

Experimenting with UD Adaptation of an Unsupervised Rule-based Approach for Sentiment Analysis of Mexican Tourist Texts

Olga Kellert^{1,*}, Mahmud Uz Zaman², Nicholas Hill Matlis³ and Carlos Gómez-Rodríguez⁴

¹Universidad de Coruña, Grupo LyS, CITIC, Depto. de Ciencias de la Computación y Tecnologías de la Información, Campus de Elviña s/n, 15071 A Coruña, Spain

²University of Göttingen, Seminar für Romanische Philologie, Humboldtallee 19, 37073 Göttingen, Germany

³Center for Free-Electron Laser Science CFEL, Deutsches Elektronen-Synchrotron DESY, Germany

⁴Universidad de Coruña, Grupo LyS, CITIC, Depto. de Ciencias de la Computación y Tecnologías de la Información, Campus de Elviña s/n, 15071 A Coruña, Spain

Abstract

This paper summarizes the results of experimenting with Universal Dependencies (UD) adaptation of an Unsupervised, Compositional and Recursive (UCR) rule-based approach for Sentiment Analysis (SA) submitted to the Shared Task at Rest-Mex 2023 (Team Olga/LyS-SALSA) (within the IberLEF 2023 conference). By using basic syntactic rules such as rules of modification and negation applied on words from sentiment dictionaries, our approach exploits some advantages of an unsupervised method for SA: (1) interpretability and explainability of SA, (2) robustness across datasets, languages and domains and (3) usability by non-experts in NLP. We compare our approach with other unsupervised approaches of SA that in contrast to our UCR rule-based approach use simple heuristic rules to deal with negation and modification. Our results show a considerable improvement over these approaches. We discuss future improvements of our results by using modality features as another shifting rule of polarity and word disambiguation techniques to identify the right sentiment words.

Keywords

Sentiment Analysis, Unsupervised, Rule-based, Sentiment dictionary

1. Introduction

The intensity of a sentiment expressed in words, sentences and paragraphs strongly depends on the syntactic context in which sentiment words like *nice* appear. For instance, *this hotel is very nice* expresses a stronger positive sentiment due to the intensifier *very* than just the statement *this hotel is nice* and the use of negation can reverse the polarity of the sentiment as in *this hotel is not nice*. Modification of sentiment words by intensifiers like *very* and negation are among the basic syntactic rules influencing the semantic interpretation of sentiment words. Our main contribution in this article is a) to exploit an unsupervised (knowledge-based) model


IberLEF 2023, September 2023, Jaén, Spain

*Corresponding author.

✉ o.kellert@udc.es (O. Kellert); mail.mahmduzzaman@gmail.com (M. U. Zaman); nicholas.matlis@desy.de (N. H. Matlis); carlos.gomez@udc.es (C. Gómez-Rodríguez)



© 2023 Copyright for this paper by its authors. Permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

 CEUR Workshop Proceedings (CEUR-WS.org)

for compositional and recursive sentiment analysis (SA) driven by basic syntactic rules (Vilares et al., 2017 [1]) and to adapt it to the formalism of Universal Dependencies (UD) which is a universal framework for annotation of grammar across different human languages [2], b) to test our adaptation on the dataset provided by the Shared Task Rest-Mex 2023 organizers [3] and finally c) to compare the results of our unsupervised approach with other unsupervised methods that use heuristic rules to address modification and negation (Hutto and Gilbert, 2014 [4]) as well as other teams that participated at the Shared Task Rest-Mex [5, 6, 3].

The remainder of this article is structured as follows. §2 reviews related work. §3 introduces the adapted formalism for syntactic operations. §4 presents experimental results of the Shared Task participation and compares the results with other unsupervised approaches that use heuristic rules. Finally, §5 concludes and discusses directions for future work.

2. Related work

The most widely used approaches to perform sentiment analysis (SA) include Machine Learning/Deep Learning based methods, lexicon-based methods and hybrid methods [7, 8].

2.1. Machine Learning/Deep Learning methods

Supervised approaches using machine learning, deep learning and pre-trained models like BERT and RoBERTa are a common practice in SA [9], [10], [11], [12], [13], [14], [15]. In machine learning and deep learning-based approaches, the dataset is separated into training and testing datasets. A training dataset is used during the training process to learn correlations between the specific input text and the sentiment polarity. The testing dataset is then used to predict sentiment polarity on the basis of learned associations during the training period. The performance of these approaches is known to be relatively high if the learning and prediction tasks are performed on a similar dataset or corpus, as in the case of the present Shared Task of the Rest-Mex 2023 [3]. The downside of these approaches is that they require a huge amount of data for learning and prediction, which makes them sub-optimal for the use in real life situations where huge amounts of data are often missing or users do not have enough expertise in training and using these models. In addition, the learned associations between input and output are often obscure to be easily understood or explained.

