.7 times more likely about a women, in the Russian Wikipedia its 4.8 time more likely about a woman and in the Spanish, Italian and French Wikipedia it is 4.2 times more likely about a women.

This example shows that a lexical bias is indeed present on Wikipedia and can be observed consistently across different language editions. This result is in line with (Bamman and Smith 2014) who also observed that in the English Wikipedia biographies of women disproportionately focus on marriage and divorce compared to those of men.

Visibility Bias

Figure 9 shows the proportion of notable men and women that showed up at the front page of the English Wikipedia in the past few years. One can see that proportions of men and women that got selected are very small and therefore also the differences are marginal. Though we observe across all years that the proportion of men that were selected and

	W_i	$p_i <$	$ks_i <$	W_k	$p_k <$	$ks_k <$
EN	-	10^{-15}	10^{-15}	-	0.03	10^{-4}
ES	+	0.17	0.02	+	10^{-4}	10^{-4}
DE	-	10^{-15}	10^{-15}	-	10^{-12}	10^{-8}
FR	-	10^{-9}	10^{-5}	-	0.07	0.09
IT	-	10^{-6}	10^{-3}	+	0.95	10^{-4}
RU	-	10^{-4}	10^{-7}	-	0.55	0.003

Table 2: **Significance of Structural Centrality Bias:** Differences between the in-degree distributions (W_i) and k-coreness distributions (W_k) of men and women. A positive difference (+) indicates that women are more central, while a negative difference (-) indicates that men are more central. The significance of the difference as suggested by the Wilcoxon test ($p_i <$) and by the Kolmogorov-Smirnov test ($ks_i <$). In some language editions like the English (EN), the Russian (RU) and the German (DE) one, men are indeed significantly more central than women according to both centrality measures.

featured at the startpage was slightly higher, the Chi-Square test suggests that the difference in proportions is not significant. Therefore, we conclude that the *selection procedure of featured articles of the Wikipedia community does not suffer from gender bias*.

Discussion

While Wikipedia’s massive reach in coverage ensures that notable women have high likelihood of being represented on

Category	Term	Female	Male
Relationship	husband	9.2	1.0
Gender	female	8.2	1.0
Relationship	aunt	6.5	1.0
Gender	women	6.4	1.0
Gender	madam	6.1	1.0
Gender	woman	5.6	1.0
Family	grandmoth	5.5	1.0
Gender	girl	5.3	1.0
Gender	mrs	4.9	1.0
Relationship	divorc	4.4	1.0
Gender	ladi	4.4	1.0
Relationship	wed	4.3	1.0
Relationship	marriag	3.8	1.0
Relationship	lover	3.8	1.0
Family	babi	3.7	1.0
Family	sister	3.5	1.0
Family	child	3.0	1.0
Family	mother	3.0	1.0

Table 3: **English Gender-specific Likelihood Ratios:** Word stems with the highest gender-specific likelihood ratio in the English Wikipedia that belong to one of the three categories (Family, Relationship and Gender).

Category	Term	Female	Male
Family	embaraz	9.6	1.0
Gender	mrs	6.1	1.0
Gender	femenin	5.3	1.0
Gender	madam	4.4	1.0
Gender	dam	4.4	1.0
Family	tia	4.4	1.0
Relationship	divorci	4.2	1.0
Relationship	bod	4.0	1.0
Gender	mujer	3.9	1.0
Gender	girl	3.9	1.0
Gender	lady	3.7	1.0
Relationship	parej	3.2	1.0
Relationship	enamor	3.0	1.0
Relationship	matrimoni	2.9	1.0
Relationship	marido	2.7	1.0
Relationship	viud	2.7	1.0
Relationship	amant	2.6	1.0
Relationship	hereder	2.5	1.0
Relationship	sexual	2.4	1.0
Family	niet	2.3	1.0

Table 4: **Spanish Gender-specific Likelihood Ratios:** Word stems with the highest gender-specific likelihood ratio in the Spanish Wikipedia that belong to one of the three categories (Family, Relationship and Gender)

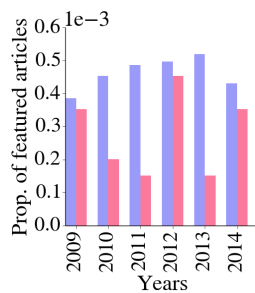


Figure 9: **Visibility Bias:** The proportion of notable men and women that were featured on the front page of the English Wikipedia in the past few years. One can see that the proportion of men is consistently higher, but the difference is marginal.

