

The Nature of Real and Perceived Bias in Chilean Media

Erick Elejalde*
Department of Computer Science
University of Concepcion
Concepción, Chile
eelejalde@udec.cl

Leo Ferres
Data Science Institute
Universidad del Desarrollo &
Telefónica R&D
Santiago, Chile
lferres@udd.cl

Eelco Herder
L3S Research Center
Leibniz-University
Hannover, Germany
herder@l3s.de

ABSTRACT

News consumers expect news outlets to be objective and balanced in their reports of events. However, there is a body of evidence of bias in the media caused by underlying political and socio-economic viewpoints. Previous studies have tried to classify the partiality of the media, sometimes giving a quantitative evaluation, but there is little reported on its nature. The vast amount of content in the social media enables us to quantify the inclination of the press to either side of the political spectrum. To describe such tendencies, we use tweets to automatically compute a news outlet's political and socio-economic orientation. We show that the media have a measurable bias, and illustrate this by showing the favoritism of Chilean media for the ruling political parties in this country. We also found that the nature of the bias is reflected in the vocabulary used and the entities mentioned by different news outlets. A survey conducted among news consumers confirms that media bias has an impact on the coverage of controversial topics and that this is perceivable by the general audience. Having a more accurate method to measure and characterize media bias will clarify to the readers where outlets stand within the socio-economic landscape, even when a self-declared position is stated. This will empower readers to better reflect on the content provided by their news outlets of choice.

KEYWORDS

Media bias, bias characterization, political quiz

ACM Reference format:

Erick Elejalde, Leo Ferres, and Eelco Herder. 2017. The Nature of Real and Perceived Bias in Chilean Media. In *Proceedings of HT '17, Prague, Czech Republic, July 04-07, 2017*, 10 pages.
DOI: <http://dx.doi.org/10.1145/3078714.3078724>

1 INTRODUCTION

The media have a strong influence on how people perceive the world that surrounds them. More and more power has been ascribed to the modern press since its inception, even calling it the “Fourth

*Funded by CONICYT doctoral scholarship (63130228)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

HT '17, Prague, Czech Republic

© 2017 ACM. 978-1-4503-4708-2/17/07...\$15.00
DOI: <http://dx.doi.org/10.1145/3078714.3078724>

State”¹ emphasizing its independence and its ability to provide strict limits to what governments may or may not do. There are well known examples of the press even toppling governments: the Washington Post in the Watergate scandal is perhaps the most resounding example.

However, as the media grows in power, the political and economic interests of news outlets and the ones who control it have grown as well, which has its impact on the news that the population of a territory gets served. Among others Herman and Chomsky [16] argue that political and doctrinal interests have penetrated the press at different stages of the news generation process, deliberately or accidental - for example through homophily effects. In certain cases the resulting bias is explicitly stated, in other cases – like FOX News – the bias is known but not communicated. People usually have some intuition of media bias. For average readers, though, it is very difficult and time-consuming to be aware or even find the bias of all media outlets, let alone quantify these biases and give them a total order in terms of the magnitude of the leaning.

Bias in the media is a global phenomenon, not exclusive to one kind of economy or particular political system. As such, there is now a quickly growing body of empirical evidence on its existence [7, 20, 24]. In previous work [29], we showed several types of bias in media coverage of ongoing news stories on crises in the world. What has not been studied as deeply, however, at least not quantitatively, is *how* outlets could be positioned in a socio-economic space. Knowing *the nature* of media bias will help individuals and organizations take actions that counteract bias. If, for example, a newspaper claims to be objective, but is in fact “right-wing, conservative” (as is the case with El Mercurio in Chile²), people should be able to recognize this and take this bias into account when reading its content. The case of El Mercurio is quite clear, and being a very old, traditional newspaper, the bias is actually known and arguably accepted. It is important to emphasize here that “bias” is not categorical, but comes embedded in a geopolitical news context determined by other outlets in the region [30]. In other, bolder words, some bias is inherent to the media, but how biased they are, depends, to an extent, upon a comparison to other media.

In this work, we automatically identify the (largely implicit) socio-economic “relative bias” of news outlets in the context of Chilean media. The value of our methodology and study here is to position those media outlets that do **not** state their socio-economic bias, or are not even aware of their bias. Socio-economic studies at this scale may help uncover patterns of editorial policies that show a systematic bias that favors governments’ propaganda or private economic interests over social welfare. Operationalizing bias is a

¹https://en.wikipedia.org/wiki/Fourth_Estate

²https://en.wikipedia.org/wiki/El_Mercurio

difficult task. It relies not only on linguistic information, but also on the actual geo-socio-economic, and even historical, context of the newspaper. We propose to automatically categorize news outlets by analyzing what they “think” about certain relevant, controversial topics using their tweet content and then map these worldviews onto a well-known political quiz: “The World’s Smallest Political Quiz” (henceforth *PolQuiz*) [27].

The *PolQuiz* is a ten-question educational quiz for an American audience³ designed by the Libertarian Advocates for Self Government⁴, created by Marshall Fritz in 1985. The quiz is based on the one proposed by David Nolan Chart in 1971 [22]⁵, which in turn can be traced back to a 2D chart proposed in 1968 [4], representing variations in political and socio-economic orientation.

In short, we use what the media say on Twitter to position them in a Cartesian plane that tells us more about their orientation based on Fritz’ *PolQuiz*. In turn, the *PolQuiz* results motivate a deeper investigation into the nature of the found bias, which we study through the vocabulary used and entities covered by the news outlets. Finally, we conducted a survey that confirm that media bias has a noticeable impact on how news related to controversial topics are presented.

2 RELATED WORK

There are several works related to the topic of media bias [7, 12, 14, 23, 31]. Some works do not try to identify bias directly, but instead try to identify and track events in order to present different points of view of the same affair to the readers in order to counteract these possible bias [23]. These are complemented by works like J. An *et.al.* [2], which create a so-called landscape of newspapers based on the similarity of their communities. They measure the exposure of Twitter users to politically diverse news. Other authors assume a certain leaning by contacts association [6]. In [24] the authors go deeper and try to identify different kinds of bias, what they term gatekeeping, coverage and statement bias, according to the stage at which the news acquire the alleged bias.

Most outlets identify themselves as unbiased free press, which makes the discussion on the direction and degree of media bias very controversial⁶. Media bias is usually found in the editorial policies that ultimately decide which stories are worth publishing and which amount and angle of coverage they get [23, 24, 26].

This bias reflects the political and socio-economic views of the institution, rather than the point of view of a particular reporter. For example, the authors in [16] use a few recent events to point out how the press applies the word “genocide” to cases of victimization in non-allied states, but almost never to similar or worse cases committed by the home state or allied regimes. In the latter case, they could use terms such as “repression of insurgency”.

In [19], the authors defined a model to predict political preference among Twitter users. Through this model they calculate, for each user, a ranking of the likelihood that they prefer a political

party over another. This model is based on the usage of *weighted words*. The words and their weights are extracted from tweets of candidates of certain political parties. Using these weights, in combination with Twitter specific features (retweets, following, etc.) the authors train classifiers that achieve a performance similar to that of human annotators. Similarly, in [12], the authors estimate the bias in newspapers according to how similar the language is to that used by congressmen for which a right/left stand is known. One interesting result is that bias in the news is found to be correlated to political inclinations of readers, showing a tendency in these news outlets to maximize profit by “catering” to a certain audience.

