

## Assessing Media Bias in Cross-Linguistic and Cross-National Populations

Allan Sales,<sup>1</sup> Albin Zehe,<sup>2</sup> Leandro Balby Marinho,<sup>1</sup> Adriano Veloso,<sup>3</sup> Andreas Hotho<sup>2</sup> and Janna Omeljanenko<sup>2</sup>

<sup>1</sup> Federal University of Campina Grande

<sup>2</sup> University of Würzburg

<sup>3</sup> Federal University of Minas Gerais

allanmelo@copin.ufcg.edu.br, zehe@informatik.uni-wuerzburg.de, lbmarinho@computacao.ufcg.edu.br,  
adrianov@dcc.ufmg.br, hotho@informatik.uni-wuerzburg.de, janna.omeljanenko@uni-wuerzburg.de

### Abstract

Media bias is a worldwide concern. Although automated methods exist for the analysis of various forms of media bias, language is still an important barrier toward spotting worldwide differences in reporting. In this paper, we propose a methodology based on word embeddings, lexicon translation, and document similarity to assess media bias in news articles published in different idioms. We model media bias under the perspective of subjective language use, i.e., the more subjective the content of a news article is, the more biased it is. Our core assumption is that news articles reporting the same events, but written in different languages, should have similar levels of subjectivity; otherwise, we may have spotted biased text. Our method consists of using translated versions of subjectivity lexicons that were originally constructed for measuring subjectivity in the Brazilian Portuguese language. We evaluate our approach on two labeled data sets to show that our method is valid and apply our methodology to analyze recent and largely resounded topics, such as the Venezuela crisis and Syrian war, on four distinct idioms: Portuguese, German, English, and Spanish.

### Introduction

In his bestseller book, “Bias: A CBS Insider Exposes How the Media Distort the News” (Goldberg 2014), Bernard Goldberg exposes how the news media industry ignored a fundamental premise in journalism: providing objective and disinterested reports. A key message from his book is that, in many cases, media is intentionally reporting facts charged with biased opinion. This may exert great influence on readers whose own biases may be reinforced or shaped. While media bias analysis has a long tradition in social sciences and communication, only in recent years it has attracted interest from the computer science and computational linguistics communities.

The literature distinguishes among different types of media bias, among which the most usual are *statement* and *framing*. The Statement Bias is the preference for expressing oneself more (or less) favourable about a certain subject (e.g., party and politician) (Saez-Trumper, Castillo, and Lalmas 2013). The *Framing effect* (or framing bias) is related

to *how* the information is conveyed to the readers, in order to influence their judgment on a given topic (Entman 2007).

Although there is a broad and interdisciplinary body of work on media bias analysis, these works are still limited in their cross-language and country compatibility. Each country has its own news outlets producing news charged with the shared qualities that define its population, such as language (or dialect) and geography. While most of the works found in the reviewed literature focus on specific countries or languages, spotting news reporting differences across different countries and languages remains an open research question.

In this paper, we introduce a new methodology based on lexicon translation, word embeddings, document similarity, and a parallel corpus for large-scale assessment of media bias in cross-language populations. We characterize media bias in terms of subjectivity bias. In order to perform frame analysis on news articles, two broad questions are usually asked (Entman 1993): (1) *What* information is conveyed? (2) *How* is that information conveyed? These questions define a *frame* (Hamborg, Donnay, and Gipp 2018), which is our unit of study. In this paper, we fix (1) to recent and largely resounded topics worldwide: the Venezuela crisis and the Syrian war. Next, we propose a new subjectivity measure to answer (2). This measure addresses the challenge of calculating comparable subjectivity scores across news written in different languages.

Our methodology relies on the translation of handcrafted subjectivity lexicons, initially constructed for the Brazilian Portuguese language, to other languages (i.e., German, English, and Spanish). The fact that subjectivity lexicons are a set of independent words that do not compose meaning enables their discrete translation. Our methodology relies on two main premises: (i) the unprecedented accuracy of current machine translators and (ii) word embeddings that can mitigate eventual translation imprecision. That is, even when the expressions in the lexicons and target textual documents do not match in a syntactic level, they shall be close to each other in the semantic space induced by the embedding.

A variety of factors might be associated with the language subjectivity, such as the structure of the linguistics (Kristiansen, Garrett, and Coupland 2005). Therefore, one issue that we need to consider in the construction of a score that is comparable across languages is that some languages may be inherently more subjective than others, and thus their subjectivity

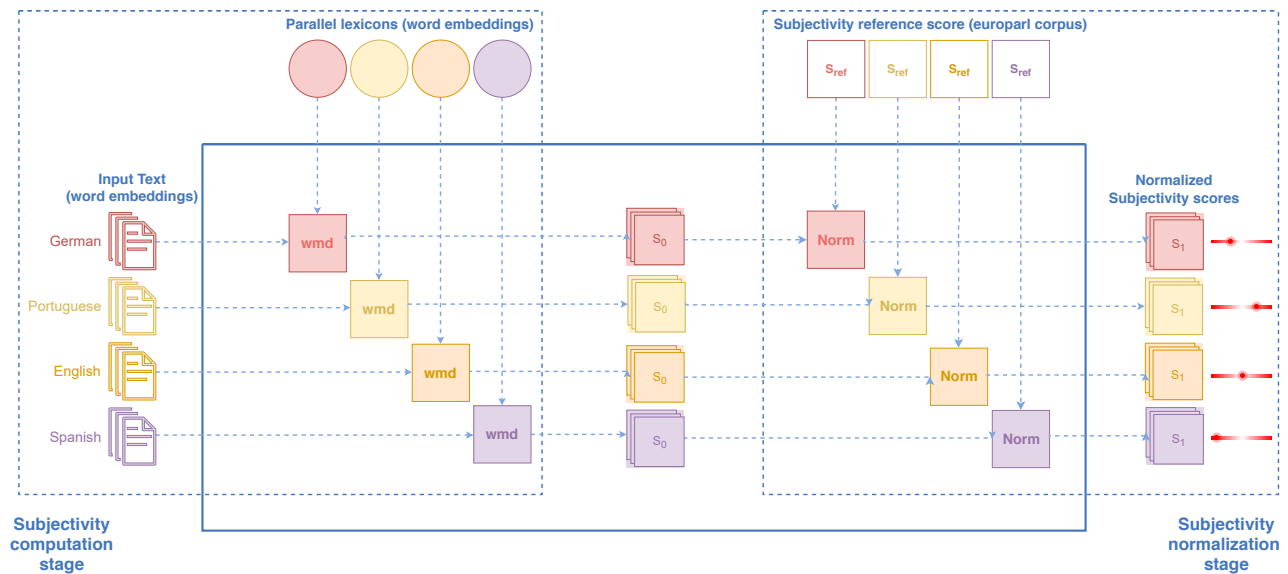


Figure 1: Task description. Given distinct sets of news articles and subjectivity lexicons in different languages, we compute the news articles’ *subjectivity bias* as the Word Mover’s Distance to their respective language subjectivity lexicon (Subjectivity computation stage). In order to remove the base subjectivity of the language, we calculate a *normalized subjectivity bias* by subtracting the values from a pre-computed subjectivity reference value of their language derived from the parallel Europarl corpus (Subjectivity Normalization Stage).

