

On being negative*

Maite Taboada

Department of Linguistics
Simon Fraser University
Canada
mtaboada@sfu.ca

Radoslava Trnavac

Department of Linguistics
Simon Fraser University
Canada
rtrnavac@sfu.ca

Cliff Goddard

School of Humanities, Languages
and Social Science
Griffith University
Australia
c.goddard@griffith.edu.au

Abstract

This paper investigates the pragmatic expressions of negative evaluation (negativity) in two corpora: (i) comments posted online in response to newspaper opinion articles; and (ii) online reviews of movies, books and consumer products. We propose a taxonomy of linguistic resources that are deployed in the expression of negativity, with two broad groups at the top level of the taxonomy: resources from the lexicogrammar or from discourse semantics. We propose that rhetorical figures can be considered part of the discourse semantic resources used in the expression of negativity. Using our taxonomy as starting point, we carry out a corpus analysis, and focus on three phenomena: adverb + adjective combinations; rhetorical questions; and rhetorical figures. Although the analysis in this paper is corpus-assisted rather than corpus-driven, the final goal of our research is to make it quantitative, in extracting patterns and resources that can be detected automatically.

Keywords: negativity, lexicogrammar, discourse semantics, rhetorical figures, corpus-assisted methodology

1. Introduction: Negativity and negative language

Our focus is the linguistic expression of negative evaluation or opinion. The general term ‘evaluation’, in this paper, refers to assessments of objects, situations, and other people, as well as the expression of emotional states. Evaluation has also been referred to as stance (Biber & Finegan, 1989), subjectivity (Wiebe et al., 2004), opinion (Pang & Lee, 2008) or appraisal (Martin & White, 2005).

In previous work, we have focused on the interpretation of evaluation in context (Taboada & Trnavac, 2013; Trnavac & Taboada, 2012), the nuances of evaluation conveyed by adjectives (Goddard et al., 2016a, 2016b), and how to extract evaluation automatically (Taboada et al., 2011). These extensive analyses of evaluative language led us to study the expression of negativity and negative evaluation, for several reasons. The first of those is the claim that humans have an inherent negative bias, that is, that

* To appear in *Corpus Pragmatics* (<http://www.springer.com/education+%26+language/linguistics/journal/41701>). Current version: February 21, 2017.

we perceive negative events and feelings as more salient and stronger. In language more specifically, previous research has shown that languages tend to contain more positive than negative words, and that negation and negativity are asymmetrical with respect to their positive counterparts. We expand on each of those aspects below.

1.1 Negativity bias

The psychological literature claims that we have a bias towards negative events and feelings, because we experience them more strongly than positive ones (Jing-Schmidt, 2007; Rozin et al., 2010; Rozin & Royzman, 2001). Rozin and Royzman (2001) describe negative events and feelings as having a 'contagion' effect. Their very vivid characterization of this phenomenon is the brief contact of a cockroach with a delicious meal, which will result in aversion. The reverse, the contact of one's favourite food with a plate of cockroaches does not render the cockroaches any more appetizing.

Negativity bias has several aspects, fundamental among them two: negative potency and negative dominance. Negative potency refers to strength of negative events and experiences: "given inverse negative and positive events of equal objective magnitude, the negative event is subjectively more potent and of higher salience than its positive counterpart" (Rozin & Royzman, 2001, p. 298). In behavioural sciences and economics, this is described as *loss aversion*, whereby a loss is perceived as stronger and more negative than the exact equivalent gain, i.e., the loss of \$100 is more painful than it is joyful to gain \$100 (Kahneman & Tversky, 1979).

Negative dominance captures the fact that "the holistic perception and appraisal of integrated negative and positive events [...] is more negative than the algebraic sum of the subjective values of those individual entities" (Rozin & Royzman, 2001, pp. 298-299). We will return to this when we discuss how algebraic sums or averages of positive and negative words in a text do not seem appropriate in the context of sentiment analysis.

This tendency of humans towards a negative bias has been explained in evolutionary terms: It is more beneficial to humans to pay attention to threatening negative events (e.g., a predator looming) than it is beneficial to attend to positive experiences (e.g., a big kill). Similar arguments have been made for the origin of disgust, as a mechanism to avoid potentially lethal pathogens, and its spread to moral disgust (Curtis, 2013).

The literary theorist and rhetorician Kenneth Burke postulated that negation is a fundamentally human trait, as part of our symbolic system and part of language.