The topology of the social network on its own has also been shown to give enough information to create classifiers concerning a user’s preference, even when the choices are very similar [11](e.g. Pepsi vs. CocaCola, Hertz vs. Avis or McDonalds vs. BurgerKing). Although we carefully select the dataset to use in our experiments to achieve extensible results [5], we notice here that in our dataset, news outlets (which may be considered the participants of our studies), regularly talk about these controversial topics, and thus, it is possible to use traditional methods to find a political stand.

Combining topological characteristics of the social networks with language features has also been tested [6], showing that users tend to interact more frequently with like-minded people. This phenomenon is known as *homophily*. As we mentioned before, our dataset is derived from a special type of users (news outlets Twitter accounts), and this method may not apply directly.

As an alternative approach, in [31] the authors propose a semi-supervised classifier for detecting political preference. They design a propagation method that, starting with a few labeled items and users, creates a graph representing the connections between users and items or even users with other users. Based on the same phenomena of homophily, they assume that users interacting with the same item, or with each other, most likely have the same political leaning. This way, they can propagate the labels from tagged users and tagged items to the rest of the graph. They report that the system achieves over 95% precision in cross-validation. In [14, 15], the authors also follow a propagation strategy to compute the political preference of Twitter users, but using Congress members as the initially tagged users.

In [18], the authors describe a framework to discover and track controversial topics that involve opposing views. They first use tags that represent each side (e.g. “#prolife” - “#prochoice”) as seeds to find an expanded set of labels to represent each side. This may also help in cases where labels may change over time as the result of new arguments for either side. With these sets of labels they identify strong partisans (anchors) that have a clear lean to one side. Having these anchors and a graph representing relationships between users (based on tweet content-similarity or based on re-tweets), they propagate the classification through the graph inferring the opinion bias of “regular” users.

Yet another approach to quantifying political leaning is presented in [20]. They based their analysis on the number of tweets and re-tweets generated about different real-life political events associated with some predefined topics. The authors developed a model that takes into account both the sentiment analysis of the tweets and

³Although we believe this does not imply a loss of generality wrt Latin American culture, at least in the topics chosen. It does, obviously, impact the polarity of attitudes towards those topics, but that is what we explore in these pages.

⁴<https://www.theadvocates.org/>

⁵Although widely cited in the literature, we could not find this manuscript online. We cite it for historical reasons here.

⁶To be fair, it is true that “bias” in journalism may arise naturally out of the interaction of reporters, rather than *a priori*, but this discussion is left for another paper.

the number of time they are re-tweeted to calculate the political leaning score of each outlet.

In [30], the authors propose an unsupervised model based on how news outlets quoted president Barack Obama's speeches. The findings suggest that quotation patterns do reveal some underlying structure in the media, and that these may be evidence of bias. They found that one of the identified dimensions roughly aligns with the traditional left(liberal)-right(conservative) political classification and the other with a mainstream/independent one. This is a strong finding. Still, we believe this is to be somewhat expected, given the selected corpus; namely, presidential speeches in the strongly bipartisan system that dominates U.S. politics. Although this model helps classify and quantify bias in the media, it does not explain the causes and nature of this bias.

In this paper, we present a new methodology that quantifies the political leaning of news outlets based on the automation of a well known political quiz. The prediction of the answers for each question for each outlet is generated based on the polarity of their tweets on subjects related to the issues addressed in the quiz. The automation of a quiz has been used before to automatically classify mood [3] but, as far as we know, this is the first attempt to quantify media bias using this approach.

3 METHODOLOGY

In this section, we describe our dataset, followed by an overview of the *PolQuiz* and an explanation on how we applied this quiz to our data. In Section 3.4, we introduce the Rank Difference method for investigating the nature of bias. We conclude with an overview of the survey that we carried out to measure perceived bias.

3.1 Data

Every news outlet, from the smallest to the largest, has some presence on the web, which opens the possibility for the automatic collection of the news stream they produce. Twitter is a prime example of a web platform that allows this. Twitter is an online social network that enables users to send and read short messages called "tweets". Twitter offers an open API to automatically access the flow of tweets and query the system for user profiles, followers and tweeting history. This makes it possible to explore the behavior and interactions of personal and institutional accounts, developing and testing social theories at a scale never seen before. This is the closest thing we have to a record of the every-day life of over 300 million people⁷. We treat every tweet as an independent document from which we can extract a statement. We assume that these reflect the ideology of the news outlet as an entity. We use Twitter as our source documents because it provides us (and any news consumer) with a standardized way to access the daily events reported by the media. Technically, tweets are much easier to collect (as opposed to, for example, scraping the content of complete articles from a wide range of newspapers and formats). A tweet from a media outlet is a man-made summary of the news, usually in the form of a *headline*. It conveys the main idea, and hence the main editorial point of view. Headlines of online news articles have shown to be a reliable source for adequately providing a high-level overview of the news

⁷Twitter reports 313 million active users, see <https://about.twitter.com/company> (accessed Aug. 2016)

events [1, 8, 28]. These summaries are expected to be representative of the newspaper's bias, but with the advantage that bias is easier to detect than in a full articles (shorter, to the point). Tweets also contain features/annotations (e.g. hashtags (#) and mentions (@)) that help to give semantic to the text.

Chile ranks among the top-10 countries regarding the average number of Twitter users per 1000 individuals [21]. This allows experiments to be comparable with other countries where a larger number of studies have been performed, such as the UK and even the United States. We find this of particular interest, since many of the studies in the literature have been conducted only on English-speaking countries, which may bias the knowledge we possess in general about these issues.

To create our database of outlets, we used different sources, with Poderopedia's "influence" database as our baseline⁸, manually adding other news outlets in Chile. Our database contains 399 *active* accounts⁹. The data set contains 1,916,709 tweets, spanning a period of 8 months - from October 6, 2015 to June 4, 2016. The accounts vary dramatically in tweet publication behavior, with some having published more than a hundred thousand tweets to others with less than a hundred in this timeframe. Out of the 399 active accounts, only 269 outlets published at least one document about the topics of interest.

3.2 PolQuiz

The *PolQuiz* has ten questions, divided into two groups: economic and personal issues, of five questions each. The answers to the questions may be *Agree*, *Maybe (or Don't Know)* or *Disagree*.

Personal issue questions:

- (1) Government should not censor speech, press, media or Internet.
- (2) Military service should be voluntary. There should be no draft.
- (3) There should be no laws regarding sex between consenting adults.
- (4) Repeal laws prohibiting adult possession and use of drugs.
- (5) There should be no National ID card.

Economic issue questions:

- (6) End corporate welfare. No government handouts to business.
- (7) End government barriers to international free trade.
- (8) Let people control their own retirement: privatize Social Security.
- (9) Replace government welfare with private charity.
- (10) Cut taxes and government spending by 50% or more.

