

Gender Bias Detection Using Facebook Reactions

Lillian Gray

Department of Computer Science

Earlham College

Richmond, Indiana

lmgray16@earlham.edu

ABSTRACT

Using data to quantify the amount of gender bias on certain websites is important to reveal discrimination on online platforms. Online platforms have to be aware of the presence of bias for them to change aspects of their platform to mitigate it. Facebook Reactions were released in 2016 and are an unexplored source to study gender bias. Related work either measures gender bias elsewhere online or analyzes Facebook Reactions for various applications. Gender bias on Facebook might be measured by analyzing the difference in reactions on posts by women or men. This project is studying bias in the Facebook Pages of United States 2020 Senate candidates using Facebook Reactions. Data was collected from Facebook Pages of politicians using a crawler and inserted into a database. The data was analyzed using an entropy function on the reactions for each post and the preliminary results were visualized.

KEYWORDS

gender bias, facebook reactions, social media, web crawling, data processing

1 INTRODUCTION

Measuring and quantifying the amount of gender bias on certain websites is important to highlight the presence of discrimination on online platforms. In order for online platforms to change features or rules of their platform to mitigate gender bias, they have to be aware of its existence. For example, the existence of gender bias across Wikipedia, which is generally viewed as an unbiased encyclopedia, can have an effect on the many users who read it [24]. Existing gaps in this topic include sites that have not been analyzed for gender bias, and quantifying the change in gender bias on a site over time. Facebook, due to its privacy restrictions, is a site that has not been studied as much for gender bias. Facebook Reactions, which were released in 2016, are an unexplored and unique source to study gender bias. By analyzing the difference in reactions on posts by or about women and men, gender bias on Facebook can be measured.

The proposed project would specifically study bias on the Facebook pages of 2020 candidates for the United States Senate. Politicians have a particularly large public presence compared to other groups of people. Furthermore, their social media presence and reception could affect their perceived eligibility for office. Mertens et al. explore gender bias in tweets to and from German politicians [18]. This project builds on that work by studying if there is similar gender bias on Facebook towards United States politicians.

Using a Facebook crawler fbcrawl built on the Scrapy framework, data was collected from Facebook pages and processed into a database. The data consists of Facebook Page data for each candidate,

and post data for each page. The data was analyzed by calculating the entropy of the reactions for each post. The data analysis of different pages was then compared to see if there were significant differences based on the gender of the politician. Then the preliminary results were visualized.

This paper will first cover related works on online gender bias and Facebook Reactions. Then it will cover the design and implementation of the project, including data collection, data storage, and analysis. Next, preliminary results will be covered. Finally, possibilities for future work will be discussed.

2 RELATED WORK

While there are no current studies that have investigated Facebook Reactions for gender bias, there are both studies that look at gender bias online and studies that research Facebook Reactions for various purposes. This section first discusses background on gender bias against political candidates. Then, studies about gender bias on websites and studies on Facebook Reactions are discussed.

2.1 Gender Bias Facing Political Candidates

There exists a variety of research that explores the effect of gender bias on elections and campaigns. Some research explores the direct effect gender bias has on voting, while other research explores how gender affects people's impression of a candidate, or the further information they search for on a candidate. Some research uses experimental data, so that data on candidates can be easily manipulated only by gender, and other research uses data from real-life candidates or votes. Current research has varying conclusions about the extent of the effect of gender bias on female candidates.

Aalberg and Jenssen studied the impressions of people after watching one of two speeches from a politician [1]. The politicians are actors, and the speeches are identical, except that one actor is male and one is female. They found that the man was found to be more knowledgeable, trustworthy, and convincing than the woman, mainly due to the scores from male participants. The study was performed in Norway, a country which, prior to the study, has done work to reach close to equal gender representation in government.

However, others argue that baseline preference in an experimental setting does not imply an effect on vote choice in the real world. Ditonto et al. argued that gender does not directly affect vote choice [6]. Instead, a candidate's gender guides the further information that voters search for over the process of a campaign. Using an environment to simulate campaigns, they found that voters would change the types of information they searched for based on the gender of the candidate. For female candidates, voters would typically search for information related to their competence and how compassionate they are. They argue that this information search

mediates the relationship between stereotypes and final vote decisions, and so the information that voters find when they search about a candidate matters greatly.