"In an age when we are told, even in song, to "accentuate the positive," and when some experts in verbalization make big money writing inspirational works that praise "the power of positive thinking," the second clause of my definition [of man/humans] must take on the difficult and thankless task of celebrating that peculiarly human marvel, the negative." (Burke, 1966, p. 9)

1.2 Higher frequency of positivity

Languages seem to have more positive than negative words¹, and seem to use positive words with higher frequency. Boucher and Osgood (1969) sampled 13 languages and cultures, and found that evaluative positive words were used more frequently across those languages. This led them to propose the Pollyanna Hypothesis, a supposedly universal tendency. Garcia et al. (2012) also found confirmation of the hypothesis when analyzing emotion words across three languages: Words with a positive emotional content were found more often in Google's *N*-gram database.

In online review texts, which we have studied in previous work, and which have been extensively studied in the field of sentiment analysis (Pang & Lee, 2008), we find more positive than negative reviews, and more positive than negative words within the reviews. For instance, in the Internet Movie Database Corpus, a collection of 1.36 movie reviews (317 million words), 73% of the reviews are positive (Potts, 2011). Similarly, the Hedonometer project (Dodds et al., 2015) has shown a higher frequency of positive terms in Twitter. In other words, in the Twitter world, the world is mostly a happy place. Perhaps because there are more positive examples in available databases, sentiment analysis methods tend to perform better on positive sentences than on negative or neutral sentences (Ribeiro et al., 2016).

This higher frequency of positivity seems contradictory to the presence of a negativity bias. They can, however, be reconciled if we view them as complementary: We discuss positive events and feelings more often, because negative events and feelings are unpleasant and have a contagion effect. As Jing-Schmidt (2007, p. 424) puts it, we “tend to talk about the bright side of life, this is not the same as *looking on the bright side of life*”.

1.3 Asymmetrical nature of negation and negativity in language

In several aspects of language, positive and negative are asymmetrical and unequal (Horn, 1989; Israel, 2004). As discussed earlier, positive words have a higher frequency than negative words. At the morphological level, negative words are often derived from a positive root plus an affix (*unhappy*, *insincere*). It is rarely the case that a positive is derived from a negative in a similar way (*sad*—**unsad*), although exceptions exist (*unselfish*, *uncontaminated*) (Greenberg, 1966; Matlin & Stang, 1978). A negated positive word is negative (*unhappy*). A negated negative word is usually neutral: *unsad* is rather neutral, and certainly not the equivalent of *happy*.

In adjective combinations, positive adjectives are more likely to occur first in conjunctions, across languages (Rozin et al., 2010). The conjuncts *more or less*, *happy or sad*, *win or lose* are more frequent and less marked than the opposite order (*less or more*).

When examining negation at the clause level, i.e., the frequency of positive clauses as compared to clauses where the main verb is negated, Halliday and colleagues have found that positive is about ten times as frequent as negative, in a wide range of discourse types (Halliday & James, 1993; Halliday & Matthiessen, 2014; Matthiessen, 2006).

¹ This refers to positive and negative words in general. When it comes to words to describe emotions more specifically, numerous sources report a higher number of words to describe negative emotions (Izard, 1971; Rozin & Royzman, 2001).

At the pragmatic level, negation has a strong effect. The negation of a positive word is not the exact polar opposite. For instance, if *good* is mildly positive, *not good* is not necessarily mildly negative, but probably more strongly negative. In a study of movie reviews, Potts (2011) found that *good* is frequently found in mildly positive reviews, whereas *not good* (and similar expressions) is more frequently found in very negative reviews. Potts postulates an emergent expressivity for negation and negative polarity: the expressivity derives not from the lexical meaning of negation alone, but from pragmatic considerations and hearer expectations.

In summary, we see a puzzle to address: how negativity is negative, but not exactly the opposite or the complement of the positive. In particular, we are interested in the linguistic resources deployed to convey negativity. Because negative expressions seem to be less frequent and more marked, they are perhaps avoided when conveying negativity, and a range of other linguistic resources, apart from and beyond negative words seem to be employed.

Our goals for this paper are twofold. First, we propose a taxonomy of the range of linguistic resources that language deploys to convey negativity, with special focus on rhetorical figures and their integration with the lexicogrammar. Secondly, we present a corpus-based analysis of some of those resources. The corpora that we studied include two main sources: (i) reviews of movies, books and consumer products posted online; (ii) comments posted online in response to a newspaper opinion article. Our approach is firmly placed in linguistic analysis, but with a pragmatic focus: We want to explore the pragmatic effects of some of the lexicogrammatical resources and the rhetorical figures that we study.