tivity biases may not be directly comparable. For overcoming that, we propose to model subjectivity bias as a combination of the interlocutor (e.g., writer or speaker) subjectivity and the language subjectivity (i.e., the level of subjectivity that is inherent to the language). We rely on a parallel corpus<sup>1</sup> from where we extract reference subjectivity bias values (or normalization factors) for each target language. Figure 1 summarizes our methodology (see Section for more details). Given a set of topic-wise similar news articles written in distinct languages, we compute their distances to their respective subjectivity lexicons in a word embedding space and normalize the results considering a precomputed subjectivity reference value from the parallel corpus. Our main contributions can be summarized as follows:

- We introduce a new media bias metric that disentangles language subjectivity from interlocutor subjectivity. This enables us to compute and compare media bias in cross-linguistic and cross-national populations.
- We conduct a thorough validation of our methodology, showing that: (i) it is able to correctly distinguish between objective and subjective text, and (ii) it is comparable to an approach that translates the whole text in another idiom to Portuguese (instead of translating the lexicons in Portuguese to the target idiom as we propose to do).
- We apply our method to analyze news in various idioms about two largely resounded topics worldwide: the Venezuela crisis and Syrian war. This leads to interesting insights about the level of bias that different countries ex-

hibit in their news articles about these topics (e.g., News articles published in Brazil addressing the Venezuela Crisis are more subjective than the ones addressing the Syrian War).

### Background: Subjectivity and Bias

In order to provide a clear basis for our analysis, this section details our understanding of what makes a text subjective and introduces the terminology that we use in this paper.

The Normative Theory of Journalism is concerned with what the media should do in society (Benson 2008). According to (Siebert et al. 1956) in their book *Four Theories of the Press*, “the press takes on the form and coloration of the social and political structures within which it operates” (pp.1-2). Given that we use news articles written in languages spoken mostly by European and American countries, we are probably transiting between two kinds of Normative theories: *Free press theory* and *Social responsibility theory*. Still according to (Siebert et al. 1956), the free press theory (mostly adopted in the US), states complete freedom of speech and economic operation of the media, dismissing any interference of the government. The social responsibility theory (mostly adopted in European countries), in turn, is similar to the free press theory but places greater emphasis on the accountability of the media.

Although there may be differences in the dominant ideas about the obligations of mass media in different societies, media organizations of different nations typically share many common ethical standards including the principles of truthfulness, accuracy, **objectivity**, impartiality, fairness, and public accountability. We propose an automatic method-

<sup>1</sup> A parallel corpus is a corpus that contains a collection of texts in one language and their translations into a set of other languages.

ology for shedding light on these differences concerning the objectivity principle.

**Dimensions of Subjectivity** Before delving into our proposed methodology, we first need to define what we regard as subjective language. Similar to (Amorim, Cançado, and Veloso 2018; Sales, Balby, and Veloso 2019), we study subjectivity under five *subjectivity dimensions*: argumentation, presupposition, sentiment, valuation, and modalization. These lexicons were constructed anchored in the pragmatics theory and (Recasens, Danescu-Niculescu-Mizil, and Jurafsky 2013). According to (Verhagen 2005), the theory of argumentation and pragmatics assume implicit markers of positioning in language as clues to the speakers subjectivity. In the following, we briefly define each dimension and give examples of expressions that characterize it.

- **Argumentation** aims to identify argumentative discourse in text, which might be indicative of an attempt to convince the reader of a specific point of view. E.g., “even” and “by the way”;
- **Presupposition** contains expressions related to a prior assumption of something’s veracity. This kind of discourse is characterized by the interlocutor’s veracity assumption of something before the current event he/she is mentioning. E.g., “nowadays” and “admit”;
- **Sentiment** implies appealing to an emotional discourse, not necessarily positive or negative. E.g., “fall in love” and “fear”;
- **Valuation** is related to giving value or intensifying something. E.g.: “a lot”, “better” and “big”;
- **Modalization** shows that the author takes an attitude towards his/her judgment. E.g., “undeniable” and “undoubted”.

**Subjectivity Bias** Amorim et al. (Amorim, Cançado, and Veloso 2018), together with Brazilian Portuguese linguists, built lexicons for each of the dimensions mentioned above in the Brazilian Portuguese language. These enable us to use several ways for measuring subjectivity of textual documents written in Portuguese, e.g., counting the occurrences of the words of each lexicon in the document, using tf-idf like scores, or calculating similarities between texts and subjectivity lexicons in some vector space. In the latter case, the closer the documents and lexicons are, the more subjective the documents are. We refer to the calculated similarities using lexicons as *subjectivity bias*.

We consider the subjectivity bias to be composed of two parts: the *language subjectivity* (caused by the language) and the *interlocutor subjectivity* (caused by the author). If we want to compare the subjectivity of news articles in different languages, we therefore need to do a normalization step to remove the language subjectivity and only compare the interlocutor subjectivity.

## Related Work

There is a large body of work in media bias analysis. While communications and social sciences have a long tradition in

this area, it has been attracting a lot of attention from the computer science and computational linguistics communities. Hamborg et al. (Hamborg, Donnay, and Gipp 2018) put together a thorough review of media bias analysis, covering the existing literature and also establishing synergy points between social and computer sciences.

On the side of news analysis, there is extensive literature on distinct media bias types, such as journalistic biases (e.g., selection and coverage), confirmation/statement bias (Lazaridou, Krestel, and Naumann 2017; Saez-Trumper, Castillo, and Lalmas 2013; Lin, Bagrow, and Lazer 2011; Nickerson 1998; Dallmann et al. 2015), psychological/cognitive biases (Recasens, Danescu-Niculescu-Mizil, and Jurafsky 2013; Caliskan, Bryson, and Narayanan 2017), and subjectivity bias (Sales, Balby, and Veloso 2019; Mihalcea, Banea, and Wiebe 2007; Chaturvedi et al. 2018).

Most of the works regarding framing bias detection are based on sentiment analysis, which aims to identify an author’s opinion toward some entity mentioned in the text (Oelke, Geisselmann, and Keim 2012; Mundim 2018; De Cock et al. 2018). However, state-of-the-art sentiment analysis methods on news articles still perform poorly, given the lack of clearly defined targets and large-scale news annotated data sets for training machine learning methods (Bahur et al. 2013; Hamborg, Donnay, and Gipp 2018).

More recently, some researchers have proposed new alternatives for assessing framing bias apart from sentiment analysis. The authors of (Card et al. 2016), for example, trained a logistic classifier using unigrams, bigrams, and personas (e.g., characterizations of entities) as features for identifying the primary frame (the emphasized dominant aspect) of news articles. The classifier received, as input, manually annotated news articles about immigration with 15 categories of framing bias, such as Economic and Politics. The authors of (Bai et al. 2018), in turn, introduced an improvement of the Matrix Factorization Method, where they apply a jointed penalty function for detecting whether frames over illegal immigration change over time. Differently from these works, we rely on an unsupervised approach given the difficulty in finding and/or assembling annotated biased data.

Sales et al. (Sales, Balby, and Veloso 2019) propose a holistic approach for analyzing the textual content of news stories about three consecutive Brazilian Presidential elections in search of three kinds of bias: coverage, association, and subjectivity. For computing subjectivity bias, they applied the Word Mover’s Distance (WMD) (Kusner et al. 2015) between each news article of interest and each of five subjectivity lexicons containing expressions associated with different subjectivity dimensions. A limitation of this approach is that it is not applicable across multiple languages. In this paper, we address this issue by proposing a new methodology that yields comparable scores for texts in different languages, enabling us to analyze the difference in subjectivity across multiple countries and languages. To this end, we (1) point out some constraints that need to be satisfied when translating the lexicons, (2) provide lexicon translations that satisfy these constraints for three additional languages, and (3) introduce a normalization stage ensuring the comparability of the scores across multiple languages.