Based on the answers to these questions, the quiz-taker is classified into one of five categories: left-liberal, libertarian, centrist, right-conservative, or statist. *Left-liberalism* is a political ideology that supports governments that take care of the welfare of vulnerable people and keeps a centralized economy, but at the same time, allows a great deal of liberties in personal matters. *Libertarians* seek freedom in both economic and personal issues, minimizing the role of the state in all matters. An extreme position in this direction would be anarchism. On the other side, *statists* - or supporters of a big government - want the state to regulate both personal and economic issues. Examples of this position would be totalitarian regimes, such as Kim Jong-Un in North Korea. *Right-wing conservatives* are more reluctant to accept changes in personal issues and want official standards on these matters (i.e. morality and family traditions), but demand economic freedom and a free market. Finally,

⁸<http://apps.poderopedia.org/mapademedios/index/>

⁹An account is *active* if it tweets at least once a month

centrists accept or even support a balance between the government reach and personal/economic freedom. They favor selective government interventions to current problems while avoiding drastic measures that may shift society to either side of the spectrum.

For each *Agree* answer, we increase the score of the quiz-taker in the corresponding dimension by 20 points. If the answer is *Maybe* (or *don't know*), we only add 10 points. Finally, if the answer is *Disagree*, no points are added. This way, if the quiz taker agrees with all the issues in one dimension, it will be in one end of that axis. In the other extreme of the axis, we will have a quiz-taker who disagrees with all issues in that dimension. In our study, we assume that news outlets are (or strive to be) unbiased, so in an ideal world, most of their comments should have no polarity toward any side of the issue and, as such, they should score as a *Maybe*. Another expected behavior would be that news outlets report on both sides of the issue to cover different points of view. Both approaches would result in the news outlet being in the center of the graph.

3.3 Operationalizing the Quiz

We filtered the collected tweets to get only those with information regarding the issues referred to in the *PolQuiz*. For this, we created a seed query for each question, containing a set of preselected keywords (see Table 1).

With the subset of documents returned by the seed queries, we then analyzed the hash-tags to find an expanded set of labels that may represent related aspects of the same issues [18]. We removed hash-tags that contain the name of a news outlet, as it is common practice in news papers accounts to use hash-tags to refer to themselves or the original source of the news (regardless of the subject). We also remove hash-tags with names of politicians: even when these politicians could potentially provide some relevant documents, they also introduce a lot of noise, mostly due to the salience of politicians who appear regularly in the news for a wide variety of issues not necessarily related to the query in question. The new labels are added as keywords to the original query. Our enriched queries give us the final set of tweets used to evaluate any possible bias of each news outlet, see Table 2.

Having the set of tweets for each question, we classified their polarity *with respect to the corresponding question*. For example, for question 7 (q7), a tweet classified as *Agree* is “TPP *abrira puertas a más de 1.600 productos chilenos no incluidos en acuerdos vigentes.*” (tr. *TPP will open doors for more than 1.600 Chilean products not included in existing agreements*). For that same question, the following tweet disagrees with it: “*El TPP: un misil contra la soberanía*” (tr. *TPP: a missile against our sovereignty*). In other words, we classify the polarity of the tweet with respect to the corresponding issue. As the number of tweets is too large to label manually, we created and trained a supervised model for each question. This approach also allows us to scale in the presence of an even larger number of resulting documents.

To create a representative sample for the training set, we randomly select, where possible, two tweets from each question from each news outlet. We took care to not include duplicate tweets (tweets with the exact same text) published by the same outlet. The training set consisted of 1916 documents (an average of about 190 documents per question). We manually classified this training set in

four groups: *Agree*, *Maybe*, *Disagree* and *Out of topic* (Not relevant). The distribution of each training set is shown in Table 2.

For the automatic classification task, we used a “Randomized Trees” model¹¹ [13]. Decision trees are less susceptible to overfitting, considering that we have relatively small training sets. Given that the classes in our training set are not evenly populated, we decided to evaluate the model using a 10 iterations stratified shuffle-split cross validation. Each fold leaves out 20% for validation. The other 80% is selected while preserving the percentage of samples for each class. The accuracy values for each model is presented in Table 2.

After the classification stage, we scored each news outlet on each question. We removed those documents classified as *off-topic* (Not relevant). We scored the remaining documents’ polarity according to the *PolQuiz* scoring system and we found the average for each question/news-outlet pair. For simplicity, in the question/news-outlet pairs for which we have no associated documents, we assume a *Maybe* (or *don't know*) answer. This assumption is the least disruptive towards the default supposition of an unbiased media¹².

In order to find out how sensitive the observed bias is to noise, we repeated the scoring steps 20 times. Each time we leave out 5% of the tweets, selected at random while maintaining the original distribution of documents per question. Each time we measure the average score of the news outlets for which we were able to answer at least one question in the corresponding dimension. We did not go over 5%, because the smallest news outlets already have a small set of documents: removing too many entries would have resulted in the elimination of an entire outlet, affecting the results.

Finally, we tested how the entire system adapts to the local environment. For a proof of concept, we introduced the subject of abortion in the personal dimension. This topic appears among the personal issues in other political quizzes (e.g Political Compass¹³) In addition, abortion was a very relevant and controversial topic in the Chilean media during this period because of a new bill presented by the president and approved by the Chamber of Deputies to legalize the abortion on three grounds: pregnancy resulting from rape, lethal fetal infeasibility or danger to the life of the pregnant women. We formulated this new question as follows:

(0) All women should be free to choose whether she wants to terminate her pregnancy or not.

Notice that the question is formulated in the same “direction” as the rest of the questions. This is, agreeing with the statement will be an indicator of a more liberal tendency by the quiz taker.

We apply the same methodology described before to the original *PolQuiz*. We named this question **q0** and the query we applied (before injecting the hash-tags) is shown in Table 3. The enriched query returned 4891 documents from our corpus. We selected two random documents for each news outlets to create a training set containing 409 tweets. We had an average precision of 0.70 (± 0.08) in the 10 iterations stratified shuffle-split cross validation.

¹¹Implemented in the python library `scikit-learn` in the module `sklearn.tree.ExtraTreeClassifier`

¹²We are aware this is too conservative. However, for the sake of space, we leave a more detailed discussion for future research

¹³<https://www.politicalcompass.org/>

Table 1: Initial set of keywords for each query.

Question	Keywords
q1	(censura libertad) & (prensa discurso expresion)
q2	(servicio reclutamiento entrenamiento reserva) & (militar ejercito armada)
q3	(ley legal legislacion regulacion penalizacion) & (sexual prostitucion sexo sodomia gay) & ¬(infantil menor niño acoso abuso agresion)
q4	(ley legal legislacion regulacion penalizacion) & (droga marihuana cannabis psicotropico cocaína)
q5	inmigracion inmigrante refugiado xenofobia
q6	(subsidio bienestar ayuda) & (corporativa empresa)
q7	(trato tratado convenio negociacion relacion) & (comercial economica) & (internacional bilateral gobierno libre liberal barrera proteccion bloque)
q8	("seguridad social" afp ¹⁰ pension jubilado prevision) & (privada gobierno estatal)
q9	("beneficio sociale" bono "ayuda sociale" "programa social") & (gobierno)
q10	(reduccion recorte aumento incremento) & (impuesto gasto) & (gobierno gubernamental)

Our actual queries are designed so they can also find variations of the keywords (such as variations in gender and number)

Table 2: Tweets extracted from our corpus after applying the enriched queries.