2. Preliminary taxonomy: Linguistic resources for negative evaluation

The expression of evaluation in general draws from different resources in the language. As the Appraisal Framework has shown, evaluation and opinion are conveyed through lexis, grammar and discourse (Martin, 2014; Martin & White, 2005). We are particularly interested in the role of phraseology, which Hunston (2011) has compellingly described as playing an important role in evaluative language. In crafting a preliminary taxonomy, we found a recurring role of resources deployed for effect, and that led us to the inclusion of rhetorical figures as central devices in evaluative language. We believe rhetorical figures are particularly important in the expression of negative evaluation, the focus of our paper. In this section, we describe how we created the taxonomy, with the help of corpora, and we list some of the resources that we believe should be included in one such classification of evaluative language.

One such taxonomy, of what kinds of phrases, expressions and linguistic devices are used when conveying negativity, could have different starting points. We could create a formal taxonomy, by categorizing different parts of speech, different syntactic structures, constructions (Goldberg, 2006) or patterns (Hunston, 2011; Hunston & Francis, 2000). The taxonomy could also start from a semantic point of view, by classifying different types of negativity, as it is done in the Appraisal framework (Martin & White, 2005). We believe, however, that what we need is a dynamic taxonomy, one which contains different linguistic devices, but which places emphasis on how those resources are deployed for effect, and which focuses on the interaction and of the effect that the speaker/writer wants to achieve in the hearer/reader. In other words, we are in search of a pragmatic taxonomy.

We have found two broad groups where we can place the phenomena under study, i.e., the linguistic devices used in the expression of negative evaluation: lexicogrammar and discourse semantics. These are planes of expression in Systemic Functional Linguistics, SFL (Halliday, 1985; Halliday & Matthiessen, 2014). Lexicogrammar includes all the grammatical devices of the language (mood, transitivity, information structure), plus the lexis. Discourse semantics (sometimes simply *semantics*) in SFL deals with meanings created through text, captured in the systems of Transitivity, Mood and Theme, and by examining the cohesion and coherence of texts (Eggins, 1994; Halliday & Matthiessen, 2014; Martin, 1992).

Such a division of labour between lexicogrammar and discourse semantics is not particularly novel. What we contribute in this paper is the addition of another plane, an expansion of the purview of discourse semantics, in the form of rhetorical figures. Our view is that in the category of discourse semantics there is additional room to explore rhetoric aspects of how language is used. One of the phenomena we are interested in is the juxtaposition of a very positive intensifier (*hilariously*) with a negative adjective (*bad*). These seem to be more frequent in certain genres (see below), and are captured under the classical rhetoric trope *oxymoron*. Most SFL descriptions of discourse semantics do not discuss these expressive aspects, leaving evaluative aspect to the Appraisal framework (Martin & White, 2005). We view lexicogrammar as the set of lexical, phraseological and grammatical resources of the language. The other plane of expression contains discourse semantics, but also includes the use of lexicogrammatical resources for effect and in the course of interaction. This is captured in Figure 1, with the two aspects interacting with each other.

The dichotomy is, fundamentally, one of structure (static) vs. rhetoric (dynamic). As with all classifications, a clear-cut separation between the two categories is not possible; discourse semantics and rhetorical figures are necessarily encoded in the lexicogrammar. These are not mutually exclusive classes, but rather categories that feed each other. The 'Rhetoric' category in Figure 1 does not comprise linguistic resources *per se*. It instead encompasses the rhetoric deployment of those resources for a particular effect.

We should point out that SFL postulates another plane, or component of language, below the lexicogrammar, to include phonology and graphology. The phonological level is not included in this taxonomy, because we are primarily dealing with text, but some phonological devices, such as rhyme and alliteration, are included in the Rhetoric level. Graphological aspects, in particular emoticons, all capitals and multiple exclamation marks, also contribute to conveying negativity. They are beyond the scope of this paper.

Our use of Rhetoric as a category draws from classical rhetoric, and in particular the set of rhetorical figures or tropes described in rhetoric as having a particular effect. Glossing over 2,500 years of scholarship, we can broadly characterize rhetoric as the study of how language is used for persuasion and for effect (Borchers, 2006; Keith & Lundberg, 2008). Rhetoric is important, because it involves an audience. Evaluative language, and especially negative evaluation, is clearly designed with audience effects in mind.

