ATC at IroSvA 2019: Shallow Syntactic Dependency-based Features for Irony Detection in Spanish Variants

Alessandra Teresa Cignarella[1,2] and Cristina Bosco[1]

1. Dipartimento di Informatica, Università degli Studi di Torino, Italy
2. PRHLT Research Center, Universitat Politècnica de València, Spain
cigna@di.unito.it, bosco@di.unito.it

Abstract. In the present paper we describe the participation of the ATC team at the IroSvA 2019 shared task at IberLEF 2019, which is focused on Irony Detection in Spanish Variants, addressing its identification as a classical binary classification task. The approach is mainly oriented in performing a preliminary test of the importance of morpho-syntactic information in the task of irony detection. For this reason, we exploited a straightforward methodology: a Support Vector Classifier with a linear kernel, combined with shallow features based on morphology and dependency syntax. For the representation of such kind of knowledge we relied on the application on the data of the well-known *Universal Dependencies* format.

Keywords: Irony · Social Media · Syntax · Universal Dependencies

1 Introduction

The recognition of irony and the identification of pragmatic and linguistic devices that activate it are known as very challenging tasks to be performed by both humans or automatic tools [14,15,16,19,20]. The presence of ironic devices in a text can change the polarity of an opinion expressed with positive words for intending a negative meaning. This can significantly undermine systems' accuracy and makes crucial the development of irony-aware systems [2,5,8,9,10,12,21,22,26,27].

The growing interest in this task is attested by the proposal of shared tasks focusing on irony detection and its impact on sentiment analysis in social media [11], in the context of periodical evaluation campaigns for NLP tools for many languages, see for instance the pilot task on irony detection proposed for Italian in SENTIPOLC at EVALITA, in the 2014 and 2016 editions [1,3] and the related task proposed for French at DEFT at TALN 2017 [4]. For what concerns English,

after a first task at *SemEval-2015* (i.e. Task 11) focusing on *Sentiment Analysis of Figurative Language in Twitter* [8], in 2018 a shared task on irony detection in tweets has been proposed (*SemEval-2018 Task 3: Irony detection in English tweets*) [25]. The setting proposed for the Semeval-2018 is an indication of the growing interest for a deeper analysis of the linguistic phenomena underlying ironic expressions. Such kind of deeper analysis naturally calls for the definition and the exploitation of schemes allowing the annotation of finer-grained features and resources in order to hopefully improve the performance of automatic systems in this especially challenging task. For instance, an especially fine-grained annotation format for irony is the one proposed in Karoui et al. [13], concerning French, Italian and English. The same scheme has later been applied on a larger Italian corpus TWITTIRÒ [6]. The resulting annotated corpus has been exploited as reference dataset within the context of the IronITA shared task at the EVALITA 2018 evaluation campaign on *Irony and Sarcasm Detection in Italian Tweets*.

In this paper we describe our submission at the IroSvA 2019 shared task [17], which focused on Irony Detection in Spanish Variants. In particular, for testing hypotheses about the involvement of morphology and syntax in figurative language phenomena, what we propose here is an investigation focused on these deep levels of representation and their usefulness for modeling irony in social media in a computational perspective.

2 Task Description

The task is structured into three subtasks, each one for predicting whether messages are ironic or not in three different variants of the Spanish language, i.e. those spoken respectively in Spain, Mexico and Cuba. The aim is that of investigating whether a short message, written in the Spanish language, is ironic or not with respect to a given context.

Table 1. Data distribution.

	Training		Test	
	ironic	not-iro	ironic	not-iro
Spain (es)	800	1,600	200	400
Mexico (mx)	800	1,600	200	400
Cuba (cu)	800	1,600	200	400
	7,200		1,800	

The organizers provided a different training and test set for each Spanish variant where the items are distributed as shown in Table 1. The Spanish and Mexican sets are composed by tweets, while the set from Cuba included news comments. As far as the distribution of irony, each training set contained 800 ironic and

1,600 not ironic texts. The same proportion of 33%-66% has been also maintained in the test set, counting respectively 200 ironic and 400 not-ironic texts in each subset.

3 Methodology

Because of our interest in features related to syntax and their contribution in figurative language detection, we didn't take into account lexical features also considering that they can be too much influenced by the involved Spanish varieties. Therefore, we trained our automatic system on the three datasets altogether (7,200 texts) and tested the same model on the three different test sets, regardless of the three variants of Spanish.