2.2. Lexicon-based approaches

The lexicon-based methodology makes use of a sentiment dictionary that contains sentiment words with corresponding polarity scores in order to find out the overall opinion or polarity of a sentence or text [16, 4, 1, 17]. These approaches fall into various branches depending on how they deal with syntactic rules that can shift or intensify the polarity. Non-compositional and non-recursive rule based approaches such as Vader [4] use heuristic rules to account for polarity changes due to negation, modification or other syntactic processes. Other unsupervised approaches use a more general architecture to account for SA by exploiting general syntactic rules that influence sentiment polarity [1]. Lexicon-based methods are considered to be simple to use and cost effective as there is no need to train huge amounts of data. However, they

must deal with syntactic rules that influence the polarity of sentiment words like negation and modification. We briefly describe here two unsupervised methods that deal with syntactic rules in a completely different way.

2.2.1. Non-compositional Rule Based Approach Vader

Vader is a rule-based model for SA which is considered one of the best sentiment classifiers of its kind (Mello et al., 2022 [18]). It is relatively easy to implement and does not require a large number of computational resources. Therefore, it is computationally efficient and is suitable for large-scale applications and real-time analysis. The creators of the Vader sentiment analysis system asked a number of human raters to provide sentiment scores for the sentiment words enlisted in the sentiment dictionary. The human ratings were then averaged for each word with the aim of creating a robust sentiment dictionary that is based on the opinions of many human raters. Vader was designed to handle punctuation, capitalization, emoticons, acronyms, and slang [4]. For instance, capitalization of a sentiment word like *NICE* instead of *nice* is used as an intensifier of the word being capitalized. To account for modification or negation of sentiment words, Vader uses the distance between the sentiment word and the modifier or negation. Farther modifying words or negation words have a relatively smaller effect on the sentiment word than words in close proximity.

Vader was primarily designed for English text analysis, but it also offers a multilingual sentiment analysis via translation [4]. Vader gives a probability measure between 0 and 1 (1=highest probability) for the positive, negative, neutral and compound or mixed sentiment of the input text as can be seen from the following examples in English and Spanish:

- English sentence : "VADER is VERY SMART, handsome, and FUNNY!!!"
- score output for English sentence : 'neg': 0.0, 'neu': 0.233, 'pos': 0.767, 'compound': 0.9342
- Spanish sentence : ¡¡¡VADER es MUY INTELIGENTE, guapo y DIVERTIDO!!!
- score output for Spanish sentence : 'neg': 0.0, 'neu': 0.27, 'pos': 0.73, 'compound': 0.9387

The performance of multilingual Vader was tested on a multilingual corpus including English and Portuguese texts, which has shown better results for English than for Portuguese SA [18].

2.2.2. Unsupervised, Compositional and Recursive rule-based approach

Vilares et al., 2017 [1] use the first unsupervised compositional and recursive rule-based approach for SA. Its main concept is introduced informally here. We refer the reader to the details and formalism in the paper [1].

In a nutshell, this approach exploits dependency relations between words in a sentence to deal with the scope of negation, modification and other syntactic processes that can shift, strengthen or weaken the polarity of sentiment words. In this approach, the first task is to search for sentiment words taken from a sentiment dictionary in the input text such as the sentiment word *handsome* in the sentence in Figure 1 and then to traverse the dependency tree and to check whether negation words or modifiers change the polarity of sentiment words.




Figure 1: A simplified version of a fragment of a dependency tree from Vilares et al., 2017

The system uses a set of compositional operations to propagate changes to the semantic orientations of the nodes in the tree in a particular order. The order is basically scope driven. First, changes to polarity due to intensification apply and only after, changes to polarity due to negation apply. As a consequence, negation has scope over an intensified sentiment word as is also reflected by the hierarchical structure in Figure 1. The node of the negation word is higher than the node of the intensifier. Once all relevant operations have been executed, the processed sentiment score of the sentence is stored at the root node, which is not represented for simplicity in Figure 1.