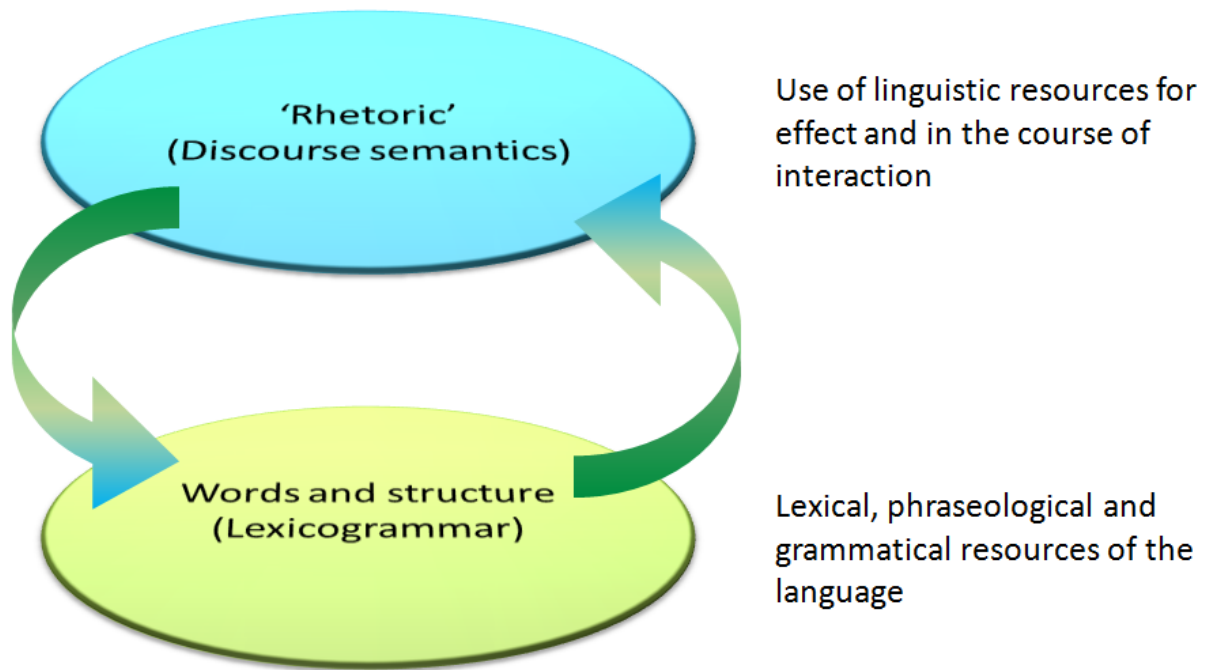


Figure 1. High-level taxonomy of resources involved in negative evaluation

The study of language from a rhetorical perspective comprises two main aspects: the structure of arguments, and the study of figures and tropes. The study of arguments refers to arrangement of parts of a discourse, the logical connections between those parts and how successful arguments are constructed. In classical terms, arrangement is one of the five canons of rhetoric (invention, arrangement, style, memory and delivery). In modern times, this is part of theories of argumentation and argumentation schemes (Toulmin, 1958; Walton et al., 2008), the New Rhetoric (Perelman & Olbrechts-Tyteca, 1969) and the emerging computational field of argumentation mining (Mochales & Moens, 2011).

Rhetorical Structure Theory (Mann & Thompson, 1988), as its name indicates, deals with rhetorical effects, mostly of the first type: How arguments are built from component parts, with the intention of creating an effect in the reader or hearer. RST deals with how propositions are combined in discourse, and it accounts for coherence through propositional relations. It is certainly part of the taxonomy, at the 'rhetoric' or discourse semantics level. In this paper, however, we focus on the second type of rhetorical devices, the use of linguistic resources to convey negative evaluation, classified as figures and tropes.

Rhetorical figures are creative arrangements of words to create an effect. Tropes are creative uses of word meaning, used for effect as well, but also with style in mind (Keith & Lundberg, 2008). Examples of such resources or rhetorical figures are exaggeration (hyperbole) or sarcasm. Rhetorical figures are deployed for the expression of negative evaluation, in what we believe are predictable and identifiable ways.

