Dependency Syntax in the Automatic Detection of Irony and Stance

Sintaxis de dependencias en la detección automática de ironía y posicionamiento

Alessandra Teresa Cignarella[1,2]
[1]PRHLT Research Center, Universitat Politècnica de València, Spain
[2]Dipartimento di Informatica, Università degli Studi di Torino, Italy
alessandrateresa.cignarella@unito.it

Abstract: PhD thesis in Computer Science written by Alessandra Teresa Cignarella under the supervision of Dr. Cristina Bosco (University of Turin) and Prof. Dr. Paolo Rosso (Universitat Politècnica de València). This thesis was developed under a cotutelle between the PRHLT Research Center of the Universitat Politècnica de València, Spain and the Computer Science Department of the University of Turin, Italy. The thesis defense was held in Torino (*online*), Italy on October 26th, 2021. The doctoral committee was composed by: Prof. Dr. Joakim Nivre (Department of Linguistics and Philology, Uppsala University, Sweden), Dr. Saif M. Mohammad (National Research Council Canada, Ottawa, Canada), and Prof. Dr. Veronique Hoste (Department of Translation, Interpreting and Communication, Ghent, Belgium). An international mention was achieved after having spent 18 months at the Università degli Studi di Torino and 18 months at the Universitat Politècnica de València.
Keywords: Universal Dependencies, Irony Detection, Stance Detection, NLP, Multilingual, User-Generated Content, Twitter, BERT.

Resumen: Tesis doctoral en Informática realizada por Alessandra Teresa Cignarella y dirigida por la Dra. Cristina Bosco (University of Turin) y el Prof. Dr. Paolo Rosso (Universitat Politècnica de València) en el marco de un convenio de cotutela entre el PRHLT Research Center de la Universitat Politècnica de València, España y el Departamento de Informatica de la Universidad de Turin, Italia. La defensa de la tesis fue en Turin (*en linea*), Italia el 26 de octubre de 2021 ante un tribunal compuesto por: Prof. Dr. Joakim Nivre (Department of Linguistics and Philology, Uppsala University, Sweden), Dr. Saif M. Mohammad (National Research Council Canada, Ottawa, Canada), and Profa. Dra. Veronique Hoste (Department of Translation, Interpreting and Communication, Ghent, Belgium). Se obtuvo la mención internacional tras una estancia de 18 meses en la Università degli Studi di Torino y 18 meses en la Universitat Politècnica de València.
Palabras clave: Dependencias universales, detección de ironía, detección de posicionamiento, PNL, multilingüe, contenido generado por el usuario, Twitter, BERT.

1 Introduction

The present dissertation is part of the broad panorama of studies of Natural Language Processing (NLP). In particular, it is a work of Computational Linguistics (CL) designed to study in depth the contribution of syntax in the field of sentiment analysis and, therefore, to study texts extracted from social media or, more generally, online content.

Furthermore, given the recent interest of the scientific community in the Universal Dependencies[1] (UD) project (De Marneffe et al., 2021), which proposes an annotation format aimed at creating a "universal" representation of the phenomena of morphology and syntax in a manifold of languages, in this work we made use of this format, thinking of a study in a multilingual perspective (Ita-

---

[1]https://universaldependencies.org/.

lian, English, French and Spanish). Although the UD format was originally conceived to be applied to texts that are more "standard" from the point of view of morphosyntactic norms and punctuation, in more recent years the same scheme has begun to be applied also to user-generated content (UGC), i.e. texts extracted from social media, blogs, forums and microblogging platforms, such as Reddit, Twitter or Wikipedia pages. Inevitably, the application of this annotation framework to such a peculiar textual genre, in which the texts are accompanied by multimedia elements such as links, photos and videos, emojis and non-standardized punctuation, has opened up several problems in the Universal Dependencies community, many of which are still the subject of open and heated debate today.

In this work we provide an exhaustive presentation of the morphosyntactic annotation format of UD, in particular underlining the most relevant issues regarding their application to UGC. Two sub-areas of NLP will be presented, and used as case studies, in order to test the research hypotheses: the first case study will be in the field of *Irony Detection* (Van Hee, Lefever, and Hoste, 2018) and the second in the area of *Stance Detection* (Mohammad et al., 2016). In both cases, historical notes are provided that can serve as a context for the reader, the problems faced are introduced and the activities proposed in the computational linguistics community are described. Furthermore, particular attention is paid to the resources currently available as well as to those developed specifically for the study of the aforementioned phenomena. Finally,through the discussion of a set of experiments performed within or outside evaluation campaigns, we describe how syntax can contribute to the resolution of such tasks.

