



Sentiment Analysis of News Articles in Spanish using Predicate Features¹

Antonio Tamayo

Julián Arias Londoño

Universidad de Antioquia

Medellín, Colombia

Diego Burgos

Wake Forest University

Winston Salem, Estados Unidos

Gabriel Quiroz

Universidad de Antioquia

Medellín, Colombia

Abstract

The automatic prediction of the course of action of agents involved in social or economic trends is an imperative challenge nowadays. However, it is a difficult task because stance or opinion is often spread throughout long, complex texts, such as news articles. The current study tests sentence predicates as features to automatically determine the writer's stance in news articles. We capture the semantics and stance of the text by encoding features such as the attribute of copulative sentences, the predicate of transitive sentences, adjectival phrases, and the section of the article. Under the assumption that these features are informative enough to model the semantics of the text, each word sequence is disambiguated and assigned a sentiment value using weighting rules. Different experiments were

¹ This study is a product of the research project "Implementación de un modelo de aprendizaje automático para análisis de sentimientos en artículos de prensa acerca de la pobreza en Colombia", approved by the University of Antioquia, with CODI record number 892 of February 2016, and of the masters thesis "Metodología para el análisis automático de sentimientos en documentos a partir de técnicas de aprendizaje automático y lingüística computacional" by Antonio Tamayo, 2018.

run using either SentiWordNet and ML-Senticon to determine words' sentiment. Feature vectors are automatically built to populate a database that is tested using two machine learning algorithms. An efficiency of 69% was achieved using a SVM with Gaussian kernel along with a feature selection strategy. This score outperformed the bag-of-words baseline in 12%. These results are promising considering that the sentiment analysis is performed on very complex texts written in Spanish.

Key words: sentiment analysis; linguistic features; socio-politics; news articles; Support Vector Machines; Naïve Bayes; dimension reduction.

Resumen

Análisis de sentimientos en artículos de prensa en español usando predicados como características

La predicción automática del modo de proceder de los agentes involucrados en las tendencias sociales o económicas es un desafío preponderante en la actualidad. Sin embargo, es una tarea difícil debido al hecho de que la postura u opinión a menudo se extiende a través de documentos extensos y complejos, como los artículos de noticias. El presente trabajo evalúa los predicados de oraciones como características para determinar automáticamente la postura del escritor en los artículos de noticias. Capturamos la semántica y la postura del texto codificando características como el atributo de oraciones copulativas, el predicado de oraciones transitivas, sintagmas adjetivales y la sección del artículo. Bajo el supuesto de que estas características son lo suficientemente informativas para modelar la semántica del texto, cada secuencia de palabras se desambigua y se le asigna un valor de sentimiento con reglas de ponderación. Se realizaron diferentes experimentos empleando SentiWordNet y ML-Senticon para determinar la opinión de las palabras. Los vectores de características se construyen automáticamente para completar una base de datos que se prueba mediante el uso de dos algoritmos de aprendizaje automático. Se logró una eficiencia del 69 % utilizando una SVM con kernel gaussiano junto con una estrategia de selección de características. Esta puntuación superó la línea de base de la técnica de "bag of words" en un 12 %. Estos resultados son prometedores si tenemos en cuenta que el análisis de sentimientos se hace en documentos muy complejos en español.

Palabras clave: análisis de sentimientos; características lingüísticas; artículos de noticias; sociopolítica; máquinas de vectores de soporte; Naïve Bayes; reducción de dimensión.

Résumé

Analyse de sentiments dans des articles de presse en espagnol en utilisant des prédicats en tant que caractéristiques

La prédiction automatique des façons d'agir des agents impliqués dans les tendances sociales ou économiques est un défi impératif de nos jours. Cependant, cette tâche s'avère difficile étant donné que les avis s'entendent dans des documents longs et complexes, tels que des articles de presse. Le présent travail évalue les prédicats de phrases en tant que caractéristiques pour déterminer systématiquement le point de vue de l'écrivain dans les articles de presse. Il s'agit de capturer la sémantique et la posture du texte en décodant des caractéristiques telles que l'attribut de phrases copulatives, le prédicat de phrases transitives, de syntagmes adjectivaux et la section de l'article. En supposant que ces fonctionnalités soient suffisamment informatives pour modéliser la sémantique du texte, chaque séquence de mots est désambiguïsée et une valeur de sentiment est attribuée à l'aide de règles de pondération. Différentes expériences ont été réalisées à l'aide de SentiWordNet et de ML-Senticon afin de déterminer l'avis des mots. Les vecteurs de fonctionnalités sont construits de façon automatique pour alimenter une base de données testée à l'aide de deux algorithmes d'apprentissage automatique. Une efficacité de 69% a été obtenue avec un SVM à noyau gaussien et une stratégie de sélection des fonctionnalités. Ce score a dépassé la base de référence de la technique «bag of words» dans 12%. Ces résultats sont prometteurs compte tenu du fait que l'analyse des sentiments est effectuée sur des documents très complexes en espagnol.

Mots-clés : analyse des sentiments ; caractéristiques linguistiques ; sociopolitique ; articles de presse ; machines à vecteurs de support ; Naïve Bayes ; réduction des dimension.

CÓMO CITAR ESTE ARTÍCULO

Tamayo. A., Arias Londoño. J., Burgos. D., & Quiroz. G. (2019). Sentiment Analysis of News Articles in Spanish using Predicate Features. *Lenguaje*, 47(2), 235-267.
doi: 10.25100/lenguaje.v47i2.7937

INTRODUCTION

Sentiment analysis (aka opinion mining) consists of analyzing people's written opinions, feelings, attitudes, and emotions with regard to a product, service, organization, individual, event, or another particular issue (Liu, 2012). From a practical perspective, the task aims at classifying a text as positive and negative, and sometimes as neutral. Sentiment analysis has gained popularity in the past fifteen years as a result of the exponential growth of digital information, much of which is of great interest to private and public organizations. The task has become especially useful for these organizations to determine people's perception of their products or services.

From a technical viewpoint, the problem of analyzing opinion or sentiment has been addressed following various approaches. The literature discussed below reports Works based on computational linguistics as well as conventional and hybrid machine learning techniques. A good deal of sentiment analysis experiments have been carried out following the Bag-of-Words approach using lexical resources and machine learning, which includes multiple adaptations of topic models (Liu, 2012; Pang & Lee, 2008). However, the literature review shows that most of these studies aim at performing sentence or aspect-based sentiment analysis of service or product reviews (Pang, Lee, & Vaithyanathan, 2002). Determining sentiment of whole complex documents such as those dealing with socio-political issues is hardly addressed.

Some results have reported the efficiency (75%-95%) of a few experiments for aspect-based sentiment analysis of web documents, but the type or domain of the documents is not specified, and testing is carried out with short texts such as product reviews (Nasukawa & Yi, 2003). The analysis of opinion in the socio-political domain has been restricted to tweets, as can be seen in Mitra, Parthasarathi and Mamani (2015) and Mohammad, Kiritchenko, Sobhani, Zhu, and Cherry (2016). The complexity of sentiment analysis in longer documents of this domain was explicitly highlighted by Liu (2015), which accounts for the scarcity of works in this direction. This is precisely the gap that we intend to contribute to fill with the present work.