3. Our Experimentation with UD Adaptation of an Unsupervised Rule-based Approach

We experimented with an adaptation of an unsupervised compositional and recursive approach in SA (Vilares et al., 2017 [1]) to the Universal Dependencies (UD) formalism [2], as it has since become the de facto standard for multilingual dependency parsing (the approach in [1] used Universal Treebanks instead, which was the latest precursor of UD available at that time). Figure 2 shows a dependency structure for an English sentence and a CoNLL-U Format which represents word lines containing the annotation of a word/token with respect to various




Figure 2: UD formalism. (<http://universaldependencies.org/eacl17tutorial/infrastructure.pdf>).

linguistic properties such as part of speech (POS), lemma, dependency relation of the word to its head, morphological features, etc. The dependency structure and linguistic properties of word/tokens as in CoNLL-U Format are an integral part of UD.

For our experimentation of an unsupervised approach to UD formalism, we used Stanza, which is a natural language toolkit based on UD-formalism that provides a basic analysis of the input text such as lemmatization, part-of-speech (POS) and dependency parsing (Peng Qi et al., 2020 ([19])). The dependency parser is based on UD parser from Qi et al. 2018 ([20]). For simplicity, we do not represent all features associated with each word and focus only on those features that are relevant for SA. As we experimented with Mexican Spanish reviews from the Shared Task of Rest-Mex 2023 for our adaptation, our examples are taken from this dataset [3]. We did very little preprocessing of the dataset (conversion of all the uppercase letters into lowercase letters) for lemmatization and parsing reasons. Let's consider a Spanish example *No es excelente* 'It is not excellent' and the associated dictionary entries with token ids, text, lemma, POS ('upos'), morphological features ('feats'), head ids and dependency relations ('deprel'):

- first word : 'id': 1, 'text': 'no', 'lemma': 'no', 'upos': 'ADV', 'feats': 'Polarity=Neg', 'head': 3, 'deprel': 'advmod'
- second word : 'id': 2, 'text': 'es', 'lemma': 'ser', 'upos': 'AUX', 'feats': 'feats': 'Mood=Ind|...', 'head': 3, 'deprel': 'cop'
- third word : 'id': 3, 'text': 'excelente', 'lemma': 'excelente', 'upos': 'ADJ', 'feats': 'Number=Sing', 'head': 0, 'deprel': 'root'

Head ids and dependency relations play an important role in our approach as they provide information about the syntactic relation of words and the hierarchical structure of the sentence. Head ids contain information about parent-child relations. Take for instance, the negation word *no* and the copular word *es* in the previous example, which have *excelente* as their head. This means that the word *excelente* is the highest node and the children *no* and *es* are the lowest

nodes in the structure. This head-child relation can be used to define the scope of negation. If the negation is a child of a sentiment word as its head, the polarity of the sentiment word needs to be shifted.

In order to be able to calculate the polarity score of a sentence, we performed several steps that can be described in a nutshell as follows:

- Step 1 : Find sentiment words in the input text and assign polarity scores to the sentiment words
- Step 2 : Create a dictionary of head ids and their correspondent children ids
- Step 3 : Identify target words that influence the sentiment word such as negation
- Step 4 : Calculate the polarity score for the input sentence

Let us illustrate these steps by looking at the given Spanish example. First, we identify the sentiment word *excelente* in the input text and add new entries to the dictionary associated with this word, namely the `elementType`: 'count' and the polarity score or `elementScore`: 5. We use the dictionaries by SO-CAL for Spanish ([16], [1]), in which the polarity score for sentiment words ranges from -5 (the most negative) to +5 (the most positive).

- Sentence : *No es excelente* 'It's not excellent'
- Step 1 : label sentiment words
- dictionary of the sentiment word : `'id': 3, 'text': 'excelente', 'lemma': 'excelente', 'upos': 'ADJ', 'feats': 'Number=Sing', 'head': 0, 'deprel': 'root', 'elementType': 'count', 'elementScore': 5`

Step 2 consists of creating a dictionary with head ids as keys and a list of children as a key value in order to find potential polarity shifters or target words such as negation and modification. Each key-value pair of this dictionary represents a head-child tree branch as represented in Figure 3.

In the UD-formalism, the head id 0 and its child represent the highest tree branch and the child of the head id 0 and its children represent the second highest branch. In our example, the second highest branch is also the lowest branch:

- Sentence : *No es excelente*
- Step 2 : Create a dictionary with heads as keys and their correspondent children as values
- head-child-dictionary : 3: [1, 2], 0: [3]

Head-child branches represent an important unit in linguistics and NLP. In linguistics, head-child branches are better known under the terms *phrases* or *maximal projections* [?, [21] and they are used to describe syntactic properties or rules. In NLP, head-child branches have been used to define the domain for Machine Translation, for example [22]. We use head-child branches for SA and more precisely to identify target words that can shift, weaken or strengthen the polarity of sentiment words.