Some works have proposed to detect subjective language in a multilingual setting (Chaturvedi et al. 2015, 2018). However, subjectivity in those works is related to the task of classifying a news article as either neutral or opinionated (e.g., positive or negative). In our work, subjectivity is defined in a broader and more transparent sense that goes way beyond positive or negative polarities.

The related works mentioned above reveal several solutions available to expose media bias. Differently from us, most of the works in this area are focused on single languages, and the ones that deal with multilingual settings, such as (Chaturvedi et al. 2015), only consider English and Spanish. With our contributions, we expect to pave the way to more effective, fast, and transparent methods for the automated detection of media bias across different languages.

## Datasets

This section describes the data collection process as well as the kind of data that was collected. It includes two data sets of news (one where we apply our methodology and other for the validation), one data set annotated with subjective/objective labels (which we have used to validate our methodology), Wikipedia corpora in all languages considered (which we use to train word embeddings and validate our methodology), and the parallel Europarl corpus (used to derive subjectivity reference values).

### Webhose News

*Webhose* is a company specialized in turning unstructured web content into structured data. They provide several news articles data sets in distinct languages crawled from several sections of different news outlets<sup>2</sup>. Nesta pesquisa nós utilizamos os datasets com notícias em alemão, inglês, português e espanhol

Nós rotulamos cada notícias de cada dataset

In order to validate our methodology, we labeled each news article in the data set as “informative” or “opinionated”. For that, we have manually defined sets of keywords that indicate opinionated news and searched whether the news article URL contains one or more of these keywords. The keywords are distinct for each language (English: blog, opinion, column; Portuguese: blog, coluna, opiniao (opinião); German: kolumne, meinung, kommentar; Spanish: blog, editor, editorial, opinion (opinión)).

For our experiments, we randomly sampled 1,200 (roughly the size of the smallest category, i.e., opinionated Portuguese news) articles from each category (opinionated and informative) and language.

### EventRegistry News

For building a data set containing news about the Venezuela Crisis and Syrian War, we used a news monitor named EventRegistry<sup>3</sup>. EventRegistry enables one to search news by keywords or by topics, such as “politics”.

We submit either the keyword “Venezuela” or “Syria” (considering their respective translations to all considered

<sup>2</sup><https://webhose.io/free-datasets/>

<sup>3</sup><http://eventregistry.org/>

idioms) to EventRegistry and filter the news published in the politics sections that contain one or more words in their body indicating that the article is addressing the “Syrian War” or the “Venezuela Crisis” topic.

For that, we select a set of keywords that are strongly related to these topics. We use the words “war” and “crisis”, respectively, as initial keywords for the two topics. These were selected as the most representative words for the topics based on the autocomplete feature of Wikipedia: when entering “Venezuela” in the search bar, Wikipedia suggests the article “Venezuelan Presidential Crisis”; for “Syria” it suggests “Syrian Civil War”. This may provide a good hint of how these topics are reported by the media. To consider alternative ways of how these topics might be reported by the media in each language, we expanded the initial set of keywords by adding the most similar words according to a word embedding. For each language, we trained a word embedding model on all news articles related to the topic (e.g., we trained a skip-gram model on all English texts returned by EventRegistry for the query “Venezuela”) and selected all words that have a cosine similarity above the (empirically defined) threshold of 0.6 to the seed words. For the English models, we end up with “conflict” and “strife” as the two most similar words to “war”, while “crises” and “turmoil” are the most similar to “crisis”.

For each article we hold information about headline, textual body, publication date, country, topic, URL and language. The resulting data set contains 13.102 news articles published in 126 distinct countries by 1.654 distinct news outlets, from which 9.004 refer to the Venezuela Crisis and 4.098 to the Syrian War. News published in Portuguese, German and English cover the period of time from 03/10/2019 to 08/26/2019, while news in Spanish range from 07/20/2019 to 08/26/2019. Figure 2 illustrates the distribution of news articles by country, language and topic.




Figure 2: Distribution of news articles by country, language and topic. The larger the cell area, the higher the number of news articles published in that country.

### Europarl Corpus

The Europarl<sup>4</sup> is a corpus containing parliamentarians’ speeches from the European Union parliament manually

<sup>4</sup><https://www.statmt.org/europarl/>

translated into all 21 European languages. The corpus is commonly used for automatic language translation tasks.

In this research, we use the German, English, Portuguese, and Spanish versions of the Europarl corpus for building a sentence aligned corpus, in which each speech in our corpus could be found translated into the other languages. We assume that, due to the careful translations that are required in the parliament, the subjectivity level in the translations should be as close as possible to the original version. We use this corpus as a reference for our normalization procedure (see Section for further details).

### Subjectivity Dataset v1.0

The Subjectivity Dataset v1.0<sup>5</sup> (Pang and Lee 2004) (SDv1) is a movie review data set commonly used for subjectivity classification tasks. It contains 5.000 sentences from IMDb plot summaries, labeled as “Objective”, and 5.000 from snippets from the Rotten Tomatoes pages, labeled as “Subjective”. We removed sentences not written in English<sup>6</sup> from the data set, ending up with 4.985 Objective and 4.963 Subjective sentences. As with the Webhose data set, in Section we show that our methodology is able to find significant differences in subjectivity between the Objective and Subjective sentences.

### Wikipedia

We downloaded the English, Portuguese, Spanish, and German Wikipedia dumps dated June 1st, 2019, and extracted the running text out of it.<sup>7</sup> Wikipedia promotes content policies to enforce a neutral point of view. Moreover, Wikipedia articles are built collaboratively and are open for review, which naturally tends to lead to unbiased articles (Greenstein and Zhu 2012). Due to this comparatively high degree of neutrality, we use the Wikipedia corpus both as a resource for training word embeddings and as an unbiased comparison corpus.

## Methodology Description

In this section, we present our methodology in more detail (depicted in Figure 1).

The methodology is composed of three stages:

- 1) Deriving parallel subjectivity lexicons for all languages in our data sets. Here, we translate the original Portuguese lexicons to all target languages, pointing out and satisfying a set of constraints that are necessary for the translation.
- 2) Computing subjectivity bias, depicted on the left-hand side of Figure 1. Here, we calculate the subjectivity bias separately for each news article in our data sets based on the subjectivity lexicons of its language.
- 3) Calculating normalized subjectivity bias, depicted on the right-hand side of Figure 1. We use language subjectivity bias to normalize the subjectivity biases calculated in the

<sup>5</sup><http://www.cs.cornell.edu/people/pabo/movie-review-data/>

<sup>6</sup>According to *langdetect*, <https://pypi.org/project/langdetect/>

<sup>7</sup><https://github.com/Kyubyong/wordvectors>




Figure 3: Example of a subset of the argumentation subjectivity lexicon in Brazilian Portuguese translated into English, German, and Spanish, respectively.

previous stage, in order to make them comparable across the distinct languages. This step is necessary because different languages may be generally more subjective than others (see Section ).