3.1 Preprocessing

We applied two preprocessing steps: the stripping of URLs from texts all normalized to lowercase letters, and the morpho-syntactic analysis. We trained indeed the *UDPipe* (which includes tokenization, Part of Speech tagging and parsing) on the UD_Spanish-GSD corpus for generating for each item of the dataset a CoNLL-U tree in Universal Dependencies (UD) format, the *de facto* standard for dependency-based syntactic representations.

3.2 The ATC System

After having performed experiments with Random Forest, Decision Trees and Support Vector Machine, we finally implemented our system with a linear kernel with the last classifier, which resulted the best performing one. We propose a straightforward approach with two different types of features: a) common *"baseline"* features, widely explored in sentiment analysis tasks and in irony detection tasks too; and b) new syntactic *"dependency-based"* features. Our novel contribution mainly consists in the exploitation of these latter features to create vectorial representations of texts; all them (listed in point b) have been made available thanks to the application of the UDpipeline and the subsequent generation of the dependency trees corresponding to the items of the datasets. Figure 1 shows a tweet where the UD format has been applied (Translation: I have launched a reporter into the air and they do not fly... 😱).

a) Baseline features

• *Bag of Words (BoW):* each tweet was pre-processed to convert it to lowercase letters. Then we extracted unigrams, bigrams and trigrams to create a binary representation
• *Bag of Char-grams (BoC)*: we considered the sequence of char-grams in a range from 2 to 5 characters.

b) Dependency-based features

• *Bag of Dependency Relations (BoDeprel):* following the approach described by Ghanem et al. [7], we considered the sets from 5 to 7 dependency relations as occurring in the linear order of the sentence from left to right.

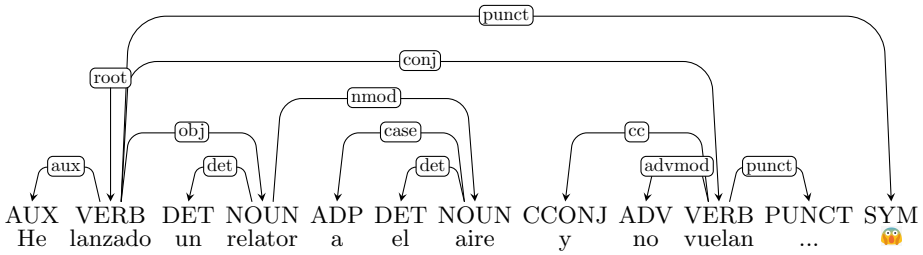

Fig. 1. Example of dependency tree.

• *Bag of SyntaxPath's Word Forms (Path_Form):* starting from the intuition of Sidorov et al. [23,24] we created a Bag of Word Forms (tokens), considering the bi-grams that can be collected following the syntactic tree structure (rather than the bi-grams that can be collected reading the sentence from left to right). For instance, the Path_Form corresponding to the sentence in Figure 1 includes ['lanzado', 'he'], ['lanzado', 'relator'], ['relator', 'un'], ['relator', 'aire'], ['aire', 'a'], ['aire', 'el'], ['lanzado', 'vuelan'], ['vuelan', 'y'], ['vuelan', 'no'], ['vuelan', '...'].

• *Bag of SyntaxPath's Deprels (Path_Deprel):* we created a Bag of Deprels, collecting the dependency relations occurring in the structure of the syntactic tree, i.e. following the syntactic paths, thus creating a vectorial space based on bi-grams, combining dependency relations in pairs. For instance, the Path_Deprel corresponding to the sentence in Figure 1 includes ['root', 'aux'], ['root', 'punct'], ['root', 'conj'], ['conj', 'cc'], ['conj', 'advmod'], ['conj', 'punct'], ['root', 'obj'], ['obj', 'det'], ['obj', 'nmod'], ['nmod', 'case'], ['nmod', 'det'].

The system is available at: `https://github.com/AleT-Cig/ATC_IroSvA_2019`.

4 Results

In Table 2 we show our official results in comparison with the four baselines proposed by the organizers. We can observe how, on average, our dependency-based approach performs better than shallow lexical techniques, such as *word n-grams*, but is not stronger than more refined approaches, such as those implemented in the neural-network based approach of *word2vec* and *LDSE* [18].

Table 2. Official results and baselines obtained over the test set.

