

A Study of Media Polarization with Stylometry Methods

Yifei Hu¹, Julia Rayz²

Purdue University

hu381, jtaylor1@purdue.edu

Abstract

This research investigated the U.S. media polarization with stylometry approaches, creating classification models to identify the political leanings of news articles based on their writing style. We tested the models of authorship attribution, while controlling for topic, stance, and style, and applied them to media companies and their identity within a political spectrum. We tested style features that could include semantic and/or sentiment-related information, such as stance taking, with features that seemingly do not capture it. We were able to successfully classify articles as left-leaning or right-leaning regardless of stance. Finally, we provide an analysis of some of the patterns that we found.

Introduction

The U.S. media polarization issue has been studied for decades.

In the media and journalism areas, researchers have been observing and reporting different forms of media bias with case studies and qualitative analysis. Most people can identify a political article as that written by media favoring one political inclination or another. Usually, the attributes by which one can tell what political side wrote a paper related to style. Style (and stylometry) are usually studied in computational work in the context of authorship attribution – finding particular features that are characteristic of a given author. Putting these two ideas together (people recognizing media leaning and computational approaches to authorship attribution), we are curious whether it is possible to detect media leaning with stylometric methods as almost none of the studies tried to measure the differences between each political inclination from the perspective of “style”.

This paper explored the stylistic difference between different authors that contribute articles to different media groups and political inclinations using computational authorship attribution approaches. The question that is addressed is it possible to learn to differentiate and explain the stylistic characteristics of news articles relative to their political spectrum?

Literature Review

In order to understand how to approach a classification of media with different political leanings, it is helpful to examine what is meant by style and stance.

Useful methods to study style can come from stylometry, which can be defined as “the use of statistical methods in the analysis of literary style.” (Holmes, 1998) Stylometry was described by Joula (2008) as the 3rd problem in authorship attribution, which he defined as the science of inferring characteristics of the author from the characteristics of documents written by that author. The other two problems of authorship attribution are: determination of an author based on a collection of texts provided as a sample and identification of an author from a list of authors based on a sample.

Joula (2008) stated “authorship attribution is a near-synonymous term of ‘stylometry.’” Some papers indeed use the term “stylometry” as a near-synonym to “authorship attribution.” For instance, Abbasi et al. (2008) described one of the stylometric analysis tasks as “compare anonymous texts against those belonging to identified entities, where each anonymous text is known to be written by one of those entities.”

From stylometric analysis, many attributes of the authors can be recovered from the texts. These attributes include age, gender (Goswami et al., 2009), native languages (Koppel et al., 2005), whether the author has dementia (Le et al., 2011), whether someone is writing deceptive online reviews (Ott et al., 2011), and whether a paper is written in the style of a conference or a workshop (Bergsma et al. 2012). As people are getting more aware of the importance of privacy, there are also applications trying to hide the authors’ attributes by obfuscating the writing style (Emmery et al., 2018). Would it be possible to add political views to this impressive list?

Computational Authorship Attribution

Computational approaches to authorship attribution range from statistical approaches to deep learning approaches.

In the research conducted by Diederich et al. (2003), support vector machine (SVM) performed well in various tests. According to the authors, SVM can handle relatively “large” sizes of input data. The dataset used by Diederich et al. consisted of 2652 documents which was larger than many other authorship attribution datasets but still relatively small compared with the datasets in deep learning experiments. In an

authorship attribution experiment on German newspapers, SVM achieved close to 100% precision for the majority of the authors. SVM also showed great robustness even when the topics were not strictly controlled.

The nature of the authorship attribution task is very similar to text classification. In authorship attribution tasks, we classify the authors based on their texts.

Zhang et al. (2015) published a famous work on general text classification using character-level convolutional neural networks. The input text is encoded into a 1 dimensional vector which is then passed to a convolutional neural network (CNN). By comparing the character-level CNN with word-level CNN, and Long Short-Term Memory (LSTM) models, Zhang et al. suggested that “the character-level CNN is an effective method.”

Shrestha et al. (2017) adopted a similar idea and used character level n-gram with convolutional neural networks for authorship attribution on short texts. They tested the approach with different n-gram sizes, different numbers of authors, and different numbers of short texts written by each author. The result indicated that the n-gram CNN approach can work on various situations. The accuracy for 50 authors with 1000 tweets each was 76.1%.