### Deriving Parallel Subjectivity Lexicons

This step is concerned with deriving parallel lexicons for all languages present in our data sets from the original Portuguese subjectivity lexicons. Parallel lexicons, in this research, can be understood as a set of lexicons translated into multiple languages<sup>8</sup>. Figure 3 depicts an example of parallel lexicons.

We perform the lexicon translation in two steps, an automatic translation step and a manual post-processing step, as shown in (Banea, Mihalcea, and Wiebe 2013):

1. *Translation*: We translate the Portuguese lexicons into the other languages using the automatic translation tool DeepL<sup>9</sup>, which has been obtaining comparable performance to the Google Translate system (Macketanz et al. 2018). Also, the company provides a BLEU score comparing itself to other companies on their website.<sup>10</sup>
2. *Post-Processing*: For our method to perform correctly, we need to perform deduplication on the translated words and ensure that all translated lexicons contain an equal number of words (see below). We do this by updating duplicated words based on an online dictionary (e.g., Pons<sup>11</sup>). We went through the translated lists of words, searched for duplicate words, and tried to find alternative translations. If we could not find an alternative translation, we randomly removed one of the words in the original lexicon that led to duplicates in the translation. As an example, consider the Portuguese words “aturdir” and “atordoar”, which translate to German as “betäuben” (to stun). Since we cannot find an alternative translation for “aturdir” or “atordoar”, then we randomly remove one of them from the lexicon set.

The post-processing step is necessary to fulfill the following requirements for creating parallel subjectivity lexicons according to our method.

First, lexicons representing the same subjectivity dimension (e.g. argumentation) must have the same length in all

<sup>8</sup>The concept is analogous to the “parallel corpus” but using lexicons instead of text documents.

<sup>9</sup><https://www.deepl.com/>

<sup>10</sup><https://www.deepl.com/quality.html>

<sup>11</sup><https://en.pons.com/translate>

languages. We show in Section that WMD (Kusner et al. 2015) (the method in charge of computing the subjectivity scores) would return slightly different results if the underlying lexicons have different sizes.

Secondly, all expressions of one lexicon must be unique for that lexicon, meaning that we will not find repeated words in one lexicon in any language. This requirement is essential because if we keep duplicated words in the lexicon, the word will influence the subjectivity score more than it should. For example, if we compute the distance between the piece of text “the book is on the table” and the lexicon “book book” the method will return a higher subjective bias than if the lexicon was “book novel.”

As a consequence of these requirements, if we cannot find a unique translation of the Portuguese word in all languages, we remove that word from our analysis. After these two steps, we end up with five equally sized, deduplicated subjectivity lexicons for each language.

After this deduplication, the argumentation lexicon decreased from 110 to 88 expressions, presenting the most significant decrease. This is somewhat expected since the argumentation set contains the more complex expressions (expressions composed of two or more words) over all lexicons, making it harder to find synonyms or alternative translations. The sentiment lexicon decreased from 153 to 138 expressions. Since many words in Portuguese are synonyms, it becomes challenging to find the same amount of synonyms in all other languages. The modalization lexicon presented the smallest decrease from 55 to 54 expressions, while the valuation and the presupposition suffered no decrease and remained with 81 and 54 expressions, respectively. While there might be a loss of detectable subjectivity in the Portuguese language, caused by the removal of some words from the original Portuguese lexicons, we show in Section that, in practice, this loss is negligible.

## Computing Subjectivity Bias

This step performs the subjectivity bias computation for each language. For that, we use the method introduced by Sales et al. (Sales, Balby, and Veloso 2019) for computing *inverse subjectivity bias* (ISB) scores for each language, which we will then normalize to make it comparable across languages in the next stage.

Given a word embedding model, five lexicons representing five subjectivity dimensions and a set of news articles, the method relies on a word embedding model for computing the Word Mover’s Distance (WMD) between each lexicon and news article. The WMD takes two documents as input, corresponding in our case to one news article and one lexicon, which are represented as a weighted set of word embeddings, and calculates the distance between them as the sum of Euclidean distance between the two documents’ words (Kusner et al. 2015). Doing this for each lexicon yields a 5-dimensional subjectivity vector representing the degree of subjectivity associated with the news article. Since the measure is based on distance rather than similarity, the resulting value is the inverse of the subjectivity bias: a higher ISB value implies lower subjectivity bias and vice versa.

For each language, we train one skip-gram word embedding (Mikolov et al. 2013b,a) over a Wikipedia dump of the respective language and then calculate the WMD between a news article and the language’s subjectivity lexicons as shown on the left side of Figure 1.

We have used the Wikipedia corpus (in each considered idiom) for training our word embeddings. As a result, we expect to have word embeddings that are mostly unbiased regarding subjectivity. This is important since we want to isolate, as much as possible, the bias caused by the interlocutor.

At the end of this stage, subjectivity bias scores are already comparable for news of the same language. However, we cannot directly compare scores across languages since similar levels of subjectivity for distinct languages might present distinct subjectivity biases.

## Subjectivity Normalization

We propose to use a parallel corpus, which we can consider to be equally biased in terms of the interlocutor across the languages, for computing ISBs reference values for all target languages and check how much the news ISB rates deviate from its respective reference values. To this end, we simply calculate the difference between the ISB of a specific news article and the median of ISBs in the reference corpus for the article’s language. This is depicted on the right hand side of Figure 1. The outcome of this stage is a normalized ISB that can take any value in the real numbers domain. The higher the score, the lesser the subjectivity associated with the news, where a value of zero means that the article is exactly as biased as the reference corpus.

We use a subset of 15,000 randomly selected speeches from the Europarl sentence aligned corpus (described in Section ) for computing the reference rates of the target languages. After computing the ISBs in the Europarl corpus, we end up with a distribution over 15,000 ISBs for each language-dimension, depicted along with their mean and median in Figure 4. We define the language’s reference rate, for each subjectivity dimension, as the median of the respective distribution. Note that the mean and median are usually close to each other for each language-dimension, which allows us to use either one. Also note that, distinct languages present different mean/median values for the same dimension. This is, per se, a clear indication that the language contains an inherent subjectivity and the normalization step is necessary. We will further substantiate the normalization step in Section section.

## Experiments

In this section, we first validate our method by applying it on corpora where we know the level of subjectivity contained in the text and showing that our method can recover this information. Next, we show an application of the proposed methodology to assess media bias on our news data set, comparing the level of subjectivity in news from different countries about two topics. In addition, we provide an analysis of how the results would change if we choose not to run the Normalization Stage and directly compared the

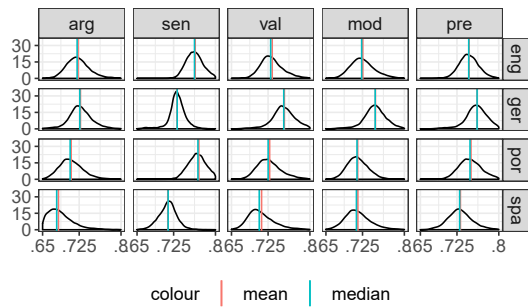


Figure 4: Density plots of ISBs by language and dimension of subjectivity calculated over the Europarl corpus. The distribution mean is shown in red and the median in blue.

subjectivity values after the subjectivity computing stage. All results are presented through confidence intervals with a confidence level of 99%. All the code, data, and resources used in this research are available at Github<sup>12</sup>.