After this description of the top levels of our taxonomy, we list below some of the resources that play a role in the expression of negative evaluation, divided into phenomena in the lexicogrammar, and

phenomena that belong in discourse semantics and rhetoric. In most cases, examples are from our corpora (see Section 3), either verbatim or adapted for readability and simplicity.

2.1 Phenomena in the lexicogrammar

2.1.1 Morphology

English does not have a rich morphological system in general, and, as a consequence, it does not count on morphology to convey evaluation. A few morphemes, however, do perform that role. One such example is *-let*, which can be used as a diminutive morpheme, but that in Ginsberg's (2011) coinage of the word *deanlet* to describe middle management at universities (Dean, Associate Dean, Deanlet) is clearly negative. Other such examples are the suffixes *-ist* or *-ism(s)*. *Chomskyism* is clearly negative, versus the more neutral *Chomskyan*. Example (1) is taken from our comments corpus. The author suggests that some comment they had or could have said would result in being labelled an *-ist* of some type, which in the author's mind would be a negative type of *-ist*.

(1) ... that would immediately brand me as an "(fill in the blank)ist"....

2.1.2 Lexis

Individual words have clear evaluative content, and many can be easily classified as positive or negative. That category of evaluative lexis is easily classified, and has been made into lists of words used in sentiment analysis and emotion detection. Additionally, we have found in our corpora certain words that do not seem to carry intrinsic polarity, but nevertheless tend to be used in negative contexts. Mentions of *evidence* or *rational* trigger negative readings of the content. Similarly, some adverbs, such as *actually* and *exactly*, seem to be used mostly when negative evaluation is intended. Example (2) contains an example of both *evidence*, as an appeal to another commenter to back their claims, and *actually*, which seems to question the veracity of the interlocutor's assertions.

(2) Is there any evidence that her server was actually hacked?

2.1.3 Person reference

The basic distinction between first and second pronouns, on the one hand, and third person pronouns, on the other hand, is relevant in the expression of negative evaluation. Direct attacks use second person pronouns, and sometimes vocatives, especially in comments, where the user's name of the person being addressed is employed as a vocative, as in (3), where the user name "bagelboy" is exploited by the user "Be responsible for YOURSELF" to mock.

(3) bagelboy: This is a start. I've always wondered why Muslims don't silence their nut-bar extremists. It would serve their own best interests - trust me.

Be responsible for YOURSELF: bagelboy, double toast mine and top it with garlic cream cheese please.

Third person reference also plays a role in mocking or derogatory terms. Consider the use of *these people* to refer to the Clintons.

(4) These "petty scandals" at a minimum should have put the Clantons [sic] in jail and at the other end of the scale probably endangered peoples lives...I will take the Donald any day over these people.

2.1.4 Adverb + adjective combinations

Adjectives tend to carry a large proportion of the evaluative load in language, and adverbs can serve to graduate their meaning, intensifying or downtoning it (*very good, somewhat good*). We have found a class of adverb+adjective combinations that seem to have a very specific evaluative load. These are adverbs which in English end in *-ly* and have more semantic content than the usual intensifiers (*very, really, so*). Examples are *hilariously, ridiculously* or *wonderfully*. These often combine with positive adjectives to intensify the positive meaning. In many cases, however, there seems to be a clash between the polarity of the modifying adverb and the adjective, as in *hilariously bad*. Such expressions seem to have a non-compositional meaning, a meaning that is more than the sum or the average of the individual polarities of the words. Consider *ridiculously complex* and *ridiculously simple*. While the first one is definitely negative, and its polarity could be inferred from the sum of its parts, the second one is more nuanced. Something *ridiculous* is often negative, but *simple* can be either positive or negative. The “Made Ridiculously Simple” line of medical books includes titles such as *Clinical Anatomy Made Ridiculously Simple* or *Ophthalmology Made Ridiculously Simple*². The titles, and the series, clearly intend for *ridiculously simple* to be a good thing. In a movie review, on the other hand, *ridiculously simple plot* will probably convey negative evaluation.

These combinations of adverb+ adjective with opposite polarities can be characterized as examples of oxymoron, a rhetorical figure. We discuss examples in Section 4.

2.1.5 Temporal adverbials

Some temporal adverbials are used negatively, to express exasperation or impatience. Examples are *at last* or *eventually* (*At last, somebody speaks the truth.*)

2.1.6 Semantic prosody, discourse prosody

Semantic prosody refers to the intrinsic positive or negative connotations that a word carries, by virtue of its association with overtly positive or negative items (Hunston, 2011; Louw, 1993). For instance, the subjects of the phrasal verb *set in* are usually unpleasant or undesirable (*reality, smog, cold weather*). Words and expressions with negative semantic prosody also include *bent on, quick to, bordering on* or *symptomatic of*.