In this study, we use Support Vector Machines (SVM) and Naïve Bayes to test the suitability of linguistic and discursive features to perform sentiment analysis of socio-political news articles in Spanish². These models were selected after a set of experiments using different algorithms, according to their

² This work is part of the research project entitled "Implementación de un modelo de aprendizaje automático para análisis de sentimientos en artículos de prensa acerca de la pobreza en Colombia", approved by the School of Languages of the University of Antioquia according to Research Approval dated February 2 of 2016-9244.

performance. The chosen classification models are also reported by other scholars as the models used for sentiment analysis classification (Pang *et al.*, 2002).

In this sense, with the choice of linguistic and discursive features along with well-established classification models, this work intended to improve the efficiency obtained with a bag-of-words approach, which is considered as the baseline. The way the feature database was automatically built and the definition and extraction of features from Colombian newspaper articles related to the phenomenon of poverty is also described here. In order to determine the features with the highest discrimination capabilities, two-dimension reduction techniques, namely principal component analysis (PCA) and sequential floating forward selection (SFFS) were also evaluated.

This paper is organized as follows. Section II discusses relevant related work. Section III presents the methodology to automatically populate a machine learning-oriented database using linguistic and discursive features designed to capture the sentiment of a document. Section IV describes the testing method and algorithms used in the classification tasks. Section V shows the experiments and results of testing the feature set using classification models and feature reduction techniques. Section VI is devoted to the discussion of the results. Section VII concludes.

240

RELATED WORK

The amount of research and experimentation on sentiment analysis has grown dramatically in the past years. The bulk of works in this field, however, have focused mainly on classifying product and service user reviews. For example, Pang *et al.* (2002) and Turney (2002), probably among the most cited works in the field of sentiment analysis, classify movie, car, and bank reviews in English. They test unigrams and bigrams as well as word association measures with a semantic orientation using SVM and Naïve Bayes. Very little and restricted work has been done, though, on classifying longer complex documents. Among them, Nasukawa and Yi (2003) experimented with aspect-based sentiment analysis in news articles, but instead of analyzing the sentiment of the whole article they classify simple sentences in the article; then they present the end user with a set of classified statements. Likewise, Balahur *et al.* (2010) work with news articles, but they analyze only quotes within the articles. Along the same lines, Q. Li *et al.* (2014), X. Li, Xie, Chen, Wang, and Deng (2014), and Schumaker, Zhang, Huang, and Chen (2012) reported experiments with whole documents to predict stock price using sentiment analysis of financial news articles. These works do not focus either on

classifying the sentiment of entire articles. Instead, they tag individual words or sentences with a sentiment value, and then they use this input as one additional factor to predict trends in the stock market. No results on the sentiment classification task itself are reported in these experiments.

With regard to features used for classification, Gamon (2004) shows an improvement for text sentiment classification using deep linguistic features, i.e., frequencies of syntactic rewrite rules and features derived from semantic graphs, which outperform n-grams. Mullen and Collier (2004) introduced a sophisticated set of features by combining n-grams, which can be categorized in three groups: (1) features extracted from the sentiment values of words or phrases, (2) adjective values for the three factors proposed by Osgood, Suci, and Tannenbaum (1957), and (3) sentiment values of words or phrases in (1) or (2) that are either close to each other or within the relevant phrase. Along the same lines, Joshi and Penstein-Rosé (2009) use syntactic dependency to capture adjectival modification of nouns. This work is highly dependent on the domain, though, which motivated a generalization of the approach using bigrams (Xia & Zong, 2010). In a similar way, Ng, Dasgupta, and Arifin (2006) utilized syntactic dependency features to capture adjective-noun, subject-verb, and verb-object relations.

Some attempts have also been made to capture discursive features, such as annotating document sections under the assumption that particular sections concentrate opinion. Both manual annotation (Zaidan, Eisner & Piatko, 2007) and automatic annotation (Yessenalina, Choi & Cardie, 2010; Yessenalina, Yue & Cardie, 2010) have been used. Another example of this can be found in Becker and Aharonson (2010), who, based on psychological and psycholinguistic experiments, suggest that sentiment analysis should focus on the last part of the text (e.g., the last sentence). No computational experiment was implemented to verify this statement, though. As for sentiment value commutators, S. Li, Lee, Chen, Huang, and Zhou (2010) explored negations and other sentiment commutators. This is done by grouping the sentences of a document into sentences with commutated sentiment and sentences without commutated sentiment. Then, these features are used to classify texts using supervised machine learning. Data was annotated manually. Lastly, Xia, Wang, Hu, Li, and Zong (2013) also present a methodology to handle negations for sentiment analysis.

METHODOLOGY

This section presents a description of the process to automatically extract predicate and discursive features from our document set. Technical details and the conceptual rationale used to define the applied characterization rules are provided here also. Using these features, a vector representation of each document is obtained, which constitutes the input for sentiment analysis using SVM and Naïve Bayes.

The document set

The document set used for validation was made up of 205 online articles taken from three of the most important Colombian newspapers (El Tiempo³ = 56, El Espectador⁴ = 65, El Colombiano⁵ = 84).

The articles were taken from the corpus of PoLaMe⁶ project following stratified random sampling.

242 This corpus was collected automatically using a web crawler (Quiroz, Tamayo & Zuluaga, 2017) according to two main criteria: 1) that the articles contained the word poverty and 2) that they had been published between 2010 and 2014. The article extension ranges from 300 to 1000 words. Thematically, the corpus compiled following these criteria contains documents related to public finance management, e.g., poverty, peace talks, armed conflict, informal economy, unemployment, affordable housing, prostitution, child adoption, education, political parties, social investment, violence-related displacement, and local and global economy, among other topics.

Document tagging

Each article was manually tagged either positive or negative by three linguist experts using joint probability of agreement as validation method. While this methodology is rather simple, we deemed it useful for these experiments with three annotators and two classes.

³ www.eltiempo.com

⁴ www.elespectador.com

⁵ www.elcolombiano.com

⁶ "Pobreza, Lenguaje y Medios- PoLaMe", funded by the Norwegian Research Council and the University of Bergen. <http://www.uib.no/en/project/polame>. This corpus was created to analyze how poverty is conceptualized in the media. Hence the use of the word poverty as a criterion to collect this corpus. Access to the corpus can be requested to Antonio Tamayo to antonio.tamayo@udea.edu.co.

The annotators were asked to:

- tag of an article positive if the newswriter's opinion is positive towards the central matter of the article.
- tag an article is negative if the newswriter's opinion is negative towards the central matter of the article.

This annotation process yielded 109 positive and 96 negative articles.

Feature extraction process

Tokenization, part-of-speech tagging, and sense disambiguation were carried out by means of an interface that we built to call Freeling's (Atserias *et al.*, 2006; Carreras, Chao, L. Padró, & M. Padró, 2004; L. Padró, 2011; L. Padró, Collado, Reese, Lloberes, & Castellón, 2010; L. Padró & Stanilovsky, 2012) modules, and taking each document as input. The output was a file per document with one word per line with its part of speech and a sense ID to be used in determining the word's sentiment value, as shown below.