Alsulami et al. (2017) used LSTM to do source code authorship attribution on 10, 25, and 70 different authors and the LSTM models outperformed the SVM and Random Forest models.

We follow these lines of research and compare statistical approaches and deep neural network approaches for our tasks.

Topic, stance, and style

Mikros et al. (2007) claimed that many stylometric variables were discriminating topics rather than authors. The conclusion from their experiments suggested that when doing authorship attribution experiments on multi-topic corpora, the researchers should be extremely careful about the impact of the topic. Sarawgi et al. (2011) performed several authorship attribution experiments on various genres and topics and ended up finding statistical evidence of gender-specific language styles beyond topics and genres, although the token-level language models had the tendency to learn topic words on top of the basic stylometric cues. Sánchez-Junquera et al. (2019) compared style-based and topic-based models on detecting hyperpartisanship in news articles. The results suggested that topic-based models tended to perform better in this task.

Shrestha et al. (2017) believed that character-level features like character-n-grams learned the style of an author (and not the topic). Applying convolutional neural network and character-n-grams can identify verbal fillers and structural features of the given authors.

Besides topic, another variable in texts is the stance. According to the Merriam-Webster dictionary, stance is defined as an intellectual or emotional attitude. Stance can be detected through quantitative methods. Lai et al. (2019) applied linear SVM to automatically detect stance from a corpus of tweets and visualized the results in different temporal phases. Among the papers talking about both stance

Table 1: Dataset Statistics - Overview

media group	# of article	political leaning	selected
Breitbart	23781	R	7076
CNN	11488	L	1306
Washington Post	11114	L	2963
New York Times	7803	L	813
Atlantic	7179	L	1579
Vox	4947	L	1236
Buzzfeed News	4854	L	586
Fox News	4354	R	1567

and style, many believed that there was no explicit boundary between the two elements. Kiesling (2009) stated that “repeatable linguistic styles emerge out of stancetaking strategies that prove repeatedly relevant and useful for particular speakers in particular kinds of interactions.” Jaffe (2009) also claimed that “a personal style is created through habitual stancetaking.”

We wonder, however, if it is possible to identify style without being affected by stance. In order to do that we go back to authorship attribution features that cannot account for stance due to their inability to detect meaning, namely part of speech. Part of speech n-grams have been successfully used for authorship attribution (Brennan et al., 2012) and we rely on them to compare classification accuracy due to stance vs style.

Methodology

The question central to this research is: what are the shared unique stylometry features within each U.S. media group or political side that can be identified by computational authorship attribution approaches?

To answer this question, we applied various authorship attribution approaches on the news articles published by media with various political leanings.

Dataset

We used “All the News” dataset which contained more than 140,000 news articles from 15 major news sources. The number of articles published by each news source (used for the experiments) and the political leaning (Jurkowitz et al., 2020 & Allsides, 2020) of each news source are shown in Table 1.

To reduce the topic’s impact on the classification models, we restricted the topic to be related to the 2016 election. We selected articles published between January and November 2016. Among those articles, the ones containing either “Donald Trump” or “Hillary Clinton” were selected (see Table 1) as the articles related to the election. If we only use the last names of the presidential candidates, their family members (husband, children, etc.) would likely to be included. Using full names might exclude some relevant articles, but can improve the precision of filtering.

We selected CNN, New York Times, Vox, Atlantic, Buzzfeed News, and Washington Post as the left-wing group.

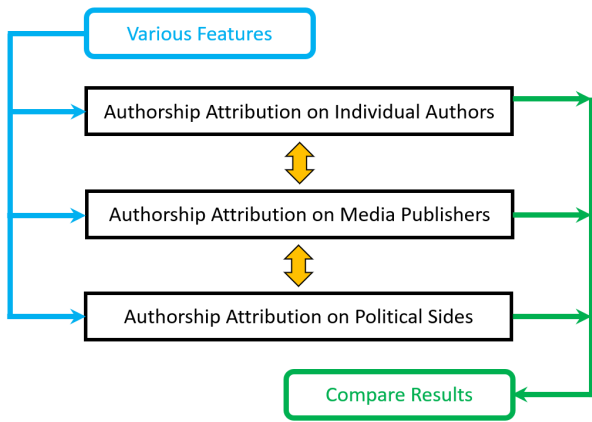


Figure 1: Study Pipeline

Fox News and Breitbart were chosen as the right-wing group (Jurkowitz et al., 2020). The final distribution of the data is perfectly balanced with 8,490 news articles from the left-wing media and 8,649 articles from the right-wing media.