Others explore the effect of stereotyping, and how it can in turn affect voting behavior. Sanbonmatsu explored the effect of gender stereotypes on voter behavior [20]. They argued that voters who have a preference for women or men are affected by stereotypes about candidate traits, beliefs, and policy strengths and that these preferences can have an effect on voting behavior. Additionally, Bauer examined automatic stereotype activation of voters from candidates for office [5]. They found that campaign communication causes automatic stereotype activation even when they are not directly proving or disproving a stereotype. This has the effect of diminishing support for female candidates. They argued that this stereotype activation could lead voters to make further stereotype-based inferences that a female candidate lacks the right traits for a leadership role.

Many studies used surveys to examine people's stereotypes of male and female politicians, and how that affects their voting. Dolan tested the ways gender stereotypes influence people's willingness to support female candidates [7]. They performed a survey examining stereotypes about male and female politicians. Their results indicated that people hold both policy and trait stereotypes about women and men. Later, Dolan and Lynch examined attitudes about women and men in politics in conjunction with vote choices in United States House elections [8]. They find in line with previous work that gender stereotypes affect people's evaluations of candidates. However, they find little evidence of stereotypes having a direct impact on the vote.

In another study using real election data as a source, Koch studied evaluations of United States Senate candidates from 1988 to 1992 [13]. They find that female candidates were evaluated more highly for their perceived ability to handle social issues, but were evaluated lower in terms of their competence. They also found that gendered evaluations were more likely from people who were highly educated. They concluded that these appraisals may have resulted from a gendered pattern of campaign messages and media coverage, and not necessarily the candidates themselves.

Finally, Anstead and Chadwick explore the relationship between technology and political institutions in the age of Web 2.0 [2]. They argued that there is a dialectical relationship between the two. Technologies have the power to reshape political institutions, but those institutions mediate the eventual results of that shaping. They discussed the large role of the Internet in the 2000s in shaping the American campaign environment. Taken today, the even larger roles Facebook, Twitter, and YouTube play in people's social lives allow the sites to play a large role in many campaigns. They also noted that most campaigning innovations in the 2000s on the Internet took place during primary campaigns. However, they found that the Internet seemed to not be as effective in reaching undecided voters.

2.2 Detecting Gender Bias on Websites

There exist a good number of articles that quantified and measured gender bias using computational methods on a website or over two similar websites. A popular site in the existing literature to analyze

for gender bias is Wikipedia. As an encyclopedia site that is edited by the public, the gender bias existing on Wikipedia is of interest, since encyclopedias should be free of bias. Graells-Garrido et al. analyzed meta-data, language, and network structure in biographies on Wikipedia because since they are each about a specific person, they can easily be grouped and compared to study gender bias [10]. Zagovora et al. analyzed German Wikipedia articles about professions to see what combination of male and/or female titles of the profession it used [26].

Since gender bias in professional interactions can affect someone's career and earnings, other researches focused on sites related to careers. Tang et al. used 10 years of job posts on LinkedIn to analyze gender bias in job postings over time [22]. On the other hand, Hannák et al. focused on two freelance sites and the gender and racial bias of users hiring workers by performing linguistic analysis on reviews and analyzing if ratings and reviews are connected to certain variables [11]. Finally, Imtiaz et al. explored interactions among software engineers and their work by analyzing a set of GitHub users [12].

Other researchers studied gender bias on social media sites. Mertens et al. used the set of tweets collected during the German federal elections in 2017 and used topic modeling and dictionary analysis to study gender bias in interactions with German politicians [18]. Their results show that for every party (excluding the Alternative for Germany party which looks to have no male data) female politicians received a higher ratio of personal to job related words than male politicians. Magno and Weber used data of users from Twitter and Google+ around the world to compare activity and followers between male and female users for each country with that of The Gender Gap Report by the World Economic Forum [16].

Since gender bias is a social issue, researchers needed to find a way to quantify this bias within their datasets. Some authors used existing gender studies research to find ways to quantify bias. For example, Mertens et al. justify using gender studies research that female politicians might receive more personal tweets than professional ones compared to male politicians [18]. Imtiaz et al. used the four effects of gender bias described by William and Dempsey to frame their research [12, 25]. They found measurements to take in their GitHub data that related to each effect and created hypotheses under the assumption that effects of gender bias would be visible on GitHub.