Qs	# tweets	Training set (TS)	% Agr (TS)	% Mb (TS)	% Dis (TS)	% Not rel (TS)	Prc ($\pm 2 * stdev$)
q1	374	179	0.486	0.167	0.178	0.167	0.76 (± 0.14)
q2	194	132	0.181	0.295	0.204	0.318	0.87 (± 0.11)
q3	144	78	0.576	0.051	0.243	0.128	0.83 (± 0.17)
q4	597	203	0.610	0.083	0.142	0.162	0.80 (± 0.10)
q5	746	219	0.351	0.127	0.159	0.360	0.73 (± 0.16)
q6	636	117	0.264	0.256	0.239	0.239	0.53 (± 0.20)
q7	1162	238	0.298	0.247	0.399	0.054	0.76 (± 0.09)
q8	251	117	0.213	0.094	0.418	0.273	0.76 (± 0.13)
q9	298	167	0.059	0.131	0.694	0.113	0.87 (± 0.09)
q10	8573	466	0.167	0.133	0.660	0.038	0.71 (± 0.06)

The last column indicates the average precision obtained by the model in cross-validation (See Section 3.3)

Table 3: Initial set of keywords for q0 query.

Question	Keywords
q0	(ley legal legislacion regulacion penalizacion despenalizacion) & (aborto interrupcion embarazo)

Our actual queries are designed so they can also find variation of the keywords (such as variations in gender and number)

3.4 Rank difference

Using the *PolQuiz*, we aim to show empirically that the news media in Chile have some socio-economic leaning. This means that news outlets tend to have a stand in at least some of the controversial topics that dominate the political landscape of the country. However, we are also interested in the *nature* of the bias regarding such controversial topics.

To do this, we use the *rank difference* method proposed in [17]. Rank difference is used to identify terms that characterize a specific domain. For example, the word *court* will be probably identified as a term if we are analyzing a corpus of legal documents. The method creates a ranking of words based on their frequency in a domain and a generic corpus. By comparing their relative position in both corpora, the algorithm identifies words that are significantly more used in a given domain. These unusual word frequencies are used as an indication of the importance of these words in the given domain. The formula for calculating rank difference is shown in Equation 1,

$$\tau(w) = \frac{r_D(w)}{\sum_{w' \in V_D} r_D(w')} - \frac{r_B(w)}{\sum_{w' \in V_B} r_B(w')} \quad (1)$$

where $r_D(w)$ and $r_B(w)$ are the ranks of word w in the domain and background corpus respectively. Rank normalization is done

against the summation of all word rankings in the corresponding vocabulary (V_D and V_B).

3.5 Survey

To investigate to what extent the bias – as measured with the *PolQuiz* and investigated using the rank difference method – is perceived by the general audience, we conducted an online survey. We chose abortion as the topic of this survey, as this is (as explained in Section 3.3) a current and controversial item in Chile that has received an important amount of coverage in the local media. This means that most people in Chile are aware of the discussion and probably have their own criteria. We also restricted our survey to the subset of news outlets who had relevant tweets for at least four questions per dimension (see Section 4) since these are the ones that we were able to position in the chart with the highest confidence.

We calculated the bi-grams' rank difference (see Section 3.4) for each news outlet. We decided to present bi-grams to users in the survey instead of words, because bi-grams offer more context, so it was easier for people to assess the connotation of a word or set of words within the selected topic. We also decided to use bi-grams over named entities because people not always recognize all the names involved in the discussion, although they do have an intuition in the discourse and the arguments used on both sides.

For each survey we presented a randomly selected and anonymised list (each list represents a news outlet) with the top-20 ranked bi-grams in one column and the bottom-20 bi-grams in another column. The top-20 list was presented as the words used with a relatively high frequency by one outlet. The bottom-20 list was presented

as words the outlet tried to avoid or used with a relatively low frequency. The user had to answer if, based on these lists, he or she considered the outlet to be “in favor” or “against” abortion. The user could also respond with an “I can’t tell” option. A user could answer the survey more than once, but the random selection was always made from the remaining lists.

We scored the “perceived bias” for each news outlet based on the answers we received in the survey. For each outlet, we calculated the percentage of users that answered “in favor” and subtracted the percentage of “against” answers. These percentages include the users that answered “I can’t tell”. So, we consider outlets with a negative score to have a conservative “perceived bias”. Equivalently, outlets with a positive score are considered as liberals in our “perceived bias”. It is worth noticing that an unbiased news outlet should be expected to score close to zero (because it should have mixed signals and, either a proportional number of user labeled in each direction or most users were unable to classify it).

4 RESULTS AND DISCUSSION

In this section, we first show how the *PolQuiz* helps to measure the bias in the media. We verified that our results are stable to small changes in the dataset and investigated the benefit of contextualizing the quiz by including new questions that fit the current political landscape in Chile. In Section 4.2, we explore the nature of the bias showed by the media by using the rank difference method. We show the differences in the type of coverage between news outlets of various leanings when we deep in the analysis of one particular topic. Finally, in Section 4.3 we investigate, using a survey, to what extent this bias is perceived by the general audience.

4.1 Measuring Bias Using the PolQuiz

For our statistical analysis we will treat each axis independently, so we can work with values in only one dimension.

We aim to make a comparative analysis between outlets, so we only report quantifications relative to the other news outlets. For this, we normalize the scores on each axis in the range [0, 100].

Using our methodology, we find that the news media do have a political bias. This corresponds with previous findings studying this subject [7, 12, 24]. We conducted a one-sample Student t-test¹⁴ (two-sided) for each dimension (economic and personal) to test if the mean score is significantly different from 50 (the assumed unbiased score). We used, for each dimension, only the scores of those news outlets for which we were able to answer at least one question on that dimension. For the economic dimension, there is a significant bias, $t(254) = -10.93, p < .001$, with a leaning to the left-wing ($M = 40.28, SD = 14.21$). In the personal issues the bias is lower, but still is statistically significant, $t(190) = -2.10, p < .05$, with a leaning to the conservative side ($M = 47.42, SD = 16.98$).