Experimental Setup

Figure 1. shows the pipeline of this study. We conducted authorship attribution experiments on individual authors, media publishers, and two sets of publishers grouped by their political inclinations. We hypothesize that the style of individuals determines the style of organizations or groups. By comparing the salient features and results we can test this hypothesis and can come closer to an understanding whether stylistometric classification must take stance into account.

We used support vector machine, logistic regression, random forest, and deep learning models to perform the authorship attribution tasks. The features were vectorized using the TF-IDF algorithm so that they could be passed into the machine learning models for further classification. To avoid selection bias during the data partition stage, we applied 5-fold cross validation on every machine learning experiment and reported the average score in the result section.

As there is no clear boundary between style, stance, or topic, in order to run a clean experiment we control for topic by selecting only articles about Trump or Clinton during a specified timeframe of 2016 election campaign, and select word-level features, character-level features, and Part-of-speech (POS) features that we believe contain different amount of non-style information. Based on common sense, word-level features can contain sentiment or stance information while POS features contain mostly style information. Character-level features break down some of the word choice information (carrying stance) but still contain more non-style information than POS features. By combining different features with the same classification models, we expected to learn how much the style differences contributed to the classification by comparing the different results generated by each type of feature.

Most of the machine learning models provided convenient ways to extract the highest weighted features. Features ac-

Table 2: Authorship Attribution with 100 Authors

Feature	Accuracy	Feature	Accuracy
Single Word	61.66%	Char-4-gram	59.43%
Word-2-gram	64.56%	POS-2-gram	29.41%
Word-3-gram	61.13%	POS-3-gram	50.74%
Char-2-gram	29.35%	POS-4-gram	56.91%
Char-3-gram	57.03%	POS-5-gram	55.62%

Table 3: Authorship Attribution on 50, 100, and 400 Authors

	chance	word-2-gram	char-4-gram
50 authors	2.00%	77.31%	72.56%
100 authors	1.00%	64.56%	59.43%
400 authors	0.25%	48.76%	44.72%

counting for the differences between the political inclinations explain the results of the machine learning models. By looking into the highest weighted features in the machine learning algorithms, we expected to learn some specific stylistometric patterns favored by each political inclination. The analysis of features could reflect the degrees of explainability of the machine learning models.

Results and Discussions

From individual to group style

We first perform the experiments on individual authors. We applied SVM models with various features to perform authorship attribution tasks on the top-100 most productive authors. Table 2 shows the results from word-level, character-level, and POS features.

The results show the word-level features generally performed better than character-level features and POS features. This is not surprising as we expect authors to be identified by certain words or phrases that are characteristic of their writing. SVM with word bi-gram feature achieved 64.56% accuracy. Char 4-gram and POS 4-gram have similar accuracy. However, these numbers may not tell a full story.

The 17,000 news articles in the dataset are written by 1677 different authors. We thus wanted to find out if these numbers hold for more authors in the dataset. The mean number of articles published by each author is 9.26, but the median is only 1, which means more than half of the authors only published 1 article. We exclude authors that do not have enough samples, and experiment with the top 400 authors.

Each of the top-400 most productive authors published at least 5 news articles, while each of the top-50 authors published at least 33 articles. We applied the SVM model with the best performing word-level and character-level features on authorship attribution tasks with different number of authors (see Table 3).

The results suggest that authorship attribution tasks with fewer authors typically produced higher accuracy, which is consistent with previous findings and matches the results

Table 4: Authorship Attribution on Media Companies

	Word-2-gram	Char-4-gram	POS-4-gram
Accuracy	80.34%	73.63%	74.61%

Table 5: Statistical models (SVM, Random Forest(RF), Logistic Regression(LR)) for word, character and POS n-grams and the classification results

	2-gram	3-gram	4-gram
Word-SVM	0.9178	0.9003	0.8751
Word-RF	0.8851	0.8732	0.8669
Word-LR	0.9087	0.9038	0.8827
Char-SVM	0.8375	0.9040	0.9052
Char-RF	0.8015	0.8071	0.8405
Char-LR	0.8101	0.8769	0.8895
POS-SVM	0.8041	0.8494	0.8657
POS-RF	0.7800	0.7763	0.7802
POS-LR	0.7737	0.8312	0.8489

from various past literature. The accuracy decreased when increasing the number of authors from 50 to 100 to 400. On the other hand, the difference between random guess and the calculated accuracy actually improved as the number of authors increased. The results also show that the word bi-grams still performs better than character 4-grams.