### Methodology Validation

We validate the proposed method from distinct perspectives. First, we show that we can find significant differences in subjectivity scores between subjective and objective texts. Secondly, we show that either translating the lexicons to other idioms or translating the texts of other idioms to Portuguese yield compatible results, indicating that both methods are valid. Thirdly, we compare the automatically translated lexicons to manually adapted lexicons produced by a linguist. Finally, we justify the need of having equal sized lexicons and show that the reduced lexicons after our translation step lead to the same conclusions as the original version on Portuguese texts.

**Subjectivity Detection** We use Wikipedia, Webhose, and SDv1 for validating our methodology, showing that it can effectively distinguish between objective and subjective text. Wikipedia and Webhose enable us to show how well our method performs in multiple languages, since they contain articles written in German, Portuguese, English, and Spanish. However, we only have weak labels for these corpora, as described in Section . SDv1, on the other hand, contains only English text, but features manually annotated sentences with objective/subjective labels. We added Wikipedia to this experiment because Wikipedia has content policies for enforcing unbiased articles, and hence serve as a good reference of objective texts. Regarding the Webhose data set, we expect that opinionated news present higher subjectivity values than informative ones, which should, in turn, be equally or more subjective than Wikipedia’s articles. For SDv1, subjective sentences should present higher subjectivity values than objective ones.

We compute the ISB for each set of textual documents, and present the confidence intervals of the mean in Figure 5. If two confidence intervals overlap, we can not infer any significant difference between them. Results show significant

<sup>12</sup><https://github.com/allansales/InternationalMediaBias>

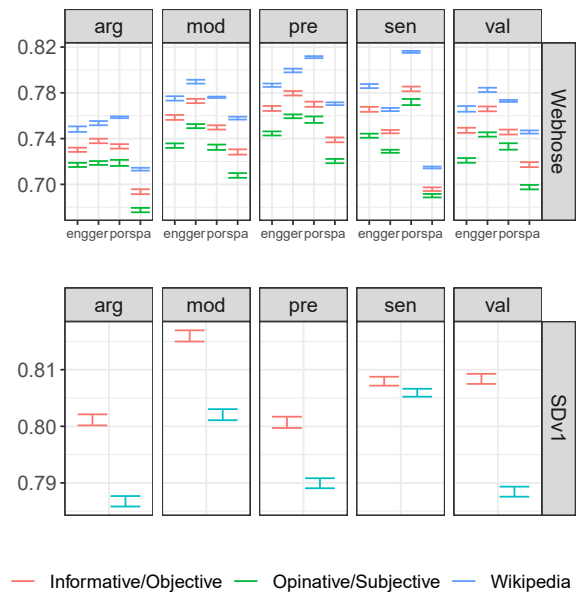


Figure 5: Confidence intervals of the mean of ISBs for informative news, opinionated news, and Wikipedia’s articles by language. If two intervals do not overlap, it implies a significant difference.

differences for all languages and dimensions across the different text sources in both Webhose and SDv1 data sets. It is noteworthy that the most subjective sources, i.e., Webhose opinionated news articles and the subjective sentences of SDv1, always present the lowest ISB scores. Moreover, the Webhose informative news present lower ISB values than Wikipedia articles, as one could expect. These results provide strong evidence that our approach can identify subjective language correctly. As an example, we provide the most subjective<sup>13</sup> and objective<sup>14</sup> news article found in our data set, according to the argumentation dimension.

**Lexicon vs News Translation** Instead of translating the lexicons to other languages, we could translate the news articles themselves to Portuguese and then apply the same methodology with the original, untranslated lexicons. This approach may have certain advantages over translating the lexicons: (i) the lexicons post-processing is no longer necessary, and; (ii) the context of the translated words in the text may be preserved. However, this approach is difficult to apply to a large number of documents in practice due to the costs of commercial machine translation services. Moreover, since lexicons are sets of words with no collective meaning, they are much easier to translate. The lack of context of the words and expressions in the lexicons, as well as eventual translation imprecision, are mitigated, to some ex-

<sup>13</sup><https://web.archive.org/web/20190920131028/https://reliefweb.int/report/syrian-arab-republic/under-secretary-general-humanitarian-affairs-and-emergency-relief-90>

<sup>14</sup><https://web.archive.org/web/20200115145321/https://www.sana.sy/en/?p=171635>

tent, by the use of the word embedding-based WMD in our method. That is, we expect that the word embedding models will place both contextual words as well as the correct translation close to the words/expressions contained in the translated lexicons in the semantic space. In order to compare the two approaches, we apply this alternative approach to a subset of our data and show that this leads to comparable conclusions, albeit with lower statistical significance due to the lower number of samples.

We test the alternative approach by translating texts from the Webhose and SDv1 data sets to Portuguese using DeepL and running the methodology as previously detailed. From the Webhose data set, we randomly select 120 news articles, out of which 60 are opinionated and 60 are informative, from the data set of each language (i.e., German, English, and Spanish). We do not include Portuguese news in this experiment since, our original lexicon is in Portuguese and does not need a translation; consequently, results would not differ from those obtained in the previous experiment. For the smaller SDv1, we were able to translate the entire data set.

Figure 6 presents the confidence intervals for differences of ISBs between objective and subjective sources for both text translation (red) and lexicon translation (blue). We can make two main observations in this plot: First, we see that translating the full text detects a significant difference of ISBs between the informative/objective and opinionated/subjective texts in all languages with 99 % confidence level, except argumentation and valuation in English in the Webhose data set. Note that, if we lower the confidence level to 90 %, the difference becomes significant in all subjectivity dimensions. We assume that the lower confidence stems mostly from the small sample size, as we only use 120 articles per language in this setting. However, the method still shows rather significant differences between the opinionated and informative news, albeit with a lower confidence.

Secondly, we can compare the confidence intervals produced by the lexicon and text translation. Two non-overlapping confidence intervals of the same language, data set, and subjectivity dimension imply a significant difference. For the Webhose data set, we cannot find a significant difference between the adopted approaches, indicating that either approach would lead to the same conclusions. As to the SDv1 data set, the lexicon translation approach presents significant differences in all subjectivity dimensions, except the sentiment one. However, there is no case where one interval includes zero and the other does not, meaning that the approaches agree with each other and our conclusions would tend to be similar using either.

Some factors might influence these results, such as the text length: news articles are commonly composed of multiple sentences while each instance of SDv1 is a single sentence. In any case, we show that our lexicons translation approach is at least as reliable as a text translation for computing the subjectivity values.

As an additional comparison, we compute the Pearson correlation between ISBs of the 120 news articles using either lexicon or news translation approach. The values produced by the approaches reach an average correlation of

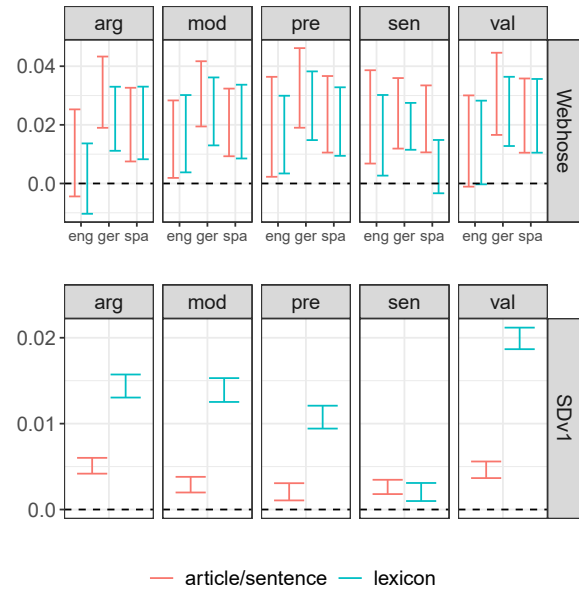


Figure 6: Confidence intervals of the difference of ISBs for objective and subjective sources. At the top: Confidence intervals computed over a subset of 60 opinionated and 60 informative news articles, for each language, from the Webhose data set. At the bottom: Confidence intervals computed over the entire set of Objective and Subjective sentences from the SDv1 data set. The color indicates the approach: lexicon (blue) and article/sentence translation (red). Confidence intervals without zero imply significant difference.

