

Mining Subjective Knowledge from Customer Reviews: A Specific Case of Irony Detection

Antonio Reyes and Paolo Rosso

Natural Language Engineering Lab - ELiRF
Departamento de Sistemas Informáticos y Computación
Universidad Politécnica de Valencia, Spain
{areyes,pross}@dsic.upv.es

Abstract

The research described in this work focuses on identifying key components for the task of irony detection. By means of analyzing a set of customer reviews, which are considered as ironic both in social and mass media, we try to find hints about how to deal with this task from a computational point of view. Our objective is to gather a set of discriminating elements to represent irony. In particular, the kind of irony expressed in such reviews. To this end, we built a freely available data set with ironic reviews collected from Amazon. Such reviews were posted on the basis of an online viral effect; i.e. contents whose effect triggers a chain reaction on people. The findings were assessed employing three classifiers. The results show interesting hints regarding the patterns and, especially, regarding the implications for sentiment analysis.

1 Introduction

Verbal communication is not a trivial process. It implies to share a common code as well as being able to infer information beyond the semantic meaning. A lot of communicative acts imply information not grammatically expressed to be able to decode the whole sense: if the hearer is not capable to infer that information, the communicative process is incomplete. Let us consider a joke. The amusing effect sometimes relies on not given information. If such information is not filled, the result is a bad, or better said, a misunderstood joke. This information, which is not expressed with “physical” words, supposes a great challenge, even from a linguistic analysis, because it points to social and cognitive layers quite difficult to be computationally represented. One of the communicative phenomena which better represents this problem is irony. According to Wilson

and Sperber (2007), irony is essentially a communicative act which expresses an opposite meaning of what was literally said.

Due to irony is common in texts that express subjective and deeply-felt opinions, its presence represents a significant obstacle to the accurate analysis of sentiment in such texts (cf. Councill et al. (2010)). In this research work we aim at gathering a set of discriminating elements to represent irony. In particular, we focus on analyzing a set of customer reviews (posted on the basis of an online viral effect) in order to obtain a set of key components to face the task of irony detection.

This paper is organized as follows. Section 2 introduces the theoretical problem of irony. Section 3 presents the related work as well as the evaluation corpus. Section 4 describes our model and the experiments that were performed. Section 5 assesses the model and presents the discussion of the results. Finally, Section 6 draws some final remarks and addresses the future work.

2 Pragmatic Theories of Irony

Literature divides two primary classes of irony: verbal and situational. Most theories agree on the main property of the former: verbal irony conveys an opposite meaning; i.e. a speaker says something that seems to be the opposite of what s/he means (Colston and Gibbs, 2007). In contrast, situational irony is a state of the world which is perceived as ironical (Attardo, 2007); i.e. situations that should not be (Lucariello, 2007). Our work focuses on verbal irony. This kind of irony is defined as a way of intentionally denying what it is literally expressed (Curc6, 2007); i.e. a kind of indirect negation (Giora, 1995). On the basis of some pragmatic frameworks, authors focus on certain fine-grained aspects of this term. For instance, Grice (1975) con-

siders that an utterance is ironic if it intentionally violates some conversational maxims. Wilson and Sperber (2007) assume that verbal irony must be understood as echoic; i.e. as a distinction between use and mention. Utsumi (1996), in contrast, suggests an ironic environment which causes a negative emotional attitude. According to these points of view, the elements to conceive a verbal expression as ironic point to different ways of explaining the same underlying concept of opposition, but specially note, however, that most of them rely on literary studies (Attardo, 2007); thus, their computational formalization is quite challenging. Furthermore, consider that people have their own concept of irony, which often does not match with the rules suggested by the experts. For instance, consider the following expressions retrieved from the web:

1. "If you find it hard to laugh at yourself, I would be happy to do it for you."
2. "Let's pray that the human race never escapes from Earth to spread its iniquity elsewhere."

These examples, according to some user-generated tags, could be either ironic, or sarcastic, or even satiric. However, the issue we want to focus does not lie on what tag should be the right for every expression, but on the fact that there is not a clear distinction about the boundaries among these terms. For Colston (2007), sarcasm is a term commonly used to describe an expression of verbal irony; whereas for Gibbs (2007), sarcasm along with jocularly, hyperbole, rhetorical questions, and understatement, are types of irony. Attardo (2007) in turn, considers that sarcasm is an overtly aggressive type of irony. Furthermore, according to Gibbs and Colston (2007), irony is often compared to satire and parody.

