

With 1 Follower I Must Be AWESOME :P. Exploring the Role of Irony Markers in Irony Recognition

Debanjan Ghosh,[§] Smaranda Muresan[‡]

[§]School of Communication & Information, Rutgers University, NJ, USA

[‡]Data Science Institute, Columbia University, NY, USA

debanjan.ghosh@rutgers.edu, smara@columbia.edu

Abstract

Conversations in social media often contain the use of irony or sarcasm, when the users say the opposite of what they really mean. Irony markers are the meta-communicative clues that inform the reader that an utterance is ironic. We propose a thorough analysis of theoretically grounded irony markers in two social media platforms: *Twitter* and *Reddit*. Classification and frequency analysis show that for *Twitter*, typographic markers such as emoticons and emojis are the most discriminative markers to recognize ironic utterances, while for *Reddit* the morphological markers (e.g., interjections, tag questions) are the most discriminative.

Introduction

With the advent of social media, irony and sarcasm detection has become an active area of research in Natural Language Processing (NLP) (Joshi, Bhattacharyya, and Carman 2016; Riloff et al. 2013; Joshi, Sharma, and Bhattacharyya 2015; Ghosh, Fabbri, and Muresan 2017). Most computational studies have focused on building state-of-the-art models to detect whether an utterance or comment is ironic/sarcastic¹ or not, sometimes without theoretical grounding. In linguistics and discourse studies, Attardo (2000) and later Burgers (2010) have studied two theoretical aspects of irony in the text: *irony factors*² and *irony markers*. Irony factors are characteristics of ironic utterances that cannot be removed without destroying the irony. In contrast, irony markers are a meta-communicative clue that “alert the reader to the fact that a sentence is ironical” (Attardo 2000). They can be removed and the utterance is still ironic.

In this paper, we examine the role of irony markers in social media for irony recognition. Although punctuations, capitalization, and hyperboles are previously used as features in irony detection (Bamman and Smith 2015; Muresan et al. 2016), here we thoroughly analyze a set of theoretically-grounded types of irony markers, such as tropes (e.g., metaphors), morpho-syntactic indicators (e.g., tag questions), and typographic markers (e.g., emoji) and their use in ironic utterances. Consider the following two irony examples from *Twitter* and *Reddit* given in Table 1.

Platform	Utterances
<i>Reddit</i>	Are you telling me iPhone 5 is only marginally better than iPhone 4S? I thought we were reaching a golden age with this game-changing device. /s
<i>Twitter</i>	With 1 follower I must be AWESOME. :P #ironic

Table 1: Use of irony markers in two social media platforms

Both utterances are labeled as ironic by their authors (using hashtags in *Twitter* and the /s marker in *Reddit*). In the *Reddit* example, the author uses several irony markers such as *Rhetorical question* (e.g., “are you telling” ...) and *metaphor* (e.g., “golden age”). In the *Twitter* example, we notice the use of capitalization (“AWESOME”) and emoticons (“:P” (tongue out)) that the author uses to alert the readers that it is an ironic tweet.

We present three contributions in this paper. First, we provide a detailed investigation of a set of theoretically-grounded irony markers (e.g., tropes, morpho-syntactic, and typographic markers) in social media. We conduct the classification and frequency analysis based on their occurrence. Second, we analyze and compare the use of irony markers on two social media platforms (*Reddit* and *Twitter*). Third, we provide an analysis of markers on topically different social media content (e.g., technology vs. political subreddits).

Data

***Twitter*:** We use a set of 350K tweets for our experiments. The ironic/sarcastic tweets are collected using hashtags, such as #irony, #sarcasm, and #sarcastic whereas the non-sarcastic tweets do not contain these hashtags, but they might include sentiment hashtags, such as #happy, #love, #sad, #hate (similar to (González-Ibáñez, Muresan, and Wacholder 2011; Ghosh, Guo, and Muresan 2015)). As pre-processing, we removed the retweets, spam, duplicates, and tweets written in languages other than English. Also, we deleted all tweets where the hashtags of interest were not located at the very end (i.e., we eliminated “#sarcasm is something that I love”). We lowercased the tweets, except the words where all the characters are uppercased.