0.76 over all dimensions and languages, showing a reasonable agreement about what is subjective and objective news. The correlations are broken down in Figure 7 per language and subjectivity dimension.

Overall, the experiments in this section show that both methods mostly agree about the subjectivity of texts, further suggesting that both are valid.

**Lexicon Translation vs Manual Adaptation** Ideally, we would have our lexicons adapted by linguists to each target language. While this is difficult to achieve and does not scale for any possible idiom, we asked a Brazilian linguist specialized in English to adapt some of our Portuguese lexicons to English. We compute the WMD between the linguist’s resulting lexicons and ours, aiming to verify how the automatic translation compares to a manual adaptation made by an expert. Given the high cognitive effort of this task, our linguist adapted only the presupposition and the modalization lexicons to English.

Note that the smaller the WMD value, the higher the semantic agreement between the lexicons. Given that WMD is defined on the  $[0, +\infty)$  interval, we normalized it for better interpretability. We use the *min-max* normalization where *min* is 0, and we define *max* as the distance of the automatically translated lexicon to a “random” lexicon. Specifically, we sampled 40 random lexicons from the vocabulary of our






Figure 7: Pearson correlation confidence intervals between ISBs taken over the parallel lexicon approach and the news articles translation approach. The darker the cell, the smaller the correlation between the approaches, and vice-versa.

English word embedding and averaged the WMD between these lexicons and our automatically translated lexicon. Now our normalized score will lie in the  $[0, 1]$  interval. Finally, to come up with a similarity metric, we subtract the normalized score from 1. So, the closer to 1, the higher the agreement between the lexicons.

For presupposition, this similarity score is 0.70, while for modalization we have 0.76. Notice that these results denote a high semantic agreement between the compared lexicons. One key strength of our approach is that even when the translation is not completely accurate, it is enough that the translated expressions are close to the correct ones in the embedding space. Moreover, even when an expression that carries subjectivity appears in a text and not in our lexicons, our method will still be able to capture this signal of subjectivity, since this expression will be probably close to the ones in our lexicons in the embedding space.

**Lexicons Size Effect** In this experiment, we investigate how much the reduction of the lexicon size in our translation process influences the results of the method, that is, whether our translation process reduces the expressiveness of the lexicon. Figure 8 shows the relative change in average ISBs from the original Portuguese lexicon and the reduced lexicon resulting from the translation process, that is  $\frac{mean(ISB_{mod}^{dim})}{mean(ISB_{orig}^{dim})}$ , where  $ISB_{orig}^{dim}$  and  $ISB_{mod}^{dim}$  are the ISB computed with the original and modified lexicon for subjectivity dimension  $dim$ , respectively. We compare the ISBs computed over a Portuguese Wikipedia sample, containing 13.000 randomly chosen articles.

The results indicate that, as expected: (i) decreasing the lexicon sizes also decreases the detectable subjectivity in the text, and (ii) the strength of the decrease is associated with the number of words removed from the lexicon. The argumentation lexicon, which had 22 words removed, presents the highest decrease in subjectivity detection; followed by the sentiment lexicon with 15 words removed and the modalization lexicon with 1 word difference from its original version. However, the loss of detected subjectivity when using




Figure 8: Relative change of ISBs computed over a sample of Portuguese Wikipedia articles through the original Portuguese lexicons and the decreased sized Portuguese lexicons after deriving parallel subjectivity lexicons.

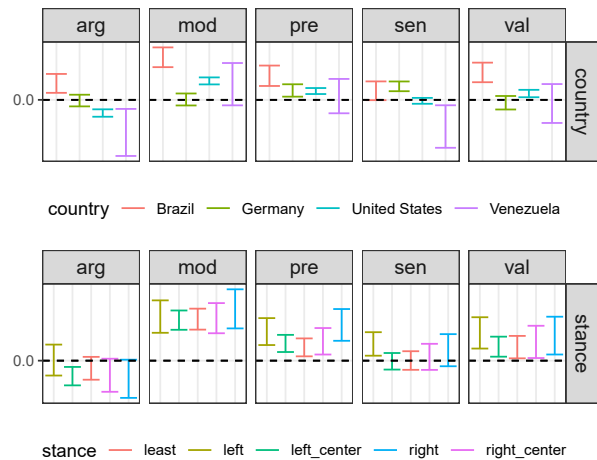


Figure 9: Confidence intervals of the difference of ISBs in news related to Venezuela Crisis and Syrian War by country, language and news outlet slant. Intervals entirely above or below zero mean higher subjectivity bias in articles about the Venezuela Crisis or Syrian War, respectively.

the smaller version of the lexicons is only a small fraction of the whole value.

Overall we can conclude that, while our translation method does reduce the expressiveness of the lexicon, it is only by a relatively small amount. Therefore, the modified lexicon should still be able to reliably detect subjective texts.

### Case Studies

After validating that our methodology can detect subjective text as expected, we apply it to news articles about the Venezuelan Crisis and the Syrian War to analyze the subjectivity regarding these topics from different points of view. Unless otherwise specified, in the following Confidence Intervals entirely above/below zero imply a higher subjectivity bias in articles about Venezuela Crisis/Syrian War topic, respectively.

**Topic Subjectivity by Countries** This experiment investigates whether the media reports in different countries are more subjective when reporting on one topic than the other. The selected countries are the ones who presented the highest number of published news for each language.

Figure 9 shows the confidence intervals built over the




Figure 10: Confidence intervals of the mean ISB in news related to Venezuela Crisis and Syrian War by country. Intervals entirely above or below zero mean higher subjectivity bias in articles about the Venezuela Crisis or Syrian War, respectively.

differences between the ISBs of news reporting on the Venezuela Crisis and the Syrian War. The dotted line represents the zero value across the boxes.

Results show that the Brazilian media is significantly more biased when reporting about the Venezuela Crisis, manifesting significant differences in all subjectivity dimensions. The United States media reports are more subjective about Venezuela, while presenting more argumentation level towards the Syrian War topic. German media is roughly equally biased when reporting about both topics. The Venezuelan media, in turn, is, surprisingly, a bit more biased when addressing the Syrian War than their own crisis situation, exhibiting two significantly different subjectivity dimensions.