Step 3 consists of identifying target words that can modify the sentiment word identified in step 1. In order to achieve this goal, we loop through branches upwards and check if we can find




Figure 3: Example of a head-child tree branch

the sentiment word, negation and/or modification in the same branch. For this, we calculate the order of branches from the lowest to the highest branch associated with a sentence. In the given sentence example *no es excelente*, the sentiment word and negation are in the same branch.

Step 4 consists of calculating the polarity score for each branch upwards by applying the formula for the calculation of the polarity score in (1) from Vilares et al. 2017 [1], where the variable a equals the elementScore of a sentiment word such as *excelente*, the variable b equals a value that depends on the strength of the intensifier such as *muy* taken from a list of intensifiers and negation has a score of -4 or +4 depending on the positive or negative value of a :

- Step 4 : Calculate the polarity score for the branch 3: [1, 2]


$$a * (1 + b) + (\text{sign}(a) * -4) = \text{polarityscore} \quad (1)$$

According to the formula in (1), the polarity score for the lowest branch 3:[1, 2] equals 1, if we calculate $5*(1+0)-4$. As the highest branch simply expresses an identity relation between the root and the head of the previous branch, the polarity score remains the same, namely 1, and the calculation finishes with the highest branch. We take the polarity score of the highest branch to be the final result for the polarity calculation.

We now discuss an example in Figure 4 with more branches.

We distinguish between qualitative and quantitative modifiers. Quantitative modifiers like *big* as in *big problem* act like intensifiers as they intensify the sentiment word *problem*, similar to *very*. In our approach, quantitative modifiers can never be sentiment words and their only

1 2 3 4 5 6
 No es una comida muy buena
 not is a food very good



Step 1 = $a*(1+b) = 2*(1+0.25) = 2.3$
 Step 2 = assign 2.3 to the head
 Step 3 = $2.3 - 4 = -1.7$
 Step 4 = label -1.7 as sentence score

Figure 4: Example with several branches

function is to contribute the value b in the formula (1). Qualitative adjectives like *good*, however, describe the quality of the referent expressed by the noun, they are not quantifying and they contribute the value a in the formula (1).

Turning to our example in Figure 4, we first compute the intensification of the qualitative adjective *buena* ‘good’ by the intensifier *muy* ‘very’. We follow the idea of assigning intensifiers like *muy* a score of 0.25 [1]. The result for this calculation is $2*(1+0.25)=2.3$. We assign the score 2.3 to the nominal head *comida* ‘food’ as the result of the nominal modification. Collecting information from the lowest branch and bringing it up to the highest branch (e.g. nominal phrase) is a common step in formal grammars such as Head-driven Phrase Structure Grammar (HPSG) [?] or Minimalist Grammar ([21]). As the negation is a child of the nominal head with a polarity score 2.3, the negation has scope over the nominal head. As a result, we subtract 4 from the polarity score 2.3 of the nominal head. The output of this calculation is -1.7.

The calculation finishes with the highest branch, which expresses an identity relation between the root and its child. The calculation steps are summarized as follows:

- Sentence : *No es una comida muy buena* ‘It’s not a very good food’
- polarity score of the lowest branch : $2 * (1 + 0.25) = 2.3$
- polarity score of the higher branch : $2.3 - 4 = -1.7$
- polarity score of the highest branch : -1.7 (final polarity score)

So far we only considered short single sentences with only one sentiment word. If the input text represents a longer sentence or several sentences with more than one sentiment word as in

the case of reviews at the Shared Task Rest-Mex 2023, the polarity score needs to be calculated for the whole review text. One possibility is to calculate the mean of all sentence scores of a review (metrics a). Another option is to provide more weight to "relevant" sentences such as the last sentence of a review under the assumption that humans often say the most relevant things at the end of a review (metrics b). The third option is to consider only extreme positive or negative values (metrics c) under the assumption that humans pay more attention to extreme sentiments in a review. We illustrate the differences between the metrics by a review example and correspondent sentence scores:

- I did not like this hotel. -1
- The room was ok. 2
- The breakfast was mediocre. -1
- The service was not so great. 1
- In sum, not recommendable! -4

According to the metrics b and c, the polarity for the review is -4. According to the metrics a, the review score is -0.6. For our experiments, we used the extreme value calculation (metrics c).