The style of each political inclination (assuming binary classification) comes from the style of each media company supporting the similar political views. The SVM model with the best performing features in previous tasks can also differentiate the media companies, as shown in Table 4.

Table 4 showed how the word-level, character-level and POS features performed on the task of identifying the publisher of a news article. There are only 8 different media companies in our dataset, thus the task is considerably easier than before. The best performing feature provided above 80% accuracy in this task, and, again, corresponds to word choices. One should notice, however, that the accuracy does not decrease by a large amount when we consider POS or character 4-grams. In other words, assuming that POS does not carry any stance information, the media companies do have differences in style.

The style of political inclinations

The previous experiments supported our assumption that the style of individual authors can shape the style of an organization. The goal of this study is to discover the style features that characterize political inclinations. Therefore, we trained multiple machine learning models to classify the political leanings of the news articles (see Table 5).

The best performing character-level and word-level models in the previous tasks also provided the best results in differentiating the political inclinations. The SVM models outperformed other machine learning models in most tests.

Word-level features were classified as stylometric features

according to Joula (2008), but Joula and other researchers did not specify whether the word-level features also contain information more than just style. One can argue that some of the semantic information was also well preserved in the word-level features, and such semantic information reflect not only the writing style but also the topic, sentiment, and other characteristics that are beyond style.

Breaking the words into character n-grams is one way to reduce some of the semantic information preserved by the words. Table 5 showed that the best performing character-level features (character 4-grams) still achieved an excellent accuracy. Shrestha et al. (2017) suggested that character-level n-grams were able to capture the stylistic features (like verbal filters) of the texts, but Shrestha et al. did not mention whether the character n-grams also stored semantic information. Zhang et al. (2015) used character-level features to classify the topics of the documents and showed that character n-grams performed just as well as word-level features in terms of identifying the topic of a given text. This makes it plausible that char level n-grams, while lower in accuracy, still carry some features of stance-taking.

Classification without semantic information Breaking down the words into character N-grams potentially decomposed some semantic information, but the results from character-level features and word-level features were still highly similar. To further explore the role of style and stance in identifying the political inclinations, we converted words to their POS and used POS n-grams in order to dissociate the semantic information from the texts. Examples below demonstrate how a regular sentence looks after such conversion.

“Democrats did what party people do.”

“PROPN VERB DET NOUN NOUN VERB PUNCT”

The results of classification using POS n-grams are shown in bottom rows of Table 5. The results strongly indicate that there was a significant stylometric difference between the media outlets. After removing the semantic (stance-carrying) information of the news articles, the classification accuracy only decreased by approximately 4.5% (comparing with the best performing word 2-gram support vector machine model). The results suggest that the semantic differences played a relatively minor role during the classification process.

Deep Learning Models We applied neural network models along with character-level N-gram features because both Zhang et al. (2015) and Shrestha et al. (2017) used the character-level features in their studies and confirmed the effectiveness of the character-level features.

Due to the complexity of the hyper-parameters only limited combinations of hyper-parameters were tested. The CNN model with the best performing hyper-parameter setting (complete results are omitted due to space limitation) only achieved 86.01% accuracy.

One of the reasons why traditional statistical machine learning methods performed better than the sophisticated deep learning methods could be attributed to the character-

Table 6: Keywords Counts for Both Political Inclinations

Keywords	Left	Right
Mr. Trump	1,201	9,010
Donald Trump	12,043	16,939
Mrs. Clinton	343	3,293
Hillary Clinton	10,193	11,499

istics of the dataset which did not favor the neural networks. Another reason could be the size of the dataset.

SVM and other statistical machine learning models are known to work well with small datasets. After a training-testing split, our training dataset only contains about 13,000 samples, which seems to be inadequate for deep learning. Although both Zhang et al. (2015) and Shrestha et al. (2017) applied CNN models with character-level features and achieved good results, there were two conditions that we did not have in our experiments.

Zhang et al. (2015) tested the CNN models on several datasets, with the smallest dataset being 10 times larger than our training dataset. Our CNN model started to overfit after 2-3 epochs of training. While our smaller neural network models contained only 1 million parameters and the larger models contained more than 4 million parameters, they still overfitted fairly quickly. This could be mitigated with an increased size of the dataset. Shrestha et al. (2017) used CNN models on Twitter texts which belonged to a different genre. The Twitter texts were informal, thus the character-level CNN models captured verbal filters (“*Uhm..*”) and internet expressions (“*XD*”) which did not exist in formal news articles. It is possible that the genre of a text also contributed to a inferior performance (if 86% can be called inferior).

