

Guidelines for Annotating Irony in Social Media Text

Version 1.1

LT3 Technical Report – LT3 16-01

Cynthia Van Hee, Els Lefever and Véronique Hoste

LT3 – Language and Translation Technology Team
Faculty of Translation Studies
Ghent University

All rights reserved.

LT3, Faculty of Arts, Humanities and Law, Ghent University, Belgium.

March 17, 2016

Contents

1	Introduction	1
2	Evaluative Expressions	3
3	Evaluation Polarity	4
4	Modifier Types	5
5	Irony Presence	6
6	Irony Harshness	9
7	Evaluation Target	10
8	Embedded Evaluations	12
9	Annotation Procedure	13
10	Brat Annotation Examples	14

Chapter 1

Introduction

The expansion of the Internet has allowed for Web 2.0 technologies like social media to be accessible to a vast amount of people. As a consequence, these Web-based technologies have become a valuable source of information about the public's opinion for interest groups like politicians, companies, researchers, trend watchers, and so on (Pang and Lee, 2008). What characterises this social media content (also referred to as *user-generated content*) is that it is often rich in **figurative language** like **irony** (Maynard and Greenwood, 2014). Modeling the underlying meaning of ironic utterances is interesting from psycholinguistic and cognitive perspectives. Moreover, it can benefit a wide range of natural language processing applications. The frequent use of irony on social media has important implications for tasks such as **sentiment analysis** or **opinion mining**, which aims to extract positive and negative opinions automatically from online texts. Since irony is commonly used in user-generated content, its presence can significantly undermine accurate analysis of opinions and sentiment in such texts (Reyes, Rosso, and Veale, 2013; Maynard and Greenwood, 2014). In order for sentiment analysis systems to perform well on ironic content, it is important to build computational models that are capable of recognising irony so one knows that the expressed sentiment is not to be understood in its literal sense. To this end, it is of key importance to understand how irony works and how it can be recognised in texts. These annotation guidelines are developed for **identifying specific aspects** and forms of irony that are susceptible to **computational analysis**.

When describing how irony works, theorists traditionally distinguish between situational irony and verbal irony. Situational irony is often referred to as situations that fail to meet some expectations (Lucariello, 1994; Shelley, 2001). Shelley (2001) illustrates this with firefighters who have a fire in their kitchen while they went out to answer a fire alarm. Verbal irony is often defined as expressions that convey an opposite meaning (Grice, 1975). There has been a large body of research in the past involving the definition of irony and the distinction between irony and sarcasm (Barbieri and Saggion (2014), Grice (1975), Kreuz and Glucksberg (1989), Wilson and Sperber (1992), amongst others). To date, however, experts do not formally agree on the distinction between irony and sarcasm. For this reason, our definition does not distinguish between both phenomena, but rather focusses on a specific linguistic form that can cover (most) cases described as verbal irony or sarcasm. We will refer to this form as *verbal irony*.

Traditionally, irony is defined as 'saying the opposite of what you mean' (e.g., Grice (1975), McQuarrie (1996), Quintilian (1959)). In accordance with this definition, we define irony as an *evaluative expression whose polarity (i.e., positive, negative) is changed between the literal and the intended evaluation, resulting in an incongruence between the literal evaluation and its context*. Our definition is similar to that of Burgers' (2010) in that verbal irony is defined as the expression of an evaluation whose polarity is changed between the literal and the intended meaning. It is different,

however, in that we do not restrict this polarity change to the use of opposition. Instead, we also consider expressions of an evaluation whose polarity is stronger (i.e., hyperbole) or less strong (i.e., understatement) than the intended evaluation polarity.

In these guidelines, no distinction is made between irony and sarcasm. However, it allows to signal variants of verbal irony that are particularly **harsh** (i.e., carrying a mocking or ridiculing tone with the intention to hurt someone). As a matter of fact, a number of previous studies found that harshness may be indicative for distinguishing between irony and sarcasm (Barbieri and Saggion, 2014; Lee, Christopher J. and Katz, Albert N., 1998).