1.1 Motivation and Objectives

My main purpose is to explore the impact of morpho-syntactic information in sentiment analysis related tasks.

Firstly, for both irony and stance, I focused on the importance of the formulation of a clear problem statement, and the subsequent computational modeling of it. Secondly, I highlighted my experience in the creation of annotated corpora for those problems, my contribution to the organization of shared tasks and the important lessons learned, in terms

of research understanding. Later on, I proposed my first approaches to solve both tasks from a shallow perspective, starting to explore the most feasible way to represent morphosyntactic information, to extract it, and exploit it for classification purposes. I ended this process by relying on the UD Dependencies annotation format. Finally, after having encountered a satisfactory combination of features, I exploited the best sets of them – some of which are encoded in UD format – and I performed a handful of experiments in a variety of settings.

In the whole thesis a multilingual scenario is kept in mind, exploring four different language settings: English, Spanish, French, and Italian, for both irony detection and stance. Furthermore, due to the availability of benchmark datasets, regarding stance detection, I also experiment on a fifth language, i.e., Catalan.

The research questions that I aimed answering to are as follows:

RQ-1: *Could features derived from morphology and syntax help to address the task of irony detection?*

RQ-2: *To what extent does using resources such as treebanks for training NLP models improve the performance in irony detection?*

RQ-3: *Could features derived from morphology and syntax help to address the task of stance detection?*

RQ-4: *To what extent does using resources such as treebanks for training NLP models improve the performance in stance detection?*

2 Thesis Overview

This thesis consists in a reorganized collection of the most relevant investigations extracted from some research projects in which I was involved during my Ph.D. studies.

A brief overview of the contents of the thesis is presented below, summarizing all the work done and resuming the results obtained in the framework of this three-year-long research path. In Chapter 3, I also show some unpublished results regarding stance detection with dependency syntax and neural networks. Lastly, I draw some conclusions and discuss future work in the final chapter.

Chapter 1 – Introduction: In this chapter I introduced the reader to the main topics that will be discussed in the thesis, starting with a broad description of Natural Language Processing and automatic text classification, followed by an introduction on Universal Dependencies, morphology and syntax. I also proposed a brief discussion on the issues that can arise while applying the UD format to social media data, mainly referring to the following work: Sanguinetti et al. (2022).

Chapter 2 – Irony Detection: In the second chapter, which deals with the topic of irony detection, I described several works regarding such topic. In particular, in Section 2.1.1, I mainly referred to Cignarella et al. (2018) in order to describe described the organization of the *IronITA 2018* shared task. In Section 2.1.2, I described the creation of the multilayered corpus TWITTIRÒ-UD, annotated both of irony and morphology and dependency syntax.

In Section 2.3, I finally described some experiments performed in which I leveraged morpho-syntactic information for irony detection, mainly referring to what was done in the participation of the IroSvA 2019 shared task (Cignarella and Bosco, 2019; Cignarella et al., 2020).

Chapter 3 – Stance Detection: In the third chapter I dealt with the task of stance detection. In Section 3.1.1, I described the organization of the *SardiStance* shared task at EVALITA 2020. In Section 3.2.1, I presented the work done in Lai, Cignarella, and Hernandez Fariás (2017), for describing the participation in the *StanceCat 2017* shared task at IberEval 2017. In Section 3.3, I described my participation in *RumorEval 2019* where I first applied a syntax-based approach to the task of stance detection.

In Section 3.3.2, I presented a completely new research, specifically done for the PhD thesis, where I introduced a BERT-based approach leveraging morphosyntactic information for the automatic detection of stance in different languages.

Chapter 4 – The Interaction of Irony and Stance: In the fourth chapter I proposed a new part of my research, in which I explored the interaction between irony and stance, through the analysis of the *Sardi-Stance* dataset, which has been annotated accordingly to both phenomena.

Chapter 5 – Conclusions and Future Work – In the last chapter, I finally summarized all the important lessons learned and I proposed new research directions for future work.

3   Conclusions

This thesis collocates within the growing trend of studies devoted to make Artificial Intelligence results more explainable, going beyond the achievement of highest scores in performing tasks, but rather making their motivations understandable and comprehensible for experts in the domain.

The novel contribution of this work mainly consists in the exploitation of features that are based on *morphology* and *dependency syntax*, which were used in order to create vectorial representations of social media texts in various languages and for two different tasks. Such features have then been paired with a manifold of machine learning classifiers, with some neural networks and also with the language model BERT.

