

Is This an Effective Way to Annotate Irony Activators?

Alessandra Teresa Cignarella^{1,2}, Manuela Sanguinetti¹, Cristina Bosco¹, Paolo Rosso²

1. Dipartimento di Informatica, Università degli Studi di Torino, Italy

2. PRHLT Research Center, Universitat Politècnica de València, Spain

{cigna|msanguin|bosco}@di.unito.it, proso@dsic.upv.es

Abstract

In this article we describe the first steps of the annotation process of specific irony activators in TWITTIRÒ-UD, a treebank of Italian tweets annotated with fine-grained labels for irony on one hand, and according to the *Universal Dependencies* scheme on the other. We discuss in particular the annotation scheme adopted to identify irony activators and some of the issues emerged during the first annotation phase. This helped us in the design of the guidelines and allowed us to draw future research directions.

1 Introduction

In the last decade, several efforts have been devoted to address the challenges of sentiment analysis and related tasks, working mainly in English and other languages such as Italian, Spanish or French. Provided that most of the existing approaches in NLP are based on supervised semantic shallow analysis and machine learning techniques, there has been a strong push towards the development of resources from where related knowledge can be learned.

In particular the detection of irony is among the tasks currently considered as especially challenging since its presence in a text can reverse the polarity of the opinion expressed, that is using positive words for intending a negative meaning or – less often – the other way around. This can significantly undermine systems' accuracy and makes it crucial to develop irony-aware systems (Bosco et al., 2013; Reyes et al., 2013; Riloff et al., 2013; Wang, 2013; Barbieri et al., 2014; Joshi et al., 2015; Hernández Farías et al.,

2015; Hernández Farías et al., 2016). Additionally, the challenge is further complicated when there is a co-occurrence with sarcasm or satire (Hernández Farías and Rosso, 2016; Joshi et al., 2017; Ravi and Ravi, 2017).

The growing interest in irony detection is also attested by the proposal of shared tasks focusing on this topic within NLP evaluation campaigns. For instance, the pilot task on irony detection proposed for Italian in SENTIPOLC at EVALITA¹, in 2014 and 2016 (Barbieri et al., 2016; Basile et al., 2014), and the related task proposed for French at DEFT at TALN 2017 (Benamara et al., 2017). For what concerns English, after a first task at *SemEval-2015* focusing on figurative language in Twitter (Ghosh et al., 2015), a shared task on irony detection in tweets has been proposed in 2018 (Van Hee et al., 2018). Concerning Spanish, the most recent shared task about irony in social media has been organized at IberLEF 2019 *Irony Detection in Spanish Variants (IroSvA 2019)*, exploring the differences among varieties of Spanish from Spain, Cuba and Mexico (Ortega et al., 2019) in which the organizers also proposed a focus on context, stressing the importance of contextual semantics in ironic productions.

While the majority of the participating systems in the above-mentioned shared-tasks are based on classical machine learning techniques (Cignarella and Bosco, 2019; Frenda and Patti, 2019), researchers have recently started to exploit approaches based on neural networks. Among these, Huang et al. (2017) applied attentive recurrent neural networks (RNNs) that capture specific words which are helpful in detecting the presence of irony in a tweet, while Wu et al. (2018) exploited densely connected LSTMs in a multi-task learning strategy, adding PoS tag features, and Zhang et al. (2019) took advantage of recent advancements in transfer learning techniques.

Copyright © 2019 for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

¹<http://www.evalita.it/>

These settings are a clear indication of the growing interest for a deeper analysis of the linguistic phenomena underlying ironic expressions. Such kind of analysis naturally calls for the exploitation of finer-grained features and resources in order to improve the performance of automatic systems. For instance, an especially fine-grained annotation format for irony is the one proposed in Karoui et al. (2017), concerning French, Italian and English. The same scheme has later been applied on a new Italian corpus: TWITTIRÒ (Cignarella et al., 2018a). The resulting annotated corpus was used as reference dataset in the *IronITA 2018* shared task² on *Irony and Sarcasm Detection in Italian Tweets* (Cignarella et al., 2018b).

1.1 Motivation and Research Questions

The present work is, indeed, part of a wider joint project with other research groups working on English and French (Karoui et al., 2015). As mentioned above, in Cignarella et al. (2018a), we created an Italian corpus of tweets, i.e. TWITTIRÒ, annotated with a fine-grained tagset for irony, and later on, we extended the same resource applying the *Universal Dependencies* (UD) scheme (Nivre et al., 2016), thus creating TWITTIRÒ-UD (Cignarella et al., 2019).