4. Experimental results

We compare our adapted system of an unsupervised rule-based approach with another lexicon-based unsupervised approach Vader, which is extensively used in research and industry ([4], Mello et al., 2020). In addition, we compare our results achieved in the Shared Task (Team Olga/LyS-SALSA) with supervised approaches on the basis of the test dataset submitted to the Shared Task Rest-Mex 2023. In both comparisons, we use accuracy as our evaluation metrics. As sentiment dictionaries, we use the dictionaries used by SO-CAL for Spanish [1]. The content of these dictionaries and their parameters are not modified or tuned. In order to calculate the sentiment score for the review, we used the extreme value metrics (metrics c).

4.1. Comparison with non-compositional rule based approaches

We used the training dataset from Rest-Mex 2023 to compare our system with Vader as we do not have access to the polarity scores of the test set. The polarity scoring system used at the Shared Task is from 1 to 5, where 1 is the lowest polarity and 5 the highest. We assumed a simple mapping between our scoring system and the one used for the Shared Task (1= >-5 and <=-3, 2= >-3 and <=-1, 3= >-1 and <=1, 4= >1 and <=3, 5= >3 and <=5). We used titles from the training dataset of Rest-Mex 2023 with positive polarity (5) and negative polarity (1) to compare the classification of our system and Vader. In order to be able to compare the polarity scores from the training dataset with Vader's scoring system that uses probability measures for polarity scores, we used three comparison metrics. In the first comparison metrics (Vader 1), we defined negative or positive scores, if the value for 'pos' or 'neg' in Vader's output was >0.7. In the second comparison metrics (Vader 2), the positive and negative values were defined more strictly ('pos' or 'neg' >0.8) and finally in the third metrics, the probability for positive and negative scores equals 1 (Vader 3). Table 1 shows that our system (Olga/Lys-SALSA) is much better at SA than the best Vader system (see 0.27 vs. 0.62 for overall accuracy).

Table 1

Comparison between unsupervised approaches for SA based on Titles from the Training dataset Rest-Mex 2023

| unsupervised | Accuracy |
|----------------|----------|
| Vader1 | 0.27 |
| Vader2 | 0.13 |
| Vader3 | 0.10 |
| Olga/LyS-SALSA | 0.62 |

Table 2

Comparison between our system and supervised systems at Shared Task Rest-Mex 2023

| Teams | Accuracy(Polarity) |
|--------------------|--------------------|
| LKE-LLMAS-Team | 0.75 |
| Javier Alonso-Team | 0.74 |
| Olga/LyS-SALSA | 0.56 |

4.2. Comparison with supervised approaches

For the participation at the Shared Task Rest-Mex 2023, participants were asked to classify the polarity of a test set. We used titles of the reviews as the input text for our polarity classification, if the title contained sentiment words such as "very good restaurant, but not so cheap". Otherwise, we used the full review text for polarity classification.

The best performing approaches used in the Shared Task of SA at Rest-Mex 2023 are supervised approaches which have been trained on the training dataset and which predicted polarity scores for the test dataset of the same domain or corpus as the training dataset [3]. It is thus not surprising to see that our results are behind the results of supervised approaches as shown in Table 2.

However, it has been shown by [1] that these methods are not so robust when compared across different domains, languages and datasets. One important note regarding the comparison of our system and supervised approaches is that we do not know from the results provided to us from the organizers of the Shared Task, whether the supervised methods used both titles and reviews or just one of the two methods as the input text for the training and/or prediction task. Given that titles are very short, it is very likely that the training period was not performed on titles alone as it would probably lead to a worse performance of supervised methods. Note that our unsupervised rule-based approach does not depend much on the shortness of the text. On the contrary, in case titles contain single sentiment words like "recommendable!", no dependency parsing is required and the word score for "recommendable" directly represents the polarity score for the whole review. Finally, we would like to mention that our participation at the Shared Task was not motivated by winning the Shared Task, but by experimenting with an unsupervised rule-based approach.

5. Discussion

There are several issues that need to be accounted for in a lexicon-based SA system. One issue is the implementation of several syntactic rules that can shift, weaken or strengthen the polarity. One important point of this implementation is the scope of negation. We have suggested to use a head-child dependency relation to control for the scope of negation. As a result, in our approach the negation word *No* in the example *No! Es excelente!* ‘No! It’s excellent!’ is ignored, although the negation word stands in very close proximity to the sentiment word. This is because the negation and the sentiment word *excelente* do not have a head-child dependency relation. In other words, *excelente* is not the head of the negation word *No*. A similar case applies to the negation *no* in the sentence *Es excelente, no?* ‘It’s excellent, isn’t it?’ Note that according to Vader, close proximity of the negation to the sentiment word counts as a trigger for shifting the sentiment polarity. As a consequence, Vader would incorrectly predict a shift in polarity for the above mentioned examples with negation.