In what follows we present the different steps for annotating irony in online texts. All annotations are done using the brat rapid annotation tool (Stenetorp et al., 2012). The example sentences in the following chapters are taken from a corpus of English Twitter messages. Note that not every element of these examples is annotated and discussed. We refer to Chapter 10 for a number of example sentences that are fully annotated.

Chapter 2

Evaluative Expressions

In this annotation scheme, irony is defined as the expression of an evaluation whose polarity is changed between the literal and the intended evaluation, resulting in an incongruence between the literal evaluation and its context. The evaluation polarity can be changed in three ways: by using **(i) opposition** (i.e., the literal evaluation is opposite to the intended evaluation, **(ii) hyperbole** (i.e., the literal evaluation is stronger than the intended evaluation) or **(iii) understatement** (i.e., the literal evaluation is less strong than the intended evaluation).

When annotating irony, annotators thus look for expressions of an **evaluation** in the text under investigation. By an evaluation, we understand the entire text span by which someone or something (e.g., a product, an event, an organisation) is evaluated, including possible modifiers (see Chapter 4). There are no restrictions as to which forms evaluations take: they can be verb phrases, predicative (adjective or nominal) expressions, and so on. Nevertheless, when possible, annotators should include the verb and its apposition in the annotated text span of the evaluation, as well as modifiers. Evaluative expressions can be found in sentences (1) to (4). As shown in example (4), one text can contain multiple evaluations.

(1) *Oh how I love working in Baltimore #not*
→ *oh how I love* = **evaluation**

(2) *What a shock. Duke Johnson is hurt in an important game. #sarcasm #canes*
→ *what a shock* = **evaluation**

(3) *So glad you'd rather read a book than acknowledge your own kid #not*
→ *so glad* = **evaluation**

(4) *Interesting visit with Terra Nova yesterday at Stoneleigh, class tent.*
→ *interesting* = **evaluation**
→ *class* = **evaluation**

Practical Remark

If an evaluative expression consists of several parts (e.g., *I love this band so much!*), these parts should be linked by means of drag and drop.

Chapter 3

Evaluation Polarity

An important subtask of annotating evaluations is **polarity classification**, which involves determining whether the expressed evaluation is **positive or negative**. Sometimes, it is not entirely clear whether a positive or a negative sentiment is expressed (due to a restricted context or ambiguity). Evaluations with an unknown polarity therefore receive the label ‘unknown’. It should be noted, however, that annotators should limit the use of the label ‘unknown’ and indicate a concrete polarity (i.e., positive or negative) when this can be inferred from the text.

(5) *I hate it when my mind keeps drifting to someone who no longer matters in life. #irony #dislike*

→ *hate* = **evaluation** [negative polarity]

→ *no longer matters in life* = **evaluation** [negative polarity]

→ *#dislike* = **evaluation** [negative polarity]

(6) *First day off for summer...kids wake up at 6:01. Love them but not Awesome. #sleepisfortheweak #not.*

→ *love* = **evaluation** [positive polarity]

→ *not Awesome* = **evaluation** [negative polarity]

→ *#sleepisfortheweak* = **evaluation** [negative polarity]

(7) *I'm surprised you haven't been recruited by some undercover agency. #sarcasm*

→ *'m surprised* = **evaluation** [unknown polarity]

As shown in the examples (4), (5) and (6), one text can contain multiple expressions of an evaluation. In Twitter data, hashtags can also contain evaluations (e.g., *#dislike* in sentence (5)). In this case, annotators should annotate the hashtag and indicate its polarity. Annotators should be careful to annotate a hashtag **as an entire unit**, even if it is a multiword expression (e.g., *#sleepisfortheweak*), and to include the hash sign (#) in the annotated text span. All hashtags that contain an evaluation should be annotated, except for the meta-hashtags *#sarcasm*, *#irony* and *#not*.

