

ValenTO at SemEval-2018 Task 3: Exploring the Role of Affective Content for Detecting Irony in English Tweets

Delia Irazú Hernández Farías

Inst. Nacional de Astrofísica,
Óptica y Electrónica (INAOE)
Mexico

dirazuherfa@hotmail.com

Viviana Patti

Dip. di Informatica
University of Turin
Italy

patti@di.unito.it

Paolo Rosso

PRHLT Research Center
Universitat Politècnica de València
Spain

proso@dsic.upv.es

Abstract

In this paper we describe the system used by the ValenTO team in the shared task on *Irony Detection in English Tweets* at SemEval 2018. The system takes as starting point *emotIDM*, an irony detection model that explores the use of affective features based on a wide range of lexical resources available for English, reflecting different facets of affect. We experimented with different settings, by exploiting different classifiers and features, and participated both to the binary irony detection task and to the task devoted to distinguish among different types of irony. We report on the results obtained by our system both in a constrained setting and unconstrained setting, where we explored the impact of using additional data in the training phase, such as corpora annotated for the presence of irony or sarcasm from the state of the art. Overall, the performance of our system seems to validate the important role that affective information has for identifying ironic content in Twitter.

1 Introduction

People use social media platforms as a forum to share and express themselves by using the language in creative ways and employing figurative language devices such as *irony* to achieve different communication purposes. Irony is closely associated with the indirect expression of feelings, emotions and evaluations, and detecting the presence of irony in social media texts is considered a challenge for research in computational linguistics, also for the impact on sentiment analysis, where irony detection is important to avoid misinterpreting the polarity of ironic statements.

Broadly speaking, under the umbrella term of *irony* two main concepts are covered: *verbal irony* and *situational irony*. Verbal irony is commonly defined as a figure of speech where the speaker intends to communicate the opposite of what is liter-

ally said (Sperber and Wilson, 1986). Situational irony, instead refers to a contradictory or unexpected outcome of events (Lucariello, 2014). In Twitter we can find many examples both of verbal irony and of posts where users describe aspects of an ironic situation. Most of the proposed approaches to the automatic detection of irony in social media (Riloff et al., 2013; Buschmeier et al., 2014; Ptáček et al., 2014) take advantage of lexical factors such as n-grams, punctuation marks, among others. Information related to affect has been also exploited (Reyes et al., 2013; Barbieri et al., 2014; Hernández Farías et al., 2015). Other scholars proposed methods exploiting the context surrounding an ironic utterance (Wallace et al., 2015; Karoui et al., 2015). Recently, also deep learning techniques have been applied (Nozza et al., 2016; Poria et al., 2016).

This paper describes our participation in the SemEval-2018 Task 3. The aim of this task is to identify ironic tweets. ValenTO exploited an extended version of *emotIDM* (Hernández Farías et al., 2016), an irony detection model based mainly on affective information. In particular, we experimented the use of a wide range of affect-related features for characterizing the presence of ironic content, covering different facets of affect, from sentiment to finer-grained emotions. Most theorist (Grice, 1975; Wilson and Sperber, 1992; Alba-Juez and Attardo, 2014) recognized, indeed, the important role of affective information for irony communication-comprehension.

2 The emotIDM model

Irony is a very subjective language device that involves the expression of affective contents such as emotions, attitudes, or evaluations towards a particular target. Attempting to take advantage of the emotionally-laden characteristics of ironic expressions, we relied on *emotIDM*, an irony detection

model that, taking advantage of several affective resources available for English (Nissim and Patti, 2016), exploits various facets of affective information from sentiment to finer-grained emotions for characterizing the presence of irony in Twitter (Hernández Farías et al., 2016).

In (Hernández Farías and Rosso, 2016) the robustness of *emotIDM* was assessed over different Twitter state-of-the-art corpora for irony detection (Reyes et al., 2013; Barbieri et al., 2014; Mohammad et al., 2015; Ptáček et al., 2014; Riloff et al., 2013). The obtained results outperform those in the previous works confirming the significance of affective features for irony detection. An additional aspect to be mentioned about *emotIDM* is that it was designed to identify ironic content in a general sense, i.e. considering irony as a broad term covering different types of irony in tweets.

emotIDM comprises two main groups of features that are described below:

Structural Features (Str). This group includes different markers that could help to identify ironic intention in tweets: punctuation marks (colon, exclamation, question marks), Part-Of-Speech labels (verbs, adverbs, nouns, adjectives), emoticons, uppercase characters, among others.