Wikipedia, evidence of gender bias surfaces from a deeper analysis of the content of those articles. Our results clearly show that subtle lexical and structural gender biases are present on Wikipedia.

Potential explanations for these biases are the following: it is possible that biases are a consequence of (i) the predominantly male editor community and the software design in general that might encourage male contributors and/or (ii) historic and present inequalities between men and women that manifest e.g. in unequal access to resources, unequal media presentation and historic documentation and implicit gender stereotyping (which has been shown to give men an unfair advantage in fame judgements (Banaji et al. 1995)). It seems to be plausible that certain biases such as the coverage or structural bias can be explained by historic inequalities and implicit cognitive biases due to gender stereotypes that may lead to the fact that notable men seem to be more present in our minds than notable women. Other biases such as the lexical bias (e.g. the fact that articles about women disproportionately focus on marriage and divorce compared to articles about men) can more likely be explained by the narrow demographics of the Wikipedia editor community and the media portrayal of men and women. We leave the question of exploring the extent to which different factors explain different biases for future research.

Implications: The low coverage and visibility bias suggest that the Wikipedia community covers notable women and men equally. However, our results highlight that editors need to pay attention to the ways women are portrayed on Wikipedia. In particular, the community needs to evaluate the gender balance of links included in articles (e.g., if an article about a woman links to the article about her husband, the husband should also link back), and to adopt a more gender-balanced vocabulary when writing articles about notable people. These existing biases might put women at a practical disadvantage: For example, because modern search and recommendation algorithms exploit both structural and textual information, women might suffer from lower visibility when it comes to ranking articles about notable people or in terms of their general visibility on Wikipedia (at least if we only take links between articles about people into

account; see Figure 6 in (Eom et al. 2014) for preliminary comparison of ranking algorithms).

Cross-lingual Analysis: We observe the strongest structural bias for the English and Russian Wikipedia. Also on the lexical dimension the strongest bias becomes visible in the English and Russian Wikipedia. Surprisingly the Spanish Wikipedia reveals the lowest structural bias. Comparing our results with the Gender Inequality Index of the World Economic Forum (WMF) (Schwab et al. 2013) shows that a positive correlation exists between the bias in the offline world and the bias on Wikipedia. However, one needs to note that it is difficult to compare our Wikipedia based gender bias rankings of languages with the ranking of countries according to the gender inequality index since countries where the same language is predominantly spoken often reveal very different positions in the WMF ranking. We use the weighted average of the WMF rank positions of countries where the same language is spoken³ and weight countries by the size of the internet population⁴. The Spearman rank correlation between the ranking of the 6 languages according to the WMF index shows a correlation of 0.89 with the coverage bias based ranking, 0.37 with the structural bias ranking and 0.09 with the lexical bias ranking. This indicates that to a certain extent gender inequalities of the real world manifest on Wikipedia. However, since the Wikipedia editor community is not representative for the larger population in a country, it is also not surprising that certain biases like the lexical bias do only reveal a very limited relation with the WMF ranking. Although Wikipedia may only reflect certain aspects of gender inequalities of the real world, gender biases that are introduced by the editor community of Wikipedia may effect the larger population and therefore it is important to investigate them.