Chapter 4

Modifier Types

Sometimes, evaluative expressions are modified. This means that their prior polarity is changed by the presence of a particular element (i.e., a *modifier*) in the text. Modifiers are lexical items that cause a **shift in the prior polarity** of other nearby lexical items (Van de Kauter, Desmet, and Hoste, 2015). Two types of modifiers are distinguished in our annotation scheme: **(i) intensifiers**, which increase the intensity of the expressed sentiment and **(ii) diminishers**, which decrease the intensity of the expressed sentiment (Kennedy and Inkpen, 2006; Polanyi and Zaenen, 2004). Besides adverbs (e.g., *absolutely*), interjections (e.g., *Wow*) and punctuation marks (e.g., *!, ???!*) can modify the evaluation polarity. Modifiers can, but are not necessarily syntactically close to the text span of the evaluation. When possible, however, they should be included in the annotated text span of the evaluation. The modifiers in the sentences (8) and (9) are bold-faced. Modifiers should be annotated only if they alter a literally expressed evaluation.

- (8) *The **most** annoying kid lives next to my door!!!*
→ *most annoying* = **evaluation** [negative polarity]
→ *most* = **intensifier** of *annoying*
→ *!!!* = **intensifier** of *most annoying*

- (9) *Throwing up at 6:00 am is **always** fun #not*
→ *is always fun* = **evaluation** [positive polarity]
→ *always* = **intensifier** of *is fun*

Practical Remark

In brat, modifiers should be linked to the evaluative expression they alter by means of drag and drop.

Chapter 5

Irony Presence

According to our definition, irony arises from an evaluative expression whose polarity is changed between the literal and the intended evaluation, resulting in an incongruence between the literal evaluation and its context. More concretely, this incongruence results from a **contrast or clash** between the intended and the literal evaluation. As is explained in Chapter 2, this contrast (or *clash*) is realised by means of an opposition, a hyperbole or an understatement between two (or more) evaluations. It should be noted, however, that all of these evaluative expressions need not be explicit. Sometimes, a literally expressed evaluation is contrasted with another explicit evaluation (see example (21)). However, as irony tends to be realised implicitly (Burgers, 2010), the literally expressed evaluation is contrasted with an **implicit evaluation** (i.e., an element in the text whose polarity can be inferred by means of contextual clues and common sense or world knowledge). Chapter 7 elaborates on the annotation of evaluations.

Annotators should carefully analyse the evaluation(s) expressed in each text and define whether the text under investigation is ironic. When an instance is considered ironic according to our definition, annotators should indicate in what way the polarity is changed between the literal and the intended evaluation (i.e., by means of an opposition, a hyperbole or an understatement). It is possible, however, that an instance is ironic although no clash can be perceived. These instances should be labeled as **possibly ironic**. They will often contain descriptions of situational irony (see Chapter 1). Instances that are not ironic should be annotated likewise.

- **Ironic by means of an opposition:** the text expresses an evaluation whose literal polarity is the opposite of the intended polarity.
- **Ironic by means of a hyperbole:** the text expresses an evaluation whose literal polarity is stronger than the intended polarity.
- **Ironic by means of an understatement:** the text expresses an evaluation whose literal polarity is less strong than the intended polarity.
- **Possibly ironic:** there is no contrast between the literal and the intended evaluation, but the text is still ironic.
- **Not ironic:** the text is not ironic.

Sentences (9) and (10) are examples of an ironic message in which the literal evaluation is opposite to the intended evaluation. In most cases, the irony will be realised by means of an opposition. Examples of ironic hyperbole and ironic understatement are generally more difficult to perceive. In these cases, the distinction between ironic and not ironic is often fuzzy. Sentences (11) and (12) illustrate cases of irony caused by an understatement and a hyperbole, respectively.

Contextual clues and common sense or world knowledge are often sufficient to understand this clash between the intended and the literal evaluation, as shown in sentence (9). In very short texts (e.g., Twitter messages), however, this might only become clear when an explicit marker such as an irony-specific hashtag (e.g., #not, #sarcasm, #irony) is present, as shown in sentence (10). In this case, annotators should mark this by indicating '**hashtag indication needed**'.