Affective Features. They are organized in three sub-groups representing different facets of affect:

Sentiment-Related Features (*Sent*). Hu&Liu (Hu and Liu, 2004), General Inquirer (Stone and Hunt, 1963), EffectWordNet (Choi and Wiebe, 2014), Subjectivity lexicon (Wilson et al., 2005), and EmoLex (Mohammad and Turney, 2013), AFINN, SWN, Semantic Orientation lexicon (Taboada and Grieve, 2004), and SenticNet (SN) (Cambria et al., 2014).

Emotional Categories (*eCat*). EmoLex, EmoSenticNet (Poria et al., 2013), SentiSense (Carrillo de Albornoz et al., 2012), and the Linguistic Inquiry and Word Count dictionary (Pennebaker et al., 2001).

Dimensional Models of Emotions (*eDim*). ANEW (Bradley and Lang, 1999), DAL (Whissell, 2009), and SN.

3 emotIDM at SemEval-2018 Task 3: Irony Detection in English Tweets

3.1 Task Description and Datasets

In the framework of SemEval-2018 was organized the *Task 3 on Irony detection in English tweets* (Van Hee et al., 2018). The main objective of this

task is to identify the presence of irony in Twitter. It was divided in two different subtasks:

1. **Task A: Ironic vs. non-ironic:** to determine whether a tweet is ironic or not.
2. **Task B: Different types of irony:** to predict one out of four labels: 0) *non-irony* (*nI*), 1) *verbal_irony_by_polarity_contrast* (*vI*), 2) *other_verbal_irony* (*oI*), 3) *situational_irony* (*sI*).

Organizers provided datasets for training and test labeled according the objectives of each subtask. The whole dataset was collected by exploiting a set of hashtags (#irony, #sarcasm and #not). Therefore, a manual annotation process was applied in order to minimize the noise in the data. For Task A, 1,911 ironic and 1,923 non-ironic tweets were provided. While for Task B, the distribution was: 1923 for *nI*, 1393 for *vI*, 213 for *oI* and 328 for *sI*. Participants were allowed to submit systems trained under two settings: *constrained* (*C*), where only the training data provided for the task should be used; *unconstrained* (*U*), where the use of additional data was permitted.

3.2 Our Proposal

We decided to participate to the shared task by using *emotIDM*. By analyzing the training data, an interesting characteristic was found: 857 out of 3,834 tweets contain an URL. From these tweets, 265 were belonging to the ironic class, while 592 were labeled as non-ironic. Notice that, in (Hernández-Farías et al., 2014), the authors found a similar behavior regarding URL information in the dataset provided by the organizers of SentiPOLC-2014 (Basile et al., 2014). Furthermore, Barbieri et al. (2014) exploited a feature for alerting the existence of an URL in a tweet; such feature was ranked among the most discriminative ones according to an information gain analysis. Since, information regarding to the presence of URL in a tweet has proven to be useful for detecting irony in Twitter, we decided to enrich *emotIDM* by adding a binary feature for reflecting the presence of URL in a tweet.

Below, we describe our participation in the task.

Task A: Ironic vs. non-ironic

We addressed this task as a binary classification by taking advantage of two of the most widely applied classifiers in irony detection: Decision Tree (DT) and Support Vector Machine

(SVM)¹. Moreover, we also included Random Forest (RF) as a classifier in our experiments². We carried out a set of experiments for assessing the performance of the original version of *emotIDM* and the one including information concerning URL (*emotIDM+URL*). Besides, to investigate the contribution of the different sets of features in *emotIDM* further experiments were performed. Several classifiers were used in order to identify the most promising setting.