In order to determine the word's sentiment value, the word's sense ID was matched with sense's sentiment value in SentiWordNet (Baccianella, Esuli & Sebastiani, 2010).

Likewise, to capture potentially useful discursive or rhetorical information, each document was splitted up into three sections. This information was later added to each linguistic feature. Thus, each article was sectioned into introduction (S1, first 25% of the tokens of the article), body (S2, 25%-75%) and "conclusions" (S3, 75%-100%). This decision was made under the assumption that some features are more discriminating depending on the section they appear. These percentages were assigned arbitrarily due to the fact that news articles do not follow an ideal rhetorical pattern or model, such as the ones described by Swales (2004) and Parodi (2010). Figure 1 illustrates the feature extraction process to build the training database.

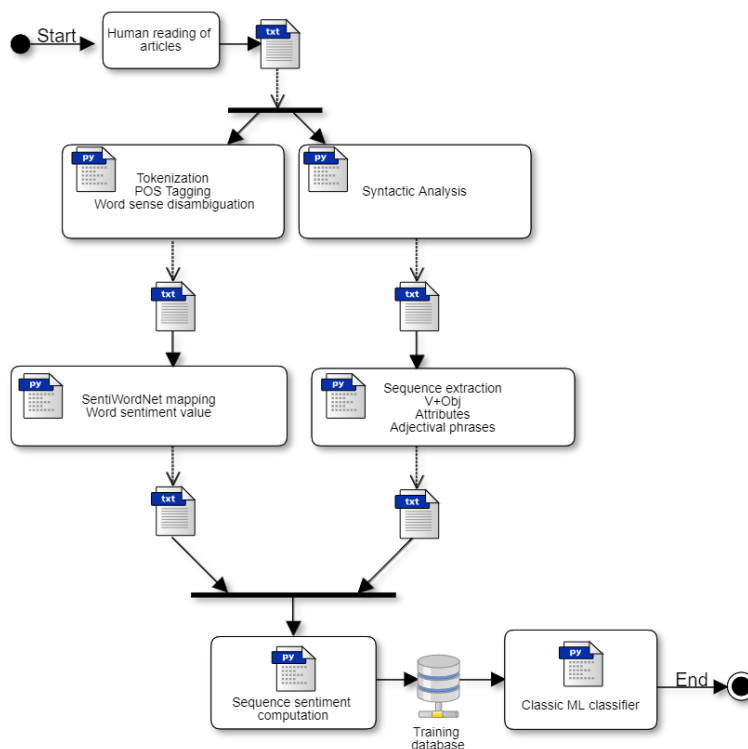


Figure 1. Linguistic feature extraction process

Extracting predicates

Feature definition was based on semantic criteria aiming at locally capturing the newswriter’s opinion in predicates. Predicates are categories designating states, actions, properties, or processes in which one or more participants take part (Real Academia Española y Asociación de Academias de la Lengua Española, 2009, pp. 63-64) [our translation]. The three linguistic categories use to extract the features are explained below. See Tables 1, 2, and 3 for some examples of these linguistic categories.

Attributes of copulative and pseudo-copulative verbs

The Spanish copulative and pseudo-copulative verbs considered here are *ser*, *parecer*, *haber*, *ponerse*, *sentirse*, *andar*, *estar*, and *tener*. As these verbs do not imply actions, their sentiment value does not affect that of the attribute.

The attribute of these verbs denotes properties or states that are told about the referent of a noun’s or sentence’s referent (Real Academia Española y Asociación de Academias de la Lengua Española, 2009, p. 2773) [our translation].

Attributes refer to characteristics of the subject, that is, adjectival concepts that can be designated with an adjective (e.g., John is tall; my children were happy), with a noun or noun phrase (John is a soldier), or with any adjectival phrase (John is from Madrid) (Real Academia Española y Asociación de Academias de la Lengua Española, 1973, p. 364) [our translation]. It can be seen in these examples that any relevant information is provided by the attribute (i.e., the aspect). This information describes the subject (i.e., the entity), and also expresses the opinion of the speaker about the entity. It is this opinion about the relevant entity that we are trying to capture here so it can be annotated with a sentiment value and contribute to classify a document as positive or negative.

Table 1. Examples of attributes

(Pseudo-)copulative	verb Attribute
to have	a better house
to be	an unacceptable alternative

Verb + direct object

The pattern transitive verb + direct object was also chosen. Many transitive verbs imply an action and require a direct object. This structure constitutes another type of highly frequent predicate. It conveys potentially substantial information about the opinion of the newswriter. A semantic analysis of transitive verbs can be done based on four criteria (Real Academia Española y Asociación de Academias de la Lengua Española, 2009, p. 2603): a) their mode of action, b) the notional class they belong to, c) the lexical nature of its direct object, and d) the semantic interpretation of its direct object [our translation]. These criteria suggest, on one hand, the important semantic connection between the constituents of the transitive predicative verbal phrase. On the other hand, the quality of the information that each phrase constituent contributes is remarkable.

245

Table 2. Examples of verb + direct object

Verb	Direct object
to affect	the living conditions of his family
to build	a better country

Adjectival phrases

Thirdly, we selected the pattern of adjectival phrases (adjective, adjective + adverb, adverb + adjective). We expected adjectives and adverbs and their syntactic dependency to convey sentiments and opinions (Liu, 2015, p. 60). It has been shown by Taboada, Brooke, Tofiloski, Voll, and Stede (2011) that the syntactic dependencies of this pattern allow to gradually work with sentiment values. For example, values from extremely negative to extremely positive can be assigned according to the adverbial modification of the adjective.

Table 3. Examples of adjectival phrases

Adjectival phrases
more vulnerable
very negative

246

Syntactic analysis is crucial for the extraction of the three linguistic characteristics described above. We used Freeling to obtain a dependency analysis, which became the input of the scripts used to extract sequences according to the defined patterns. These sequences represent the semantics of the documents and are used to build the features that populate our training database.

Notice that we are not keeping copulative verbs as features because they end up not contributing to the meaning of the text; we only keep their attributes. For example, in the sentence poverty creates much violence and is an unfortunate event, the difference in the contribution of meaning of create compared to is can be noted. It is clear that in the sequence “creates much violence” the verb plays an important role with regard to the analysis of opinion; it reinforces the negative value of the direct object “much violence”. The verb “is”, however, does not have any effect on the value of the attribute in the sequence “is an unfortunate event”; the copulative verb neither reinforces nor weakens the attribute sentiment value.

Representing documents as vectors

Each document is modeled as a feature vector and then it is passed to the machine learning algorithms to be classified as positive or negative. We start obtaining feature vectors by converting each extracted sequence into a real number that informs about its sentiment value. For each word of the extracted sequences, two

real numbers are generated, one represents the positive sentiment of the word and the other one represents its negative sentiment.

Then, two additional steps take place: a) determining an overall sentiment value for multiword sequences, and b) factoring the effect of sentiment or polarity commutators, such as negations.

Sentiment weighting rules for multiword sequences

The features defined in this study mostly represent multiword sequences. Therefore, it is necessary to find a way to compute the overall sentiment value of the sequence.

This overall value was achieved by defining a weighting rule for each syntactic pattern. These rules are described below.