In accordance with these statements, the limits among these figurative devices are not clearly differentiable. Their differences rely indeed on matters of usage, tone, and obviousness, which are not so evident in ordinary communication acts. Therefore, if there are no formal boundaries to separate these concepts, even from a theoretical perspective, people will not be able to produce ironic expressions as the experts suggest. Instead, there will be a mixture of expressions pretending to be ironic but being sarcastic, satiric, or even humorous. This get worse

when dealing with non prototypical examples. Observe the following fragment from our corpus:

3. "I am giving this product [a t-shirt] 5 stars because not everyone out there is a ladies' man. In the hands of lesser beings, it can help you find love. In the hands of a playa like me, it can only break hearts. That's why I say use with caution. I am passing the torch onto you, be careful out there folks."

In this text irony is perceived as a mixture of sarcasm and satire, whose effect is not only based on expressing an opposite or negative meaning, but a humorous one as well.

Taking into account these assumptions, we begin by defining irony as a *verbal subjective expression whose formal constituents attempt to communicate an underlying meaning, focusing on negative or humorous aspects, which is opposite to the one expressed*. Based on this definition, we consider sarcasm, satire, and figures such as the ones suggested in (Gibbs, 2007), as specific extensions of a general concept of irony, and consequently, we will not make any fine-grained distinction among them; i.e. irony will include them.

3 Approaching Irony Detection

As far as we know, very few attempts have been carried out in order to integrate irony in a computational framework. The research described by Utsumi (1996) was one of the first approaches to computationally formalize irony. However, his model is too abstract to represent irony beyond an idealized hearer-listener interaction. Recently, from a computational creativity perspective, Veale and Hao (2009) focused on studying irony by analyzing humorous similes. Their approach gives some hints to explain the cognitive processes that underly irony in such structures. In contrast, Carvalho et al. (2009) suggested some clues for automatically identifying ironic sentences by means of identifying features such as emoticons, onomatopoeic expressions, punctuation and quotation marks. Furthermore, there are others approaches which are focused on particular devices such as sarcasm and satire, rather than on the whole concept of irony. For instance, Tsur et al. (2010) and Davidov et al. (2010) address the problem of finding linguistic elements that mark the use of sarcasm in online product reviews and tweets, respectively. Finally, Burfoot and Baldwin (2009) explore the task of automatic satire

detection by evaluating features related to headline elements, offensive language and slang.

3.1 Evaluation Corpus

Due to the scarce work on automatic irony processing, and to the intrinsic features of irony, it is quite difficult and subjective to obtain a corpus with ironic data. Therefore, we decided to rely on the wisdom of the crowd and use a collection of customer reviews from the Amazon web site. These reviews are considered as ironic by customers, as well as by many journalists, both in mass and social media. According to such means, all these reviews deal with irony, sarcasm, humor, satire and parody (hence, they are consistent with our definition of irony). All of them were posted by means of an online viral effect, which in most cases, increased the popularity and sales of the reviewed products. The *Three Wolf Moon T-shirt* is the clearest example. This item became one of the most popular products, both in Amazon as well as in social networks, due to the ironic reviews posted by people¹.

Our positive data are thus integrated with reviews of five different products published by Amazon. All of them were posted through the online viral effect. The list of products is: i) *Three Wolf Moon T-shirt* (product id: B002HJ377A); ii) *Tuscan Whole Milk* (product id: B00032G1S0); iii) *Zubaz Pants* (product id: B000WVXM0W); iv) *Uranium Ore* (product id: B000796XXM); and v) *Platinum Radiant Cut 3-Stone* (product id: B001G603AE). A total of 3,163 reviews were retrieved. Then, in order to automatically filter the ones more likely to be ironic without performing a manual annotation (which is planned to be carried out in the near future), we removed the reviews whose customer rating, according to the Amazon rating criteria, was lesser than four stars. The assumptions behind this decision rely on two facts: i) the viral purpose, and ii) the ironic effect. The former caused that people to post reviews whose main purpose, and perhaps the only one, was to exalt superficial properties and non-existent consequences; thus the possibilities to find *real* reviews were minimal. Considering this scenario, the lat-

¹According to results obtained with Google, apart from the more than one million of results retrieved when searching this product, there are more than 10,000 blogs which comment the effect caused by these reviews.

ter supposes that, if someone ironically wants to reflect properties and consequences such as the previous ones, s/he will not do it by rating the products with one or two stars, instead, s/he will rate them with the highest scores.