This new corpus collocates in the panorama of treebanks with data extracted from social media, such as those recently developed for Italian and released in the UD repository³, and to the best of our knowledge it is one of the few linguistic resources where sentiment analysis and syntactic annotation are applied within the same framework. The main research question that we want to address is:

RQ1. Is there any syntactic pattern that can help us to automatically detect irony?

The intuition that we follow in this work is that if such “syntactic patterns” which activate irony do actually exist, therefore, they should be particularly evident in the syntactic context of certain lexical elements that create a semantic clash in a text.

For this reason, in the present article, we describe the first steps of the annotation process

²<http://di.unito.it/ironita18>.

³https://github.com/UniversalDependencies/UD_Italian-PoSITITA.

of specific irony activators in the TWITTIRÒ-UD corpus, taking advantage of the fact that the annotation format we adopted for the syntactic annotation allows us also to label specific activators at token level and retrieve dependency relations connected to them. In doing so, we are led to the following research questions, anticipated by the title of the paper:

RQ2. Is there an effective way to annotate irony activators?

RQ3. If so, is the one we propose valid?

The paper is organized as follows. In Section 2 the novel dataset TWITTIRÒ-UD and its annotation layers are presented. In Section 3 we describe the annotation process concerning irony activators, and we comment the inter-annotator agreement showing some examples. Finally, in Section 4 and Section 5 we discuss some difficult cases and we conclude the paper.

2 Corpus Description

The current version of TWITTIRÒ-UD comprises 1,424 tweets, annotated at multiple levels: a pragmatic level that attempts to model irony (see Section 2.1) and a syntactic level based on the UD scheme that represents the underlying syntactic structure of the tweets in the corpus (Section 2.2). In addition, we have recently introduced a further level that tries to act as an interface between the previous two (Section 3).

2.1 Annotating Irony

As far as the annotation for irony is concerned, the data of this corpus were manually annotated according to a multi-layered annotation scheme described in Karoui et al. (2017), which in turn includes 4 different levels.⁴ Beyond the annotation of irony vs non-irony (henceforth level 1), the multifaceted annotation scheme is organized in three further layers, namely the *activation type* (level 2), the *categories* (level 3) and the *clues* (level 4).

Irony is often activated by the presence of a clash or a contradiction between two elements (also called P1 and P2). This motivates the annotation of the two different *activation types* at level 2: explicit when both these elements are lexicalized in the message, implicit otherwise.

⁴See annotation guidelines at <https://github.com/IronyAndTweets/Scheme>.

```

# sent_id = -----
# twittiro = EXPLICIT EX:OXYMORON PARADOX
# activators = 3 12
# text = Il Pd diviso in due. Non è mai stato così unito. [@user]

1 Il il DET RD Definite=Def|Gender=Masc|Number=Sing|PronType=Art 2 det _ _
2 Pd Pd PROPN SP _ 3 nsubj _ _
3 diviso diviso ADJ A Gender=Masc|Number=Sing 0 root _ _
4 in in ADP E _ 5 case _ _
5 due due NUM N NumType=Card 3 obl _ SpaceAfter=No
6 . . PUNCT FS _ 3 punct _ _
7 Non non ADV BN PronType=Neg 12 advmod _ _
8 è essere AUX VA Mood=Ind|Number=Sing|Person=3|Tense=Pres|VerbForm=Fin 12 cop _ _
9 mai mai ADV B _ 12 advmod _ _
10 stato essere AUX VA Gender=Masc|Number=Sing|Tense=Past|VerbForm=Part 12 aux _ _
11 così così ADV B _ 12 advmod _ _
12 unito unito ADJ A Gender=Masc|Number=Sing 3 parataxis _ SpaceAfter=No
13 . . PUNCT FS _ 12 punct _ _
14 [ [ PUNCT FB _ 15 punct _ SpaceAfter=No
15 @user @user SYM SYM _ 12 vocative:mention _ SpaceAfter=No
16 ] ] PUNCT FB _ 15 punct _ SpaceAfter=\n

```

Figure 1: Example of tweet in CoNLL-U format.