- **Weighting rule for verb + direct object:** a) Out of the two sentiment values of the verb (i.e., negative and positive), take the highest value as the sentiment of the verb; b) out of the two sentiment values of each word in the direct object, take the highest value as the sentiment of the word; c) add the sentiment values of the words in the direct object and take this number as the sentiment value of the object; d) apply the following rules:

Rule 1: $V_p + Obj_p = P$

Where V_p stands for verb with positive value, Obj_p stands for object with positive value, and P means final positive value of the sequence. This rule determines that if the final value of both the verb and the object is **positive**, then the overall sentiment value of the sequence is positive. See Table 4 for an example of this rule.

Table 4. Example of weighting rule 1

Verb	Pos. Value	Neg. Value	Direct object	Pos. Value	Neg. Value
improve	1	0	quality	0.375	0
Verb + direct object				Pos. Value	Neg. Value
Improve quality				1.375	0

Rule 2: $V_p + Obj_n = P$

Where V_p stands for verb with positive value, Obj_n stands for object with negative value, and P means final positive value of the sequence. This rule determines that if the final value of the verb is positive but the final value of the object is negative, then the overall sentiment value of the sequence is **positive**. See Table 5 for an example of this rule.

Table 5. Example of weighting rule 2

Verb	Pos. Value	Neg. Value	Direct object	Pos. Value	Neg. Value
meet	1	0	the requirements	0	0.125
Verb + direct object				Pos. Value	Neg. Value
meet the requirements				1.125	0

248

Rule 3: $V_n + Obj_p = N$

Where V_n stands for verb with negative value, Obj_p stands for object with positive value, and N means final negative value of the sequence. This rule determines that if the final value of the verb is negative but the final value of the object is positive, then the overall sentiment value of the sequence is **negative**. See Table 6 for an example of this rule.

Table 6. Example of weighting rule 3

Verb	Pos. Value	Neg. Value	Direct object	Pos. Value	Neg. Value
lose	0	1	the job	0.625	0
Verb + direct object				Pos. Value	Neg. Value
lose the job				0	1.625

Rule 4: $V_n + Obj_n = P$

Where V_n stands for verb with negative value, Obj_n stands for object with negative value, and P means final positive value of the sequence. This rule determines that if the final value of both the verb and the object is negative, then the overall sentiment value of the sequence is **positive**. See Table 7 for an example of this rule.

Table 7. Example of weighting rule 4

Verb	Pos. Value	Neg. Value	Direct object	Pos. Value	Neg. Value
reduce	0	1	the huge inequality in the country	0.125	0.625
Verb + direct object				Pos. Value	Neg. Value
reduce the huge inequality in the country				1.625	0

- **Weighting rule for attributes and adjectival phrases:** the positive value of the sequence is obtained by adding the positive values of the words, and the negative value of the sequence is obtained by adding the negative values of the words. See Table 8 for an example of this rule.

249

Table 8. Computing adjectival phrases

Verb	Pos. Value	Neg. Value	Direct object	Pos. Value	Neg. Value
much	0	0	worse	0.25	0.75
Verb + direct object				Pos. Value	Neg. Value
much worse				0.25	0.75

Polarity commutators

- **Negations:**

Negations are an important factor for sentiment analysis because they often commute the sentiment. Table 9 shows an example of how the negation

“no” before a sequence of the type verb + object commutates the sequence’s polarity that was initially positive.

Table 9. Example of the negation “no” commutating a verb + object sequence

Verb	Pos. Value	Neg. Value	Direct object	Pos. Value	Neg. Value
meet	1	0	the requirements	0	0.125
Verb + direct object				Pos. Value	Neg. Value
meet the requirements				1.125	0
Negation + verb + direct object				Pos. Value	Neg. Value
(no) do not meet the requirements				0	1.125

250

In this study, we considered the Spanish negations “no” and “ni”. We dealt with non-adjacent negations with a simple rule: if there is negation anywhere preceding a relevant syntactic pattern in the same sentence, the negation works and commutates the polarity of the sequence. The effect of the negation stops after commutating the sentiment of the first relevant pattern found, or when the sentence ends. Table 10 shows an interesting case of non-adjacent negation.

Table 10. Example of non-adjacent negation

Verb	Pos. Value	Neg. Value	Direct object	Pos. Value	Neg. Value
disminuido (decreased)	0	1	la enorme desigualdad que hay en el país (the huge inequality in the country)	0.125	0.625
Verb + direct object				Pos. Value	Neg. Value
disminuido la enorme desigualdad que hay en el país (decreased the huge inequality in the country)				1.625	0
Negation + verb + direct object				Pos. Value	Neg. Value
No se ha disminuido la enorme desigualdad que hay				0	1.625

en el país (has not been decreased the huge inequality in the country)		
--	--	--

- **Within-sentence commutators:** Besides negations, other words are also considered polarity commutators. Although the rules for these commutators do not always apply, due to the unpredictable and dynamic nature of language, we did observe trends that we tried to capture. Within-sentence commutators commute the sentiment of all the syntactic patterns found in the sentence. The commutators considered in this category are *pese a* (despite), *poco*, (little), *ningún* (none), *nunca* (never), and *pero* (but).
- **Out-of-sentence commutators:** These commutators change the polarity of the preceding sentence. Automatically determining the scope of a polarity commutator is not a trivial task (Liu, 2015, pp. 122-123). Sometimes, having a negation nearby a sequence does not require a sentiment change. This is why the assigned sentiment value of a sequence is not always right. However, we used a simple strategy that proved to be useful for out-of-sentence commutators. The sentiment value of a sequence was inverted if it preceded a commutator such as *sin embargo* (however/notwithstanding) and *aunque* (although).

Although the literature reports on the use of other commutators such as *except*, *instead of*, *on the other hand*, etc. (Liu, 2015, p. 129), we do not include them in our rules because we observed they are less frequent (see Figure 2) and less predictable; only the commutators mentioned above were used in the present study.

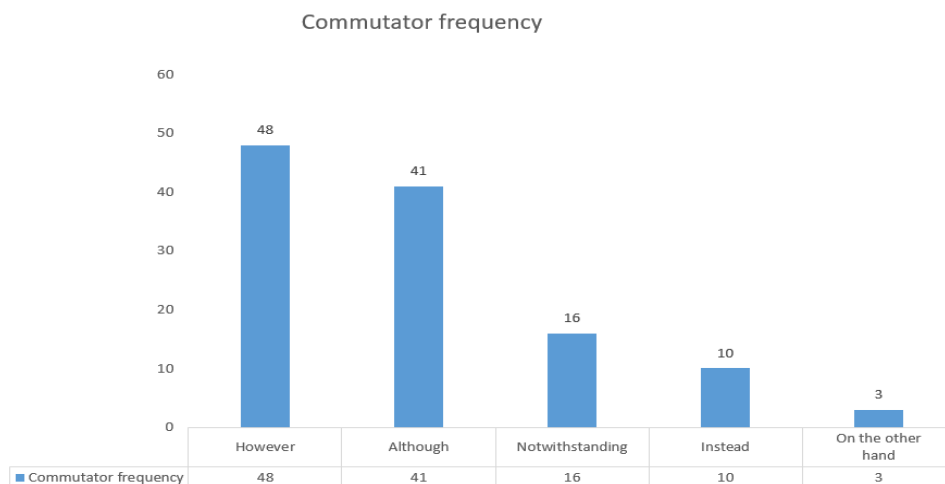


Figure 2. Frequency of commutators

Feature vectors then encode information related to the positive and negative sentiment of the relevant sequences, to the effect of the polarity commutators, and to the section of the document where a sequence appear. They constitute a training database with 18 features per vector. A schematic view of the database is shown in Table 11.