As mentioned before, exploiting external data was allowed in the *unconstrained* setting. We took advantage of a set of corpora previously used in the state of the art in irony detection. We exploited data from a set of corpora collected exploiting different approaches: *self-tagging* or *manual annotation* or *crowd-sourcing*³. We exploited the corpora developed by (Reyes et al., 2013), (Barbieri et al., 2014), (Mohammad et al., 2015), (Ptáček et al., 2014), (Riloff et al., 2013), (Ghosh et al., 2015), (Karoui et al., 2017), and (Sulis et al., 2016). Besides, we also take advantage of an in-house collection of tweets containing the hashtags #irony and #sarcasm⁴.

Table 1 shows the obtained results during the developing phase for Task A. We experimented with different sets of features and classifiers considering a five fold-cross validation setting.

Features	Classifiers					
	DT		RF		SVM	
	C	U	C	U	C	U
emotIDM	0.57	0.70	0.64	0.71	0.64	0.79
emotIDM + URL	0.56	0.74	0.62	0.70	0.64	0.81
Str + Sent	0.59	0.69	0.60	0.70	0.63	0.78
Str+eCat+eDim	0.58	0.69	0.62	0.70	0.65	0.77
Sent+eCat+eDim	0.52	0.61	0.54	0.62	0.57	0.70

Table 1: Training phase: results for Task A with different experimental settings in *C* and *U* scenarios.

SVM emerges as the classifier with the best performance in both *C* and *U* scenarios. We noticed that, when using SVM, adding the URL feature

¹For experimental purposes we used Scikit-learn: <http://scikit-learn.org/>. The default configuration of parameters in the classifiers was applied.

²This was motivated by the fact that it demonstrated a competitive performance for classifying tweets with #irony, #sarcasm, and #not hashtags in (Sulis et al., 2016).

³Further details on the approaches for collecting corpora for irony detection can be found in (Hernández Farías and Rosso, 2016).

⁴The tweets were retrieved during the 2016 US Elections period from 8th up to 18th November 2016.

to *emotIDM* helps to improve the overall performance of our system. When we experimented by removing a set of features in *emotIDM*, a drop in the performance (in most of the cases) is observed. The results of the experiments with external data are higher than those using only the training data. The last row in Table 1 shows the obtained results when only affect-related features were used; even though there is a drop in the performance respect to the experiments using structural features, it seems that affective features on their own provide useful information for irony detection.

We participated in the subtask A by submitting two runs (constrained and unconstrained) exploiting the experimental setting with the best performance: *emotIDM+URL* with a SVM as classifier.

Task B: Different types of irony

Distinguishing between different kinds of ironic devices is still a controversial issue. In computational linguistics, only few research works have attempted to address such a difficult task (Wang, 2013; Barbieri et al., 2014; Sulis et al., 2016; Van Hee et al., 2016). We are interested in assessing the performance of *emotIDM* when it deals with different types of irony, in order to test if a wide variety of affective features can help in discriminating also in the finer-grained classification task here proposed. This could give some insights on the role of affective content among ironic devices having different communication purposes.

emotIDM+URL was trained with the dataset provided for Task B (constrained setting) to test the effectiveness of affective features in such finer-grained task. We exploited the same classifiers than in Task A attempting to evaluate their performance when different classes of irony should be classified. Overall, the best performance was achieved by SVM (see Table 2). However, when the performance of each single class was considered, the best results were those obtained with DT. For this reason, we decided to combine two classifiers with the following criterion: the *sI* and *oI* classes are assigned by the DT; while *irony* and *non-irony* are assigned by SVM or RF. Table 2 shows the obtained results of the experiments carried out over the dataset for Task B. A five fold cross-validation was applied. From the results in Table 2, it can be noticed that when two classifiers are combined the performance of our model improves. The DT + SVM was selected as the system for participating in the Task B.

Classifier (s)	Macro F-measure
DT	0.31
RF	0.30
SVM	0.31
DT + RF	0.33
DT + SVM	0.34

Table 2: Training phase: obtained results for Task B.