Scope of negation is not a trivial issue and has caused some headache to (computational) linguists, because the scope of negation does not only depend on the dependency relation, but also on semantic features of the elements negation has scope over. Take for instance, the difference between definite and indefinite nouns. Negation has usually scope over indefinite nouns, but not over definite nouns. This is because definite nouns express presuppositions and presuppositions are preserved under negation [23]. Consider the relevant examples: *It’s not this amazing hotel, it’s the other one* vs. *It’s not an amazing hotel*. Negation has scope over the adjective *amazing* in the latter, but not in the former example, because in the former example it is presupposed to be true that the hotel is amazing. We will review various cases of scope of negation in the future.

Another issue is the verb modality. Take for instance the sentiment word *great* which has a positive polarity in a sentiment dictionary. However, if the sentiment word occurs with counterfactual verbs like *would have been* as in *Everything in this hotel would have been great if it only were not so far from the center*, the verb modality weakens the polarity of the sentiment word *great*. To account for the modality, we will experiment with verbal features represented in ‘feats’ in UD such as feats: Mood=xx, VerbForm=yyy in future research. We assume that capturing the verb modality of a sentence will considerably improve unsupervised SA.

Another issue is agent-based modality. Current NLP tasks express the need for a more structured sentiment analysis that accounts for the agent-based modality (SemEval 2022, Task 10 [24]). Sentiment words can be associated with different agents or opinion holders as in *According to Trip-Advisor, this hotel is great, but I don’t think so.* Quotation is another example which shows that the opinion expressed inside quotations does not need to correspond to the opinion of the reviewer as in *"this restaurant is the best" says Trip Advisor*. We will look into modalities in more detail in the future research.

One important issue is word ambiguity. Take for instance the word *vieja* ‘old’. It has a negative score in our sentiment dictionary. However, if it is used in a proper name as in *Havana vieja* it does not have a negative sentiment. Depending on the domain, *old* can have a positive value as in *old tradition*. Several approaches have been dealing with Word Sense Disambiguation (WSD) including WSD for lexicon-based approaches for SA [25, 17], most of which are exploring WordNet [26]). Other more current approaches use neural language models for WSD ([27]). We

will deal with WSD in SA in the future.

The final issue is the calculation of the polarity score of a whole text or review that contains several sentences. We have shown that the correct prediction of sentiment scores is not only a matter of syntactic rules, but also a matter of finding the right scoring metrics that optimally represents how humans generally write reviews. As with syntactic rules that determine the syntactic context of polarity, we can assume that there are stylistic rules that determine the stylistic context of polarity. Such stylistic rules can reflect the optimal organization or structure of a text [28]. Marking relevant sentences by sentence or word order, key words like "in sum", capitalization, repetition, etc. can be part of these stylistic rules. Determining these rules will considerably improve unsupervised rule-based approaches and is therefore reserved for the future research.

6. Acknowledgments

We acknowledge the European Research Council (ERC), which has funded this research under the Horizon Europe research and innovation programme (SALSA, grant agreement No 101100615), ERDF/MICINN-AEI (SCANNER-UDC, PID2020-113230RB-C21), Xunta de Galicia (ED431C 2020/11), and Centro de Investigación de Galicia "CITIC", funded by Xunta de Galicia and the European Union (ERDF - Galicia 2014–2020 Program), by grant ED431G 2019/01.

7. Appendices

References

- [1] D. Vilares, C. Gómez-Rodríguez, M. A. Alonso, Universal, unsupervised (rule-based), uncovered sentiment analysis, *Knowledge-Based Systems* 118 (2017) 45–55. URL: <https://doi.org/10.10162Fj.knosys.2016.11.014>. doi:10.1016/j.knosys.2016.11.014.
- [2] D. Zeman, J. Nivre, M. Abrams, E. Ackermann, N. Aeppli, H. Aghaei, Ž. Agić, A. Ahmadi, L. Ahrenberg, C. K. Ajede, S. F. Akkurt, G. Aleksandravičiūtė, I. Alfina, A. Algom, K. Alnajjar, C. Alzetta, E. Andersen, et al., Universal dependencies 2.12, 2023. URL: <http://hdl.handle.net/11234/1-5150>, LINDAT/CLARIAH-CZ digital library at the Institute of Formal and Applied Linguistics (ÚFAL), Faculty of Mathematics and Physics, Charles University.
- [3] M. Á. Álvarez-Carmona, Á. Díaz-Pacheco, R. Aranda, A. Y. Rodríguez-González, L. Bustio-Martínez, V. Muñis-Sánchez, A. P. Pastor-López, F. Sánchez-Vega, Overview of rest-mex at iberlef 2023: Research on sentiment analysis task for mexican tourist texts, *Procesamiento del Lenguaje Natural* 71 (2023).
- [4] C. Hutto, E. Gilbert, Vader: A parsimonious rule-based model for sentiment analysis of social media text, *Proceedings of the International AAAI Conference on Web and Social Media* 8 (2014). URL: <https://ojs.aaai.org/index.php/ICWSM/article/view/14550>. doi:10.1609/icwsml.v8i1.14550.
- [5] M. Á. Álvarez-Carmona, R. Aranda, S. Arce-Cardenas, D. Fajardo-Delgado, R. Guerrero-Rodríguez, A. P. López-Monroy, J. Martínez-Miranda, H. Pérez-Espinosa, A. Y. Rodríguez-