***Reddit*:** Khodak, Saunshi, and Vodrahalli (2018) introduced an extensive collection of sarcastic and non-sarcastic

Copyright © 2018, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

¹We treat irony and sarcasm similarly in this paper.

posts collected from different subreddits. In Reddit, authors mark their sarcastic intent of their posts by adding “/s” at the end of a post/comment. We collected 50K instances from the corpus for our experiments (denoted as *Reddit*), where the sarcastic and non-sarcastic replies are at least two sentences (i.e., we discard posts that are too short).

For brevity, we denote ironic utterances as *I* and non-ironic utterances as *NI*. Both *Twitter* and *Reddit* datasets are balanced between the *I* and *NI* classes. We use 80% of the datasets for training, 10% for development, and the remaining 10% for testing.

Irony Markers

Three types of markers — tropes, morpho-syntactic, and typographic are used as features.

Tropes:

Tropes are figurative use of expressions.

- **Metaphors** - Metaphors often facilitate ironic representation and are used as markers. We have drawn metaphors from different sources (e.g., 884 and 8,600 adjective/noun metaphors from (Tsvetkov et al. 2014) and (Gutiérrez et al. 2016), respectively, and used them as binary features. We also evaluate the metaphor detector (Rei et al. 2017) over *Twitter* and *Reddit* datasets. We considered metaphor candidates that have precision ≥ 0.75 (see Rei et al. (2017)).
- **Hyperbole** - Hyperboles or intensifiers are commonly used in irony because speakers frequently overstate the magnitude of a situation or event. We use terms that are denoted as “strong subjective” (positive/negative) from the MPQA corpus (Wilson, Wiebe, and Hoffmann 2005) as hyperboles. Apart from using hyperboles directly as the binary feature we also use their sentiment as features.
- **Rhetorical Questions** - Rhetorical Questions (for brevity *RQ*) have the structure of a question but are not typical information seeking questions. We follow the hypothesis introduced by Oraby et al. (2017) that questions that are in the middle of a comment are more likely to be RQ since questions followed by text cannot be typical information seeking questions. Presence of *RQ* is used as a binary feature.

Morpho-syntactic (MS) irony markers:

This type of markers appear at the morphologic and syntactic levels of an utterance.

- **Exclamation** - Exclamation marks emphasize a sense of surprise on the literal evaluation that is reversed in the ironic reading (Burgers 2010). We use two binary features, single or multiple uses of the marker.
- **Tag questions** - We built a list of tag questions (e.g., “didn’t you?”, “aren’t we?”) from a grammar site and use them as binary indicators.²

²<http://www.perfect-english-grammar.com/tag-questions.html>

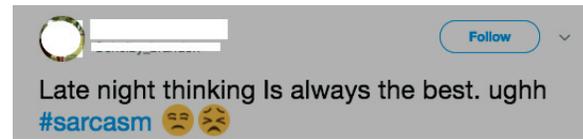


Figure 1: Utterance with emoji (best in color)

- **Interjections** - Interjections seem to undermine a literal evaluation and occur frequently in ironic utterances (e.g., “yeah”, “wow”, “yay”, “ouch” etc.). Similar to tag questions we assembled interjections (a total of 250) from different grammar sites.

Typographic irony markers:

- **Capitalization** - Users often capitalize words to represent their ironic use (e.g., the use of “GREAT”, “SO”, and “WONDERFUL” in the ironic tweet “GREAT i’m SO happy shattered phone on this WONDERFUL day!!!”).
- **Quotation mark** - Users regularly put quotation marks to stress the ironic meaning (e.g., “great” instead of GREAT in the above example).
- **Other punctuation marks** - Punctuation marks such as “?”, “:”, “;” and their various uses (e.g., single/multiple/mix of two different punctuations) are used as features.
- **Hashtag** - Particularly in *Twitter*, hashtags often represent the sentiment of the author. For example, in the ironic tweet “nice to wake up to cute text. #suck”, the hashtag “#suck” depicts the negative sentiment. We use binary sentiment feature (positive or negative) to identify the sentiment of the hashtag, while comparing against the MPQA sentiment lexicon. Often multiple words are combined in a hashtag without spacing (e.g., “fun” and “night” in #fun-night). We use an off-the-shelf tool to split words in such hashtags and then checked the sentiment of the words.³
- **Emoticon** - Emoticons are frequently used to emphasize the ironic intent of the user. In the example “I love the weather :) #irony”, the emoticon “:)” (wink) alerts the reader to a possible ironic interpretation of weather (i.e., bad weather). We collected a comprehensive list of emoticons (over one-hundred) from Wikipedia and also used standard regular expressions to identify emoticons in our datasets.⁴ Beside using the emoticons directly as binary features, we use their sentiment as features as well (e.g., “wink” is regarded as positive sentiment in MPQA).
- **Emoji** - Emojis are like emoticons, but they are actual pictures and recently have become very popular in social media. Figure 1 shows a tweet with two emojis (e.g., “unassumed” and “confounded” faces respectively) used as markers. We use an emoji library of 1,400 emojis to identify the particular emoji used in irony utterances and use them as binary indicators.⁵