We found that the bias is more evident in economic issues than in personal issues. This can be seen in the number of news outlets that comment on those issues and the average score for that dimension. The slight left-wing bias in the economic issues might be explained by the political context of Chile during the observed period: the president, Michelle Bachelet, is affiliated to the socialist

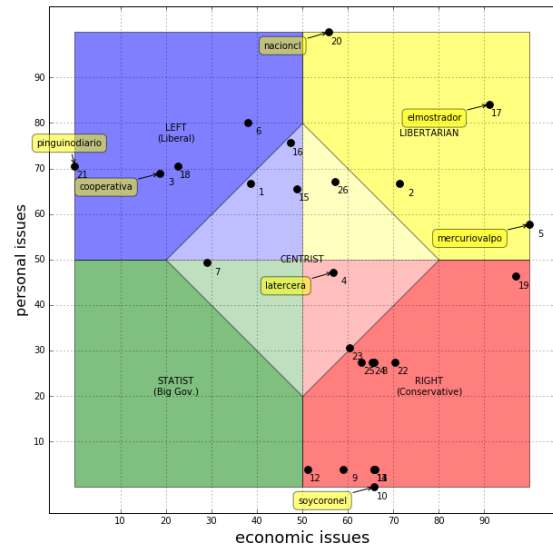


Figure 1: Relative position of the 26 news outlets who had relevant tweets for at least four questions per dimension (the 26ers): 1. adnradiochile, 2. biobio, 3. cooperativa, 4. latercera, 5. mercuriovalpo, 6. publicmetrochile, 7. emol, 8. soyarauco, 9. soyconcepcion, 10. soyconel, 11. soyquillota, 12. soysanantonio, 13. soyaltcahuano, 14. soytome, 15. dfnanciero, 16. el_ciudadano, 17. elmostrador, 18. tele13_radio, 19. el_dinamo, 20. nacioncl, 21. pinguinodiaro, 22. soyhillan, 23. soycopiapo, 24. soyvaldivial, 25. soyvalparaiso, 26. t13

party. The ruling coalition is “Nueva Mayoría”, which mainly consists of center-left to left-wing parties, but with a strong component of Christian democracy. Christian democracy is still a center-left party, but conservative, specially in personal issues. So, in this case, the observed bias has a similar tendency to the political alignment of the ruling coalition. This result is also in correspondence with the theory postulated in the *Propaganda Model* [16]. On the personal issues dimension, we can also see some bias, although less prominent, tending to the conservative end of the spectrum.

We show in Figure 1 the 26 news outlets for which we were able to answer at least four questions on each dimension (we will call this subset of news outlets the **26ers**). This represents 10% of our database and 13% of those that regularly report on economics and politics. Even though we are showing relative values with respect to the other outlets, we show them on the original Nolan’s chart, to give an intuition on their political tendencies. We tagged some of the most prominent ones to help to understand the landscape.

In the rightmost cell, we have *mercuriovalpo*¹⁵ (*El Mercurio de Valparaíso*), one of the oldest newspapers in Chile currently in circulation. This newspaper is part of a big conglomerate (*El Mercurio S.A.P*) that owns more than 20 news papers and several radio stations, among other broadcast media (such as magazines, TV cable, etc.). The regional newspaper *Soy Coronel* (*soyconel*), on the bottom, is also part of this group. In fact, 11 regional newspapers of *El Mercurio S.A.P* are within this 26 and are all clustered bottom-right, with scores under 31 on personal issues. As we mentioned earlier,

¹⁴The QQplot and the histogram suggested normality was a reasonable assumption.

¹⁵Tags in Figure 1 are the corresponding Twitter accounts (e.g. <https://twitter.com/mercuriovalpo>)

the *El Mercurio*'s newspapers are popularly perceived as right-wing conservative, which coincides with the scores they received.

La Tercera (*latercera*), is owned by *Copesa S.A.*, which is *El Mercurio*'s closest competitor. These two companies have a so-called news media duopoly. *La Tercera*, also in the lower-right quadrant but closer to the center of the chart, is thought to be moderate-conservative¹⁶. *El Mostrador* (*elmostrador*) is an on-line newspaper with a perceived orientation to progressivism¹⁷, which corresponds with its position in the chart.

Finally, *La Nación* (*nacioncl*) is a newspaper that currently only publishes its online edition and is partially controlled by the government. This newspaper appears in the topmost region of the *personal* dimension. Compared to the other 25 news outlets, this one appears as the most progressive on personal issues. This score is due to a series of populist reforms promoted by the government during the observed period (i.e. therapeutic marijuana legalization, decriminalization of abortion, anti-xenophobic campaigns, promote voluntary enlistment of women to the military service, etc.)

To summarize, using the PolQuiz we were able to position Chilean newspapers on a chart with respect to their bias with regard to economic and personal issues. In the next sections we will investigate the nature of the bias and to what extent these positions are in line with common perceptions. However, before doing so, we first need to know how stable the results are.

4.1.1 Stability of the results. In order to find out the stability of the observed bias with respect to changes in the obtained evidence (i.e. the collected tweets), we repeated the scoring steps 20 times. Each time we leave out 5% of the tweets selected at random, while maintaining the original distribution of documents per question. Each time, we measure the average score of the news outlets for which we were able to answer at least one question in the corresponding dimension. In the *economic* issues, we could observe a consistent bias to the left ($M = 40.45$, 95% CI[36.91, 43.99]). On the other hand, the *personal* dimension, although it is also leaning to one side, is much closer to the center of the spectrum ($M = 46.89$, 95% CI[43.99, 49.79]). Figure 2 shows a similar analysis, but at an individual level in the *26ers*. The mean for each individual score stays close to its original position, and each newspaper can be located in a relatively small neighborhood with high confidence, meaning that there are not any drastic changes in the previous classification.

The relatively low impact of leaving out data in the positioning process indicates that the results are not very sensitive to change and not influenced by only a small number of tweets.

4.1.2 Contextualizing the PolQuiz. We noticed that some of our queries, particularly in the *personal* issues dimension, returned only a small number of documents (e.i. **q2** and **q3**). This is because of lack of interest or too few relevant events related to the corresponding topics during the observed period. We think that a way to counteract this environmental/circumstantial effect is to substitute the respective questions or to increase the number of questions.

We repeated our analysis using **q0** as a replacement for question **q3** (related to laws concerning sex between consenting adult, see

¹⁶https://en.wikipedia.org/wiki/La_Tercera

¹⁷https://es.wikipedia.org/wiki/El_Mostrador

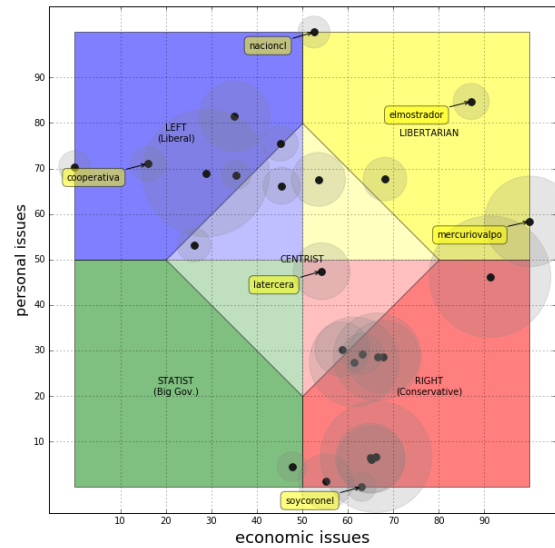


Figure 2: Relative position of the *26ers*. The score on each dimension is the average over 20 repetitions, leaving out each time a random 5% of the documents. Gray shade around outlets is its 95% confidence interval.

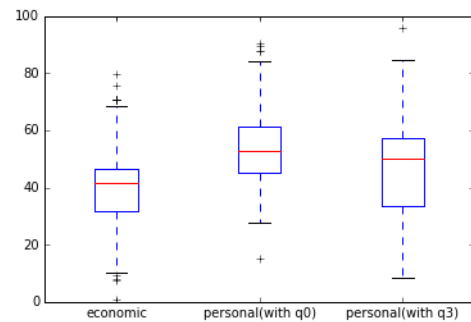


Figure 3: New scores after replacing **q3 by **q0**.** These are the scores of news outlets for which we were able to answer at least one question in the corresponding dimension.