2.3 Facebook Reactions

Despite Facebook Reactions being first released in 2016, there exist several papers already that study them. While no papers found currently have studied gender bias, the methods used by some authors could be adapted to study bias.

Matamoros-Fernandez studied how groups of Facebook users used Facebook Reactions for racist means [17]. While the study is more qualitative than quantitative, it demonstrates how Facebook Reactions can have a negative impact on minority groups.

Authors performed a variety of analyses on the reactions data. Basile et al. used entropy functions on reactions to measure the controversy of posts on different pages of news sites [4]. Their code is published, and can be used to see if posts by or about women are more "controversial" than posts by or about men. Freeman et al.

measured the Pearson r correlation coefficients for all reaction pairs in their data [9]. Sandoval-Almazan and Valle-Cruz grouped negative and positive Facebook Reactions and measured the sentiment index of a post based on number of positive and negative reactions [21]. Tian et al. applied a K-means clustering algorithm to these Reaction proportions for all posts to investigate which Reactions were most likely to be seen together on a post [23].

Some authors also performed analyses on the comments of posts. Krebs et al. used an emotion lexicon to perform emotion mining on comments [14]. Kuo et al. collected comments by users who reacted on the same post and found lexical patterns in those comments for each reaction [15]. This could be interesting to see if the patterns for a reaction are different based on gender, meaning the emotional usage of the reaction is generally different when reacting to different genders.

Multiple authors stated they used the Facebook Graph API for data collection [3, 9, 14, 19]. However, the Graph API cannot be employed for this project due to changes Facebook made to the API to protect data privacy after the Cambridge Analytica scandal in 2018.¹ This project used a crawler to collect data, but that data will be handled as privately as possible and only consolidated and anonymized data will be published in the Results.

3 DESIGN AND IMPLEMENTATION

The design for the project is shown in Figure 1.

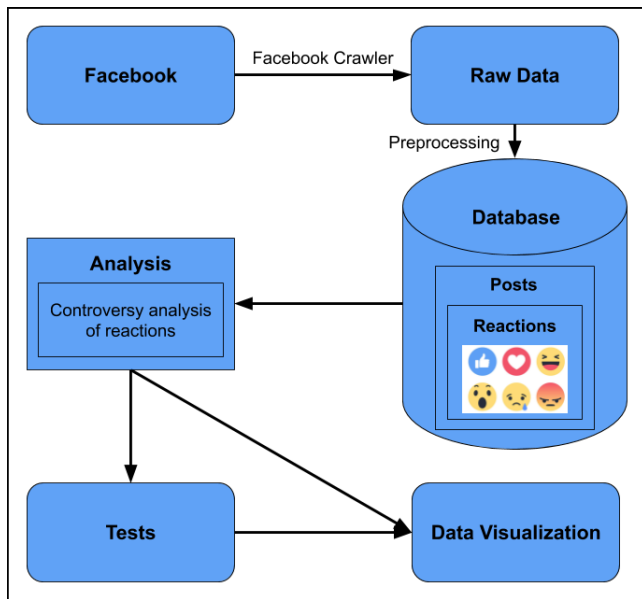


Figure 1: Project Framework

Data was collected from pages using a web crawler. The data was then be processed into a database. The data included basic information on each candidate and their Facebook Page, and data for each post on each page, including the numbers for each reaction. The data was analyzed by calculating the entropy for each post and updating the post data with each post’s entropy in the database.

¹<https://about.fb.com/news/2018/04/restricting-data-access/>

Finally, the average entropy was calculated for all posts, and posts by women or men, with different constraints. While the preliminary results were visualized, there was not time to run extensive tests on the results.