The results, in some cases, reflect these countries' current positioning regarding the Venezuelan Crisis and the Syrian War. Regarding Brazil, one could expect the country to present a significant difference in subjectivity towards Venezuela Crisis topic, since Brazil has currently a far right-wing government that is often conflicting with Venezuela<sup>15</sup> but has little involvement with Syria's current war<sup>16</sup>. Germany, in its turn, has taken a position regarding Venezuela's situation<sup>17</sup> and also has some involvement with Syria. The United States took part in both events. Last, Venezuela's case is intriguing since it is naturally expected that its media would be more biased towards Venezuela's than other situations. However, one possible reason is the decreasing of press freedom during the Chavéz and, afterward, Maduro's government (Hawkins 2016).

A more detailed version of the countries subjectivities split by topic is given in Figure 10.

**Topic Subjectivity by Political Stance** This experiment aims to investigate whether the media reports, split by their political stance, are more subjective when reporting on one

<sup>15</sup><https://www.bbc.com/news/world-latin-america-47300962>

<sup>16</sup>[https://en.wikipedia.org/wiki/Foreign\\_involvement\\_in\\_the\\_Syrian\\_Civil\\_War](https://en.wikipedia.org/wiki/Foreign_involvement_in_the_Syrian_Civil_War)

<sup>17</sup><https://www.bbc.com/news/world-latin-america-47115857>

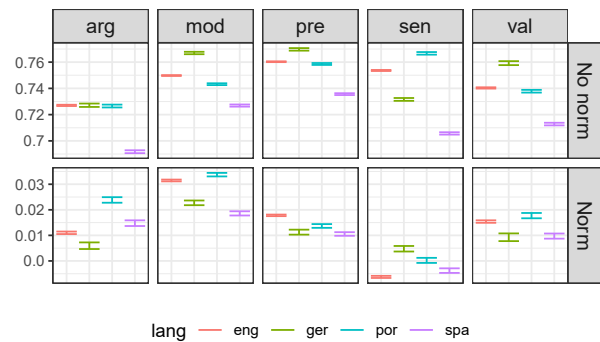


Figure 11: Subjectivity confidence intervals of Syrian War and Venezuela Crisis computed before (No norm) and after the Normalization Stage (Norm).

topic or the other. For example, are news reported by the right-wing news outlets more subjective than the left-wing subjective when reporting about Venezuela Crisis in comparison with the Syrian War?

For finding the political stance of each news outlet, we gather information from Media Bias Fact Check<sup>18</sup> (MBFC). It is important to remark that we run this experiment only with the 2,351 articles published by the 219 news outlets (covering 37 countries) in our data sets mapped by MBFC.

Results, depicted in Figure 9, show that news outlets are more subjective when reporting about the Venezuela Crisis, regardless of the political stance. Also, right-wing news outlets present the highest difference of subjectivity between the topics. In some cases (e.g., presupposition), the difference is significantly higher than the other political stances.

The right-wing results might be related to their rivalry with left-wing governments and the fact that far right-wing representatives often associate Venezuela's current situation with socialism.<sup>19</sup>

### Normalization Stage Effect

Our last experiment aims to attest how the Normalization Stage affects results and how not using it would lead to different conclusions. For doing this, we compute the ISBs for all news in our database a) including the Normalization stage and b) excluding the Normalization stage. We show the confidence intervals for the mean of these values for each language and subjectivity dimension in Figure 11.

We can point out some differences in results obtained from the different subjectivity computation approaches (with and without the Normalization Stage):

- The ISBs values decrease when running the Normalization Stage (in the presented example, in a scale range of about 0.7). This decreasing effect reflects what is intended when applying normalization: removing the presence of the language subjectivity in the final computed value. Before normalization, each subjectivity value aggregates the

<sup>18</sup><https://mediabiasfactcheck.com/>

<sup>19</sup><https://www.theguardian.com/world/2018/dec/16/liberate-venezuela-from-maduro-urges-bolsonaro-ally>

language subjectivity itself summed up with the interlocutor's subjectivity; after normalization, each value represents only the interlocutor's subjectivity.

- The distances between confidence intervals for different languages in each particular subjectivity dimension are smaller. Focusing on one subjectivity dimension (e.g., argumentation), it is noticeable that the distances between confidence intervals inside a box increase from the normalized to the unnormalized values;
- Some conclusions would change without eliminating the language subjectivity: For example, taking a look at the argumentation dimension, without normalization one would conclude that English, German and Portuguese news articles do not present significant subjectivity differences, as their confidence intervals overlap.

## Conclusion

Media Bias is an important research topic due to its influence on people's personal decisions on important issues. The languages singularities are a barrier for automatically assessing media bias in cross-linguistic and cross-national scenario, where the language is charged with the qualities that characterize each country and its people.

In this paper, we present a methodology for assessing media bias, instantiated as subjectivity analysis, in cross-linguistic scenarios, and on a large scale. The methodology requires parallel lexicons, subjectivity reference values, and a word embedding model representing the vocabulary for each language. We use machine translation for creating the parallel lexicons and compute subjectivity based on the distance between the input text and its respective language lexicon. The subjectivity reference values are computed over a parallel corpus, serving as an equally biased corpus in terms of interlocutor's subjectivity, making it possible to estimate the language bias differences. Finally, we applied the methodology over the news of two recent and resounded topics. Among our main findings we can highlight:

- Different languages exhibit different "base levels" of subjectivity, that is, one language may be generally more subjective than another;
- Taking into account the language bias is important in order to isolate the interlocutor's bias, which is the measure we are really interested in;
- We find subjectivity in news about the Venezuela Crisis significantly higher than in news about the Syrian War, mainly in Portuguese and English languages, and their most publishing countries, Brazil and the United States;
- Right-wing news outlets showed a higher subjectivity bias towards the crisis in Venezuela than news outlets following other political ideologies.

**Limitations and Future Work** Some limitations of our work should be noted. First, regarding the labelling of our validation data set, the Webhose News (Section ), we use a rather small number of keywords for detecting opinionated news with high precision, albeit potentially low recall.

However, leaving some "false negatives" in the informative corpus will not heavily influence the overall subjectivity distribution in the informative news. This means that the differences should still be significant if we do not detect all opinionated news with our keywords (as confirmed in Section ). Last, factors other than the interlocutor and the language subjectivity, such as cultural or regional biases, might have influenced our obtained subjectivity bias values. Also, as a threat to validity we point out that we do not remove quotes or anything else from news articles. We assume that the selection of quotations is a conscious decision made by the author and all such decisions must be taken into account for our subjectivity calculation.

As future work we plan to measure whether the results change if we normalize the subjectivity biases after the subjectivity computing stage, instead of forcing the lexicon sizes to be equal by manually removing some translations. Also, we plan to include more languages into the analysis, measuring the impact of this addition to the lexicon sizes and allowing to draw even more global conclusions about the addressed as well as other topics.