The main linguistic devices reported in literature as irony triggers are described instead at level 3 by the *categories* of the scheme (i.e. analogy, euphemism, false assertion, oxymoron/paradox, context shift, hyperbole, rhetorical question and other). Table 1 shows the distribution of ironic categories throughout the corpus.

	n#	%
ANALOGY	261	18%
EUPHEMISM	84	6%
EX:CONTEXT SHIFT	185	13%
EX:OXYMORON PARADOX	277	19%
HYPERBOLE	81	6%
IM:FALSE ASSERTION	117	8%
OTHER	198	14%
RHETORICAL QUESTION	221	16%
TOTAL	1,424	

Table 1: Ironic categories in TWITTIRÒ-UD.

Finally the *clues* of level 4 are lexical or morpho-syntactic signals of the *activation types* and *categories* that can be found in a given ironic tweet, such as the preposition “like” or the presence of comparative structures in the *analogy* type, or the adverb “very” for *hyperbole*. For more details about this annotation scheme, see Karoui et al. (2017).

2.2 Annotating Universal Dependencies

The availability of social media data annotated also at syntactic level is a prerequisite for our study and for the kind of annotation we intend to perform; as a dependency-based representation was deemed to be more suitable for our purposes, Universal Dependencies became our natural choice.

To obtain the data thus annotated, we ran UD-Pipe (Straka and Straková, 2017) for tokenization, PoS tagging, lemmatization and dependency parsing, using a model trained on two Italian resources available in the UD repository, the ISDT (Simi et al., 2014) and PoSTWITA-UD (Sanguinetti et al., 2018) treebanks⁵. The former includes multiple text genres (legal texts, news, Wikipedia articles, among others), but it mostly deals with well-edited texts and a standard language. The latter is made up of so-called user-generated contents, an in particular of Twitter posts in Italian. As using both resources for training proved to give better results when analyzing Italian tweets (Sanguinetti et al., 2018), we used the same approach in this work.

Figure 1 shows an example from the TWITTIRÒ-UD corpus⁶ in CoNLL-U format: along with the typical fields indicating the sentence id and the raw text, two resource-specific fields have been introduced, to encode the information on irony categories (described in Section 2.1) and irony activators (see Section 3).

As also described in Cignarella et al. (2019), and as expected, the main critical issues in applying the UD scheme to our corpus namely consisted in finding the proper tags and coding conventions for those linguistic phenomena typically occurring in Italian tweets. The guidelines provided in Sanguinetti et al. (2018) represented a helpful ground-

⁵More details in Cignarella et al. (2019).

⁶The id of the tweet and the user mention are encrypted due to privacy regulations. – Translation: The Democratic Party is split in two. It has never been so united. [@user].

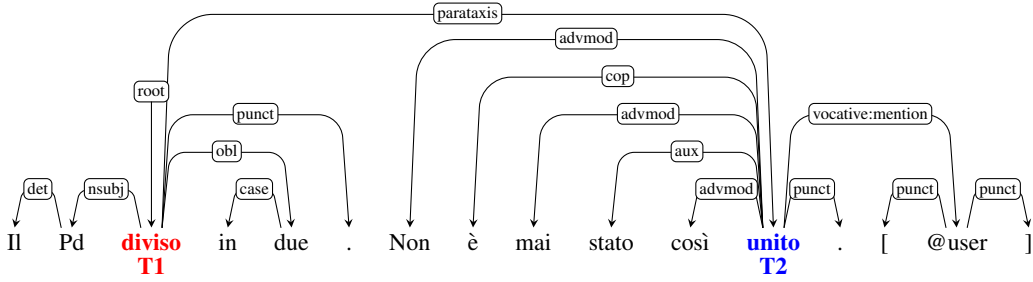


Figure 2: Dependency graph of the tweet in Figure 1 with irony activators T1 and T2 highlighted in red and blue, respectively.

work in this respect.

The fully-annotated treebank, including the annotation of irony categories, is going to be made available with the release of UD version 2.5. Due to its preliminary nature, however, the annotation of irony activators will be included in the resource at a later stage.

3 Annotating Irony Activators

As previously mentioned, irony is activated by the presence of a clash or a contradiction between two elements or two propositions (P1 and P2), which are indeed the triggers of the activation of irony. According to the scheme proposed by Karoui et al. (2017) there are two kinds of *activation types*: EXPLICIT when both these elements are lexicalized in the message, IMPLICIT otherwise.