3.1 Data Collection

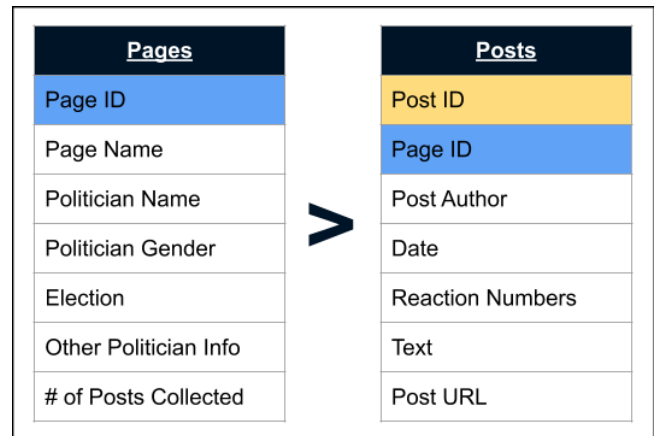


Figure 2: Database Schema

Data was collected from pages using an open source Facebook crawler built on the Scrapy framework.² The program crawls a Facebook page given a page name, and dumps the collected posts into a CSV file. The post data for each page includes the link to the post. The crawler, given a post link, can also collect data for the comments on that post. The crawler also requires a Facebook login to make the request from.

The crawler initially had issues running on the Earlham Computer Science servers, with Facebook blocking requests. After creating a new Facebook account for the sole purpose of this project, and setting up the crawler on the Windows Subsystem for Linux (WSL), the program works a majority of the time. This is reasonably because the request through WSL had the same IP address as when the new account was created.

The page names, along with other information on each candidate running for US Senate in 2020, were collected manually into a spreadsheet. With the wide variability in page names and popularity of candidates, having people fill out the spreadsheet was a simpler solution than trying to build or find a program to search for candidate information. Recruitment for help in filling out the spreadsheet occurred to cut down on the time to collect candidate data.

Many members of Congress have both a Facebook Page for their candidacy and a page for interacting with their constituents as a government official, since this work is required to be separate by law. Because this project is focusing on candidates for US Senate, only the candidate pages are being collected for incumbent candidates.

A few candidates have opted to use their Facebook profile, rather than a Facebook page, for their candidacy. While the structure of posts, reactions, and comments on profiles is the same as a page

²<https://github.com/rugantio/fbcrawl>

from a user perspective, the crawler being used is unable to collect data from profiles. Therefore, since few cases of profiles have been found, they were ignored.

3.2 Data Conversion and Database

The data fetched using the crawler was stored in a Postgres database on the cluster. The generated CSV files were copied into the database to both help manage the data, and allow for simpler comparisons of subsets of the data. The database includes a table generated from the spreadsheet of candidate data, and a table of all crawled posts from pages.

The design for the database is shown in Figure 2. The Page IDs collected manually were also stored in each row in the Posts table for the page a post is on in order to create an index. Because the Page ID is not stored in the collected posts by the crawler, the program that runs the crawler had to keep track of that information until it was entered into the database. To do so, each csv of collected posts from a page was named after that page's Page ID.

Preprocessing on the CSV files for both pages and posts had to be performed so that the converted files could be easily copied into the database. Two Python scripts were written, one to convert a CSV of candidate data, and one to convert a CSV of posts from a page. The CSV of candidate data was converted to only include candidates that had a Facebook Page. This was so that an index for Page ID could be created on the posts and pages tables. Each CSV file of posts from a page was converted to filter out identical rows generated by the crawler. Numbers that contained characters were also converted so that they could be inserted into the database as integers (eg. 15K was converted to 15000). During the conversion, the Page ID for the page a post was from was also added to each row for the index to be created. The two Python conversion scripts were run using a bash script, using a for loop for the post files.

Two bash scripts were written to make the database tables, one for the posts table and one for the pages table. For the posts table, the converted CSV files were copied in a for loop, and then an entropy column was added for later updates. For the pages table, the converted CSV file was copied, the index created on Page ID, and the column for number of posts collected was updated with the count from the posts table for each Page ID.

3.3 Analysis

Preliminary analysis was performed on the reactions for each post. These analyses were compared between posts by female politicians and those by male politicians.

For each post collected, an entropy function was performed on the reactions [4]. The calculations were performed in a Python script using `psycpg2` library to retrieve the data from the database, and the entropy function from the SciPy stats library.³ The script also updated the database for each post with the entropy value for that post's reactions.