After applying this filter, we obtained an ironic set integrated with 2,861 documents. On the other hand, two negative sets were automatically collected from two sites: Amazon.com (AMA) and Slashdot.com (SLA). Each contains 3,000 documents. The products selected from AMA were: Bananagrams (toy), The Help by Kathryn Stockett (book), Flip UltraHD Camcorder (camera), I Dreamed A Dream (CD), Wii Fit Plus with Balance Board (Videogame console). Finally, the data collected from SLA contain web comments categorized as funny in a community-driven process. The whole evaluation corpus is integrated with 8,861 documents. It is available at <http://users.dsic.upv.es/grupos/nle>.

4 Model

We define a model with six categories which attempts to represent irony from different linguistic layers. These categories are: *n-grams*, *POS n-grams*, *funny profiling*, *positive/negative profiling*, *affective profiling*, and *pleasantness profiling*.

4.1 N-grams

This category focuses on representing the ironic documents in the simplest way: with sequences of n-grams (from order 2 up to 7) in order to find a set of recurrent words which might express irony. Note that all the documents were preprocessed. Firstly, the stopwords were removed, and then, all the documents were stemmed. The next process consisted in removing irrelevant terms by applying a *tf - idf* measure. This measure assesses how relevant a word is, given its frequency both in a document as in the entire corpus. Irrelevant words such as *t-shirt*, *wolf*, *tuscan*, *milk*, etc., were then automatically eliminated. The complete list of filtered words, stopwords included, contains 824 items. Examples of the most frequent sequences are given in Table 1.

4.2 POS n-grams

The goal of this category is to obtain recurrent sequences of morphosyntactic patterns. According to

Table 1: Statistics of the most frequent word n-grams.

Order	Sequences	Examples
2-grams	160	opposit sex; american flag; alpha male
3-grams	82	sex sex sex; fun educ game
4-grams	78	fun hit reload page; remov danger reef pirat
5-grams	76	later minut custom contribut product
6-grams	72	fals function player sex sex sex
7-grams	69	remov danger reef pirat fewer shipwreck surviv

our definition, irony looks for expressing an opposite meaning; however, the ways of transmitting that meaning are enormous. Therefore, we pretend to symbolize an abstract structure through sequences of POS tags (hereafter, POS-grams) instead of only words. It is worth highlighting that a statistical substring reduction algorithm (Lü et al., 2004) was employed in order to eliminate redundant sequences. For instance, if the sequences “he is going to look so hot in this shirt” and “he is going to look hot in this shirt” occur with similar frequencies in the corpus, then, the algorithm removes the last one because is a substring of the first one. Later on, we labeled the documents employing the FreeLing resource (Atserias et al., 2006). The N-best sequences of POS-grams, according to orders 2 up to 7, are given in Table 2.

4.3 Funny profiling

Irony takes advantage of humor aspects to produce its effect. This category intends to characterize the documents in terms of humorous properties. In order to represent this category, we selected some of the best humor features reported in the literature: *stylistic features*, *human centeredness*, and *keyness*. The stylistic features, according to the experiments reported in (Mihalcea and Strapparava, 2006), were obtained by collecting all the words labeled with the tag “sexuality” in WordNet Domains (Bentivogli et al., 2004). The second feature focuses on social relationships. In order to retrieve these words, the elements registered in WordNet (Miller, 1995), which belong to the synsets *relation*, *relationship* and *relative*, were retrieved. The last feature is represented by obtaining the keyness value of the words (cf. (Reyes et al., 2009)). This value is calculated comparing the word frequencies in the ironic documents against their frequencies in a reference corpus. Google N-grams (Brants and Franz, 2006) was

Table 2: Statistics of the most frequent POS-grams.

Order	Sequences	Examples
2-grams	300	dt nn; nn in; jj nn; nn nn
3-grams	298	dt nn in; dt jj nn; jj nn nn
4-grams	282	nn in dt nn; vb dt jj nn
5-grams	159	vbd dt vbg nn jj
6-grams	39	nnp vbd dt vbg nn jj
7-grams	65	nns vbd dt vbg nn jj fd

used as the reference corpus. Only the words whose keyness was ≥ 100 were kept.