	ES	MX	CU	AVG
LDSE	0.6795	**0.6608**	**0.6335**	**0.6579**
W2V	**0.6823**	0.6271	0.6033	0.6376
Word nGrams	0.6696	0.6196	0.5684	0.6192
MAJORITY	0.4000	0.4000	0.4000	0.4000
ATC	0.6512	0.6454	0.5941	0.6302

A fine-grained observation of results shows that our system performs better on the Spanish and Mexican datasets (F_{avg} ES = 0.6512, F_{avg} MX = 0.6454) and slightly worse in the Cuban variety (F_{avg} CU = 0.5941). We recall this is connected to the nature of the datasets, in fact, the first two are composed by tweets while the Cuban dataset is composed by news comments which may be featured by a slightly different syntactical structure.

5 Conclusions

In this paper we presented an overview of the ATC submission for the IroSvA 2019 shared task on Irony Detection in Spanish Variants at IberLEF 2019. We participated to the shared task by submitting one single system for the detection of irony in Spanish, Mexican and Cuban texts with a view to testing the suitability of the features we engineered across different datasets of Spanish variants based on syntactic knowledge.

Our approach, chiefly based on morphological and dependency-based syntactic features, proved to perform better than straightforward baselines, such as *word n-grams* and *majority voting*, and also to be able to provide in-line results with respect to stronger baselines based on *word2vec* semantic representation and the *LDSE* system [18].

The results show that focusing on syntactic features, namely *Bag of SyntaxPath's Word Forms* and *Bag of SyntaxPath's Deprels* (i.e. the novel contribution of our work), produced a good contribution to the Irony Detection task in Spanish Variants. Considering that the results seem quite promising, and that the dependency-based features deserves a finer-grained study, in future work we will further investigate them observing their behavior in other tasks related to sentiment analysis. In particular, thanks to the great adaptability of the UD format across different languages, we plan to test these new features in a multilingual scenario too.

Acknowledgements

The work of Cristina Bosco was partially funded by Progetto di Ateneo/CSP 2016 (*Immigrants, Hate and Prejudice in Social Media*, S1618L2BOSC01).

References

1. Barbieri, F., Basile, V., Croce, D., Nissim, M., Novielli, N., Patti, V.: Overview of the EVALITA 2016 SENTIment POLarity Classification Task. In: Proceedings of Third Italian Conference on Computational Linguistics (CLiC-it 2016) & Fifth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2016). vol. 1749. CEUR-WS.org (2016)
2. Barbieri, F., Saggion, H., Ronzano, F.: Modelling sarcasm in Twitter, a novel approach. In: Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis. pp. 50–58. ACL (June 2014)
3. Basile, V., Bolioli, A., Nissim, M., Patti, V., Rosso, P.: Overview of the EVALITA 2014 Sentiment Polarity Classification Task. In: Proceedings of the 4th Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2014) (2014)
4. Benamara, F., Grouin, C., Karoui, J., Moriceau, V., Robba, I.: Analyse d'opinion et langage figuratif dans des tweets: présentation et résultats du Défi Fouille de Textes DEFT2017. In: Actes de l'atelier DEFT2017 associé à la conférence TALN. Association pour le Traitement Automatique des Langues (ATALA) (2017)
5. Bosco, C., Patti, V., Bolioli, A.: Developing Corpora for Sentiment Analysis: The Case of Irony and Senti-TUT. IEEE Intelligent Systems **28**(2), 55–63 (2013)
6. Cignarella, A.T., Bosco, C., Patti, V., Lai, M.: Application and Analysis of a Multi-layered Scheme for Irony on the Italian Twitter Corpus TWITTIRÒ. In: Proceedings of the 11th Language Resources and Evaluation Conference (LREC) (2018)
7. Ghanem, B., Cignarella, A.T., Bosco, C., Rosso, P., Rangel Pardo, F.M.: UPV-28-UNITO at SemEval-2019 Task 7: Exploiting Post's Nesting and Syntax Information for Rumor Stance Classification. In: Proceedings of the 12th International Workshop on Semantic Evaluation (SemEval-2019). pp. 1121–1127. ACL (2019)
8. Ghosh, A., Li, G., Veale, T., Rosso, P., Shutova, E., Barnden, J., Reyes, A.: Semeval-2015 Task 11: Sentiment Analysis of Figurative Language in Twitter. In: Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval-2015). pp. 470–478. ACL (2015)
9. González-Ibáñez, R., Muresan, S., Wacholder, N.: Identifying sarcasm in Twitter: A closer look. In: Proceedings of the 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies. pp. 581–586. ACL (2011)
10. Hernández Farías, D.I., Benedí, J.M., Rosso, P.: Applying basic features from sentiment analysis for automatic irony detection. In: Pattern Recognition and Image Analysis, vol. 9117, pp. 337–344. Springer (2015)
11. Hernańdez Farías, D.I., Patti, V., Rosso, P.: Irony detection in twitter: The role of affective content. ACM Transactions on Internet Technology (TOIT) **16**(3), 19 (2016)
12. Joshi, A., Sharma, V., Bhattacharyya, P.: Harnessing context incongruity for sarcasm detection. In: Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 2: Short Papers). pp. 757–762. ACL (July 2015)
13. Karoui, J., Benamara, F., Moriceau, V., Patti, V., Bosco, C.: Exploring the Impact of Pragmatic Phenomena on Irony Detection in Tweets: A Multilingual Corpus Study. In: Proceedings of the European Chapter of the Association for Computational Linguistics (EACL 2017). pp. 262–272 (2017)