Results suggest that fine-grained dependency-based syntactic information is more informative for the detection of irony, and less informative for what concerns stance detection. Nonetheless, dependency syntax might still prove useful in the task of stance detection if, firstly, irony detection is considered as a pre-processing step. I also believe that the approach based on dependency syntax that I proposed could help in understanding and explaining a such a complex phenomenon like irony.

In fact, the several studies presented here allowed to investigate whether syntactic structures, independently from the target language, may provide information useful to understand whether a message is ironic or not.

Although it has been duly noted that syntax does not seem to be particularly informative regarding directly the task of stance detection (the second case study, presented in Chapter 3). On the other hand, also supported by some previous linguistic studies, syntax seems to play an important role in the detection of irony. Therefore, a new speculation that comes to mind, is that it could be more useful to perform a *"cascade task"*. Meaning that, firstly it might be useful to predict irony, with the help of morphosyntactic cues (step 1), and only then (as step 2),

proceeding in the detection of stance. In general, my assumption, is that predicting irony could be the first step in numerous other tasks, even shallow sentiment analysis, or the identification of fake news.

This outcome is something that should not be ignored, but obviously carrying out supervised studies in this sense would also mean dedicating a great effort and consuming much time in the creation of annotated datasets (that ought to be annotated on various layers, for different dimensions and phenomena). In fact, to further study this line of investigation, in Chapter 4, I proposed a shallow analysis of the Italian dataset regarding the Sardines Movement, which is only a small and limited beginning, but it is also certainly opening a new research perspective.

My work has certainly many limitations. Firstly, I needed to deal with the scarcity of data annotated in some adequate way and with the reduced size of the few datasets that are indeed available or those I helped to develop. Furthermore, this kind of investigation, mostly based on morphosyntactic cues that are applied to NLP tasks, is a rather new one. In fact, there are very few studies going towards this direction up to these days.

By having looked at some of the results obtained in the wide variety of experiments performed in this thesis, it is fundamental to stress that we did not solely want to appreciate the outcomes in terms of numerical performances, but rather being more focused in the more profound linguistic reasons behind them. And the same is valid also for why sometimes results are poorer and why features do not make improvements on a certain task.

I am positive that if we manage to understand what is the linguistic knowledge a certain approach, or a group of features, leverages when it produces good (or poor) results, among many possible approaches, it could allow us to make more mature choices for following work. Indeed, the future of NLP research needs to go towards approaches that better integrate different types of knowledge (such as syntactic knowledge, for once) and that manage to be more versatile for certain types of data and in different application contexts.

Acknowledgments

References

Cignarella, A. T., V. Basile, M. Sanguinetti, C. Bosco, F. Benamara, and P. Rosso. 2020. Multilingual Irony Detection with Dependency Syntax and Neural Models. In *Proceedings of the 28th International Conference on Computational Linguistics (COLING 2020)*. ACL.

Cignarella, A. T. and C. Bosco. 2019. ATC at IroSvA 2019: Shallow Syntactic Dependency-based Features for Irony Detection in Spanish Variants. In *Proceedings of IberLEF 2019*. CEUR-WS.org.

Cignarella, A. T., S. Frenda, V. Basile, C. Bosco, V. Patti, and P. Rosso. 2018. Overview of the EVALITA 2018 Task on Irony Detection in Italian Tweets (IronITA). In *Proceedings of EVALITA 2018*. CEUR-WS.org.

De Marneffe, M.-C., C. D. Manning, J. Nivre, and D. Zeman. 2021. Universal Dependencies. *Computational Linguistics*, 47(2):255–308.

Lai, M., A. T. Cignarella, and D. I. Hernandez Farías. 2017. iTACOS at IberEval2017: Detecting Stance in Catalan and Spanish Tweets. In *Proceedings of IberEval 2017*. CEUR-WS.org.

Mohammad, S., S. Kiritchenko, P. Sobhani, X. Zhu, and C. Cherry. 2016. SemEval-2016 Task 6: Detecting Stance in Tweets. In *Proceedings of SemEval 2016*. ACL.

Sanguinetti, M., C. Bosco, L. Cassidy, Ö. Çetinoğlu, A. T. Cignarella, T. Lynn, I. Rehbein, J. Ruppenhofer, D. Seddah, and A. Zeldes. 2022. Treebanking User-generated Content: A UD based Overview of Guidelines, Corpora and Unified Recommendations. *Language Resources and Evaluation*, pages 1–52.

Van Hee, C., E. Lefever, and V. Hoste. 2018. SemEval-2018 Task 3: Irony Detection in English Tweets. In *In Proceedings of SemEval 2018*. ACL.