## References

- Amorim, E.; Caçado, M.; and Veloso, A. 2018. Automated Essay Scoring in the Presence of Biased Ratings. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, volume 1, 229–237.
- Bai, Q.; Wei, K.; Chen, M.; Hu, Q.; and He, L. 2018. Mining Temporal Discriminant Frames via Joint Matrix Factorization: A Case Study of Illegal Immigration in the US News Media. In *International Conference on Knowledge Science, Engineering and Management*, 260–267. Springer.
- Balahur, A.; Steinberger, R.; Kabadjov, M. A.; Zavarella, V.; der Goot, E. V.; Halkia, M.; Pouliquen, B.; and Belyaeva, J. 2013. Sentiment Analysis in the News. *CoRR* abs/1309.6202. URL <http://arxiv.org/abs/1309.6202>.
- Banea, C.; Mihalcea, R.; and Wiebe, J. 2013. Porting multilingual subjectivity resources across languages. *IEEE Transactions on Affective Computing* 4(2): 211–225.
- Benson, R. 2008. Journalism: normative theories. *The international encyclopedia of communication* .
- Caliskan, A.; Bryson, J. J.; and Narayanan, A. 2017. Semantics derived automatically from language corpora contain human-like biases. *Science* 356(6334): 183–186. doi:10.1126/science.aal4230. URL <http://opus.bath.ac.uk/55288/>.
- Card, D.; Gross, J.; Boydston, A.; and Smith, N. A. 2016. Analyzing framing through the casts of characters in the news. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, 1410–1420.
- Chaturvedi, I.; Cambria, E.; Zhu, F.; Qiu, L.; and Ng, W. K. 2015. Multilingual subjectivity detection using deep multiple kernel learning. *Proceedings of Knowledge Discovery and Data Mining, Sydney* .

- Chaturvedi, I.; Ragusa, E.; Gastaldo, P.; Zunino, R.; and Cambria, E. 2018. Bayesian network based extreme learning machine for subjectivity detection. *Journal of The Franklin Institute* 355(4): 1780–1797.
- Dallmann, A.; Lemmerich, F.; Zoller, D.; and Hotho, A. 2015. Media bias in german online newspapers. In *Proceedings of the 26th ACM Conference on Hypertext & Social Media*, 133–137. ACM.
- De Cock, R.; Mertens, S.; Sundin, E.; Lams, L.; Mistiaen, V.; Joris, W.; and d’Haenens, L. 2018. Refugees in the news: Comparing Belgian and Swedish newspaper coverage of the European refugee situation during summer 2015. *Communications* 43(3): 301–323.
- Entman, R. 1993. Framing: Toward Clarification of A Fractured Paradigm. *The Journal of Communication* 43: 51–58. doi:10.1111/j.1460-2466.1993.tb01304.x.
- Entman, R. 2007. Framing Bias: Media in the Distribution of Power. *Journal of Communication* 57: 163 – 173. doi:10.1111/j.1460-2466.2006.00336.x.
- Goldberg, B. 2014. *Bias: A CBS insider exposes how the media distort the news*. Regnery Publishing.
- Greenstein, S.; and Zhu, F. 2012. Collective Intelligence and Neutral Point of View: The Case of Wikipedia. Working Paper 18167, National Bureau of Economic Research. doi:10.3386/w18167. URL <http://www.nber.org/papers/w18167>.
- Hamborg, F.; Donnay, K.; and Gipp, B. 2018. Automated identification of media bias in news articles: an interdisciplinary literature review. *International Journal on Digital Libraries* doi:10.1007/s00799-018-0261-y.
- Hawkins, K. A. 2016. Responding to radical populism: Chavismo in Venezuela. *Democratization* 23(2): 242–262.
- Kristiansen, T.; Garrett, P.; and Coupland, N. 2005. Introducing subjectivities in language variation and change. *Acta Linguistica Hafniensia* 37(1): 9–35.
- Kusner, M.; Sun, Y.; Kolkin, N.; and Weinberger, K. 2015. From word embeddings to document distances. In *International Conference on Machine Learning*, 957–966.
- Lazaridou, K.; Krestel, R.; and Naumann, F. 2017. Identifying Media Bias by Analyzing Reported Speech. In *Data Mining (ICDM), 2017 IEEE International Conference on*, 943–948. IEEE.
- Lin, Y.-R.; Bagrow, J. P.; and Lazer, D. 2011. More Voices Than Ever? Quantifying Media Bias in Networks. *ICWSM 1*(arXiv: 1111.1227): 1.
- Macketanz, V.; Ai, R.; Burchardt, A.; and Uszkoreit, H. 2018. TQ-AutoTest—An Automated Test Suite for (Machine) Translation Quality. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*.
- Mihalcea, R.; Banea, C.; and Wiebe, J. 2007. Learning Multilingual Subjective Language via Cross-Lingual Projections. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, 976–983. Association for Computational Linguistics. URL <http://www.aclweb.org/anthology/P07-1123>.
- Mikolov, T.; Chen, K.; Corrado, G.; and Dean, J. 2013a. Efficient Estimation of Word Representations in Vector Space. *CoRR* abs/1301.3781. URL <http://dblp.uni-trier.de/db/journals/corr/corr1301.html#abs-1301-3781>.
- Mikolov, T.; Sutskever, I.; Chen, K.; Corrado, G. S.; and Dean, J. 2013b. Distributed Representations of Words and Phrases and their Compositionality. In Burges, C. J. C.; Bottou, L.; Welling, M.; Ghahramani, Z.; and Weinberger, K. Q., eds., *Advances in Neural Information Processing Systems* 26, 3111–3119. Curran Associates, Inc. URL <http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.pdf>.
- Mundim, P. S. 2018. O viés da cobertura política da imprensa nas eleições presidenciais brasileiras de 2002, 2006 e 2010. *Revista Brasileira de Ciência Política* 7 – 46. ISSN 0103-3352. URL <http://www.scielo.br/scielo.php?script=sci.arttext&pid=S0103-33522018000100007&nrm=iso>.
- Nickerson, R. S. 1998. Confirmation bias: A ubiquitous phenomenon in many guises. *Review of general psychology* 2(2): 175.
- Oelke, D.; Geisselmann, B.; and Keim, D. A. 2012. Visual Analysis of Explicit Opinion and News Bias in German Soccer Articles. In *EuroVis Workshop on Visual Analytics, EuroVA 2012, Vienna, Austria, June 4-5, 2012*. doi:10.2312/PE/EuroVAST/EuroVA12/049-053. URL <https://doi.org/10.2312/PE/EuroVAST/EuroVA12/049-053>.
- Pang, B.; and Lee, L. 2004. A Sentimental Education: Sentiment Analysis Using Subjectivity Summarization Based on Minimum Cuts. In *Proceedings of the ACL*.
- Recasens, M.; Danescu-Niculescu-Mizil, C.; and Jurafsky, D. 2013. Linguistic models for analyzing and detecting biased language. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, 1650–1659.
- Saez-Trumper, D.; Castillo, C.; and Lalmas, M. 2013. Social media news communities: gatekeeping, coverage, and statement bias. In *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, 1679–1684. ACM.
- Sales, A.; Balby, L.; and Veloso, A. 2019. Media Bias Characterization in Brazilian Presidential Elections. In *Proceedings of the 30th ACM Conference on Hypertext and Social Media*, HT ’19, 231–240. New York, NY, USA: ACM. ISBN 978-1-4503-6885-8. doi:10.1145/3342220.3343656. URL <http://doi.acm.org/10.1145/3342220.3343656>.
- Siebert, F.; Siebert, F.; Peterson, T.; Peterson, T.; Schramm, W.; of the Churches of Christ in the United States of America. Department of the Church, N. C.; and Life, E. 1956. *Four Theories of the Press: The Authoritarian, Libertarian, Social Responsibility, and Soviet Communist Concepts of What the Press Should Be and Do*. An Illini book ; 72421-6. University of Illinois Press. ISBN 9780252724213. URL <https://books.google.com.br/books?id=4Q-oePDdc8C>.
- Verhagen, A. 2005. *Constructions of intersubjectivity: Discourse, syntax, and cognition*. Oxford University Press on Demand.