Reference datasets: Our findings with regard to coverage bias are effected by the (unknown) biases inherent in the reference datasets used. Due to this, we can not make any absolute statements about coverage inequality on Wikipedia. However, regardless of this problem, we can assert that *Wikipedia covers women and men from our reference datasets better equally well*. Using external reference datasets that represent collections of notable people to prune down the number of biographies in Wikipedia rather than studying all of them further helps to uncouple lexical bias and structural bias from coverage bias and ensures that only people that are notable from a global perspective become the subject of study. An alternative would be to select all people from Wikipedia using category pages such as “Births by Year”⁵ or “Deaths by Year”⁶ as starting point. However, these category pages do not exist in all language editions and therefore the selection would be based on the categories of the English Wikipedia only, which introduces a bias since every language editions tends to focus on their “local heros” (Callahan and Herring 2011;

³<http://www.infoplease.com/ipa/A0855611.html>

⁴http://en.wikipedia.org/wiki/List_of_countries_by_number_of_Internet_users

⁵http://en.wikipedia.org/wiki/Category:Births_by_year

⁶http://en.wikipedia.org/wiki/Category:Deaths_by_year

Related Work

Gender Inequalities in Traditional Media: Feminist often claim that news is not simply mostly about men, but overwhelmingly seen through the eyes of men. In (Ross and Carter 2011) the authors analyze longitudinal data from the GMMP (Global Media Monitoring Project) which spans over 15 years. The authors conclude that the role of women as a producer and subject of news has seen a steady improvement, but the relative visibility of women compared to men has stuck at 1:3 which means that the world's new agencies still consider the life of men three time more worth to write about it as those of women. Gender inequalities also manifest in films that are used for education purposes, as revealed by the application of the Bechdel test to teaching content (Scheiner-Fisher and Russell 2012). In (Sugimoto et al. 2013) the authors present a cross-disciplinary, global, bibliometric analysis of the relation between gender and scientific output (i.e., number of papers, citations per paper and internationality of collaborations) using data from more than 5 million scientific publications. They find that the research output in most countries is dominated by males and that the few countries that are dominated by females have lower research output which indicates that barriers are present.

Gender Inequalities on Wikipedia: Our work is not the first work which recognises the importance of understanding gender biases on Wikipedia (Reagle and Rhue 2011; Eom et al. 2014; Callahan and Herring 2011; Aragón et al. 2012). In (Reagle and Rhue 2011) thousands of biographical subjects from six reference sources (e.g., The Atlantic's 100 most influential figures in American history, TIME Magazine's list of 2008's most influential people) are compared against the English-language Wikipedia and the online Encyclopedia Britannica with respect to coverage and article length. The authors do not find gender-specific differences in the coverage and article length on Wikipedia, but Wikipedia's missing articles are disproportionately female relative to those of Britannica. Our findings on the coverage dimension confirm their findings and further we also analyze the content of articles on Wikipedia which they left for future work.

In (Bamman and Smith 2014) the authors present a method to learn biographical structures from text and observe that in the English Wikipedia biographies of women disproportionately focus on marriage and divorce compared to those of men, which is in line with our findings on the lexical dimension. Recent research showed that most important historical figures across Wikipedia language editions are born in Western countries after the 17th century, and are male (Eom et al. 2014). On average only 5.2 female historic figures are observed among the top 100 persons. The authors use different link-based ranking algorithms and focus on the top 100 figures in each language edition. Their results clearly show that very few women are among the top 100 figures in all language editions, but since the authors do not use any external reference lists it remains unclear how many women we would expect to see among the top 100

figures.

Previous research has also explored gender inequalities in the editor community of Wikipedia and potential reasons for it (cf. (Lam et al. 2011; Collier and Bear 2012; Hill and Shaw 2013)). Also among Wikipedians, the importance of this issue has been acknowledged for example through the initiation of the "Countering Systemic Bias" WikiProject⁷ in 2004.