252

Table 11. Schematic view of the training database

Introduction (0-25% of the text)		Body (25%-75%)		Conclusions (75%-100%)	
positive	negative	positive	negative	positive	Negative
V+O_posS1	V+O_negS1	V+O_posS2	V+O_negS2	V+O_posS3	V+O_negS3
Atr_posS1	Atr_negS1	Atr_posS2	Atr_negS2	Atr_posS3	Atr_negS3
SAdj_posS1	SAdj_negS1	SAdj_posS2	SAdj_negS2	SAdj_posS3	SAdj_negS3

Where V+O_posS1 is a feature obtained by adding the positive sentiment values of the sequences of the type verb + object in section S1 (i.e., “introduction”) of a document; Atr_posS1 is obtained by adding the positive sentiment values of the attribute sequences in section S1, and SAdj_posS1 is obtained by adding the positive sentiment values of the adjectival phrases in section S1. And so on to obtain the other 15 features.

VALIDATION AND ALGORITHMS

Validation

All the machine learning algorithms in this study were tested following a bootstrapping validation methodology. This consists of using 70% of the data for training and 30% for testing. This training/testing cycle was iterated 20 times for each experiment, taking different data each time for training and testing. During each iteration, each algorithm was trained and tested, and then the average of the 20 iterations was used to determine error, efficiency, sensitivity, and specificity.

Error metric

The set of measures included to evaluate the performance of the proposed approach consists of:

- $Sensitivity = \frac{TP}{TP+FN}$ (*Recall*)
- $Specificity = \frac{TN}{TN+FP}$
- $Efficiency = \frac{TP+TN}{TP+TN+FP+FN}$
- $Error = 1 - Efficiency$

Where TP = true positives, TN = true negatives, FP = false positives and FN = false negatives.

In this study, we use Support Vector Machines (SVM) and Naïve Bayes to test the suitability of linguistic and discursive features. Beyond that these algorithms are used by other scholars for tasks of sentiment analysis classification (Pang *et al.*, 2002), with the choice of linguistic and discursive features, we intend to improve the accuracy obtained with a bagof- words approach, which we consider the baseline.

EXPERIMENTS AND RESULTS

This section presents the results of nine experiments. We start reporting the results of a bag-of-words approach (BOW), which we consider the baseline. The efficiency reached with each approach is reported in Table 12.

The baseline: Classification with Bag-of-words

This is the simplest approach for text classification. It does not require feature extraction and uses as input the frequency of words with their positive and negative sentiment values. Being simple and commonly used, it can be set as the baseline to compare the performance of other techniques.

We used this approach to classify our news articles following the process below:

- Tokenization
- Part-of-speech tagging
- Word sense disambiguation
- Automatic sentiment annotation using SentiWordNet
- Positive and negative word count
- Classification according to the most voted sentiment value

254

Here also, out of the two sentiment values of each word in SentiWordNet (i.e., negative or positive), the highest value was taken as the sentiment of the word.

Although this strategy reaches a sensitivity of 94%, it presents low efficiency. This is because the model classifies badly the negative documents, which is reflected on a specificity of only 12%.

We aim at outperforming this bag-of-words approach using a vector space with the linguistic features described above, and using SVM and Naïve Bayes algorithms. The algorithms and the strategies used are described below:

Classification with Support Vector Machines

Table 12 shows the best run to classify the news articles using a SVM with Gaussian kernel. The parameters γ and C are tuned through a grid search method, which searches in a two-dimension space, $\gamma \in \{0.001, 0.01, 0.1, 1\}$ and $C \in \{1, 10, 100, 1000\}$, which total $4 \times 4 = 16$ combinations. This same grid search method was used for all the experiments with Gaussian Kernel SVM.

Table 12. Results using the whole feature set

Method	C	γ	Efficiency	Error	Sensitivity	Specificity
BOW	-	-	0.56	0.44	0.94	0.12
SVM	1	0.1	0.63+/-0.05	0.37+/-0.05	0.69+/-0.1	0.57+/-0.11
Naïve Bayes	-	-	0.5+/-0.08	0.5+/-0.08	0.35+/-0.21	0.69+/-0.14
PCA-SVM	100	0.01	0.64+/-0.05	0.36+/-0.05	0.7+/-0.09	0.57+/-0.11
SFFS-SVM- SWN⁷	1000	0.001	0.67+/-0.05	0.33+/-0.05	0.72+/-0.08	0.61+/-0.1
SFFS-SVM- MLS⁸	10	0.01	0.69+/-0.05	0.31+/-0.05	0.72+/-0.08	0.66+/-0.09

255

Compared to the baseline (bag-of-words), these results show an improvement in the classification of negative documents with an SVM using predicate-based features (specificity = 57%). However, the classification of positive documents decreases (sensitivity= 69%). The efficiency with this approach outperforms the baseline in 7%.

⁷ Sequential Floating Forward Selection with Support Vector Machines using SentiWordNet

⁸ Sequential Floating Forward Selection with Support Vector Machines using ML-Senticon

Classification using Naïve Bayes

The Naïve Bayes model shows a noticeable improvement in the classification of negative documents in comparison to the baseline. Specificity reaches 69%, but sensitivity is low, which affects its efficiency (50%). These results disagree in a certain way with other reports about this model (Pang *et al.*, 2002). It is important to consider, however, that the literature reports experiments with Naïve Bayes for sentiment analysis in documents of less complexity than the socio-political news articles that we are dealing with here.

As Naïve Bayes worked better in the classification of negative documents, we tried combining both SVM and Naïve Bayes together. Indeed, the classification of negatives improved that of the SVM alone, but efficiency decreased due to the poor classification of positives.

In addition to this, we tested the predictive power of each linguistic category separately using a SVM. That is, we first classified only with the pattern verb + object, then with attributes only, and then with adjectival phrases only. Table 13 shows the best runs for this strategy per linguistic category.

Table 13. Results using a SVM with verb + object, attributes and adjectival phrases independently

Pattern	C	γ	Efficiency	Error	Sensitivity	Specificity
Verb + Object	10	0.01	0.63+/-0.05	0.37+/-0.05	0.74+/-0.08	0.51+/-0.1
Attributes	10	0.1	0.57+/-0.05	0.43+/-0.05	0.68+/-0.1	0.44+/-0.12
Adjectival Phrases	10	0.01	0.56+/-0.06	0.44+/-0.06	0.81+/-0.01	0.28+/-0.1

It can be seen in Table 13 that classification with the pattern verb + direct object yields an efficiency of 63%. This equals the efficiency obtained with the three linguistic categories together. On the other hand, classification with adjectival phrases and attributes separately does not perform well.

Alternatively, although there seems not to be curse of dimensionality in this feature space, we tried reduction dimension with principal component analysis (PCA) and sequential floating forward selection (SFFS). The details and results of this strategy are shown below.