Section 3.2). We replaced **q3**, because it was the one with the lowest number of retrieved documents. This substitution increased the number of news outlets with at least one answer. There is now a stronger statistical effect for the *personal* issues dimension, $t(239) = 3.54$, $p < .001$. Interestingly, this dimension now leans to the more liberal end of the spectrum ($M = 53.57$, $SD = 15.63$) (see Figure 3).

In Figure 4 we plot the scores of the *26ers* in the original quiz (dots) and the adapted quiz (diamonds). Note that the difference in scores between the quiz with **q3** and the quiz with **q0** is considerably larger (with a negative difference) for outlets in the right/conservative quadrant. This is expected and validates the model.

4.2 Investigating the nature of bias using rank difference

The PolQuiz showed the existence of bias in Chilean media. In this section, we investigate the nature of this bias in terms of vocabulary used and entities mentioned in the different newspapers' tweets (see Section 3.4). We focused on the *26ers* and the topic of abortion.

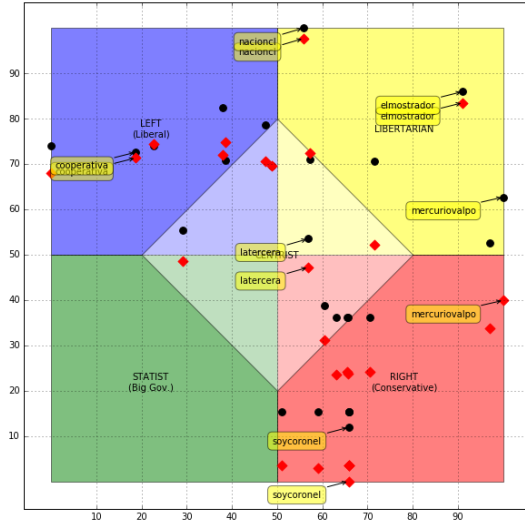


Figure 4: Relative position of the 26ers. Dots represent the scores with q3. Diamonds represent the scores with q0.

4.2.1 Topic bias based on named entities. We used the Stanford’s NE recognizer system [9] to extract the entities mentioned in the tweets related to the abortion issue. We compare the extracted entities against a list of politicians, public personalities and activist groups. For the list of politicians and their position in the abortion issue, we use the vote sessions in the house of representatives¹⁸ and in the senate¹⁹. We manually labeled another 53 personalities and groups according to comments and events reported in the local news. The complete list L_E has 199 labeled entities. We labeled with -1 the politicians who voted against the abortion bill, and the public figures that were openly against the issue. Equivalently, we use $+1$ for politicians and personalities in favor of the subject. We assign a 0 to the entities not included in our list. We will refer to these labels as the leaning of the entities (e.g. $leaning(entity)$)

After applying the rank difference method to the NE mention counts, we calculated a score for each outlet in function of the $\tau(entity)$ and the leaning of $entity$ in the issue (for every $entity$ mentioned more than once in the news). This final score of the outlet o_i is found using the equation 2.

$$score(o_i) = \sum_{e \in L_E} (\tau(e) * leaning(e)) / size_of(L_E) \quad (2)$$

A low value in this score indicates that this outlet tends to mention with relatively high frequency entities with a conservative leaning and/or it tends to ignore those with a more liberal view.

As expected, the outlets tagged as independent, libertarians and classical liberalism are within the higher values. According to our scores, this means they have comparably more mentions of entities with a liberal leaning than the rest of the outlets. Interestingly, the lower values are occupied by the outlets linked to parties in the ruling coalition (Christian democracy and Left-Liberal(*nacioncl*)). Apparently these outlets focus their tweets in negative reports of the opposition. For example, when we look at the rank-difference

results for *nacioncl*, within the top-20 entities, only two refer to entities with a liberal leaning (‘President Michelle Bachelet’ and ‘Government’). To investigate more on this, we run a sentiment analysis on the most used bi-grams. The results are presented in the next section.

4.2.2 Topic bias based on bi-grams. We again apply the rank difference method, this time using the bi-gram counts in the tweets relevant to the subject of abortion. Following the same strategy as before, we calculated a score for each outlet in function of the $\tau(bigram)$ and the sentiment calculated for $bigram$ (for every $bi-gram$ mentioned more than once in the news). For determining the sentiment of words and bi-grams we use the Spanish lexicon from [25]. To calculate $\tau(bigram)$ we use a formula equivalent to that shown in Equation 2. Accordingly, we give a similar interpretation to these scores. That is, a high value indicates that this outlet tends to convey mostly positive sentiments with the bi-grams used with relatively high frequency and/or avoid using negative sentiments when referring to the issue of abortion. For example, *elmostrador*, with the highest score, has as a frequently use bi-gram “proyecto aprobado” (tr. “project approved” - referring to the bill). This bi-gram is classify as positive by the sentiment analyzer, so it will add to the score. On the other hand, this same outlet has “injusticia gobierno” (tr. “government injustice”) as a totally ignored bi-gram. Since the bi-gram is assigned a negative sentiment and the rank-difference is also negative, the bi-gram will also add to the score of the outlet, pushing it to the liberal side. Following the same reasoning, an outlet with a very low score can be understood as an outlet that uses predominantly negative words with relatively high frequency.

When we analyze the scores of the 26ers, we notice that *nacioncl* (controlled by the government) has the lowest score. This, together with the previous NE analysis, confirms the theory that this outlet focuses in tweeting negative reports of the opposition, at least for the abortion issue. Most of the others outlets show the expected behavior, with conservative in the lower half of the ranking (i.e. lower scores) and liberals in the higher positions.

4.3 Survey

For the survey described in Section 3.5, we collected 372 answers from 54 unique Chilean²⁰ users on how they perceive the bias on the topic of abortion in the different Chilean newspapers. We received between 11 and 19 answers for each of the 26ers (M: 14.31, SD: 2.07). We carried out 10 Fleiss’ kappa measurements; each time we selected 10 ratings at random per outlet (subject). This shows a fair agreement in the answers (M: 0.2253, SD: 0.0167). In Table 4 we show the 26ers and their corresponding “Perceived bias” (see Section 3.5). The political alignment information shown in the table was extracted from Wikipedia²¹, the official web site of the news outlet or the political alignment known for the owners. Note that the Christian Democracy party is part of the center-left coalition that was ruling in Chile during the observed period, so is generally in favor of the social changes promoted by the government.

¹⁸https://www.camara.cl/trabajamos/sala_votacion_detalle.aspx?prmID=23099

¹⁹<http://www.senado.cl/appsenado/index.php?mo=sesionessala&ac=detalleVotacion&votaid=6668>

²⁰The IP addresses indicate we have representation of different regions of the country

²¹Since Wikipedia pages are a crowdsourced content, we consider the political alignment extracted from there as either self-declared or a popular perception

Table 4: Results from popular survey for the 26ers.