(9) *Exams start tomorrow. Yay, can't wait!*

→ the message is **ironic by means of an opposition**: the polarity of the literally expressed evaluation (i.e., *Yay, can't wait!*) is positive, whereas the intended evaluation is negative (having exams is generally experienced as unpleasant, hence the implicit evaluation of *exams* is negative).

(10) *My little brother is absolutely awesome! #not.*

→ the message is **ironic by means of an opposition**: the polarity of the literal evaluation (*is absolutely awesome!*) is positive whereas the intended evaluation is negative.

In sentence (10), the contrast between the literal and the intended evaluation is only clear by the presence of the hashtag #not.

(11) *A+? So you did quite well..*

→ the message is **ironic by means of an understatement**: the polarity of the literally expressed evaluation *did quite well* is slightly positive, whereas the intended evaluation is strongly positive (given that A+ is the highest grade).

(12) *58 degrees and a few sunbeams breaking through the clouds. Now could the weather be any better for a picnic?*

→ the message is **ironic by means of a hyperbole**: the polarity of the literal evaluation (*could ... be any better*) is strongly positive whereas the intended evaluation is slightly positive (*58 degrees and a few sunbeams*).

Descriptions of ironic situations (i.e., situations that do not meet a certain norm or expectation) are examples of **situational irony** (Shelley, 2001) and should be annotated as **possibly ironic**. Consider for instance the sentences (13) to (15).

(13) *"@Buchinator.: Be sure you get in all those sunset instagrams before the sun explodes in 4.5 billion years." Look at your next tweet #irony*

(14) *Just saw a non-smoking sign in the lobby of a tobacco company #irony*

(15) *My little sister ran away from me throwing a water balloon at her and fell into the pool... #irony.*

Examples of **non-ironic** messages are presented in sentences (16) to (19). As non-ironic we consider instances that do not contain any indication of irony (sentence (16)) or instances where there is not sufficient context available to perceive the irony (sentence (17)). Additionally, this category encompasses tweets in which an irony-related hashtag (e.g., *#not*, *#sarcasm*) is used in a self-referential meta-sentence (sentence (18)), or functions as a negator (sentence (19)).

(16) *Drinking a cup of tea in the morning sun, lovely!*

(17) *@GulfNewsTabloid Wonder why she decided to cover her head though! #Irony*

(18) *Those that are #Not #BritishRoyalty should Not presume #Titles or do any #PublicDuties*

(19) *@TheSunNewspaper Missed off the #irony hashtag?*

Practical Remark

Whether an instance is ironic, possibly ironic or not ironic should be annotated on the dummy token ¶ preceding each text.

Chapter 6

Irony Harshness

Some previous studies claim that sarcasm can be considered to be a **sharp form of irony** that is meant to ridicule or hurt a specific target (Lee, Christopher J. and Katz, Albert N., 1998; Attardo, 2000; Barbieri and Saggion, 2014). To gain insight into this possibly distinctive feature, annotators should indicate the **harshness** of each ironic evaluation (i.e., whether the irony is used to ridicule or hurt a person/a company,...) on a **two-point scale (0-1)** where 0 means that the evaluation is not harsh and 1 that the evaluation is harsh. Additionally, a confidence score (i.e., low, medium, high) should be given for this annotation.

(20) Well this exam tomorrow is gonna be a bunch of laughs #not

→ the message is ironic by means of an opposition: the polarity of the literal evaluation is opposite to that of the intended evaluation

→ the ironic evaluation is **not harsh** (remember that a harshness score of 0 need not be annotated explicitly)

(21) Yeah you sure have great communication skills #not

→ the message is ironic by means of an opposition

→ the ironic evaluation is **harsh** (score 1), the evaluation is aimed at a person and is ridiculing

Practical Remark

For convenience and to speed up the annotation, a harshness score of 0 need not be annotated explicitly. When there is no harshness score indicated, the message is assumed to be not harsh.