Entropy can be used as a measure of controversy, since the more variety of reactions to a post there are, the higher the entropy will be [4]. For example, the entropy of a post with 10 likes and no other reactions equals 0, while the entropy for a post with 5 likes, 3 ahah, 1 wow, and 1 grr equals 1.69. The entropy equation as it

is calculated by the SciPy stats library is shown below, where x_i is the number of each reaction for a post, and $p(x_i)$ is the ratio of that reaction to the total reactions on a post.

$$H(X) = - \sum_{i=1}^n p(x_i) \cdot \log(p(x_i)) \quad (1)$$

Finally, two bash scripts were written to generate preliminary results. The first script averaged the entropy over all posts in total, and for male and female candidates. The averages were calculated with certain constraints to see how much the output was affected. The first constraint was date. Because of the timeline for each campaign was different, and the crawler was run at different dates for different pages, the date was set between September 2019 and April 2020, which was found to contain the majority of posts. Another constraint used was to not use posts with under twenty reactions total. This is because posts with very few reactions were thought to be more likely to have outlying entropy values [4].

If the average entropy of female politician's pages is much higher than that of male politician's pages, then entropy of Facebook Reactions will be shown to be a valid measure of bias. Figure 3 shows the output from the various averages in the Results section. The second script generates the average entropy for each page, using the same constraints as noted above. The output of that script is included with the software for this project.⁴

4 RESULTS

The preliminary results for average entropy over posts are shown in Figure 3. The figure shows that entropy has no clear difference across gender. While performing statistical analysis over different elections, or controlling for other variables may show otherwise, these results show that entropy is probably not a good indicator of gender bias.

However, initial analysis of average entropy over individual pages may continue to support entropy of reactions as a good measurement of controversy. Among the male politicians with the highest average entropy are Mitch McConnell, Lindsey Graham, and Cory Booker. These politicians are all high profile and therefore have been under a fair amount of public scrutiny. Without controlling for date or minimum reactions, the female politician with the highest average entropy is Susan Collins, a candidate currently in a tight election and under much scrutiny. At first glance, the majority of pages with the highest entropy are from male politicians. It is unclear whether this is due to their gender, their behavior, and/or having more years in office.

5 FUTURE WORK

There are many possibilities for future work in this project. The data collected from this project and comments data that could be collected has many possibilities for further analysis, results, and testing. Some possibilities for future work within this project are outlined below.

³<https://github.com/scipy/scipy>

⁴https://gitlab.cluster.earlham.edu/senior-capstones-2020/lmgray16-senior-capstone/blob/master/storage_machine/generateResultsByPage_out.txt