³<https://github.com/matchado/HashTagSplitter>

⁴<http://sentiment.christopherpotts.net/code-data/>

⁵<https://github.com/vdurmont/emoji-java>

Features	Category	<i>P</i>	<i>R</i>	<i>F1</i>
all	<i>I</i>	66.93	77.32	71.75
	<i>NI</i>	73.13	61.78	66.97
- tropes	<i>I</i>	67.70	48.00	56.18
	<i>NI</i>	59.70	77.09	67.29
- MS	<i>I</i>	63.59	78.09	70.10
	<i>NI</i>	71.59	55.27	62.38
- typography	<i>I</i>	57.30	77.95	66.05
	<i>NI</i>	65.49	41.86	51.07

Table 2: Ablation Tests of irony markers for *Twitter*. **bold** are best scores (in %).

Features	Category	<i>P</i>	<i>R</i>	<i>F1</i>
all	<i>I</i>	73.16	48.52	58.35
	<i>NI</i>	61.49	82.20	70.35
- tropes	<i>I</i>	71.45	50.36	59.08
	<i>NI</i>	61.67	79.88	69.61
- MS	<i>I</i>	58.37	49.36	53.49
	<i>NI</i>	56.13	64.8	60.16
- typography	<i>I</i>	73.29	48.52	58.39
	<i>NI</i>	61.52	82.32	70.42

Table 3: Ablation Tests of irony markers for *Reddit* posts. **bold** are best scores (in %).

Classification Experiments and Results

We first conduct a binary classification task to decide whether an utterance (e.g., a tweet or a *Reddit* post) is ironic or non-ironic, exclusively based on the irony marker features. We use Support Vector Machines (SVM) classifier with linear kernel (Fan et al. 2008). Table 2 and Table 3 present the results of the ablation tests for *Twitter* and *Reddit*. We report Precision (*P*), Recall (*R*) and *F1* scores of both *I* and *NI* categories.

Table 2 shows that for ironic utterances in *Twitter*, removing tropes have the maximum negative effect on Recall, with a reduction on *F1* score by 15%. This is primarily due to the removal of hyperboles that frequently appear in ironic utterances in *Twitter*. Removing typographic markers (e.g., emojis, emoticons, etc.) have the maximum negative effect on the Precision for the irony *I* category, since particular emojis and emoticons appear regularly in ironic utterances (Table 4). For *Reddit*, Table 3 shows that removal of typographic markers such as emoticons does not affect the *F1* scores, whereas the removal of morpho-syntactic markers, e.g., tag questions, interjections have a negative effect on the *F1*.

Table 4 and Table 5 represent the *top* most discriminative features for both categories based on the feature weights learned during the SVM training for *Twitter* and *Reddit*, respectively. Table 4 shows that for *Twitter*, typographic features such as emojis and emoticons have the highest feature weights for both categories. Interestingly, we observe that for ironic tweets users often express negative sentiment directly via emojis (e.g., angry face, rage) whereas for non-ironic utterances, emojis with positive sentiments (e.g., hearts, wedding) are more familiar. For *Reddit* (Table 5), we observe that instead of emojis, other markers such as ex-

Category	Top features
<i>I</i>	emoticons: annoyed (";-"), perplexed (":-/"); emojis: angry face/monster, unamused, expressionless, confounded, rage, neutral_face, thumbsdown; negative_tag questions ("is n't it?", "don't they?")
<i>NI</i>	emojis: birthday, tophat, hearts, wedding, rose, ballot_box_with_check; quotations, hashtags (positive sentiment), emoticons: happy ("😊"), overjoyed ("😄")