In this step of our work, we focused our attention on the manual annotation of irony activators and on providing annotation guidelines that could be useful also for other datasets in different languages, within the same multilingual project. Indeed, the starting point of the present work is connected to the work of Karoui (2017), on a French dataset, in which the author tried to annotate at tweet level some elements that are responsible for the activation of irony. In that approach, each tweet had to be annotated using the Glozz tool (Widlöcher and Mathet, 2009), in terms of units and relationships between units (if the relationship existed). Three types of relationship were taken into account: 1) relation of comparison, 2) relation of explicit contradiction, and 3) relation of cause/consequence.

With respect to this work we opted for a finer-grained annotation also taking advantage from the availability of tokenized data and a full syntactic analysis in UD format.

3.1 Our approach

Our aim is to annotate irony activators in the whole TWITTER-UD corpus. Differently from what proposed in Karoui (2017), in which the elements creating an ironic contrast (P1 and P2) could be words, phrases or even full sentences; in this work, since we want to highlight the interaction between the pragmatic phenomenon of irony and its syntactic representation, we define as irony activators a pair of words T1 and T2 that must correspond to nodes of the syntactic dependency tree.

Given an ironical utterance (in our case a tweet) and its dependency-based syntactic representation, where each node in the tree structure represents a word, T1 and T2 is thus a pair of words – regardless of their grammatical category – such that:

- either they are both lexicalized (in explicit irony) or one of them is left unspecified (implicit irony);
- they act as triggers by signaling the presence of an ironic device.

The intuition behind this choice is inspired by the work of Saif et al. (2016), in which the authors underline the importance of contextual and conceptual semantics of words when calculating their sentiment, which in turn comes from the popular dictum “You shall know a word by the company it keeps!” (Firth, 1957). Our idea is, in fact, to proceed in two steps: firstly, to annotate irony triggers at token level, and subsequently to retrieve the other tokens that “keep company” to them by means of the dependency relations available from the UD annotation.

Therefore, as we have already highlighted in Section 1.1, if any kind of “syntactic pattern” that can help us to automatically detect irony does exist, we assume this will be particularly evident in

the “syntactic circle” around the lexical elements that create a contradiction and are the lexical activators of the ironic realization, namely T1 and T2.

In the present research, being a preliminary study, and in order to validate the strengths and weaknesses of annotation guidelines for irony activators, two skilled annotators (A1 and A2) annotated a first sample of 277 tweets, focusing on the most frequent category: EX:OXYMORON PARADOX, which covers almost 20% of the whole corpus, as it is shown in Table 1 in Section 2.1. In the following sections we will describe the guidelines that emerged throughout the discussion between A1 and A2, we will discuss the most relevant comments reported by the annotators and we will comment on some examples, thus providing an evaluation and the measures of inter-annotator agreement.

3.2 Annotation process

A sample of 277 tweets, from the ironic category EX:OXYMORON PARADOX, was annotated in parallel by two skilled annotators (A1 and A2), experts both in sentiment analysis annotations and also familiar with the CoNLL-U format.

Both of them were asked, given a tweet, to annotate two words T1 and T2 that are responsible for the activation of irony, bearing in mind these basic guiding principles:

- T1 and T2 can be nodes of any type: no specific constraints are given on the morpho-syntactic category;
- the identification of the proper T1 and T2 is guided by the irony category: for example, if the ironic tweet fits the category *oxymoron/paradox*, select the activators so that the type of relation triggered will be a contrast or a contradiction:


 la cosa bella del governo Monti è che ha acceso^{T1} le speranze di tutti e le spegnerà^{T2} pure ...
→ *the good thing about the Monti government is that it has kindled everyone's hopes and it will stifle them as well*

Figure 2 provides an example of annotated tweet, where the words *diviso* (divided) and *unito* (united) have been annotated as T1 and T2, respectively. From a procedural perspective, since the

tokens “diviso” and “unito” are respectively at position 3 and 12 in the CoNLL-U format (cfr. Figure 1), annotators were asked to add a line in the header of the annotation file, such as this one:

activators = 3 12

Furthermore, the annotators were asked to annotate any kind of doubt it might occur to them in order to provide material to a discussion about the efficacy of the guidelines.

3.3 Evaluation and Agreement

In a first phase, the annotators sketched a draft of the guidelines for the annotation of ironic activators T1 and T2, and, as a pilot experiment, they tested their efficacy on a sample of 50 tweets. Discussing the uncertain cases and the instances in disagreement helped to significantly improve the quality of the annotation choices between A1 and A2. In fact, after the first “training phase”, the guidelines were cleared up, and the annotators could proceed to annotate all the 277 OXYMORON PARADOX tweets. The inter-annotator agreement (IAA) on the 277 tweets was later calculated by means of simple observed agreement (expressed in percentage).