2.1.7 Phraseology

Certain combinations of words, phrases and expressions tend to be used either with higher frequency in negative evaluation (thus becoming ‘lexical bundles’), or to be fixed in their structure or meaning (i.e., phrases). As with all other phenomena, we concentrate on those that seem to be used for negative evaluation. We have found instances of idioms used flexibly, that is, used by appealing to a fixed expression, but in slightly different formulations. An example is Example (5), which reuses the expression *muddy the waters*.

(5) The waters get muddy when the drug trade becomes a third party.

² <http://www.medmaster.net/>

Lexical bundles are usually defined with respect to their frequency in a particular register. They are phraseological units that occur with a high frequency in a register (Biber & Barbieri, 2007; Biber et al., 1999). In our analyses, we used AntConc to find four-gram sequences that appear at least twice in each of the corpora (we set a low limit of two, because our corpora are small in size). We found that there were not enough clear lexical bundles in the data. We did find, however, a few expressions of variable size that seem to be used frequently to convey opinion, and in particular negative opinion: *nothing but*, *nothing short of*, *nothing more than*, *anything but*. These all contain instances of negative words. Other expressions that seem relevant include: *leave little doubt*, *believe it or not*, and *yet, so much for*.

More revealing were lexicalized sentence stems (Hunston, 2011) that point to a negative opinion in the clause: *The trouble/issue with X*, *It's hard to believe that X*. Questions in particular seem to make frequent use of lexicalized stems: *Isn't it time that X*, *Why would you X*, *Why wouldn't you X*, *Do you know nothing about X*, *Don't you know that X*.

2.2 Phenomena within discourse semantics

The phenomena and resources described in the previous section can be described as belonging to the lexicogrammar, to the area of language within the sentence. In this section, we review resources that take place within the discourse and pragmatic realms, i.e., that either take place above the sentence level, or that are interpretable only when context is taken into account. This includes rhetorical resources, because their reliance on the effect being created presupposes an interlocutor, and thus a context.

2.2.1 Direct and indirect speech

Quoting, whether an authority, a portion of an article under discussion, or the words of other online commenters, seems to often have a negative effect. In particular in online comments, commenters quote the article or other comments in order to highlight that they said something wrong. Graphically, there is also some overlap with the use of quotes as a distancing mechanism, a device to signal that the wording or the concept is not the writer's own, the same effect that is achieved with the expression *so called*, as in *the so called liberal elites (the "liberal" elites, if quotes are used)*.

As an example of quoting for negative effect, in this case to contradict, is in (6). The author quotes from the article, and then questions the basis in truth of the statement.

(6) "we have seen an outpouring of public support." Not sure where the author is from but I live in Canada and I certainly haven't seen much support for more refugees.

2.2.2 Rhetorical questions

Rhetorical questions affirm or deny a point strongly by asking it as a question. The rhetorical question includes an emotional dimension and is usually defined as any question asked for a purpose other than to obtain the information the question asks. In negative contexts they usually start with lexicalized sentence stems.

(7) Why are you so mean?

(8) Can that really be boiled down to a "civil war?"

2.2.3 Rhetorical figures

As mentioned in Section 2, in the broad division of rhetoric into arrangement of ideas versus use of figures for effect, we concentrate on the latter. By rhetorical figures we mean the creative use of language for effect, whether in the arrangement of the words (figures), or in the manipulation of meanings (tropes).

We have identified, from the vast literature on rhetorical figures and tropes, a few that we believe play an important role in negative evaluation, such as antithesis and hyperbole, which we describe in Section 4.3. They include: metaphors, hyperbole, euphemism and litotes, sarcasm and irony, puns, antithesis, repetition and parallelism, and ellipsis.

Our definitions of rhetorical figures are culled from different sources, mostly Keith and Lundberg (2008) and Gideon Burton's comprehensive website³.

We summarize In Figure 2 the resources discussed in this section. These are just a sampling of the many lexicogrammar and discourse resources that we believe play a role in the expression of (negative) evaluation.