Chapter 7

Evaluation Target

According to our definition, verbal irony arises from a contrast (or *clash*) between two evaluation polarities. As irony tends to be realised implicitly (Burgers, 2010), one of the opposite evaluations may be expressed in an **implicit** way. The evaluation has to be inferred from the context or by world knowledge/common sense (see examples (22), (23), and (24)) and is referred to as the **evaluation target**. The evaluation target is the textual element of which the implicit polarity **contrasts with the literally expressed evaluation**. The implicit polarity of an evaluation target can be **positive, negative, neutral or unknown**. The polarity of an evaluation target is neutral when it refers to another target (see sentence (25)). The evaluation target can, but is not necessarily the syntactic target of the expressed evaluation (see for instance sentence (24)). The evaluation target is often expressed by a complement to a verb phrase (i.e., verb + verb, verb + adverb, verb + noun), but it can also be a noun (e.g., *Christmas Day, school*). If two targets are connected by means of a conjunction (e.g., *and*), they should be annotated separately (see for example sentence (22)). In brat, the annotation of a target should always be linked to the corresponding evaluative expression. It cannot cross sentence boundaries.

(22) *I just love when it's extremely hot and I can't sleep #not*

→ *just love* = **evaluation** [positive polarity]

→ *just* = **intensifier** of *love*

→ *it's extremely hot* = **target** of *love*

→ *can't sleep* = **target** of *love*

(23) *@jessebatters @MarkDice already did this segment a while ago. your so clever and original! #not #foxsucks*

→ *so clever and original!* = **evaluation** [positive polarity]

→ *so* = **intensifier** of *clever and original*

→ *!* = **intensifier** of *so clever and original*

→ *#foxsucks* = **evaluation** [negative polarity]

→ In this sentence, there is a contrast between two literally expressed evaluations (*so clever and original* and *#foxsucks*).

(24) *I did so well on my history test that I got an F-!*

→ *did so well* = **evaluation** [positive polarity]

→ *F-* = **target** of *did so well*

In sentence (22), the syntactic target of the evaluation *did so well* is *I*, but *F-* is the target of which the polarity contrasts with the expressed evaluation.

(25) *Ahh 7 a.m. bedtimes, how I've missed you #not #examproblems*

→ *'ve missed* = **evaluation** [positive polarity]

→ *you* = **target** of *(I)'ve missed* that refers to the actual target **7 a.m. bedtimes**

As shown in sentence (25), an evaluation target sometimes refers to another target. In that case, they should be linked by means of a coreferential relation. Note that, in sentence the implicit polarity of the referring element (in this case *you*) is neutral, whereas the implicit polarity of the actual target is positive or negative (in this case, the implicit polarity of *7 a.m. bedtimes* is negative).

Brat howto

In brat, a coreferential relation between evaluation targets can be added by means of drag and drop. In the same way, evaluation targets can be linked to the corresponding evaluation.

Chapter 8

Embedded Evaluations

Sometimes an evaluation is contained by another evaluation or by the target of another evaluation (see sentences (26) and (27)). This is called an *embedded* evaluation and needs to be annotated. Similarly to evaluative expressions, the polarity of embedded evaluations can be **positive**, **negative** (or **unknown** in the case there is not sufficient context) and its prior polarity can be changed by a modifier. If the embedded evaluation is not ironic, it has no evaluation target.