Table 4: Irony markers based on feature weights for *Twitter*

Category	Top features
<i>I</i>	exclamation (single, multiple), negative_tag questions ("is n't it?", "don't they?"), interjections, presence of metaphors, positive sentiment hyperbolic words (e.g., "notably", "goodwill", "recommendation")
<i>NI</i>	negative sentiment hyperbolic words (e.g., "vile", "lowly", "fanatic"), emoticon: laugh ("😂"), positive_tag questions ("is it?", "are they?"), punctuations such as periods/multiple periods

Table 5: Irony markers based on feature weights for *Reddit*

clamation marks, negative tag questions, and metaphors are discriminatory markers for the irony category. In contrary, for the non-irony category, positive tag questions and negative sentiment hyperboles are influential features.

Frequency analysis of markers

We also investigate the occurrence of markers in the two platforms via frequency analysis (Table 7). We report the mean of occurrence per utterance and the standard deviation (SD) of each marker. Table 7 shows that markers such as hyperbole, punctuations, and interjections are popular in both platforms. Emojis and emoticons, although the two most popular markers in *Twitter* are almost unused in *Reddit*. Exclamations and *RQs* are more common in the *Reddit* corpus. Next, we combine each marker with the type they belong to (i.e., either trope, morpho-syntactic and typographic) and compare the means between each pair of types via independent t-tests. We found that the difference of means is significant ($p \leq 0.005$) for all pair of types across the two platforms.

Irony markers across topical subreddits

Finally, we collected another set of irony posts from (Khodak, Saunshi, and Vodrahalli 2018), but this time we collected posts from specific topical subreddits. We collected irony posts about politics (e.g., subreddits: politics, hillary, the_donald), sports (e.g., subreddits: nba, football, soccer), religion (e.g., subreddits: religion) and technology (e.g., subreddits: technology). Table 6 presents the mean and SD for each genre. We observe that users use tropes such as hyperbole and *RQ*, morpho-syntactic markers such as exclamation and interjections and multiple-punctuations more in politics and religion than in technology and sports. This is expected since subreddits regarding politics and religion are often more controversial than technology and sports and the users might want to stress that they are ironic or sarcastic using the markers.

Conclusion

We provided a thorough investigation of irony markers across two social media platforms: Twitter and Reddit. Classification experiments and frequency analysis suggest that

Type	Irony Markers		Genres		
	Marker	Technology (a)	Sports (b)	Politics (c)	Religion (d)
Trope	Metaphor	0.01 (0.06)	0.002 (0.05)	0.02 (0.12)	0.01 (0.10)
	Hyperbole	0.19 (0.39)	0.34 (0.48) ^{a**}	0.74 (0.44) ^{(a,b)**}	0.76 (0.43) ^{(a,b)**} , ^{c*}
	RQ	0.06 (0.23)	0.11 (0.32) ^{a**}	0.22 (0.41) ^{(a,b)**}	0.2 (0.4) ^{(a,b)**}
MS	Exclamation	0.09 (0.29)	0.14 (0.34) ^{a**}	0.42 (0.49) ^{(a,b)**}	0.37 (0.48) ^{(a,b,c)**}
	Tag Question	0.03 (0.16)	0.05 (0.23) ^{a**}	0.11 (0.32) ^{(a,b)**}	0.1 (0.30) ^{(a,b)**}
	Interjection	0.13 (0.34)	0.23 (0.42) ^{a**}	0.45 (0.50) ^{(a,b)**}	0.52 (0.5) ^{(a,b,c)**}
Typographic	Capitalization	0.04 (0.19)	0.08 (0.27) ^{a**}	0.20 (0.40) ^{(a,b)**}	0.1 (0.31) ^{(a,b,c)**}
	Punctuations	0.23 (0.42)	0.45 (0.50) ^{a**}	0.84 (0.36) ^{(a,b)**}	0.89 (0.31) ^{(a,b,c)**}

Table 6: Frequency of irony markers in different genres (subreddits). The mean and the SD (in bracket) are reported. ^{a**} and ^{x*} depict significance on $p \leq 0.005$ and $p \leq 0.05$, respectively.