Words and structure: Resources in the lexicogrammar	Rhetoric: Resources in discourse semantics
<ul style="list-style-type: none">• Morphology• Person reference• Lexis<ul style="list-style-type: none">• Evaluative words• Temporal adverbials• Semantic & discourse prosody• Phraseology<ul style="list-style-type: none">• Lexical bundles• Sentence stems• Adverb+adjective combinations	<ul style="list-style-type: none">• Direct and indirect speech• Rhetorical questions• Rhetorical figures<ul style="list-style-type: none">• Metaphor• Hyperbole• Euphemism & litotes• Sarcasm & irony• Puns• Antithesis• Repetition & parallelism• Ellipsis

Figure 2. Summary, resources for (negative) evaluation

³ <http://rhetoric.byu.edu/>

3. Corpora

The taxonomy presented in the previous section was guided by our qualitative analysis of two main corpora. In the remainder of the paper, we concentrate on more quantitative studies of specific phenomena. We first describe the two corpora.

3.1 The SFU Review Corpus

The SFU Review Corpus (Taboada, 2008) is a collection of 400 texts, reviews posted online on the website Epinions.com. It contains reviews of eight different types of works of art (books, movies, music), consumer products (cars, computers, cookware, phones), or tourist destinations (hotels). In each of those categories, equal numbers of positive and negative reviews were collected (25 each per category). Positive and negative reviews were labelled as such based on the number of stars assigned by the reviewer: Reviews of 1 or 2 stars were negative; reviews of 4 or 5 stars were positive. No neutral reviews (i.e., 3 stars) were collected. We show words and sentence counts in Table 1, also broken down into positive and negative reviews.

Total number of words	283,508
Number of sentences	16,736
Words in positive reviews	164,715
Words in negative reviews	118,793

Table 1. Review corpus statistics

Table 2 shows raw frequency (n) of positive and negative expressions in positive and negative reviews. Positive and negative expressions were extracted with SO-CAL (Taboada et al., 2011), a sentiment analysis system. They include individual words (*interesting*), and such words in the context of intensifiers (*very interesting*) and negation (*not interesting*).

The normalized frequency (per 1,000 words) shows that positive and negative words have similar frequency in negative reviews, but that positive words are much more frequent in positive reviews. That is, what distinguishes positive from negative reviews is that positive reviews contain much higher numbers of positive expressions. Negative reviews are not disproportionately negative, however. They are simply more balanced.

	Positive reviews		Negative reviews	
	n	Per 1,000 words	n	Per 1,000 words
Positive expressions	7,955	48.30	4,294	36.15
Negative expressions	3,849	23.37	4,378	36.85

Table 2. Frequency and normalized frequency of positive and negative expressions in reviews

3.2 Globe and Mail corpus

The second corpus is a collection of comments posted in response to editorials on the website of the Canadian newspaper The Globe and Mail. We have a standing interest in online comments, because they seem to be particularly aggressive, and perceived to be so. Media organizations devote a large number of resources to managing their social media presence, and their interactions with authors of comments. Abuse, of other commenters, of authors of articles, and of particular groups, seems widespread. As an example, the British newspaper The Guardian conducted an analysis of over 70m comments posted to their website in the last 10 years, in particular of those comments that had been blocked by moderators. They found that some topics were more controversial than others. They also found that some authors (of the original articles) seemed to attract a higher level of abuse: Of the top 10 more abused authors, eight were female, and the other two were black men (Gardiner et al., 2016).

With this background in mind, we set out to analyze a set of comments posted to the Globe and Mail in August and September 2016. They are comments in response to 10 different editorial pieces (columns and op-eds, i.e., pieces with a by-line). We do not have yet a classification of comments into positive or negative, so our distribution (Table 3) is for the entire content of the comments, regardless of their polarity. What we find is that, just like the reviews, they contain similar frequencies of positive and negative words and expressions. Further analysis would show whether there is an imbalance once the polarity of the comment is taken into account, as there is for the reviews.

	<i>n</i>	Per 1,000 words
Positive expressions	763	33.24
Negative expressions	796	34.68
Total number of words	22,955	
Number of sentences	1,473	

Table 3. Frequency and normalized frequency of positive and negative expressions in comments

4. Detailed analyses

4.1 Adverb adjective combinations

Evaluative language is generally considered to reside, to a very large extent, in adjectives (Pang & Lee, 2008). Similarly, it is well known that adjectives may be graduated (Martin & White, 2005), that is, they can be intensified and downtoned through different linguistic devices, such as diminutive morphemes in languages that employ those, or degree words (*very*, *somewhat*).