(26) *The most hideous spider, that makes me feel sooo much better. #not*

→ *makes me feel sooo much better* = **evaluation** [positive polarity]

→ *sooo much* = **intensifier** of *makes me feel better*

→ *that* = **target** of *makes me feel sooo much better* that refers to the actual target **the most hideous spider**

→ *most hideous* = (embedded) **evaluation** of *spider* [negative polarity]

→ *most* = **intensifier** of *hideous*

(27) *I'm really looking forward to the awful weather that is coming this week-end.*

→ *'m really looking forward to* = **evaluation** [positive polarity]

→ *really* = **intensifier** of *looking forward to*

→ *the awful weather that is coming this week-end* = **target** of *really looking forward*

→ *awful* = (embedded) **evaluation** [negative polarity]

Chapter 9

Annotation Procedure


In what follows are described the different steps in the annotation procedure. Chapter 10 presents some annotation examples in brat. Note that, even if a message is not ironic or possibly ironic, annotators should indicate all evaluations that are expressed in the text under investigation.

1. **Based on the definition, indicate for each text whether it is ironic, possibly ironic or not ironic**
 - **Ironic by means of an opposition:** the text expresses an evaluation whose literal polarity is the **opposite** of the intended polarity.
 - **Ironic by means of a hyperbole:** the text expresses an evaluation whose literal polarity is **stronger** than the intended polarity.
 - **Ironic by means of an understatement:** the text expresses an evaluation whose literal polarity is **less strong** than the intended polarity.
 - **Possibly ironic:** there is no contrast between the literal and the intended evaluation, however, the text contains another form of irony (e.g., situational irony).
 - **Not ironic:** the text is not ironic.
2. **If the text is ironic:**
 - Indicate whether an **irony-related hashtag** (*#not, #sarcasm, #irony*) is required for recognising the irony (in the case of Twitter data).
 - Indicate the **harshness** of the irony on a two-point scale (0-1) and give a confidence score for this annotation.
3. **Annotate all evaluations contained by the text**
 - Indicate the polarity of the evaluation.
 - If present, annotate modifiers and link them to the corresponding evaluation.
 - If present, annotate the evaluation target(s) and link it/them to the evaluation with which it is in contrast.
 - * If the target refers to another target, link them by means of a coreferential relation.
 - * Indicate the implicit polarity of the target based on context, world knowledge or common sense.
4. **If present, annotate embedded evaluations.**
5. **Proceed with the next text.**

Chapter 10


Brat Annotation Examples

(28)



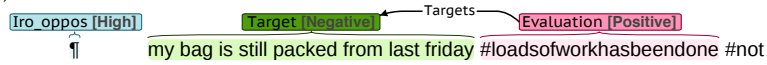
- the message is ironic by means of an opposition
→ the irony is not harsh
- *'s nothing better than* = **evaluation** [positive polarity]
- *there* = **refers to the target**
- *being squeezed between too many people on the train in the morning* = **target** of *nothing better than* → **implicit polarity** = negative

(29)



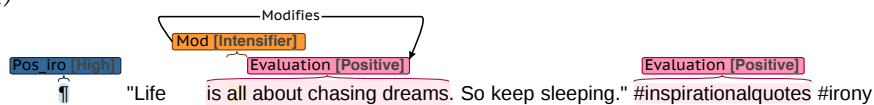
- the text is ironic by means of an opposition
→ the irony is not harsh
- *just love* = **evaluation** [positive polarity]
- *just* = **intensifier** of *love*
- *the best!* = **evaluation** [positive polarity]
- *!* = **intensifier** of *the best*
- *it* = **refers to the target**
- *being ignored* = **target** of *just love* and *the best!* → **implicit polarity** = negative

(30)



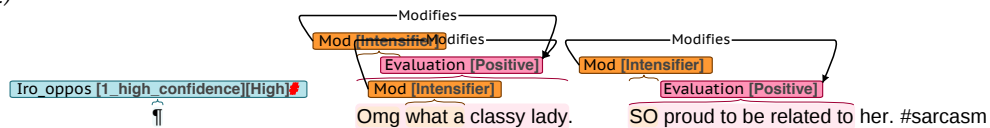
- the text is ironic by means of an opposition
→ the irony is not harsh
- #loadsofworkhasbeendone = **evaluation** [positive polarity]
- my bag is still packed from last friday = **target** of #loadsofworkhasbeendone
→ **implicit polarity** = negative