Gender inequalities in Social Media: In (Szell and Thurner 2013) the author study a communication network in a MMOG and find a similar effect as (Smoreda and Licoppe 2000). Female players send about 25% more messages (0.74 per day) than males (0.60 per day). Consequently, females show a significantly higher average degree in their communication networks, however, the communication partners of females have a significantly lower average degree than those of males, i.e. females have more communication partners, while males tend to have better connected ones. Recent research (Magno and Weber 2014) suggests that in Twitter and Google+ online inequality is strongly correlated to offline inequality, but the directionality can be counter-intuitive. In particular, they consistently observe women to have a higher online status, as defined by a variety of measures, compared to men in countries such as Pakistan or Egypt, which have one of the highest measured gender inequalities. In (Garcia, Weber, and Garimella 2014) the authors show that subconscious biases which contribute to the creation of inequality are not only present in movie scripts but also in Twitter conversations. Also the viewing and sharing patterns of youtube videos reveal differences in which content is consumed and discussed by different genders (Abisheva et al. 2014). This kind of differences also manifest in wall discussions in MySpace, where emotional expression patterns differ across genders (Thelwall, Wilkinson, and Uppal 2009).

Conclusions

Wikipedia seems to have successfully established processes that ensure that notable women have a high likelihood of being portrayed on Wikipedia. At the same time, our work surfaces evidence of more subtle forms of gender inequality. In particular, women on Wikipedia tend to be more linked to men than vice versa, which can put women at a disadvantage in terms of - for example - visibility or reachability on Wikipedia. In addition, we find that women's romantic relationships and family-related issues are much more frequently discussed in their Wikipedia articles than in men's articles. This suggests that there are gender differences w.r.t. how the Wikipedia community conceptualizes notable men/women. Because modern search and recommendation algorithms exploit both, structure and content, women may suffer from lower visibility in social networks (or article networks) where men (or articles about men) are more central and include more links to other men than to other women. To reduce such effects, the editor community needs to evaluate the gender balance of links included in

⁷http://en.wikipedia.org/wiki/Wikipedia:WikiProject_Countering_systemic_bias

articles (e.g., if an article about a woman links to the article about her husband, the husband should also link back), and to adopt a more gender-balanced vocabulary when writing articles about notable people. Further, engineers and researchers need to develop a deeper understanding of how different types of search and recommendation algorithms impact the visibility of minorities.

In summary, the contributions of this work are twofold: (i) we present a computational method for assessing gender bias on Wikipedia *along multiple dimensions* and (ii) we apply this method to several language editions of Wikipedia and share empirical insights on observed gender inequalities. We translate our findings into some potential actions for the Wikipedia editor community to reduce gender biases in the future. We hope our work contributes to increasing awareness about gender biases online, and in particular to raising attention to the different levels in which these biases can manifest themselves. The methods presented in this work can be used to assess, monitor and evaluate these issues on Wikipedia on an ongoing basis.