Classification using dimension reduction techniques

- **Principal Components Analysis**

This approach with PCA and SVM increases efficiency in 1% compared with the 63% reached by the classification with SVM and Gaussian kernel only. This is an improvement in the classification of positive documents, which can be seen in the sensitivity increasing in 1%, while specificity did not change.

- **Sequential Floating Forward Selection**

We implemented SFFS with a wrapper objective function and an SVM.

The results improved with 2 features selected by the technique, out of the original 18 features. These two features are:

- **SAdjPosS1**: positive sentiment of adjectival phrases in the “introduction” (section S1) of the documents.
- **V+ObjNegS3**: negative sentiment of sequences of the type verb + direct object in the “conclusions” (section S3) of the documents.

Table 12 shows the best classification run using these two characteristics and an SVM.

257

A combination of SFFS and SVM with Gaussian kernel yields an increase of 4% in the efficiency of classification of news articles, and a 67% is reached. Likewise, sensitivity reaches 72% and specificity 61%. The techniques stated before outperformed other five models that we also tested, namely, ANN Multilayer Perceptron (Efficiency = 51%), KNN (Efficiency = 61%), Random Forests (Efficiency = 61%), Gradient Boosting (Efficiency = 60%), and SVM with linear kernel (Efficiency = 59%).

DISCUSSION

Even though 67% is not an optimal efficiency and there is still room for improvement, the experiments and strategies described here proved to be useful toward making progress for document sentiment analysis. An important achievement of this work is the automatic construction of the feature vector based on linguistic information; no human annotation was involved.

Likewise, the efficiency of 63% obtained with an SVM with Gaussian kernel, predicate-based features, and no dimension reduction outperformed the bag-of-words baseline also. This suggests a higher power of the features defined here to capture the sentiment of a document.

In addition to this, the document representation proposed in this paper improved the results of experiments that we carried out following the methodologies reported in mainstream works in the field such as Pang *et al.* (2002), Turney (2002), and Nasukawa and Yi (2003).

The efficiency of the various experiments described here seems to have recurrently been affected by a difficulty in classifying the negative documents. We devote here some lines to discuss the possible causes of this as well as other related issues.

Issues related to news article classification

From a qualitative standpoint, we hypothesize that the low specificity of the models is caused by the journalistic style of the texts. It is not uncommon that a negative article contains a great deal of sequences with positive sentiment. However, using rhetorical or discursive strategies, the newswriter makes the article negative with the use of a few words or expressions sometimes at the end of the text. There are also other important challenges related not only to the semantics of the text but also to pragmatic issues that hinder a higher efficiency in a classification task of socio-political texts. We discuss here some specific key issues to consider for further work with document sentiment analysis.

258

- **Co-reference resolution and named-entity recognition:** Let us consider this example: *“...to keep his great promise about completely eradicating absolute poverty. It is a false argument. If we analyze his fancy speech, it can be realized that this is impossible, blamable and gullible.”* This text is clearly negative but it contains positive sequences such as:
 - keep his great promise (verb + direct object)
 - eradicating poverty (verb + direct object)
 - analyze his fancy speech (verb + direct object)

Entity recognition and relation is key here to determine that both expressions with negative sentiment such as “impossible”, “blamable” and “gullible” and expressions with positive sentiment such as “keep his great promise”, “eradicating poverty” and “fancy speech” refer to the same entity,

i.e., “his argument”. Therefore, the great challenge here is to determine that the latter group’s polarity needs to be commutated.

- **Aspect-based sentiment analysis (ABSA):** This approach to sentiment analysis may enable a more finegrained analysis of the document. This makes sense because most of the news articles mention multiple entities, but there is only one central matter towards which the newswriter has a general sentiment. A correct identification of these entities and their aspects may lead to detect the general sentiment of the document.

- **The scope of polarity commutators:** In this work, we managed to resolve the effect of some of the polarity commutators (e.g., negations) in evaluating sentiment.

However, much work still needs to be done to accurately model the scope and effects that commutators have in a document. Other tasks such as entity recognition and relation or ABSA may collaterally help with this issue. The problem is not trivial, though, as a single commutator may well change the polarity of just one word, but also of the whole document.

- **The socio-political domain:** The news articles classified in this study belong to the socio-political context. This domain entails complex semantic and even pragmatic issues that may not be typical of other documents, such as tweets or reviews. For example, many articles contain references to facts and figures in the form of percentages such as extreme poverty went from 15% to 12% under this government. World or context knowledge is necessary to determine if such figures carry a positive or a negative sentiment.

- **The linguistic features:** While the linguistic features defined here proved to be promising, other linguistic patterns are still to be defined. For example, rhetorical questions and exclamations may express sarcasm or even commute other sequences’ sentiment. The following question illustrates this: *Did the government even keep their wonderful promises?* The sequence *keep their wonderful promises* has an apparent positive sentiment, but there are chances that it may be negative in this context.

Inherited error

Inherited error is an error carried forward from a previous step in a sequential process. In this work, by inherited error we mean the error reported by the authors of the processing tools used in this research. These tools are Freeling's UKB for word sense disambiguation in Spanish, SentiWordNet for sentiment value annotation, and Freeling's pos tagging and syntactic parsing for annotating and extracting the syntactic patterns for the features defined here. We value the fact that these tools made possible these experiments. However, we also recognize that the inherited error from these tools may directly affect both the overall efficiency of our models and the power of the features. Let us use Table 14 as an illustration of one aspect of the problem. The table illustrates the sentiment value assigned by SentiWordNet. The third column shows the actual sentiment of the word, manually determined, in the context of the article the word was taken from. These discrepancies may be due to the complexity of the documents, but it is also likely that the values in SentiWordNet⁹ and the word sense disambiguation error play a role here.

Table 14. Discrepancies between SentiWordNet and in-context sentiment

Word	Pos. Value	Neg. Value	In-context sentiment
illegal	0.625	0	Negative
evasion	0.375	0.25	Negative
concern	0.375	0.25	Negative
natural	0.75	0.125	Neither positive nor negative
corruption	0.5	0.25	Negative

260

Therefore, in the light of the discrepancies between SentiWordNet sentiment and in-context sentiment described above, we ran a final experiment using ML-Senticon (Cruz, Troyano, Pontes & Ortega, 2014), a sentiment lexicon for Spanish. Although this resource has only 11,542 entries (cf. SentiWordNet with 117,659), it is relevant for the present work as it was built for the Spanish language.

The huge difference in the number of entries between these two resources may be due to the fact that SentiWordNet is synset-oriented, that is, the same word

⁹ As there were verbs in SentiWordNet with 0 in both their negative and positive values, we enriched SentiWordNet by manually annotating 959 verbs with sentiment value. This resource can be obtained by emailing Antonio Jesús Tamayo Herrera to antonio.tamayo@udea.edu.co.

form may appear several times depending on the number of different senses the word has. ML-Senticon, on the other hand, is lemma-oriented and it only includes words that are clearly negative or positive with independence of the context. Therefore, by using ML-Senticon, there is no need of word sense disambiguation and sense mapping. We just match lemmas in each document with lemmas in ML-Senticon and take the sentiment value when a matching occurs.