(31)



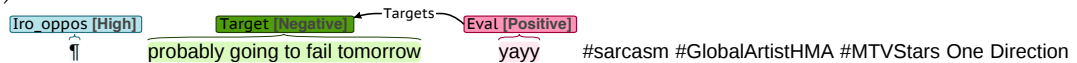
- the text is possibly ironic, it describes an ironic situation
- is all about chasing dreams = **evaluation** [positive polarity]
- all = **intensifier** of is about chasing dreams
- #inspirationalquotes = **evaluation** [positive polarity]

(32)




- the text is ironic by means of an opposition
→ the hashtag #sarcasm is required to perceive the contrast
→ the irony is harsh
- omg what a classy lady = **evaluation** [positive polarity]
- what a = **intensifier** of omg (...) a classy lady
- omg = **intensifier** of what a classy lady
- SO proud to be related to = **evaluation** [positive polarity]
- SO = **intensifier** of proud to be related to

(33)



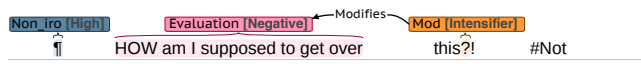
- the text is ironic by means of an opposition
→ the irony is not harsh
- yayy = **evaluation** [positive polarity]
- probably going to fail tomorrow = **target** of yayy
→ **implicit polarity** = negative

(34)



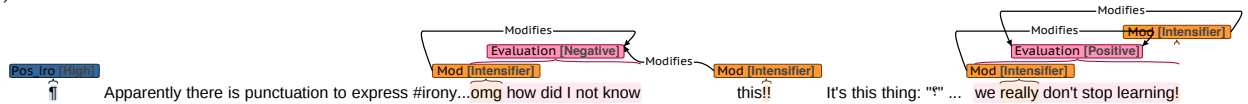
- the text is ironic by means of an opposition
→ the irony is harsh
- *love* = **evaluation** [positive polarity]
- *it* = **refers to the target**
→ **implicit polarity** = neutral
→ refers to the actual target *my roommate uses my textbook to roll joints on and now it wrecks* [implicit polarity = negative]

(35)



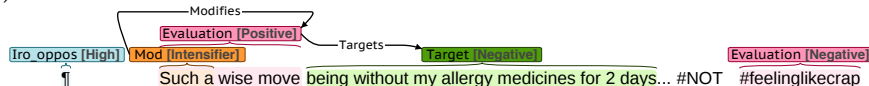
- the text is not ironic
- *HOW am I supposed to get over* = **evaluation** [negative polarity]
- *?! = intensifier* of *HOW am I supposed to get over*

(36)



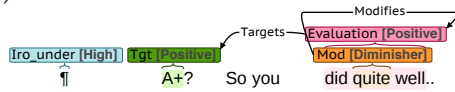
- the text is possibly ironic
- *omg how did I not know* = **evaluation** [negative polarity]
- *omg* = **intensifier** of *how did I not know*
- *!! = intensifier* of *omg how did I not know*
- *we really don't stop learning!* = **evaluation** [positive polarity]
- *really* = **intensifier** of *we don't stop learning!*
- *!* = **intensifier** of *we really don't stop learning*

(37)



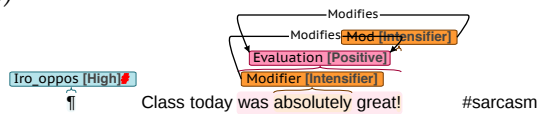
- the text is ironic by means of an opposition
→ the irony is not harsh
- *such a wise move* = **evaluation** [positive polarity]
- *such a* = **intensifier** of *wise move*
- *being without my allergy medicines for 2 days* = **target** of *such a wise move* [implicit polarity = negative]
- *#feelinglikecrap* = **evaluation** [negative polarity]

(38)



- the text is ironic by means of an understatement
→ the irony is not harsh
- *did quite well* = **evaluation** [positive polarity]
- *quite* = **diminisher** of *did quite well*
- *A+* = **target** of *did quite well* [implicit polarity = positive]