4.4 Positive/Negative Profiling

As we have already pointed out, one of the most important properties of irony relies on the communication of negative information through positive one. This category intends to be an indicator about the correlation between positive and negative elements in the data. The Macquarie Semantic Orientation Lexicon (MSOL) (Saif et al., 2009) was used to label the data. This lexicon contains 76,400 entries (30,458 positive and 45,942 negative ones).

4.5 Affective Profiling

In order to enhance the quality of the information related to the expression of irony, we considered to represent information linked to psychological layers. The affective profiling category is an attempt to characterize the documents in terms of words which symbolize subjective contents such as emotions, feelings, moods, etc. The WordNet-Affect resource (Strapparava and Valitutti, 2004) was employed for obtaining the affective terms. This resource contains 11 classes to represent affectiveness. According to the authors, these classes represent how speakers convey affective meanings by means of selecting certain words and not others.

4.6 Pleasantness Profiling

The last category is an attempt to represent ideal cognitive scenarios to express irony. This means that, like words, the contexts in which irony appears are enormous. Therefore, since it is impossible to make out all the possibilities, we pretend to define a schema to represent favorable and unfavorable ironic contexts on the basis of pleasantness values. In order to represent those values, we used the Dictionary of Affect in Language (Whissell, 1989). This dictionary assigns a score of pleasantness to

~ 9,000 English words. The scores were obtained from human ratings. The range of scores goes from 1 (unpleasant) to 3 (pleasant).

5 Evaluation

In order to verify the effectiveness of our model, we evaluated it through a classification task. Two underlying goals were analyzed: a) feature relevance; and b) the possibility of automatically finding ironic documents.

The classifiers were evaluated by comparing the positive set against each of the two negative subsets (AMA and SLA, respectively). All the documents were represented as frequency-weighted term vectors according to a representativeness ratio. This ratio was estimated using Formula 1:

$$\delta(d_k) = \frac{\sum_{i,j} fdf_{i,j}}{|d|} \quad (1)$$

where i is the i -th conceptual category ($i = 1 \dots 6$); j is the j -th feature of i ; $fdf_{i,j}$ (*feature dimension frequency*) is the frequency of features j of category i ; and $|d|$ is the length of the k -th document d_k . For categories funny, positive/negative, affective, and pleasantness, we determined an empirical threshold of representativeness ≥ 0.5 . A document was assigned the value = 1 (presence) if its δ exceeded the threshold, otherwise a value = 0 (absence) was assigned. A different criterion was determined for the n-grams and POS-grams because we were not only interested in knowing whether or not the sequences appeared in the corpus, but also in obtaining a measure to represent the degree of similarity among the sets. In order to define a similarity score, we used the Jaccard similarity coefficient.

The classification accuracy was assessed employing three classifiers: Naïve Bayes (NB), support vector machines (SVM), and decision trees (DT). The sets were trained with 5,861 instances (2,861 positive and 3,000 negative ones). 10-fold cross validation method was used as test. Global accuracy as well as detailed performance in terms of *precision*, *recall*, and *F – measure*, are given in Table 3.

5.1 Discussion

Regarding the first goal (feature relevance), our a-priori aim of representing some irony features in

Table 3: Classification results.

		Accuracy	Precision	Recall	F-Measure
NB	AMA	72,18%	0,745	0,666	0,703
	SLA	75,19%	0,700	0,886	0,782
SVM	AMA	75,75%	0,771	0,725	0,747
	SLA	73,34%	0,706	0,804	0,752
DT	AMA	74,13%	0,737	0,741	0,739
	SLA	75,12%	0,728	0,806	0,765

terms of six general categories seems to be acceptable. According to the results depicted in Table 3, the proposed model achieves good rates of classification which support this assumption: from 72% up to 89%, whereas a classifier that labels all texts as non-ironic would achieve an accuracy around 54%. Moreover, both precision and recall, as well as F-measure rates corroborate the effectiveness of such performance: most of classifiers obtained scores > 0.7 . This means that, at least regarding the data sets employed in the experiments, the capabilities for differentiating an ironic review from a non-ironic one, or a web comment, are satisfactory.

With respect to the second goal, an information gain filter was applied in order to verify the relevance of the model for finding ironic documents regarding the different *discourses* profiled in each negative subset. In Table 4 we detailed the most discriminating categories per subset according to their information gain scores. On the basis of the results depicted in this table, it is evident how the relevance of the categories varies in function of the negative subset. For instance, when classifying the AMA subset, it is clear how the POS-grams (order 3), pleasantness and funny categories, are the most informative ones; in contrast, the pleasantness, n-grams (order 5) and funny categories, are the most relevant ones regarding the SLA subset. Moreover, it is important to note how the negative words, without being the most differentiable ones, function as discriminating elements.