Thus, we followed the methodology described above (without sense disambiguation and mapping) to build the feature vector using ML-Senticon. Then we classified using SFFS and an SVM with Gaussian kernel. The SFFS technique selected three features:

- **SAdjPosS1:** positive sentiment of adjectival phrases in the “introduction” (section S1) of the documents.
- **SAdjNegS2:** negative sentiment of adjectival phrases in the “body” (section S2) of the documents.
- **V+VObjNegS3:** negative sentiment of sequences of the type verb + direct object in the “conclusions” (section S3) of the documents.

Table 12 shows the best classification run using these three characteristics. Efficiency in this experiment increases 2% (0.69) compared to the same experiment with SentiWordNet and 12% above the bag-of-word baseline with ML-Senticon. It can be noticed that negative documents are better classified here (Specificity= 0.66), and there is more of a balance with the classification of positive documents (Sensitivity= 0.72). This suggests that the subtask of sense disambiguation and mapping might be affecting specificity.

Likewise, the fact that SentiWordNet gives positive or negative values to words with neutral sentiment (e.g., territory, district) may distort the computation of weights in predicates and hence affect the classification outcome. Lastly, although ML-Senticon still seems a limited resource, we anticipate that greater entry coverage by this or other similar opinion mining resources for Spanish will contribute to improve precision in sentiment-oriented classification tasks.

CONCLUSIONS

This work presented a methodology to test linguistic features for sentiment analysis in socio-political news articles in Spanish. Classification of texts was carried out using natural language processing, Support Vector Machines, Naïve Bayes models, and dimension reduction techniques.

Determining the sentiment of this type of documents is a novelty, as the vast majority of works have been done in texts such as tweets and movie, product or service reviews. We proved that successful methodologies reported in widely cited works (Nasukawa & Yi, 2003; Pang *et al.*, 2002; Turney, 2002) do not perform well in the documents explored here. The present work contributes to the resolution of the sentiment analysis problem by characterizing complex documents in Spanish and combining linguistic and computational techniques. This challenging configuration of factors has not been addressed sufficiently by the scientific community in the field.

The whole characterization and classification process presented here was completely automatic. This makes a difference with most of the reported works, which include some part of manual annotation. We also managed to devise robust features that include features used by benchmark researchers in other projects.

262

With all the complexity of the document set, our best experiment reached 69% of efficiency in the classification by sentiment. This efficiency was obtained using sequential forward floating selection and a SVM with Gaussian kernel. Our approach to classification using sentence predicates to build the feature vector outperformed the bag-of-words baseline in 12%.

Of course, there are also limitations in the selected linguistic features. The syntactic rules not always capture all the modification of the relevant syntactic patterns. Sometimes preceding or following phrases are left out, which may be determining to assign a more accurate sentiment value. Extending the rule to include other phrases may have an effect contrary to the expected one, though, as irrelevant sequences might be captured and undermine their efficiency. With regard to polarity commutators, a more sophisticated approach is needed to better define their scope and effects.

REFERENCES

- Atserias, J., Casas, B., Comelles, E., González, M., Padró, L., & Padró, M. (2006). Freeling 1.3: Syntactic and semantic services in an open-source NLP library. In *Proceedings of LREC 2006* (pp. 2281-2286). Retrieved from <https://pdfs.semanticscholar.org/60a6/9dab3b7d53ba1032bd698ee62935d3d56921.pdf>.
- Baccianella, S., Esuli, A., & Sebastiani, F. (2010). SentiWordNet 3.0: An Enhanced Lexical Resource for Sentiment Analysis and Opinion Mining. In *Proceedings of LREC 2010* (pp. 2200-2204). Retrieved from https://www.researchgate.net/publication/220746537_SentiWordNet_30_An_Enhanced_Lexical_Resource_for_Sentiment_Analysis_and_Opinion_Mining.
- Balahur, A., Steinberger, R., Kabadjov, M., Zavarella, V., Van Der Goot, E., Halkia, M., ... Belyaeva, J. (2010). Sentiment analysis in the news. In *Proceedings of the 7th International Conference on Language Resources and Evaluation (LREC'2010)* (pp. 2216-2220). Retrieved from <https://arxiv.org/abs/1309.6202>.
- Becker, I., & Aharonson, V. (2010). Last but *definitely* not Least: On the Role of the Last Sentence in Automatic Polarity-Classification. In *Proceedings of the ACL 2010 Conference Short Papers* (pp. 331-335). Retrieved from <https://www.aclweb.org/anthology/P/P10/P10-2061.pdf>.
- Carreras, X., Chao, I., Padró, L., & Padró, M. (2004). Freeling: An Open-Source Suite of Language Analyzers. In *Proceedings of LREC 2004* (pp. 239-242). Retrieved from <http://nlp.lsi.upc.edu/papers/carreras04.pdf>.
- Cruz, F. L., Troyano, J. A., Pontes, B., & Ortega, F. J. (2014). ML-SENTICON: Un lexicón multilingüe de polaridades semánticas a nivel de lemas. *Procesamiento del Lenguaje Natural*, (53), 113-120.
- Gamon, M. (2004). Linguistic correlates of style: authorship classification with deep linguistic analysis features. In *Proceedings of the 20th International Conference on Computational Linguistics* (Article No. 611). Retrieved from <https://dl.acm.org/citation.cfm?id=1220443>.
- Joshi, M., & Penstein-Rosé, C. (2009). Generalizing dependency features for opinion mining. In *Proceedings of the ACL-IJCNLP 2009 Conference short papers* (pp. 313-316). Retrieved from <https://dl.acm.org/citation.cfm?id=1667680>.
- Li, Q., Wang, T., Li, P., Liu, L., Gong, Q., & Chen, Y. (2014). The effect of news and public mood on stock movements. *Information Sciences*, 278, 826-840. doi : 10.1016/j.ins.2014.03.096.
- Li, S., Lee, S. Y. M., Chen, Y., Huang, C.-R., & Zhou, G. (2010). Sentiment classification and polarity shifting. In *Proceedings of the 23rd International*

Conference on Computational Linguistics (pp. 635-643). Retrieved from <https://dl.acm.org/citation.cfm?id=1873853>.