Type	Irony Markers	Corpus	
		Twitter	Reddit
Trope	Metaphor	0.02 (0.16)	0.01 (0.08)
	Hyperbole	0.45 (0.50)	0.43 (0.50)
	RQ	0.01 (0.08)	0.15 (0.36)
MS	Exclamation	0.02 (0.16)	0.19 (0.39)
	Tag Question	0.02 (0.10)	0.08 (0.26)
	Interjection	0.22 (0.42)	0.32 (0.46)
Typographic	Capitalization	0.03 (0.16)	0.10 (0.30)
	Quotation	0.01 (0.01)	-
	Punctuations	0.10 (0.29)	0.47 (0.50)
	Hashtag	0.02 (0.14)	-
	Emoticon	0.03 (0.14)	0.001 (0.03)
	Emoji	0.05 (0.22)	-

Table 7: Frequency of irony markers in two platforms. The mean and the SD (in bracket) are reported.

typographic markers such as emojis and emoticons are most frequent for *Twitter* whereas tag questions, exclamation, metaphors are frequent for *Reddit*. We also provide an analysis across different topical subreddits. In future, we are planning to experiment with other markers (e.g., ironic echo, repetition, understatements).

References

- Attardo, S. 2000. Irony markers and functions: Towards a goal-oriented theory of irony and its processing. *Rask* 12(1):3–20.
- Bamman, D., and Smith, N. A. 2015. Contextualized sarcasm detection on twitter. In *ICWSM*.
- Burgers, C. F. 2010. *Verbal irony: Use and effects in written discourse*. [SI: sn].
- Fan, R.-E.; Chang, K.-W.; Hsieh, C.-J.; Wang, X.-R.; and Lin, C.-J. 2008. Liblinear: A library for large linear classification. *JMLR* 9(Aug):1871–1874.
- Ghosh, D.; Fabbri, A. R.; and Muresan, S. 2017. The role of conversation context for sarcasm detection in online interactions. *arXiv preprint arXiv:1707.06226*.
- Ghosh, D.; Guo, W.; and Muresan, S. 2015. Sarcastic or not: Word embeddings to predict the literal or sarcastic meaning of words. In *EMNLP*, 1003–1012.
- González-Ibáñez, R.; Muresan, S.; and Wacholder, N. 2011. Identifying sarcasm in twitter: A closer look. In *ACL*.
- Gutiérrez, E. D.; Shutova, E.; Marghetis, T.; and Bergen, B. 2016. Literal and metaphorical senses in compositional distributional semantic models. In *ACL (1)*.
- Joshi, A.; Bhattacharyya, P.; and Carman, M. J. 2016. Automatic sarcasm detection: A survey. *arXiv preprint arXiv:1602.03426*.
- Joshi, A.; Sharma, V.; and Bhattacharyya, P. 2015. Harnessing context incongruity for sarcasm detection. In *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics 2015 (Volume 2: Short Papers)*.
- Khodak, M.; Saunshi, N.; and Vodrahalli, K. 2018. A large self-annotated corpus for sarcasm. In *Proceedings of the 11th edition of the Language Resources and Evaluation Conference (LREC 2018)*.
- Muresan, S.; Gonzalez-Ibanez, R.; Ghosh, D.; and Wacholder, N. 2016. Identification of nonliteral language in social media: A case study on sarcasm. *JASIST*.
- Oraby, S.; Harrison, V.; Misra, A.; Riloff, E.; and Walker, M. 2017. Are you serious?: Rhetorical questions and sarcasm in social media dialog. *arXiv preprint arXiv:1709.05305*.
- Rei, M.; Bulat, L.; Kiela, D.; and Shutova, E. 2017. Grasping the finer point: A supervised similarity network for metaphor detection. *arXiv preprint arXiv:1709.00575*.
- Riloff, E.; Qadir, A.; Surve, P.; De Silva, L.; Gilbert, N.; and Huang, R. 2013. Sarcasm as contrast between a positive sentiment and negative situation. In *EMNLP*, 704–714.
- Tsvetkov, Y.; Boytsov, L.; Gershman, A.; Nyberg, E.; and Dyer, C. 2014. Metaphor detection with cross-lingual model transfer.
- Wilson, T.; Wiebe, J.; and Hoffmann, P. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT-EMNLP*, 347–354.