14. Kouloumpis, E., Wilson, T., Moore, J.D.: Twitter sentiment analysis: The good the bad and the omg! In: Proceedings of the ICWSM: International AAAI Conference on Web and Social Media. pp. 538–541. AAAI (2011)
15. Maynard, D., Funk, A.: Automatic Detection of Political Opinions in Tweets. In: Proceedings of the ESWC: Extended Semantic Web Conference. pp. 88–99. Springer (2011)
16. Mihalcea, R., Pulman, S.: Characterizing Humour: An Exploration of Features in Humorous Texts. In: International Conference on Intelligent Text Processing and Computational Linguistics. pp. 337–347. Springer (2007)
17. Ortega-Bueno, R., Rangel, F., Hernández Farías, D.I., Rosso, P., Montes-y-Gómez, M., Medina Pagola, J.E.: Overview of the Task on Irony Detection in Spanish Variants. In: Proceedings of the Iberian Languages Evaluation Forum (IberLEF 2019), co-located with 34th Conference of the Spanish Society for Natural Language Processing (SEPLN 2019). CEUR-WS.org (2019)
18. Rangel, F., Rosso, P., Franco-Salvador, M.: A low dimensionality representation for language variety identification. In: 17th International Conference on Intelligent Text Processing and Computational Linguistics, CICLing'16. Springer-Verlag, LNCS(9624), pp. 156-169 (2018)
19. Reyes, A., Rosso, P., Buscaldi, D.: Finding humour in the blogosphere: the role of wordnet resources. In: Proceedings of the 5th Global WordNet Conference. pp. 56–61 (2010)
20. Reyes, A., Rosso, P., Buscaldi, D.: From Humor Recognition to Irony Detection: The Figurative Language of Social Media. Data & Knowledge Engineering **74**, 1–12 (2012)
21. Reyes, A., Rosso, P., Veale, T.: A Multidimensional Approach for Detecting Irony in Twitter. Language Resources and Evaluation **47**(1), 239–268 (2013)
22. Riloff, E., Qadir, A., Surve, P., Silva, L.D., Gilbert, N., Huang, R.: Sarcasm as Contrast between a Positive Sentiment and Negative Situation. In: Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing, (EMNLP 2013). pp. 704–714. ACL (2013)
23. Sidorov, G., Velasquez, F., Stamatatos, E., Gelbukh, A., Chanona-Hernández, L.: Syntactic dependency-based n-grams: More evidence of usefulness in classification. In: International Conference on Intelligent Text Processing and Computational Linguistics. pp. 13–24. Springer (2013)
24. Sidorov, G., Velasquez, F., Stamatatos, E., Gelbukh, A., Chanona-Hernández, L.: Syntactic n-grams as machine learning features for natural language processing. Expert Systems with Applications **41**(3), 853–860 (2014)
25. Van Hee, C., Lefever, E., Hoste, V.: SemEval-2018 Task 3: Irony detection in English Tweets. In: In Proceedings of the International Workshop on Semantic Evaluation (SemEval-2018). ACL (2018)
26. Wang, A.P.: #Irony or #Sarcasm — A Quantitative and Qualitative Study Based on Twitter. In: Proceedings of the PACLIC: the 27th Pacific Asia Conference on Language, Information, and Computation. pp. 349–356. ACL (2013)
27. Zhang, S., Zhang, X., Chan, J., Rosso, P.: Irony detection via sentiment-based transfer learning. Information Processing & Management **56**(5), 1633–1644 (2019)