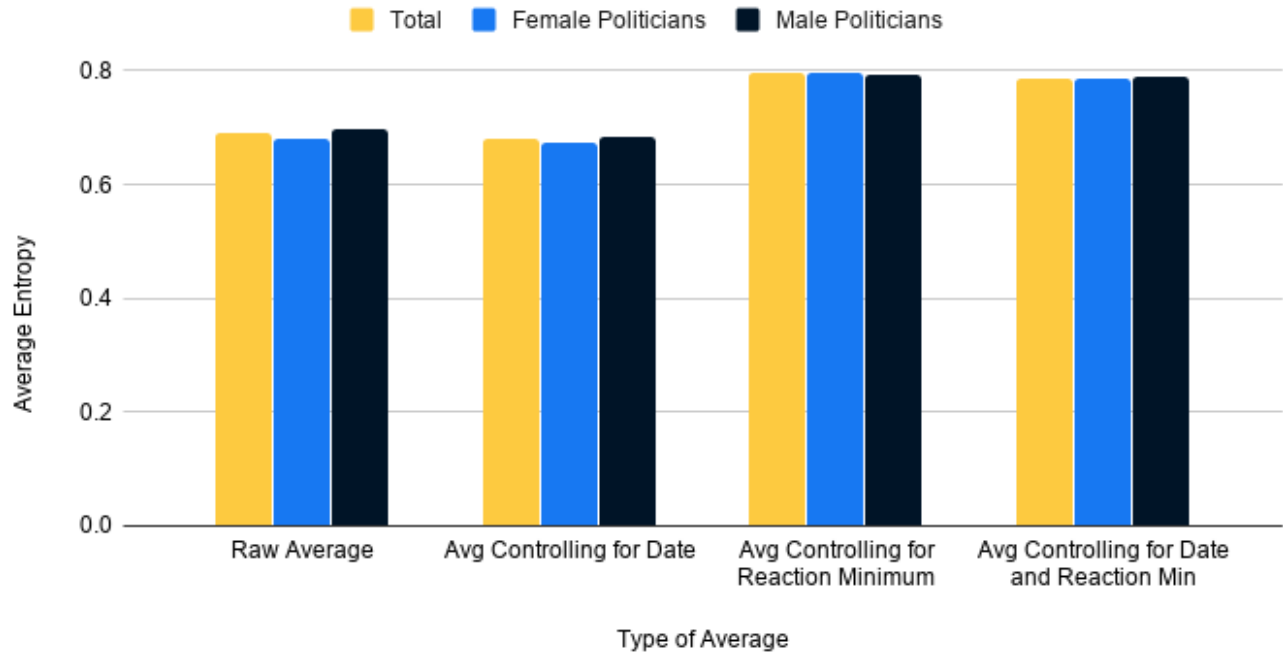


Figure 3: Average Entropy Across Posts

5.1 Comments Analysis

One large possibility for future work within this project would be inspecting the comments on collected posts. Using the already collected post urls, the crawler can collect the comments on a post. The comments on each post could be analyzed for gender bias using Linguistic Inquiry and Word Count (LIWC) dictionaries. Due to gender stereotypes, it is feasible that female politicians receive more personal comments rather than professional ones [18]. Therefore, the log of the ratio of personal vs. job related words could be calculated for each comment. Personal words would be calculated from the LIWC family and friends dictionaries, and job related words from the LIWC work dictionary. The equation for personal vs. job related words is shown below.

$$personal\ vs.\ job = \log\left(\frac{personal + .5}{job + .5}\right) \quad (2)$$

5.2 Testing

Since there are no closely related works to this project, multiple tests could be run to see if the analysis of reactions and/or comments can measure gender bias.

First, a sanity check should be performed on the results. The entropy results could be compared to the work of Basile et al. to see if they are all reasonable values [4]. Similarly, the comment analysis could be compared to the results of Mertens et al. to see if the values are reasonable. This would also see if the gender bias against German politicians on Twitter is any similar to the gender bias against United States politicians on Facebook [18].

Next, the average entropy and personal vs. job related words in comments could be calculated over the collected posts for each page. If outlier posts are found with very high entropy, they may not be calculated in the average if they were inspected and seem to be a controversial news topic rather than an average post by the page. Then it would be checked whether the difference in these averages for male and female politicians is statistically significant. The averages could also be inspected over other categories such as party, age, race, or political status. Party, election, years in office for incumbents, and whether the candidate is challenged in the primary and/or general election is already included in the pages table. If the measurements have a clearer difference over any of these categories rather than gender, then they might not be as good of an indicator of gender bias than they are of something else.

If both the reactions and comment analysis are good indicators of bias, then the measurements should correlate in some way post to post, and/or page to page. Therefore a graph with a point for each post could be plotted with one axis being the entropy and the other being the personal vs. job ratio in comments. Then measurements such as the Pearson correlation coefficient could be calculated on the data to see if entropy correlates with personal vs. job related language.

Finally, another test that could be performed is measuring if the reactions and/or comment analysis correlates with existing analysis of bias against U.S. politicians. As described in Related Work, Magno and Weber test if their measurements of bias are valid by comparing their results to The Global Gender Gap Report [16]. If existing work from other disciplines can be found rating

bias against U.S. politicians, this could be compared to the analysis to test if it is measuring gender bias well.

6 CONCLUSION

Using software and big data to research gender bias online is also still an emerging field with many areas to explore. Additionally, because Facebook Reactions were first released in 2016, research on them and their effect is still limited. The methods used in this project could be used for a variety of research on Facebook and for other categories of pages. While social media data has the potential for privacy issues beyond the scope of this paper, it is clear that it can be used to uncover large trends and learn more about how people communicate and interact with one another.