Table 4: The 5 most discriminating categories regarding information gain results.

AMA	POS 3-grams	Pleasantness	Funny	POS 2-grams	POS 4-grams
SLA	Pleasantness	5-grams	Funny	Affectiveness	6-grams

Taking into consideration all previous remarks, we would like to stress some observations with re-

spect to each category. Regarding the **n-grams**, it is important to note the presence of some interesting sequences which are not common to the three subsets. For instance: *pleasantly surprised*. However, we cannot define irony only in terms of these sequences because they might represent domain-specific information such as the bigram: *customer service*.

With respect to the **POS-grams**, the fact of focusing on morphosyntactic templates instead of only on words seem to be more affective. For instance, the sequence *noun + verb + noun + adjective* would represent more information than the sum of simple words: *[grandpa/hotel/bed] + [looks/appears/seems] + [years/days/months] + [younger/bigger/dirtier]*. These sequences of POS tags show how an abstract representation could be more useful than a simple word representation.

The **funny** category seems to be a relevant element to express irony. However, its relevance might be supported by the kind of information profiled in the positive set. Considering the comic trend in the reviews posted by Amazon's customers, it is likely that many of the words belonging to this category appeared in such reviews. For instance, in the following example the words in italics represent funny elements: "I am an attractive *guy*. Slender, weak, and I have never shaved in my 19 years, but *sexy* as hell, and I cannot tell you how many *women* have flocked to me since my purchase". Regardless, it is important to stress that this category is equally discriminating for all sets, funny web comments included.

Concerning the **positive/negative profiling**, it is necessary to emphasize that, despite the greater number of negative words in the MSOL (more than 15,000 words of difference; cf. Section 4.4), the positive elements are the most representative in the ironic documents. This fact corroborates the assumption about the use of positive information in order to express an underlying negative meaning: "The cool_{POS}, refreshing_{POS} taste_{POS} of the milk_{POS} washed away my pain_{NEG} and its kosher_{POS} source_{POS} of calcium_{POS} wash away my fear_{NEG}".

Regarding the **affective** category, its relevance is not as important as we have a-priori considered, despite it is one of the categories used to discriminate the SLA subset: "Man, that was *weird* . . . I think is

funny, because there's a *good* overlap". However, if we take into account the whole accuracy for this subset, then we can conclude that its relevance is minor. Nonetheless, we still consider that the affective information is a valuable factor which must be taken into account in order to provide rich knowledge related to subjective layers of linguistic representation.

The role played by the **pleasantness** category on the classifications is significant. Despite the category is not the most discriminating, its effectiveness for increasing the classification accuracy is remarkable. For instance, consider the following ironic sentence: "I *became* the *man* I *always* dreamed I *could be* all those *nights* *staying up* late watching *wrestling*", where most of its constituents are words whose pleasantness score is ≥ 2.5 ; i.e. these words (in italics) should communicate information related to favorable pleasant contexts.

6 Conclusions and Future Work

Irony is one of the most subjective phenomena related to linguistic analysis. Its automatic processing is a real challenge, not only from a computational perspective but from a linguistic one as well. In this work we have suggested a model of six categories which attempts to describe salient characteristics of irony. They intend to symbolize low and high level properties of irony on the basis of formal linguistic elements. This model was assessed by creating a freely available data set with ironic reviews. The results achieved with three different classifiers are satisfactory, both in terms of classification accuracy, as well as precision, recall, and F-measure. Further work consists of improving the quality of every category, as well as of identifying new ones in order to come up with an improved model capable to detect better ironic patterns in different kinds of texts.

Acknowledgments

The National Council for Science and Technology (CONACyT - México) has funded the research of the first author. This work was carried out in the framework of the MICINN Text-Enterprise (TIN2009-13391-C04-03) research project and the Microcluster VLC/Campus (International Campus of Excellence) on Multimodal Intelligent Systems.