- González, Overview of rest-mex at iberlef 2021: Recommendation system for text mexican tourism, *Procesamiento del Lenguaje Natural* (2021).
- [6] M. Á. Álvarez-Carmona, Á. Díaz-Pacheco, R. Aranda, A. Y. Rodríguez-González, D. Fajardo-Delgado, R. Guerrero-Rodríguez, L. Bustio-Martínez, Overview of rest-mex at iberlef 2022: Recommendation system, sentiment analysis and covid semaphore prediction for mexican tourist texts, *Procesamiento del Lenguaje Natural* 69 (2022) 289–299.
- [7] M. A. Álvarez Carmona, R. Aranda, A. Y. Rodríguez-Gonzalez, D. Fajardo-Delgado, M. G. Sánchez, H. Pérez-Espinosa, J. Martínez-Miranda, R. Guerrero-Rodríguez, L. Bustio-Martínez, Ángel Díaz-Pacheco, Natural language processing applied to tourism research: A systematic review and future research directions, *Journal of King Saud University - Computer and Information Sciences* 34 (2022) 10125–10144. URL: <https://www.sciencedirect.com/science/article/pii/S1319157822003615>. doi:<https://doi.org/10.1016/j.jksuci.2022.10.010>.
- [8] A. Diaz-Pacheco, M. A. Álvarez Carmona, R. Guerrero-Rodríguez, L. A. C. Chávez, A. Y. Rodríguez-González, J. P. Ramírez-Silva, R. Aranda, Artificial intelligence methods to support the research of destination image in tourism. a systematic review, *Journal of Experimental & Theoretical Artificial Intelligence* 0 (2022) 1–31. URL: <https://doi.org/10.1080/0952813X.2022.2153276>. doi:10.1080/0952813X.2022.2153276. arXiv:<https://doi.org/10.1080/0952813X.2022.2153276>.
- [9] R. Socher, A. Perelygin, J. Wu, J. Chuang, C. D. Manning, A. Ng, C. Potts, Recursive Deep Models for Semantic Compositionality Over a Sentiment Treebank, in: *EMNLP 2013. 2013 Conference on Empirical Methods in Natural Language Processing. Proceedings of the Conference, ACL, Seattle, Washington, USA, 2013*, pp. 1631–1642.
- [10] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, BERT: Pre-training of deep bidirectional transformers for language understanding, in: *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, Association for Computational Linguistics, Minneapolis, Minnesota, 2019, pp. 4171–4186. URL: <https://aclanthology.org/N19-1423>. doi:10.18653/v1/N19-1423.
- [11] S. M. Rezaeinia, A. Ghodsi, R. Rahmani, Improving the accuracy of pre-trained word embeddings for sentiment analysis, 2017. arXiv:1711.08609.
- [12] J. Hong, M. Fang, Analysis with deeply learned distributed representations of variable length texts, 2015.
- [13] B. S. Ainapure, R. N. Pise, P. Reddy, B. Appasani, A. Srinivasulu, M. S. Khan, N. Bizon, Sentiment analysis of covid-19 tweets using deep learning and lexicon-based approaches, *Sustainability* 15 (2023). URL: <https://www.mdpi.com/2071-1050/15/3/2573>.
- [14] J. Chun, Sentimentarcs: A novel method for self-supervised sentiment analysis of time series shows SOTA transformers can struggle finding narrative arcs, *CoRR abs/2110.09454* (2021). URL: <https://arxiv.org/abs/2110.09454>. arXiv:2110.09454.
- [15] J. Huang, Y. Meng, F. Guo, H. Ji, J. Han, Weakly-supervised aspect-based sentiment analysis via joint aspect-sentiment topic embedding, in: *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Association for Computational Linguistics, Online, 2020, pp. 6989–6999. URL: <https://aclanthology.org/2020.emnlp-main.568>. doi:10.18653/v1/2020.emnlp-main.568.