- Li, X., Xie, H., Chen, L., Wang, J., & Deng, X. (2014). News impact on stock price return via sentiment analysis. *Knowledge- Based Systems*, 69, 14-23. doi: 10.1016/j.knosys.2014.04.022.
- Liu, B. (2012). Sentiment Analysis and Opinion Mining. *Synthesis Lectures on Human Language Technologies*, 5(1), 1-167. doi: 10.2200/S00416ED1V01Y201204HLT016.
- Liu, B. (2015). *Sentiment analysis: Mining opinions, sentiments, and emotions*. Cambridge: University Press.
- Mitra, A., Parthasarathi, V., & Mamani, S. (2015). Using narbs for political prediction. In P. Kommers & P. Isaías (Eds.), *Proceedings of the 13th International Conference on e-Society* (pp. 274-278). Retrieved from <http://www.iadisportal.org/digital-library/using-narbs-for-political-prediction>.
- Mohammad, S., Kiritchenko, S., Sobhani, P., Zhu, X., & Cherry, C. (2016). SemEval-2016 task 6: Detecting stance in tweets. In S. Bethard, M. Carpuat, D. Cer, D. Jurgens, P. Nakov, & T. Zesch (Eds.), *Proceedings of the 10th International Workshop on Semantic Evaluation (semEval-2016)* (pp. 31-41). Stroudsburg: Association for Computational Linguistics.
- Mullen, T., & Collier, N. (2004). Sentiment analysis using support vector machines with diverse information sources. In *Proceedings of the 2004 Conference on Empirical Methods in Natural Language Processing*. Retrieved from <http://www.aclweb.org/anthology/W04-3253>.
- Nasukawa, T., & Yi, J. (2003). Sentiment analysis: Capturing favorability using natural language processing. In *Proceedings of the 2nd International Conference on Knowledge Capture* (pp. 70-77). Retrieved from <https://dl.acm.org/citation.cfm?id=945658&dl=ACM&coll=DL>.
- Ng, V., Dasgupta, S., & Arifin, S. (2006). Examining the role of linguistic knowledge sources in the automatic identification and classification of reviews. In *Proceedings of the COLING/ACL on main conference poster sessions* (pp. 611-618). Retrieved from <https://dl.acm.org/citation.cfm?id=1273073&picked=prox>.
- Osgood, C. E., Suci, G. J., & Tannenbaum, P. H. (1957). *The Measurement of Meaning*. Illinois: University of Illinois Press.
- Padró, L. (2011). Analizadores Multilingües en FreeLing. *Linguamática*, 3(2), 13-20.
- Padró, L., Collado, M., Reese, S., Lloberes, M., & Castellón, I. (2010). FreeLing 2.1: Five years of open-source language processing tools. In N. Calzolari, K. Choukri, B. Maegaard, J. Mariani, J. Odijk, S. Piperidis, M. Rosner, & D.

- Tapias (Eds.), *Proceedings of the Seventh International Conference on Language Resources and Evaluation: LREC 2010* (pp. 931-936). Valletta, Malta: European Language Resources Association.
- Padró, L., & Stanilovsky, E. (2012). FreeLing 3.0: Towards Wider Multilinguality. In *Proceedings of LREC 2012* (pp. 2473-2479). Retrieved from <https://pdfs.semanticscholar.org/b6de/1f38185454dd32e6ff4ac86405d46681d0ee.pdf>.
- Pang, B., & Lee, L. (2008). Opinion mining and sentiment analysis. *Foundations and Trends in Information Retrieval*, 2(1-2), 1-135.
- Pang, B., Lee, L., & Vaithyanathan, S. (2002). Thumbs up? Sentiment Classification using Machine Learning Techniques. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP 2002)* (pp. 79–86). Retrieved from <https://aclanthology.info/papers/W02-1011/w02-1011>.
- Parodi, G. (2010). La organización retórica del género Manual a través de cuatro disciplinas: ¿cómo se comunica y difunde la ciencia en diferentes contextos universitarios? *Boletín de Lingüística*, 22(33), 43-69.
- Quiroz, G., Tamayo, A., & Zuluaga, J. F. (2017). Cuestiones metodológicas y técnicas en la recolección de un corpus de prensa con la palabra “pobreza”. In A. Chiquito & G. Quiroz (Eds.), *Pobreza, lenguaje y medios en América Latina* (pp. 21-43). Bern: Peter Lang.
- Real Academia Española y Asociación de Academias de la Lengua Española. (1973). *Esbozo de una nueva gramática de la lengua española*. Madrid: Espasa Calpe.
- Real Academia Española y Asociación de Academias de la Lengua Española. (2009). *Nueva gramática de la lengua española*. Madrid: Espasa.
- Schumaker, R. P., Zhang, Y., Huang, C.-N., & Chen, H. (2012). Evaluating sentiment in financial news articles. *Decision Support Systems*, 53(3), 458-464. doi: 10.1016/j.dss.2012.03.001.
- Swales, J. M. (2004). *Research genres: Explorations and applications*. Cambridge: Cambridge University Press.
- Taboada, M., Brooke, J., Tofiloski, M., Voll, K., & Stede, M. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2), 267-307.
- Turney, P. D. (2002). Thumbs up or thumbs down?: Semantic orientation applied to unsupervised classification of reviews. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics (ACL)* (pp. 417–424). Retrieved from <https://dl.acm.org/citation.cfm?id=1073153>.
- Xia, R., Wang, T., Hu, X., Li, S., & Zong, C. (2013). Dual Training and Dual Prediction for Polarity Classification. In *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics (Vol. 2: Short papers)*, (pp. 521-525). Retrieved from <http://aclweb.org/anthology/P13-2093>.

- Xia, R., & Zong, C. (2010). Exploring the use of word relation features for sentiment classification. In *Proceedings of the 23rd International Conference on Computational Linguistics: Posters* (pp. 1336-1344). Retrieved from <https://dl.acm.org/citation.cfm?id=1944719>.
- Yessenalina, A., Choi, Y., & Cardie, C. (2010). Automatically Generating Annotator Rationales to Improve Sentiment Classification. In *Proceedings of the ACL 2010 Conference Short Papers* (pp. 336-341). Retrieved from <https://aclanthology.info/papers/P10-2062/p10-2062>.
- Yessenalina, A., Yue, Y., & Cardie, C. (2010). Multi-level Structured Models for Document-level Sentiment Classification. In *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing* (pp. 1046-1056). Retrieved from <https://www.aclweb.org/anthology/D/D10/D10-1102.pdf>.
- Zaidan, O., Eisner, J., & Piatko, C. (2007). Using "Annotator Rationales" to Improve Machine Learning for Text Categorization. In *Proceedings of NAACL HLT 2007* (pp. 260-267). Retrieved from <http://www.aclweb.org/anthology/N07-1033>.

SOBRE LOS AUTORES

266

Antonio Tamayo

Profesor de cátedra de ingeniería de sistemas de la Universidad de Antioquia y la Universidad de Medellín. Magíster en ingeniería. Desarrolla su docencia e investigación en las áreas de aprendizaje automático, procesamiento del lenguaje natural, estadística y lingüística computacional.

Correo electrónico: antonio.tamayo@udea.edu.co

Orcid: 0000-0002-5984-7463.

Julián Arias Londoño

Profesor asociado de la Universidad de Antioquia. PhD en Ingeniería de sistemas y automatización. Desarrolla su docencia e investigación en las áreas de analítica de datos, aprendizaje automático y procesamiento de señales.

Correo electrónico: julian.ariasl@udea.edu.co

Orcid: 0000-0002-1928-773X.

Diego Burgos

Profesor asociado de español y lingüística de Wake Forest University. PhD en Ciencias del lenguaje y lingüística aplicada. Desarrolla su docencia e investigación en las áreas de estudios y tecnologías de la traducción, terminología y procesamiento del lenguaje natural.

Correo electrónico: burgosda@wfu.edu

Orcid: 0000-0002-5784-3952.

Gabriel Quiroz

Profesor titular de traducción y lingüística de la Universidad de Antioquia. PhD en Lingüística. Desarrolla su docencia e investigación en las áreas de traducción, terminología y tecnologías del lenguaje.

Correo electrónico gabriel.quiroz@udea.edu.co

Orcid: 0000-0003-4940-3022.