Analysis on Features

Word and word N-grams Even though the classification results on POS tags indicated the existence of non-stance stylistic differences between the media outlets, we wanted to analyze what specific stylistic differences were important.

We start with word frequency as it is one of the features that can be easily interpreted. Table 6 shows the word frequency differences we manually selected from the results. The table shows that the right-leaning media tends to use titles more frequently to address the presidential candidates.

The word N-gram results in Table 7 show that among the 10 highest-weighted features, half of them (bold entries) seemed to be purely format features. Since the news articles were crawled from the web pages, features like “*read rest*” and “*follow twitter*” were treated as part of the content of the news articles. Such minor format artifacts were not removed during text pre-processing, but they did play a very important role in the classification process. It can be argued that they should be removed if one wants to test the effects of the text on classification. A counter argument can be made that classification is done on text that appear on a website and not on paper, and thus, these artifacts are legitimate document characteristics.

Table 7: Highest weighted word bi-grams

Feature	Weight
read rest	4.795745
follow twitter	4.397071
news com	4.057425
frontrunner politician	3.526376
news daily	3.503085
press contributed	3.300977
prediction map	3.056871
democrat party	3.039265
nominee politician	2.958779
illegal alien	2.821554

Table 8: POS-N-Gram Coefficients from Linear SVM

Feature	Weight
NUM PUNCT NUM PUNCT	5.91
VERB ADJ NOUN SYM	4.62
VERB PROPN PROPN ADP	4.62
VERB PROPN PUNCT NOUN	3.92
PRON ADP PROPN ADP	3.81

Further analysis shows that the bi-gram “*follow twitter*” occurred in 829 left leaning news articles, vs. 316 articles from the right leaning media. The similar features (related to Twitter) were also the highly weighted features in authorship attribution tasks on individual authors. The possible explanation could be the authors put their own Twitter account in the web pages and the web crawlers (spider) treated them as part of the content. Besides the format features, some keywords (like “*illegal alien*”) strongly indicate the political lean of the new articles due to particular word choices that are characteristic of a single party.

Part-of-speech N-grams Part-of-speech annotations do not contain word choices that can serve as a signature for a particular party, and thus their analysis is more difficult and interesting. Table 8 shows the highest-weighted POS N-grams from a linear support vector machine model.

We cannot precisely interpret the corresponding semantic meanings behind the Part-of-speech N-grams without manually locating their positions in the raw text, but some of them still point to the format artifacts. For example, the feature “*VERB ADJ NOUN SYM*” represents “verb adjective noun symbol” which is likely to be a specific format feature used by Fox News exclusively since this feature appeared 333 times in Fox News articles and only 8 times in other articles. In 259 of the Fox News articles, “*VERB ADJ NOUN SYM*” simply means “*See Latest Coverage →*” - an HTML element that is not part of the content. Other features are more likely to be universal but are preferred by one political inclination. “*VERB PROPN PUNCT NOUN*” (which represents “verb proper-noun punctuation noun”) appeared in articles from all media outlets. However, the right-leaning media uses this feature 47% more frequently than the left-

leaning media.

It appears that both the highly weighted word-level features and part-of-speech features were mainly format specific, and had limited correlation with any political views per se. However, as they serve as classification features, we wanted to test whether a previously unseen media outlets could be correctly classified. We thus used the models that we learned on articles from Reuters, New York Post, and NPR (see Table 1). Our models (word, character, and POS) indicated that Reuters was a right-leaning media outlet, while NPR was a left-leaning one. The results for New York Post were less conclusive, with word-level and POS models indicating that it is a right-leaning venue, while char model pointed that it should be in the center. Interestingly, Pew Research and AllSides also do not agree on whether New York Post should be a center or a right media outlet.

Conclusion

This study applied stylometry techniques to analyze the political leanings of the media outlets. We hypothesized that the style of a particular political inclination comes from the shared style of media companies supporting similar political views, which in turn can be traced to the styles of individual authors working for the same company. The results confirmed this hypothesis.