- [16] M. Taboada, J. Brooke, M. Tofiloski, K. Voll, M. Stede, Lexicon-based methods for sentiment analysis, *Computational Linguistics* 37 (2011) 267–307.
- [17] S. Vashishtha, S. Susan, Fuzzy rule based unsupervised sentiment analysis from social media posts, *Expert Systems with Applications* 138 (2019) 112834. URL: <https://www.sciencedirect.com/science/article/pii/S0957417419305366>. doi:<https://doi.org/10.1016/j.eswa.2019.112834>.
- [18] G. T. Caio Mello, Gullal S. Cheema, Combining sentiment analysis classifiers to explore multilingual news articles covering London 2012 and Rio 2016 Olympics, *Int J Digit Humanities* (2022). doi:<https://doi.org/10.1007/s42803-022-00052-9>.
- [19] P. Qi, Y. Zhang, Y. Zhang, J. Bolton, C. D. Manning, Stanza: A python natural language processing toolkit for many human languages, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, Association for Computational Linguistics, Online, 2020, pp. 101–108. URL: <https://aclanthology.org/2020.acl-demos.14>. doi:10.18653/v1/2020.acl-demos.14.
- [20] P. Qi, T. Dozat, Y. Zhang, C. D. Manning, Universal Dependency parsing from scratch, in: *Proceedings of the CoNLL 2018 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, Association for Computational Linguistics, Brussels, Belgium, 2018, pp. 160–170. URL: <https://aclanthology.org/K18-2016>. doi:10.18653/v1/K18-2016.
- [21] T. Graf, B. Marcinek, Evaluating evaluation metrics for minimalist parsing, in: V. Demberg, T. O’Donnell (Eds.), *Proceedings of the Fifth Workshop on Cognitive Modeling and Computational Linguistics, CMCL@ACL 2014*, Baltimore, Maryland, USA, June 26, 2014, Association for Computational Linguistics, 2014, pp. 28–36. URL: <https://doi.org/10.3115/v1/W14-2004>. doi:10.3115/v1/W14-2004.
- [22] F. Meng, Z. Lu, M. Wang, H. Li, W. Jiang, Q. Liu, Encoding source language with convolutional neural network for machine translation, *CoRR abs/1503.01838* (2015). URL: <http://arxiv.org/abs/1503.01838>. arXiv:1503.01838.
- [23] J. D. Atlas, *On the Semantics of Presupposition and Negation: An Essay in Philosophical Logic and the Foundations of Linguistics*, Ph.D. thesis, Princeton University, 1976.
- [24] J. Barnes, L. Oberlaender, E. Troiano, A. Kutuzov, J. Buchmann, R. Agerri, L. Øvrelid, E. Velldal, SemEval 2022 task 10: Structured sentiment analysis, in: *Proceedings of the 16th International Workshop on Semantic Evaluation (SemEval-2022)*, Association for Computational Linguistics, Seattle, United States, 2022, pp. 1280–1295. URL: <https://aclanthology.org/2022.semeval-1.180>. doi:10.18653/v1/2022.semeval-1.180.
- [25] D. Vilares, M. A. Alonso, C. Gómez-Rodríguez, Supervised polarity classification of Spanish tweets based on linguistic knowledge, in: *DocEng’13. Proceedings of the 13th ACM Symposium on Document Engineering*, ACM, Florence, Italy, 2013, pp. 169–172.
- [26] C. Hung, S.-J. Chen, Word sense disambiguation based sentiment lexicons for sentiment classification, *Knowledge-Based Systems* 110 (2016) 224–232. URL: <https://www.sciencedirect.com/science/article/pii/S0950705116302453>. doi:<https://doi.org/10.1016/j.knosys.2016.07.030>.
- [27] M. Giulianelli, M. Del Tredici, R. Fernández, Analysing lexical semantic change with contextualised word representations, in: *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Association for Computational Linguistics, Online, 2020, pp. 3960–3973. URL: <https://aclanthology.org/2020.acl-main.365>. doi:10.

18653/v1/2020.acl-main.365.

[28] D. Wilson, Relevance theory, Oxford Research Encyclopedia of Linguistics (2019).

A. Online Resources

The files used for the experiment are available on GitHub.

- [GitHub](#)