Figure 3: Observed IAA on 277 tweets.

As we can see from Figure 3 a complete agreement was immediately reached on 113 tweets (40.9%), other 94 tweets (34.1%) were in partial agreement (meaning that the annotators agreed only on T1 or T2), while 69 (25%) presented a complete disagreement.


After the first annotation step was completed and the agreement was calculated, the annotators tried to solve the partial disagreement. As a result, the percentage of T1-T2 pairs where agreement has been reached went up to approximately 69.2% (191 tweets), while the proportion of complete disagreement rose to approximately 30.8% (85 tweets).

4 Discussion


Overall, the outcome of the experimental annotation of irony activators is rather encouraging. Not only from a quantitative perspective (see Section 3.3), but also from a qualitative point of view. In fact, annotators pointed out several difficult cases, but in general they were able to find an agreement discussing the possibilities within the few restrictions posed by the guidelines.

Among the unresolved cases of disagreement (difficult cases) we were able to find recurring patterns, that need to be addressed adding new specific rules before continuing with the annotation on the rest of the dataset. Below we provide a short description.

More than two irony activators For instance, in the following tweet a list of names is presented. The contrast is created with *migliori* (best) and all three entities, but it is difficult to only choose one.


 Fantagoverno. **Fabio Volo**^{T1}, **Giovanni Sartori**^{T1}, **Roberto Saviano**^{T1}: ecco il governo dei **Migliori**^{T2} Mario Monti ... URL
→ *Fantagovernment. Fabio Volo, Giovanni Sartori, Roberto Saviano: here is the government of the best Mario Monti... URL*

Multiple categories There is more than one ironic category (e.g. overlap between an ANALOGY and a PARADOX). Such as in the tweet below, in which there is a clear analogy between Superman and Mario Monti; but also the paradoxical sentence “if you didn’t exist you should be invented!” referred to a country (Italy), which, of course already exists.


 E vai adesso con **Mario Monti**^{T1}/**Superman**^{T2}, crisi finita, stipendi in aumento, e riforme. Grazie **Stato**^{T1}! **Se non ci fossi bisognerebbe inventarti!**^{T2}
→ *And now let's go with Mario Monti/Superman, the crisis is over, the salaries are raising, and there are reforms. Thank you country! If you didn't exist you should be invented!*

Paraprosdokian There is a peculiar kind of ironic production, known in literature as “*paraprosdokian*”, in which the latter part of a sentence is surprising or unexpected in a way that causes the reader or listener to reinterpret the first part. This kind of ironic production is not specif-

ically taken into account in the annotation scheme.

 I Soliti Idiotti in scena a **Sanremo**^{T1}. Ieri erano alla **Camera**^{T2}. [@user] #dopofestival
→ *The Usual Idiots on Sanremo's stage. Yesterday there were at the Chamber of Deputies. [@user] #afterfestival*

Different activation type The tweet has been annotated as EXPLICIT, but the elements that create the ironic clash are to be found in the outer world (world knowledge is needed).

 #labuonascuola è avere una scuola.
→ *#thegoodschool is to have a school.*

5 Conclusion

In this article we described the preliminary steps of the annotation process of irony activators in the TWITTIRÒ-UD corpus, a novel Italian treebank of ironic tweets. In particular, we described the problems that emerged during the first annotation phase, the strengths and weaknesses of the scheme itself, in order to highlight future research directions. Being a preliminary study, and having no benchmark to compare with, the results obtained in the observed agreement are rather promising; moreover, the tweets included in TWITTIRÒ were retrieved from different pre-existing Italian corpora (as described in Cignarella et al. (2017)): the heterogeneous sources the data were gathered from thus represents a signal of the potential portability of the scheme and paves the way for a more systematic annotation process of the whole dataset. The next steps will then consist in the guidelines improvement and the annotation of the remaining part of TWITTIRÒ-UD accordingly.

Furthermore, the availability of English and French datasets annotated with the same scheme described in Section 2.1 (see Karoui et al. (2017)) allows the direct applicability of the annotation of irony activators in other languages than Italian. While this can be considered a further validation step to test the overall validity and portability of the scheme, it may also provide useful insights into the linguistic mechanisms underlying verbal irony in different languages.