References

- J. Atserias, B. Casas, E. Comelles, M. González, L. Padró, and M. Padró. 2006. Freeling 1.3: Syntactic and semantic services in an open-source nlp library. In *Proceedings of the 5th International Conference on Language Resources and Evaluation*, pages 48–55.
- S. Attardo. 2007. Irony as relevant inappropriateness. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 135–174. Taylor and Francis Group.
- L. Bentivogli, P. Forner, B. Magnini, and E. Pianta. 2004. Revising the wordnet domains hierarchy: semantics, coverage and balancing. In Gilles Sérasset, editor, *Multilingual Linguistic Resources*, pages 94–101.
- T. Brants and A. Franz. 2006. Web 1t 5-gram corpus version 1.
- C. Burfoot and T. Baldwin. 2009. Automatic satire detection: Are you having a laugh? In *ACL-IJCNLP '09: Proceedings of the ACL-IJCNLP 2009 Conference Short Papers*, pages 161–164.
- P. Carvalho, L. Sarmento, M. Silva, and E. de Oliveira. 2009. Clues for detecting irony in user-generated contents: oh...!! it's "so easy" ;-). In *TSA '09: Proceeding of the 1st international CIKM workshop on Topic-sentiment analysis for mass opinion*, pages 53–56.
- H. Colston and R. Gibbs. 2007. A brief history of irony. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 3–24. Taylor and Francis Group.
- H. Colston. 2007. On necessary conditions for verbal irony comprehension. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 97–134. Taylor and Francis Group.
- I. Councill, R. McDonald, and L. Velikovich. 2010. What's great and what's not: learning to classify the scope of negation for improved sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing*, pages 51–59, July.
- C. Curcó. 2007. Irony: Negation, echo, and metarepresentation. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 269–296. Taylor and Francis Group.
- D. Davidov, O. Tsur, and A. Rappoport. 2010. Semi-supervised recognition of sarcastic sentences in Twitter and Amazon. In *Proceeding of the 23rd international conference on Computational Linguistics*, July.
- R. Gibbs and H. Colston. 2007. The future of irony studies. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 339–360. Taylor and Francis Group.
- R. Gibbs. 2007. Irony in talk among friends. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 339–360. Taylor and Francis Group.
- R. Giora. 1995. On irony and negation. *Discourse Processes*, 19(2):239–264.
- H. Grice. 1975. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and semantics*, volume 3, pages 41–58. New York: Academic Press.
- X. Lü, L. Zhang, and J. Hu. 2004. Statistical substring reduction in linear time. In *Proceedings of IJCNLP-04, HaiNan island*.
- J. Lucariello. 2007. Situational irony: A concept of events gone away. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 467–498. Taylor and Francis Group.
- R. Mihalcea and C. Strapparava. 2006. Learning to Laugh (Automatically): Computational Models for Humor Recognition. *Journal of Computational Intelligence*, 22(2):126–142.
- G. Miller. 1995. Wordnet: A lexical database for english. *Communications of the ACM*, 38(11):39–41.
- A. Reyes, P. Rosso, and D. Buscaldi. 2009. Humor in the blogosphere: First clues for a verbal humor taxonomy. *Journal of Intelligent Systems*, 18(4):311–331.
- M. Saif, D. Cody, and D. Bonnie. 2009. Generating high-coverage semantic orientation lexicons from overtly marked words and a thesaurus. In *Proceedings of the 2009 Conference on EMNLP*, pages 599–608, Morristown, NJ, USA. Association for Computational Linguistics.
- C. Strapparava and A. Valitutti. 2004. WordNet-affect: an affective extension of WordNet. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, volume 4, pages 1083–1086.
- O. Tsur, D. Davidov, and A. Rappoport. 2010. {ICWSM} — a great catchy name: Semi-supervised recognition of sarcastic sentences in online product reviews. In William W. Cohen and Samuel Gosling, editors, *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, pages 162–169, Washington, D.C., 23–26 May. The AAAI Press.
- A. Utsumi. 1996. A unified theory of irony and its computational formalization. In *Proceedings of the 16th conference on Computational Linguistics*, pages 962–967, Morristown, NJ, USA. Association for Computational Linguistics.
- T. Veale and Y. Hao. 2009. Support structures for linguistic creativity: A computational analysis of creative irony in similes. In *Proceedings of CogSci 2009, the 31st Annual Meeting of the Cognitive Science Society*, pages 1376–1381.
- C. Whissell. 1989. The dictionary of affect in language. *Emotion: Theory, Research, and Experience*, 4:113–131.
- D. Wilson and D. Sperber. 2007. On verbal irony. In R. Gibbs and H. Colston, editors, *Irony in Language and Thought*, pages 35–56. Taylor and Francis Group.