7 ACKNOWLEDGEMENTS

Thank you to Dr.Charles Peck for advising me on my project and my initial proposal ideas and for leading the capstone course. Thank you to Dr.Xunfei Jiang for her advice on my proposal and for leading the proposal course. Thank you to everyone in the Department of Computer Science at Earlham for teaching me the skills to implement this project. Thank you to Sari Stissi, Laurence Ruberl, and Porter Libby for their support and feedback throughout the year.

REFERENCES

- [1] Toril Aalberg and Anders Todal Jenssen. 2007. Gender Stereotyping of Political Candidates. *Nordicom Review* 28, 1 (May 2007), 17–32. <https://doi.org/10.1515/nor-2017-0198>
- [2] Nick Anstead and Andrew Chadwick. 2008. Parties, Election Campaigning, and the Internet: Toward a Comparative Institutional Approach. (2008), 21.
- [3] Ismail Badache and Mohand Boughanem. 2017. Emotional Social Signals for Search Ranking. In *Proceedings of the 40th International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR '17)*. ACM, 1053–1056. <https://doi.org/10.1145/3077136.3080718> event-place: Shinjuku, Tokyo, Japan.
- [4] Angelo Basile, Tommaso Caselli, and Malvina Nissim. 2017. *Predicting Controversial News Using Facebook Reactions*. Accademia University Press, 12–17. <https://doi.org/10.4000/books.aaccademia.2370>
- [5] Nichole M. Bauer. 2015. Emotional, Sensitive, and Unfit for Office? Gender Stereotype Activation and Support Female Candidates. *Political Psychology* 36, 6 (2015), 691–708.
- [6] Tessa M. Ditonto, Allison J. Hamilton, and David P. Redlawsk. 2014. Gender Stereotypes, Information Search, and Voting Behavior in Political Campaigns. *Political Behavior* 36, 2 (2014), 335–358.
- [7] Kathleen Dolan. 2010. The Impact of Gender Stereotyped Evaluations on Support for Women Candidates. *Political Behavior* 32, 1 (2010), 69–88.
- [8] Kathleen Dolan and Timothy Lynch. 2014. It Takes a Survey: Understanding Gender Stereotypes, Abstract Attitudes, and Voting for Women Candidates. *American Politics Research* 42, 4 (Jul 2014), 656–676. <https://doi.org/10.1177/1532673X13503034>
- [9] Cole Freeman, Mrinal Kanti Roy, Michele Fattoruso, and Hamed Alhoori. 2019. Shared Feelings: Understanding Facebook Reactions to Scholarly Articles. *arXiv:1905.10975 [cs, stat]* (May 2019). <http://arxiv.org/abs/1905.10975> arXiv: 1905.10975.
- [10] Eduardo Graells-Garrido, Mounia Lalmas, and Filippo Menczer. 2015. First Women, Second Sex: Gender Bias in Wikipedia. In *Proceedings of the 26th ACM Conference on Hypertext Social Media (HT '15)*. ACM, 165–174. <https://doi.org/10.1145/2700171.2791036> event-place: Guzelyurt, Northern Cyprus.
- [11] Anikó Hannák, Claudia Wagner, David Garcia, Alan Mislove, Markus Strohmaier, and Christo Wilson. 2017. Bias in Online Freelance Marketplaces: Evidence from TaskRabbit and Fiverr. In *Proceedings of the 2017 ACM Conference on Computer Supported Cooperative Work and Social Computing (CSCW '17)*. ACM, 1914–1933. <https://doi.org/10.1145/2998181.2998327> event-place: Portland, Oregon, USA.
- [12] Nasif Imtiaz, Justin Middleton, Joymallya Chakraborty, Neill Robson, Gina Bai, and Emerson Murphy-Hill. 2019. Investigating the Effects of Gender Bias on GitHub. In *Proceedings of the 41st International Conference on Software Engineering (ICSE '19)*. IEEE Press, 700–711. <https://doi.org/10.1109/ICSE.2019.00079> event-place: Montreal, Quebec, Canada.
- [13] Jeffrey W. Koch. 1999. Candidate Gender and Assessments of Senate Candidates. *Social Science Quarterly* 80, 1 (1999), 84–96.