References

- Abisheva, A.; Garimella, V. R. K.; Garcia, D.; and Weber, I. 2014. Who Watches (and Shares) What on YouTube ? And When ? Using Twitter to Understand YouTube Viewership. In *WSDM'14*.
- Amy Yu, Kevin Hu, S. R. D. G., and Hidalgo, C. A. 2013. The Pantheon Multilingual Wikipedia Expression Dataset. MIT project.
- Aragón, P.; Kaltenbrunner, A.; Laniado, D.; and Volkovich, Y. 2012. Biographical social networks on wikipedia - A cross-cultural study of links that made history. *Proceedings of WikiSym*.
- Bamman, D., and Smith, N. 2014. Unsupervised discovery of biographical structure from text. *Transactions of the Association for Computational Linguistics* 2:363–376.
- Banaji, M. R.; Greenwald, A. G.; Kelley, C. M.; Brown, J.; Jasechko, J.; and Experiments, I. 1995. Implicit gender stereotyping in judgments of fame. *Journal of Personality and Social Psychology* 68:181–198.
- Callahan, E. S., and Herring, S. C. 2011. Cultural bias in wikipedia content on famous persons. *Journal of the American Society for Information Science and Technology* 62(10):1899–1915.
- Collier, B., and Bear, J. 2012. Conflict, criticism, or confidence: An empirical examination of the gender gap in wikipedia contributions. In *Proceedings of the ACM 2012 Conference on Computer Supported Cooperative Work, CSCW '12*, 383–392. New York, NY, USA: ACM.
- Eom, Y.-H.; Aragón, P.; Laniado, D.; Kaltenbrunner, A.; Gigna, S.; and Shepelyansky, D. L. 2014. Interactions of culture and top people of wikipedia from ranking 24 language editions. *Plos One*.
- Finkbeiner, A. 2013. What I'm not going to do: Do media have to talk about family matters? DoubleXScience.
- Fox, M. F.; Johnson, D. G.; and Rosser, S. V. 2006. *Women, gender, and technology*. University of Illinois Press.
- Garcia, D.; Weber, I.; and Garimella, V. R. K. 2014. Gender asymmetries in reality and fiction: The bechdel test of social media. In *International AAAI Conference on Weblogs and Social Media (ICWSM)*.
- Giatsidis, C.; Thilikos, D. M.; and Vazirgiannis, M. 2013. Dcores: measuring collaboration of directed graphs based on degeneracy. *Knowledge and information systems* 35(2):311–343.
- Hecht, B., and Gergle, D. 2010. The tower of babel meets web 2.0: user-generated content and its applications in a multilingual context. In *Conference on Human Factors in Computing Systems*, 291–300. ACM.
- Hill, B. M., and Shaw, A. 2013. The wikipedia gender gap revisited: Characterizing survey response bias with propensity score estimation. *PLoS One* 6(9).
- Lam, S. K.; Uduwage, A.; Dong, Z.; Sen, S.; Musicant, D. R.; Terveen, L. G.; and Riedl, J. 2011. Wp: clubhouse?: an exploration of wikipedia's gender imbalance. In Ortega, F., and Forte, A., eds., *Int. Sym. Wikis*, 1–10. ACM.
- Magno, G., and Weber, I. 2014. International gender differences and gaps in online social networks. In *The 6th International Conference on Social Informatics (SocInfo)*.
- Murray, C. 2003. *Human Accomplishment. The pursuit of excellence in the Arts and Sciences*.
- Newman, M. E. J. 2003. Mixing patterns in networks. *Phys. Rev. E* 67(2):026126.
- Pollitt, K. 1991. Hers; The Smurfette Principle. The New York Times.
- Reagle, J., and Rhue, L. 2011. Gender bias in Wikipedia and Britannica. *International Journal of Communication* 5.
- Ross, K., and Carter, C. 2011. Women and news: A long and winding road. *Media, Culture and Society* 33(8):1148–1165.
- Scheiner-Fisher, C., and Russell, W. B. 2012. Using Historical Films to Promote Gender Equity in the History Curriculum. *The Social Studies* 103(6):221–225.
- Schich, M.; Song, C.; Ahn, Y.-Y.; Mirsky, A.; Martino, M.; Barabási, A.-L.; and Helbing, D. 2014. A network framework of cultural history. *Science* 345(6196):558–562.
- Schwab, K.; Brende, B.; Zahidi, S.; Bekhouche, Y.; Guinault, A.; Soo, A.; Hausmann, R.; and Tyson, L. D. 2013. The Global Gender Gap Report 2013. World Economic Forum.
- Schwartz, H. A.; Eichstaedt, J. C.; Kern, M. L.; Dziurzynski, L.; Ramones, S. M.; Agrawal, M.; Shah, A.; Kosinski, M.; Stillwell, D.; Seligman, M. E.; and Ungar, L. H. 2013. Personality, gender, and age in the language of social media: the open-vocabulary approach. *PLoS One* 8(9).
- Smoreda, Z., and Licoppe, C. 2000. Gender-Specific Use of the Domestic Telephone. *Social Psychology Quarterly* 63(3):238–252.
- Sugimoto, Cassidy R. and Larivière, V.; Ni, C.; Gingras, Y.; and Cronin, B. 2013. Bibliometrics: Global gender disparities in science. *Nature* 504(7479):211–213.
- Szell, M., and Thurner, S. 2013. How women organize social networks different from men. *Scientific Reports* 3(1214).
- Thelwall, M.; Wilkinson, D.; and Uppal, S. 2009. Data mining emotion in social network communication: Gender differences in MySpace. *Journal of the American Society for Information Science and Technology* 61(1):1–10.
- Thomas, G. 1992. *A position to command respect : women and the eleventh Britannica*. Scarecrow Press.