Moreover, we tested style features that could include semantic and/or sentiment-related information, such as stance taking, with features that seemingly do not capture it. We were able to successfully classify articles as left-leaning or right-leaning regardless of stance. Further analysis, however, showed that the pattern of some of the non-stance features can be linked to particular phrases that are used by individual media outlets as they format their articles. While this can still be counted as "style" and not "stance", this does not relate to style of content, but rather to the style of formatting. We thus tested our models on previously unseen media outlets and reported the results.

References

- Holmes, D. I. (1998). The evolution of stylometry in humanities scholarship. *Literary and linguistic computing*, 13(3), 111-117.
- Diederich, J., Kindermann, J., Leopold, E., & Paass, G. (2003). Authorship attribution with support vector machines. *Applied intelligence*, 19(1-2), 109-123.
- Koppel, M., Schler, J., & Zigdon, K. (2005, August). Determining an author's native language by mining a text for errors. In *Proceedings of the eleventh ACM SIGKDD international conference on Knowledge discovery in data mining* (pp. 624-628).
- Mikros, G. K., & Argiri, E. K. (2007, July). Investigating Topic Influence in Authorship Attribution. In *PAN*.
- Juola, P. (2008). *Authorship attribution* (Vol. 3). Now Publishers Inc.
- Abbasi, A., & Chen, H. (2008). Writeprints: A stylometric approach to identity-level identification and similarity detection in cyberspace. *ACM Transactions on Information Systems (TOIS)*, 26(2), 1-29.
- Koppel, M., Schler, J., & Argamon, S. (2009). Computational methods in authorship attribution. *Journal of the American Society for information Science and Technology*, 60(1), 9-26.
- Goswami, S., Sarkar, S., & Rustagi, M. (2009, March). Stylo-metric analysis of bloggers' age and gender. In *Third international AAAI conference on weblogs and social media*.
- Le, X., Lancashire, I., Hirst, G., & Jokel, R. (2011). Longitudinal detection of dementia through lexical and syntactic changes in writing: a case study of three British novelists. *Literary and Linguistic Computing*, 26(4), 435-461.
- Ott, M., Choi, Y., Cardie, C., & Hancock, J. T. (2011). Finding deceptive opinion spam by any stretch of the imagination. *arXiv preprint arXiv:1107.4557*.
- Sarawgi, R., Gajulapalli, K., & Choi, Y. (2011, June). Gender attribution: tracing stylometric evidence beyond topic and genre. In *Proceedings of the fifteenth conference on computational natural language learning* (pp. 78-86).
- Brennan, M., Afroz, S., Greenstadt, R. (2012). Adversarial stylometry: Circumventing authorship recognition to preserve privacy and anonymity. *ACM Transactions on Information and System Security (TISSEC)*, 15(3), 1-22.
- Bergsma, S., Post, M., & Yarowsky, D. (2012, June). Stylo-metric analysis of scientific articles. In *Proceedings of the 2012 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies* (pp. 327-337). Association for Computational Linguistics.
- Zhang, X., Zhao, J., & LeCun, Y. (2015). Character-level convolutional networks for text classification. In *Advances in neural information processing systems* (pp. 649-657).
- Shrestha, P., Sierra, S., González, F., Montes, M., Rosso, P., & Solorio, T. (2017). Convolutional Neural Networks for Authorship Attribution of Short Texts. *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, 669-674. <https://www.aclweb.org/anthology/E17-2106>
- Emmery, C., Manjavacas, E., & Chrupala, G. (2018). Style obfuscation by invariance. *arXiv preprint arXiv:1805.07143*.
- Lai, M., Tambuscio, M., Patti, V., Ruffo, G., & Rosso, P. (2019). Stance polarity in political debates: A diachronic perspective of network homophily and conversations on Twitter. *Data Knowledge Engineering*, 124, 101738.
- Sánchez-Junquera, J., Rosso, P., Montes-y-Gómez, M., & Ponzetto, S. P. (2019). Unmasking Bias in News. *arXiv preprint arXiv:1906.04836*.
- Jurkowitz, M., Mitchell, A., Shearer, E., & Walker, M. (2020, May 08). 1. Democrats report much higher levels of trust in a number of news sources than Republicans. Retrieved July 04, 2020, from <https://www.journalism.org/2020/01/24/democrats-report-much-higher-levels-of-trust-in-a-number-of-news-sources-than-republicans/>
- AllSides Media Bias Ratings. (2020). Retrieved July 04, 2020, from <https://www.allsides.com/media-bias/media-bias-ratings>