- [14] Florian Krebs, Bruno Lubascher, Tobias Moers, Pieter Schaap, and Gerasimos Spanakis. 2017. Social Emotion Mining Techniques for Facebook Posts Reaction Prediction. *arXiv:1712.03249 [cs]* (Dec 2017). <http://arxiv.org/abs/1712.03249> arXiv: 1712.03249.
- [15] Po Chen Kuo, Fernando H. Calderon Alvarado, and Yi-Shin Chen. 2018. Facebook Reaction-Based Emotion Classifier as Cue for Sarcasm Detection. *arXiv:1805.06510 [cs]* (May 2018). <http://arxiv.org/abs/1805.06510> arXiv: 1805.06510.
- [16] Gabriel Magno and Ingmar Weber. 2014. *International Gender Differences and Gaps in Online Social Networks*. Vol. 8851. Springer International Publishing, 121–138. https://doi.org/10.1007/978-3-319-13734-6_9
- [17] Ariadna Matamoros-Fernandez. 2018. Inciting anger through Facebook reactions in Belgium: The use of emoji and related vernacular expressions in racist discourse. *First Monday* 23 (Sep 2018). <https://eprints.qut.edu.au/122413/>
- [18] Armin Mertens, Franziska Pradel, Ayjeren Rozyjumayeva, and Jens Wäckerle. 2019. As the Tweet, So the Reply?: Gender Bias in Digital Communication with Politicians. In *Proceedings of the 10th ACM Conference on Web Science (WebSci '19)*. ACM, 193–201. <https://doi.org/10.1145/3292522.3326013> event-place: Boston, Massachusetts, USA.
- [19] Bin Tareaf Raad, Berger Philipp, Hennig Patrick, and Meinel Christoph. 2018. ASEDs: Towards Automatic Social Emotion Detection System Using Facebook Reactions. In *2018 IEEE 20th International Conference on High Performance Computing and Communications; IEEE 16th International Conference on Smart City; IEEE 4th International Conference on Data Science and Systems (HPCC/SmartCity/DSS)*. IEEE, 860–866. <https://doi.org/10.1109/HPCC/SmartCity/DSS.2018.00143>
- [20] Kira Sanbonmatsu. 2002. Gender Stereotypes and Vote Choice. *American Journal of Political Science* 46, 1 (2002), 20–34. <https://doi.org/10.2307/3088412>
- [21] Rodrigo Sandoval-Almazan and David Valle-Cruz. 2018. Facebook Impact and Sentiment Analysis on Political Campaigns. In *Proceedings of the 19th Annual International Conference on Digital Government Research: Governance in the Data Age (dg.o '18)*. ACM, 56:1–56:7. <https://doi.org/10.1145/3209281.3209328> event-place: Delft, The Netherlands.
- [22] Shiliang Tang, Xinyi Zhang, Jenna Cryan, Miriam J. Metzger, Haitao Zheng, and Ben Y. Zhao. 2017. Gender Bias in the Job Market: A Longitudinal Analysis. *Proc. ACM Hum.-Comput. Interact.* 1, CSCW (Dec 2017), 99:1–99:19. <https://doi.org/10.1145/3134734>
- [23] Ye Tian, Thiago Galery, Giulio Dulcinati, Emilia Molimpakis, and Chao Sun. 2017. Facebook sentiment: Reactions and Emojis. In *Proceedings of the Fifth International Workshop on Natural Language Processing for Social Media*. Association for Computational Linguistics, 11–16. <https://doi.org/10.18653/v1/W17-1102>
- [24] Claudia Wagner, David Garcia, Mohsen Jadidi, and Markus Strohmaier. 2015. It's a Man's Wikipedia? Assessing Gender Inequality in an Online Encyclopedia. In *Ninth International AAAI Conference on Web and Social Media*. <https://www.aaai.org/ocs/index.php/ICWSM/ICWSM15/paper/view/10585>
- [25] Joan C Williams and Rachel Dempsey. 2014. *What Works for Women at Work: Four Patterns Working Women Need to Know*. NYU Press.
- [26] Olga Zagovora, Fabian Flöck, and Claudia Wagner. 2017. “(Weitergeleitet Von Journalistin)”: The Gendered Presentation of Professions on Wikipedia. In *Proceedings of the 2017 ACM on Web Science Conference (WebSci '17)*. ACM, 83–92. <https://doi.org/10.1145/3091478.3091488> event-place: Troy, New York, USA.