The actual usability of this kind of resources will be finally tested when training NLP tools for irony detection, in both mono- and multi-lingual settings.

Acknowledgments

The work of C. Bosco and M. Sanguinetti was partially funded by Progetto di Ateneo/CSP 2016 (*Immigrants, Hate and Prejudice in Social Media*, S1618L2BOSC01). The work of P. Rosso was partially funded by the Spanish MICINN under the research project MISMIS-FAKEHATE on MIS-information and MIScommunication in social media: FAKE news and HATE speech (PGC2018-096212-B-C31).

References

- Francesco Barbieri, Horacio Saggion, and Francesco Ronzano. 2014. Modelling sarcasm in Twitter, a novel approach. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 50–58. ACL.
- Francesco Barbieri, Valerio Basile, Danilo Croce, Malvina Nissim, Nicole Novielli, and Viviana Patti. 2016. 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)*, volume 1749. CEUR-WS.org.
- Valerio Basile, Andrea Bolioli, Malvina Nissim, Viviana Patti, and Paolo Rosso. 2014. 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)*. ELRA.
- Farah Benamara, Cyril Grouin, Jihen Karoui, Véronique Moriceau, and Isabelle Robba. 2017. 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).
- Cristina Bosco, Viviana Patti, and Andrea Bolioli. 2013. Developing Corpora for Sentiment Analysis: The Case of Irony and Senti-TUT. *IEEE Intelligent Systems*, 28(2):55–63.
- Alessandra Teresa Cignarella and Cristina Bosco. 2019. ATC at IroSvA 2019: Shallow Syntactic Dependency-based Features for 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.
- Alessandra Teresa Cignarella, Cristina Bosco, and Viviana Patti. 2017. TWITTIRÒ: a Social Media Corpus with a Multi-layered Annotation for Irony. In *Proceedings of the Fourth Italian Conference on Computational Linguistics (CLiC-it 2017)*, volume 2006, pages 101–106. CEUR-WS.org.
- Alessandra Teresa Cignarella, Cristina Bosco, Viviana Patti, and Mirko Lai. 2018a. Application and Analysis of a Multi-layered Scheme for Irony on the Italian Twitter Corpus TWITTIRÒ. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC-2018)*, pages 4204–4211. ELRA.
- Alessandra Teresa Cignarella, Simona Frenda, Valerio Basile, Cristina Bosco, Viviana Patti, Paolo Rosso, et al. 2018b. Overview of the EVALITA 2018 task on Irony Detection in Italian Tweets (IronITA). In *Sixth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian (EVALITA 2018)*, volume 2263, pages 1–6. CEUR-WS.org.
- Alessandra Teresa Cignarella, Cristina Bosco, and Paolo Rosso. 2019. Presenting TWITTIRÒ-UD: An Italian Twitter Treebank in Universal Dependencies. In *Proceedings of SyntaxFest 2019*.
- John R Firth. 1957. A synopsis of linguistic theory, 1930-1955. *Studies in linguistic analysis*.
- Simona Frenda and Viviana Patti. 2019. Computational models for irony detection in three 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.
- Aniruddha Ghosh, Guofu Li, Tony Veale, Paolo Rosso, Ekaterina Shutova, John Barnden, and Antonio Reyes. 2015. Semeval-2015 Task 11: Sentiment Analysis of Figurative Language in Twitter. In *Proceedings of the 9th International Workshop on Semantic Evaluation (SemEval-2015)*, pages 470–478. ACL.
- Delia Irazú Hernández Farías, José-Miguel Benedí, and Paolo Rosso. 2015. Applying basic features from sentiment analysis for automatic irony detection. In *Pattern Recognition and Image Analysis*, volume 9117, pages 337–344. Springer.
- Delia Irazú Hernández Farías, Viviana Patti, and Paolo Rosso. 2016. Irony detection in Twitter: The role of affective content. *ACM Transactions on Internet Technology (TOIT)*, 16(3):19.
- Delia Irazú Hernández Farías and Paolo Rosso. 2016. Irony, Sarcasm, and Sentiment Analysis. In *Sentiment Analysis in Social Networks*, pages 113–128. Elsevier Science and Technology.
- Yu-Hsiang Huang, Hen-Hsen Huang, and Hsin-Hsi Chen. 2017. Irony detection with attentive Recurrent Neural Networks. In *European Conference on Information Retrieval*, pages 534–540. Springer.