Id	Name	Owner	Political alignment	Perceived bias	PolQuiz Personal issues
21	pinguinodiario	Patagónica Publicaciones	—	-66.67	39.18
24	soyvaldiviacl	El Mercurio	Right-wing, conservative	-66.67	-50.49
22	soychillan	El Mercurio	Right-wing, conservative	-57.14	-50.55
25	soyvalparaiso	El Mercurio	Right-wing, conservative	-43.75	-51.81
8	soyarauco	El Mercurio	Right-wing, conservative	-42.86	-51.27
12	soysanantonio	El Mercurio	Right-wing, conservative	-30.77	-92.98
13	soyतालcahuano	El Mercurio	Right-wing, conservative	-30.77	-92.98
18	tele13_radio	Grupo Luksic & PUC	—	-28.57	52.42
9	soyconcepcion	El Mercurio	Right-wing, conservative	-25.00	-94.09
14	soytome	El Mercurio	Right-wing, conservative	-25.00	-92.92
7	emol	El Mercurio	Right-wing, conservative	-25.00	-0.59
10	soycoronel	El Mercurio	Right-wing, conservative	-23.53	-100
11	soyquillota	El Mercurio	Right-wing, conservative	-18.18	-92.92
15	dfinanciero	Grupo Claro	Right-wing, conservative	0.00	42.57
5	mercuriovalpo	El Mercurio	Right-wing, conservative	21.43	-51.81
2	biobio	BÑo-BÑo Comunicaciones	Independent	23.53	6.91
6	publimetrochile	Grupo metro	International	25.00	47.50
17	elmostrador	La Plaza	Libertarian	26.32	70.95
19	el_dinamo	Ediciones Giro Pais	Christian democracy	29.41	-30.79
4	latercera	Copesa	Classical liberalism	33.33	-3.38
1	adnradiochile	Grupo Prisa	International	37.50	52.98
16	el_ciudadano	Red de medios de los pueblos	Libertarian	37.50	44.54
23	soycopiapo	El Mercurio	Right-wing, conservative	38.46	-36.21
3	cooperativa	Co. Chilena de Comunicaciones	Christian democracy	57.14	46.04
26	t13	Grupo Luksic & PUC	—	57.14	48.45
20	nacioncl	Estado de Chile	Left, Liberal	63.64	100

The list is sorted by the perceived bias. Outlets with an unclear Political Alignment (shadowed rows in the table) were left out of the analysis.

Results show that there is a perceivable difference in the language used by the outlets in both sides of the spectrum. Note that, based on the rank difference of bi-grams, the users were able to collectively classify the outlets with over 90% precision²². Our positioning of these outlets in the adapted *PolQuiz* has also a high correspondence in the bias direction (90%) with the political alignment and a good agreement with the direction of the Perceived bias (80%).

To evaluate the relative positions of the outlets in our *PolQuiz*, we calculated the number of inversions with respect to the ranking of the outlets in the perceived bias. The Kendall's Tau-b coefficient between the two rankings is $\tau_b(21) = 0.4203$ ($z = 2.66, p < .01$). Even though the popular perception resulting from the survey can not be seen as ground-truth for the relative positioning of the outlets, it is important to notice that our results show a good correlation with the intuition of the public. As a future work, we aim to add some other content features (e.g. leaning of the named entities) to the polarity classification of the tweets as these may help to refine the relative positioning found by our model.

To summarize, we have shown that reported political alignment is highly correlated with the *PolQuiz* results as well as with the bias, as perceived by the general audience. This implies that existing bias has a noticeable influence on how controversial issues such as abortion are reported in the media.

5 CONCLUSIONS

In this paper, we presented an automatic approach for estimating the political bias of news outlets in Chile, exploiting the well-known and widely used “The World’s Smallest Political Quiz”. We empirically confirmed the estimation results and showed that they are stable with respect to evolving data. Furthermore, we have demonstrated the benefits of adapting questions to the local context. Building upon the *PolQuiz* results, we investigated the nature of this political bias and found this to exist in the chosen vocabulary and the entities covered by the newspaper. Finally, we conducted a survey, of which the results confirm that political bias in newspapers has an impact on how controversial topics are covered and that the general audience does notice this bias. Our methodology does not make too many assumptions about the underlying system. The way it is designed could be applied to any Western culture. Our system can deal with any number of outlets, can compare relative quantitative positions, can show empirical evidence of consistent bias, and can partially explain the source of these tendencies.

As a future work we are interested to see what is the most accurate way to score the missing answers. Since “coverage” is a form of bias [24], perhaps the outlet is not being *neutral* by not mentioning a specific subject. Even when the decision of which stories/events are newsworthy is subjective and depends on the editorial strategy [29], there are some events that are very relevant in the national context and are covered for the majority of the

²²We are not taking into account those for which we could not find a political alignment or those that belong to international groups

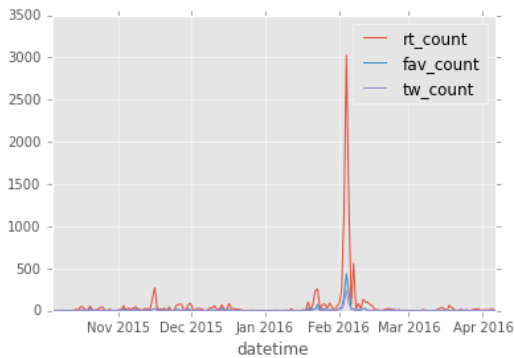


Figure 5: Counting tweets and re-tweets for q7.

media. So, a complete silence of a news outlet on such an event may be interpreted as something other than neutrality.

For example, question q7 is about international free trade. Taking the number of tweets and re-tweet as an indicator of important events [20], we can see in Figure 5 that this topic has had at least one major event during this period. This event was the adscription of Chile to the Trans-Pacific Partnership (TPP) signed by the country on Feb 3th, 2016. Despite the magnitude of the event, only 135 out of 198 newspapers with a section on politics mentioned it. A plausible cause is that the other news outlets decided not to report about this event, in other words ‘bias by omission’.

We show that a careful selection or update of the questions may lead to a significant improvement in the results. If we have an inside understanding of the socio-economic environment from where the news are being collected, then we could replace the questions to capture more relevant topics. In this sense, we could benefit from advances in systems that focus on identifying controversial topics in social media [10]. On the other hand, if we do not have any intuition on the news collected, then we can accumulate the new questions so we can widen the spectrum of topics and have a better chance of capturing relevant events/discussions with our queries.

For individuals as well as for society as a whole it is important to recognize and understand media bias that are shaped through underlying general political or socio-economic orientations. As we have shown in this paper, these general tendencies have a clear and noticeable effect on the way concrete topics are covered and commented upon, and therefore should be investigated and published.