3.3 Results

The results of ValenTO participation in the shared task are summarized in Table 3. In Task A, on the official CodaLab ranking, we ranked in the 16th position with the unconstrained version of our submission. When comparing our official result with the one obtained by the best-ranked system (0.7054), it can be noticed that the difference is lower than 0.1 in F-score terms. It is an interesting result considering that our system relies mainly on features covering different facets of affect in ironic tweets, and confirms the key role that such kind of affective information plays for detecting irony in Twitter. In addition, the organizers also provided separate rankings for constrained and unconstrained submissions. Our system ranked in the 17th position in the constrained setting, while in the unconstrained one we ranked as 4th. Moreover, the performance of our system seems to be stable in the two C and U settings. Concerning Task B, our system performed relatively well, considering that we did not apply further tuning to capture different ironic devices. We ranked in the 17th position of 31 submissions in the Official ranking at CodaLab.

Run	Accuracy	Precision	Recall	F1-score
Task A				
C	0.6709	0.5764	0.6431	0.6079
U	0.5982	0.4959	0.7814	0.6067
Task B				
C	0.5599	0.3534	0.3521	0.3520

Table 3: Official results of ValenTO team in both tasks

3.4 Discussion and Error Analysis

Data provided for the task were retrieved by exploiting hashtags #irony, #sarcasm and #not, which according to (Sulis et al., 2016) seems to label different kinds of ironic phenomena. We analyzed the gold standard labels provided by the organizers (where the ironic hashtags were also included in the tweets) in order to see the performance of our model for recognizing tweets labeled with distinct hashtags. Considering the results in Task A, we noticed that our system was able to identify all the three kinds of tweets without any kind of skew towards a particular hashtag.

It somehow confirms the robustness of emotIDM for recognizing irony in a broad sense.

Our system was able to correctly identify instances expressing an apparent positive emotional distress with an ironic intention, such as: *Sunday is such a fun day to study #ew #saywhat* and *Yay I just love this time of the month...!*. A special mention is for tweets labeled with #not. This hashtag is not always used for highlighting a figurative meaning. Our system was able to correctly identify instances containing #not when it was used for figurative meaning such as: *Yay for Fire Alarms at 3AM #not*, and also when it was used as part of the text in a tweet: *#Myanmar #men #plead #not #guilty to #murder of #British #tourists. http://t.co/flrKr3H6Kl via @reuters*.

For what concerns the performance of emotIDM in Task B, Table 4⁵ shows that our model performed better in identifying tweets where verbal irony was expressed by means of a polarity contrast. Moreover, it was recognizing better “situational irony” than “other irony”.

	vI	oI	sI	nI	tT	Correct (%)
vI	75	10	14	66	166	45
oI	6	4	11	25	62	9.2
sI	16	9	13	35	85	40
nI	67	39	47	347	473	79

Table 4: Confusion Matrix: Task B.

Since our model relies mainly on affective information, ironic instances lacking of subjective-related content are hard to recognize, as in: *Being a hipster now is so mainstream. Oh, the irony. #hipster #irony*. Moreover, we found some tweets where context information is crucial for capturing the ironic sense, like in: *So there used to be a crossfit place here.... #irony #pizzawins http://t.co/9BDkxT9GFJ*; or where the hashtag is the only signal for ironic intention.

4 Conclusions

In this paper, we described our participation at SemEval-2018 Task 3. We exploited an enhanced version of emotIDM. In our experiments, SVM emerges as the classifier with the best performance. The obtained results serve to validate the usefulness of affect-related features for distinguishing ironic tweets. As future work, it could be interesting to enrich emotIDM with features for capturing other kinds of information such as common-knowledge and semantic incongruity.

⁵The column “tT” refers to the amount of tweets in the test set. The column “Correct” refers to the percentage of instances correctly classified per class.

Acknowledgments

The work of D. I. Hernández Farías was funded by CONACYT project FC-2016/2410. The work of P. Rosso has been funded by the SomEMBED TIN2015-71147-C2-1-P MINECO project. The work of V. Patti was partially funded by the IHatePrejudice project (S1618_L2_BOSC_01).