(39)



- the text is ironic by means of an opposition
→ the hashtag *#sarcasm* is required to perceive the contrast
→ the irony is not harsh
- *was absolutely great!* = **evaluation** [positive polarity]
- *absolutely* = **intensifier** of *was great!*
- *!* = **intensifier** of *was absolutely great*

References

- Attardo, Salvatore. 2000. Irony as relevant inappropriateness. *Journal of Pragmatics*, 32(6):793–826.
- Barbieri, Francesco and Horacio Saggion. 2014. Modelling Irony in Twitter. *Proceedings of the Student Research Workshop at the 14th Conference of the European Chapter of the Association for Computational Linguistics*, pages 56–64.
- Burgers, Christian. 2010. *Verbal irony: Use and effects in written discourse*. UB Nijmegen [Host].
- Grice, H. Paul. 1975. Logic and conversation. In Peter Cole and Jerry L. Morgan, editors, *Syntax and Semantics: Vol. 3: Speech Acts*. Academic Press, New York, pages 41–58.
- Kennedy, Alistair and Diana Inkpen. 2006. Sentiment Classification of Movie Reviews Using Contextual Valence Shifters. *Computational Intelligence*, 22(2):110–125.
- Kreuz, Roger J. and Sam Glucksberg. 1989. How to be sarcastic: The echoic reminder theory of verbal irony. *Journal of Experimental Psychology: General*, 118(4):374.
- Lee, Christopher J. and Katz, Albert N. 1998. The Differential Role of Ridicule in Sarcasm and Irony. *Metaphor and Symbol*, 13(1):1–15.
- Lucariello, Joan. 1994. Situational Irony: A Concept of Events Gone Awry. *Journal of Experimental Psychology: General*, 123(2):129–145.
- Maynard, Diana and Mark A. Greenwood. 2014. Who cares about Sarcastic Tweets? Investigating the Impact of Sarcasm on Sentiment Analysis. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation (LREC-2014), Reykjavik, Iceland, May 26-31, 2014.*, pages 4238–4243.
- McQuarrie, Edward F. and David Glen Mick. 1996. Figures of Rhetoric in Advertising Language. *The Journal of Consumer Research*, 22(4):424–438.
- Pang, Bo and Lillian Lee. 2008. Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*, 2(1-2):1–135, January.
- Polanyi, Livia and Annie Zaenen. 2004. Contextual Valence Shifters. In *Working Notes — Exploring Attitude and Affect in Text: Theories and Applications (AAAI Spring Symposium Series)*.
- Quintilian. 1959. *The Institutio Oratoria of Quintilian: With an English Translation by H. E. Butler, MA (Vol. 3)*. William Heinemann (Original work published around AD 95).
- Reyes, Antonio, Paolo Rosso, and Tony Veale. 2013. A Multidimensional Approach for Detecting Irony in Twitter. *Language Resources and Evaluation*, 47(1):239–268, March.
- Shelley, Cameron. 2001. The bicoherence theory of situational irony. *Cognitive Science*, 25(5):775–818.
- Stenetorp, Pontus, Sampo Pyysalo, Goran Topić, Tomoko Ohta, Sophia Ananiadou, and Jun'ichi Tsujii. 2012. BRAT: A Web-based Tool for NLP-assisted Text Annotation. In *Proceedings of the Demonstrations at the 13th Conference of the European Chapter of the Association for Computational Linguistics, EACL '12*, pages 102–107, Stroudsburg, PA, USA. Association for Computational Linguistics.
- Van de Kauter, Marjan, Bart Desmet, and Véronique Hoste. 2015. The Good, the Bad and the Implicit: A Comprehensive Approach to Annotating Explicit and Implicit Sentiment. *Language Resources and Evaluation*.
- Wilson, Deirde and Dan Sperber. 1992. On verbal irony. *Lingua*, 87(12):53 – 76.