REFERENCES

- [1] Scott L. Althaus, Jill A. Edy, and Patricia F. Phalen. 2001. Using Substitutes for Full-Text News Stories in Content Analysis: Which Text Is Best? *American Journal of Political Science* 45, 3 (2001), 707–723. <http://www.jstor.org/stable/2669247>
- [2] Jisun An, Meeyoung Cha, Krishna Gummadi, and Jon Crowcroft. 2011. Media Landscape in Twitter: A World of New Conventions and Political Diversity. In *Proceedings of the Fifth International Conference on Weblogs and Social Media*. AAAI, Menlo Park, CA, USA.
- [3] Johan Bollen, Alberto Pepe, and Huina Mao. 2009. Modeling public mood and emotion: Twitter sentiment and socio-economic phenomena. *CoRR* abs/0911.1583 (2009). <http://arxiv.org/abs/0911.1583>
- [4] Maurice Bryson and William McDill. 1968. The political spectrum: A bi-dimensional approach. *Rampart Journal of Individualist Thought* 4, 2 (1968), 19–26.
- [5] Raviv Cohen and Derek Ruths. 2013. Classifying Political Orientation on Twitter: It’s Not Easy!. In *JCWSM*, Emre Kiciman, Nicole B. Ellison, Bernie Hogan, Paul Resnick, and Ian Soboroff (Eds.). The AAAI Press. <http://dblp.uni-trier.de/db/conf/jcwsml/jcwsml2013.html#CohenR13>
- [6] M Conover, B Gonçalves, J Ratkiewicz, A Flammini, and F Menczer. Predicting the Political Alignment of Twitter Users. In *Proceedings of 3rd IEEE Conference on Social Computing (SocialCom)*.
- [7] Alexander Dallmann, Florian Lemmerich, Daniel Zoller, and Andreas Hotho. 2015. Media Bias in German Online Newspapers. In *Proceedings of the 26th ACM Conference on Hypertext & Social Media (HT ’15)*. ACM, 133–137.
- [8] Daniel Dor. 2003. On newspaper headlines as relevance optimizers. *Journal of Pragmatics* 35, 5 (2003), 695 – 721. DOI : [http://dx.doi.org/10.1016/S0378-2166\(02\)00134-0](http://dx.doi.org/10.1016/S0378-2166(02)00134-0)
- [9] Jenny Rose Finkel, Trond Grenager, and Christopher Manning. 2005. Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics (ACL ’05)*. 363–370.
- [10] Kiran Garimella, Gianmarco De Francisci Morales, Aristides Gionis, and Michael Mathioudakis. 2015. Quantifying Controversy in Social Media. *CoRR* abs/1507.05224 (2015). <http://arxiv.org/abs/1507.05224>
- [11] Venkata Rama Kiran Garimella and Ingmar Weber. 2014. Co-following on Twitter. In *Proceedings of the 25th ACM Conference on Hypertext and Social Media (HT ’14)*. 249–254.
- [12] Matthew Gentzkow and Jesse M. Shapiro. 2010. What Drives Media Slant? Evidence From U.S. Daily Newspapers. *Econometrica* 78, 1 (2010), 35–71.
- [13] Pierre Geurts, Damien Ernst, and Louis Wehenkel. 2006. Extremely randomized trees. *Machine Learning* 63, 1 (2006), 3–42.
- [14] Jennifer Golbeck and Derek Hansen. 2011. Computing Political Preference Among Twitter Followers. In *Proceedings of the SIGCHI Conference on Human Factors in Computing Systems (CHI ’11)*. ACM, 1105–1108.
- [15] Y. Gu, T. Chen, Y. Sun, and B. Wang. 2016. Ideology Detection for Twitter Users with Heterogeneous Types of Links. *ArXiv e-prints* (Dec. 2016). arXiv:1612.08207
- [16] E.S. Herman and N. Chomsky. 1988. *Manufacturing consent*. Pantheon Books. <https://books.google.cl/books?id=Up5sAAAIAAJ>
- [17] Chunyu Kit and Xiaoyue Liu. 2008. Measuring mono-word termhood by rank difference via corpus comparison. *Terminology. International Journal of Theoretical and Applied Issues in Specialized Communication* 14, 2 (2008), 204–229.
- [18] Haokai Lu, James Caverlee, and Wei Niu. 2015. BiasWatch: A Lightweight System for Discovering and Tracking Topic-Sensitive Opinion Bias in Social Media. In *Proceedings of the 24th ACM International Conference on Information and Knowledge Management*. 213–222.
- [19] Aibek Makazhanov and Davood Rafiei. 2013. Predicting Political Preference of Twitter Users. In *Proceedings of the 2013 IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM ’13)*. ACM, 298–305.
- [20] Felix Ming, Fai Wong, Chee Wei Tan, Soumya Sen, and Mung Chiang. 2013. Quantifying Political Learning from Tweets and Retweets. In *Proceedings of the 7th International AAAI Conference on Web and Social Media*. AAAI Press, Boston, Massachusetts, USA.
- [21] Delia Moccanu, Andrea Baronchelli, Nicola Perra, Bruno Gonçalves, Qian Zhang, and Alessandro Vespignani. 2013. The Twitter of Babel: Mapping World Languages through Microblogging Platforms. *PLoS ONE* 8, 4 (04 2013), e61981.
- [22] DF Nolan. 1971. Classifying and Analysing Politico-Economic Systems. *The Individualist* (Jan. 1971), 5–11.
- [23] Sounel Park, Seungwoo Kang, Sangyoung Chung, and Junehwa Song. 2012. A Computational Framework for Media Bias Mitigation. *ACM Trans. Interact. Intell. Syst.* 2, 2, Article 8 (2012), 32 pages.
- [24] Diego Saez-Trumper, Carlos Castillo, and Mounia Lalmas. 2013. Social Media News Communities: Gatekeeping, Coverage, and Statement Bias. In *Proceedings of the 22nd ACM International Conference on Information & Knowledge Management (CIKM ’13)*. ACM, 1679–1684.
- [25] Hans Stadthagen-Gonzalez, Constance Imbault, Miguel A. Pérez Sánchez, and Marc Brysbaert. 2016. Norms of valence and arousal for 14,031 Spanish words. *Behavior Research Methods* (2016), 1–13.
- [26] Saatviga Sudhahar, Thomas Lansdall-Welfare, Ilias Flaounas, and Nello Cristianini. 2012. ElectionWatch: Detecting Patterns in News Coverage of US Elections. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics (EACL ’12)*. 82–86.
- [27] The Advocates for Self Government. 1985. The World’s Smallest Political Quiz. (1985). [Online <http://www.theadvocates.org/quiz/quiz.php>; 12-August-2016].
- [28] Giang Tran, Mohammad Alrifai, and Elco Herder. 2015. *Timeline Summarization from Relevant Headlines*. Springer International Publishing, Cham, 245–256. DOI : http://dx.doi.org/10.1007/978-3-319-16354-3_26
- [29] Giang Binh Tran and Elco Herder. 2015. Detecting Filter Bubbles in Ongoing News Stories. In *UMAP Extended Proceedings*.
- [30] Justine Zhang Cristian Danescu-Niculescu-Mizil Jure Leskovec Vlad Niculae, Caroline Suen. 2015. QUOTUS: The Structure of Political Media Coverage as Revealed by Quoting Patterns. In *Proceedings of WWW 2015*.
- [31] Daniel Xiaodan Zhou, Paul Resnick, and Qiaozhu Mei. 2011. Classifying the Political Learning of News Articles and Users from User Votes. In *Proceedings of the International AAAI Conference on Web and Social Media*.