References

- Laura Alba-Juez and Salvatore Attardo. 2014. The Evaluative Palette of Verbal Irony. In Geoff Thompson and Laura Alba-Juez, editors, *Evaluation in Context*, pages 93–116. John Benjamins Publishing Company, Amsterdam/Philadelphia.
- Jorge Carrillo de Albornoz, Laura Plaza, and Pablo Gervás. 2012. SentiSense: An Easily Scalable Concept-based Affective Lexicon for Sentiment Analysis. In *Proceedings of the Eight International Conference on Language Resources and Evaluation (LREC'12)*, pages 3562–3567. European Language Resources Association (ELRA).
- 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. Association for Computational Linguistics.
- 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 Fourth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian EVALITA 2014*, pages 50–57.
- Margaret M Bradley and Peter J Lang. 1999. Affective Norms for English Words (ANEW): Instruction Manual and Affective Ratings. Technical report, Center for Research in Psychophysiology, University of Florida, Gainesville, Florida.
- Konstantin Buschmeier, Philipp Cimiano, and Roman Klinger. 2014. An Impact Analysis of Features in a Classification Approach to Irony Detection in Product Reviews. In *Proceedings of the 5th Workshop on Computational Approaches to Subjectivity, Sentiment and Social Media Analysis*, pages 42–49.
- Erik Cambria, Daniel Olsher, and Dheeraj Rajagopal. 2014. SenticNet 3: A Common and Common-Sense Knowledge Base for Cognition-Driven Sentiment Analysis. In *Proceedings of AAAI Conference on Artificial Intelligence*, pages 1515–1521. AAAI.
- Yoonjung Choi and Janyce Wiebe. 2014. +/-EffectWordNet: Sense-level Lexicon Acquisition for Opinion Inference. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1181–1191.
- 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*, pages 470–478. Association for Computational Linguistics.
- H. P. Grice. 1975. Logic and Conversation. In P. Cole and J. L. Morgan, editors, *Syntax and Semantics: Vol. 3: Speech Acts*, pages 41–58. Academic Press.
- Delia Irazú Hernández Farías, Viviana Patti, and Paolo Rosso. 2016. Irony Detection in Twitter: The Role of Affective Content. *ACM Trans. Internet Technol.*, 16(3):19:1–19:24.
- Delia Irazú Hernández Farías and Paolo Rosso. 2016. Irony, Sarcasm, and Sentiment Analysis. Chapter 7. In Federico A. Pozzi, Elisabetta Fersini, Enza Messina, and Bing Liu, editors, *Sentiment Analysis in Social Networks*, pages 113–127. Morgan Kaufmann.
- 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 of *Lecture Notes in Computer Science*, pages 337–344. Springer International Publishing.
- Irazú Hernández-Farías, Davide Buscaldi, and Belém Priego-Sánchez. 2014. IRADABE: Adapting English Lexicons to the Italian Sentiment Polarity Classification Task. In *Proceedings of the First Italian Conference on Computational Linguistics (CLiC-it 2014) & the Fourth Evaluation Campaign of Natural Language Processing and Speech Tools for Italian EVALITA 2014*, pages 75–81.
- Minqing Hu and Bing Liu. 2004. Mining and Summarizing Customer Reviews. In *Proceedings of the Tenth ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '04*, pages 168–177, Seattle, WA, USA. ACM.
- 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 *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 644–650. Association for Computational Linguistics.
- Jihen Karoui, Farah Benamara, Veronique Moriceau, Viviana Patti, Cristina Bosco, and Nathalie Aussenac-Gilles. 2017. . In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 1, Long Papers*, Valencia, Spain. Association for Computational Linguistics.