- Aditya Joshi, Vinita Sharma, and Pushpak Bhattacharyya. 2015. 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)*, pages 757–762. ACL, July.
- Aditya Joshi, Pushpak Bhattacharyya, and Mark J. Carman. 2017. Automatic sarcasm detection: A survey. *ACM Computing Surveys (CSUR)*, 50(5):73.
- Jihen Karoui, Farah Benamara, Véronique Moriceau, Nathalie Aussenac-Gilles, and Lamia Hadrich Belguith. 2015. Towards a contextual pragmatic model to detect irony in tweets. In *53rd Annual Meeting of the Association for Computational Linguistics (ACL 2015)*.
- Jihen Karoui, Farah Benamara, Véronique Moriceau, Viviana Patti, and Cristina Bosco. 2017. 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*.
- Jihen Karoui. 2017. *Détection automatique de l'ironie dans les contenus générés par les utilisateurs*. Ph.D. thesis, Université Toulouse 3 Paul Sabatier; Faculté des Sciences Economiques et de Gestion, Université de Sfax (Tunisie).
- Joakim Nivre, Marie-Catherine De Marneffe, Filip Ginter, Yoav Goldberg, Jan Hajic, Christopher D. Manning, Ryan T. McDonald, Slav Petrov, Sampo Pyysalo, Natalia Silveira, et al. 2016. Universal Dependencies v1: A Multilingual Treebank Collection. In *Proceedings of the 10th International Conference on Language Resources and Evaluation (LREC-2016)*.
- Reynier Ortega, Francisco Rangel, Irazú Hernández, Paolo Rosso, Manuel Montes, and José E. Medina. 2019. 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.
- Kumar Ravi and Vadlamani Ravi. 2017. A novel automatic satire and irony detection using ensemble feature selection and data mining. *Knowledge-Based Systems*, 120:15–33.
- Antonio Reyes, Paolo Rosso, and Tony Veale. 2013. A Multidimensional Approach for Detecting Irony in Twitter. *Language Resources and Evaluation*, 47(1):239–268.
- Ellen Riloff, Ashequl Qadir, Prafulla Surve, Lalin-dra De Silva, Nathan Gilbert, and Ruihong Huang. 2013. 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)*, pages 704–714. ACL.
- Hassan Saif, Yulan He, Miriam Fernandez, and Harith Alani. 2016. Contextual Semantics for Sentiment Analysis of Twitter. *Information Processing & Management*, 52(1):5–19.
- Manuela Sanguinetti, Cristina Bosco, Alberto Lavelli, Alessandro Mazzei, Oronzo Antonelli, and Fabio Tamburini. 2018. PoSTWITA-UD: an Italian Twitter Treebank in Universal Dependencies. In *Proceedings of the 11th International Conference on Language Resources and Evaluation (LREC-2018)*.
- Maria Simi, Cristina Bosco, and Simonetta Montemagni. 2014. Less is more? Towards a reduced inventory of categories for training a parser for the Italian Stanford Dependencies. In *Language Resources and Evaluation 2014*, pages 83–90. European Language Resources Association (ELRA).
- Milan Straka and Jana Straková. 2017. Tokenizing, POS Tagging, Lemmatizing and Parsing UD 2.0 with UDPipe. In *Proceedings of the CoNLL 2017 Shared Task: Multilingual Parsing from Raw Text to Universal Dependencies*, pages 88–99. ACL.
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2018. SemEval-2018 Task 3: Irony detection in English Tweets. In *In Proceedings of the International Workshop on Semantic Evaluation (SemEval-2018)*. ACL.
- Angela P. Wang. 2013. #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*, pages 349–356. ACL.
- Antoine Widlöcher and Yann Mathet. 2009. La plate-forme glozz: environnement d’annotation et d’exploration de corpus. In *Actes de TALN 2009*.
- Chuhan Wu, Fangzhao Wu, Sixing Wu, Junxin Liu, Zhigang Yuan, and Yongfeng Huang. 2018. Thu_ngn at semeval-2018 task 3: Tweet irony detection with densely connected LSTM and multi-task learning. In *Proceedings of The 12th International Workshop on Semantic Evaluation*, pages 51–56.
- Shiwei Zhang, Xiuzhen Zhang, Jeffrey Chan, and Paolo Rosso. 2019. Irony detection via sentiment-based transfer learning. *Information Processing & Management*, 56(5):1633–1644.