- Joan Lucariello. 2014. Situational Irony: A Concept of Events Gone Awry. *Journal of Experimental Psychology: General*, 123(2):129–145.
- Saif M. Mohammad and Peter D. Turney. 2013. Crowdsourcing a Word–Emotion Association Lexicon. *Computational Intelligence*, 29(3):436–465.
- Saif M. Mohammad, Xiaodan Zhu, Svetlana Kiritchenko, and Joel Martin. 2015. [Sentiment, Emotion, Purpose, and Style in Electoral Tweets](#). *Information Processing & Management*, 51(4):480–499.
- Malvina Nissim and Viviana Patti. 2016. Semantic aspects in sentiment analysis. In Fersini Elisabetta, Bing Liu, Enza Messina, and Federico Pozzi, editors, *Sentiment Analysis in Social Networks*, chapter 3, pages 31–48. Elsevier.
- Debora Nozza, Elisabetta Fersini, and Enza Messina. 2016. [Unsupervised Irony Detection: A Probabilistic Model with Word Embeddings](#). In *Proceedings of the 8th International Joint Conference on Knowledge Discovery, Knowledge Engineering and Knowledge Management*, pages 68–76.
- James W. Pennebaker, Martha E. Francis, and Roger J. Booth. 2001. Linguistic Inquiry and Word Count: LIWC 2001. *Mahway: Lawrence Erlbaum Associates*, 71.
- Soujanya Poria, Erik Cambria, Devamanyu Hazarika, and Prateek Vij. 2016. [A Deeper Look into Sarcastic Tweets Using Deep Convolutional Neural Networks](#). *CoRR*, abs/1610.08815.
- Soujanya Poria, Alexander Gelbukh, Amir Hussain, Newton Howard, Dipankar Das, and Sivaji Bandyopadhyay. 2013. [Enhanced SenticNet with Affective Labels for Concept-Based Opinion Mining](#). *IEEE Intelligent Systems*, 28(2):31–38.
- Tomáš Ptáček, Ivan Habernal, and Jun Hong. 2014. Sarcasm Detection on Czech and English Twitter. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics*, pages 213–223. Dublin City University and Association for Computational Linguistics.
- 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, Lalindra 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*, pages 704–714. Association for Computational Linguistics.
- Dan Sperber and Deirdre Wilson. 1986. *Relevance: Communication and Cognition*. Harvard University Press, Cambridge, MA, USA.
- Philip J. Stone and Earl B. Hunt. 1963. [A Computer Approach to Content Analysis: Studies Using the General Inquirer System](#). In *Proceedings of the May 21-23, 1963, Spring Joint Computer Conference*, AFIPS '63 (Spring), pages 241–256. ACM.
- Emilio Sulis, Delia Irazú Hernández Farías, Paolo Rosso, Viviana Patti, and Giancarlo Ruffo. 2016. [Figurative Messages and Affect in Twitter: Differences between #irony, #sarcasm and #not](#). *Knowledge-Based Systems*, 108:132–143.
- Maite Taboada and Jack Grieve. 2004. Analyzing Appraisal Automatically. In *Proceedings of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*, pages 158–161, Stanford, US. AAAI.
- Cynthia Van Hee, Els Lefever, and Veronique Hoste. 2016. Exploring the realization of irony in Twitter data. In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC 2016)*. European Language Resources Association (ELRA).
- Cynthia Van Hee, Els Lefever, and Véronique Hoste. 2018. SemEval-2018 Task 3: Irony Detection in English Tweets. In *Proceedings of the 12th International Workshop on Semantic Evaluation, SemEval-2018*, New Orleans, LA, USA. Association for Computational Linguistics.
- Byron C. Wallace, Do Kook Choe, and Eugene Charniak. 2015. [Sparse, Contextually Informed Models for Irony Detection: Exploiting User Communities, Entities and Sentiment](#). In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 1035–1044. Association for Computational Linguistics.
- Angela P. Wang. 2013. #irony or #sarcasm — A Quantitative and Qualitative Study Based on Twitter. In *Proceedings of the 27th Pacific Asia Conference on Language, Information, and Computation*, pages 349–356. National Chengchi University.
- Cynthia Whissell. 2009. Using the Revised Dictionary of Affect in Language to Quantify the Emotional Undertones of Samples of Natural Languages. *Psychological Reports*, 2(105):509–521.
- Deirdre Wilson and Dan Sperber. 1992. [On Verbal Irony](#). *Lingua*, 87(1-2):53–76.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffmann. 2005. [Recognizing Contextual Polarity in Phrase-level Sentiment Analysis](#). In *Proceedings of the Conference on Human Language Technology and Empirical Methods in Natural Language Processing, HLT '05*, pages 347–354. Association for Computational Linguistics.