We are interested in graduation through adverbs which are not prototypical degree words. These can be described as content adverbs, and typically have an *-ly* suffix in English. Examples of these content adverbs are *hilariously*, *ridiculously* or *spectacularly*.

Our interest arose out of observations that some adverb-adjective combinations present special challenges when performing automatic sentiment analysis. We found that expressions such as *hilariously*

bad did not seem to have a clear compositional analysis. To illustrate this point, consider the evaluative load given to individual words in lexicon-based sentiment analysis systems (Liu, 2015; Pang & Lee, 2008). In one of them, SO-CAL (Taboada et al., 2011), *hilariously* is included in the dictionary as a positive word, with a score of 4 (the dictionaries include words in a range between +5 and -5). The adjective *bad* has a dictionary score of -3. If we perform arithmetic operations on them, we can obtain a sum of +1, or an average of +0.5. Neither value seems to capture well the complex meaning conveyed by the combination of a negative adjective and a positive modifying adverb.

We then extracted all the “Adverbly Adjective” combinations in our respective corpora. We found the following.

Reviews		Comments	
Tokens	Types	Tokens	Types
782	684	60	56

Table 4. Tokens and types for "Adverbly Adjective" combinations

By contrast, a similar search in the Corpus of Contemporary American English, COCA (Davies, 2008-) yields a different type of adverb-adjective combination. The first 1,000 results (types) have a token frequency of 276,065, i.e., a much lower type-token ratio, indicating that the combination is not as creative, and the same types are more frequently re-used. More interestingly, the double evaluation (evaluation both from the adverb and the adjective) seems to be less common. Typical modifying adverbs are *absolutely*, *really* or *highly*, which are closer to intensifiers. Another category of adverbs can be considered descriptive rather than evaluative: *genetically*, *politically* or *statistically*. Therefore, we postulate that the oxymoronic “Adverbly adjective” combination may be characteristic of these evaluative genres, and plan to study the phenomenon further.

Liebrecht et al. (2017) have also found that there is a difference in intensity differential when adjectives are modified by either a pure intensifier, such as *very*, or by a content adverb, such as *extraordinarily*. Their study (with Dutch texts) shows that pure intensifiers have a different contribution depending on whether the word they modify is positive or negative: Positive words receive higher intensification than negative words (*very good* is more positively intensified than *very bad* is negatively intensified). When content adverbs are used, however, there is no difference in the level of intensification.

4.2 Questions

Questions are often used for rhetorical purposes, and seem to be a tool of choice in particular when the author attacks an interlocutor, real or abstract. Real interlocutors may be other participants in the comments thread, often addressed by name (or screen name). Abstract interlocutors are the entire audience who may be reading the comment or review, and can be seen as rhetorical questions, another rhetorical figure.

We extracted all the questions in the two corpora, focusing on root questions (i.e., no embedded questions, as in *The question is why did it take Obama 3 years to show his birth certificate*).

We found 334 questions in the reviews, and 66 in the comments. Although the numbers are small, questions seem to be a more prominent feature of the comments (4.5% of the sentences) than for the reviews (2%).

Questions do seem to perform a rhetorical function. Examples from the review corpus (9) and the comments corpus (10), respectively, show some common lexical patterns and sentence stems deployed to convey evaluation.

- (9) a. Can you believe a computer company is still installing 250W power supplies in their computers?
b. Does this make it a genuine bargain among sports sedans?
c. What does his appearance have to do with The Last Samurai?
d. Why would you buy that freaky looking computer?
- (10) a. Do you have evidence that supports your claim that most patients are happy to participate?
b. Does the Globe know nothing about who this fellow is?
c. What good is it to "own the podium" when fewer and fewer young people are participating?
d. Why should someone in Toronto care what a hack from St. Louis thinks about anything.

The sample questions in the examples above show a use of yes-no interrogative forms to engage the reader (*Can you believe...*) and to engage a particular reader (*Do you have evidence...*). Wh-interrogatives perform similar functions, but in particular *why* questions seem to be particularly negative, questioning some underlying assumption, or pointing out that there are no possible good answers to the question, i.e., that no good reasons can be provided as to the *why* being posed. For instance, in (9d), the author is trying to persuade the reader (and the commenter from St. Louis) that there are no possible good reasons why somebody in Toronto could care about the opinion of somebody from St. Louis.

Our analysis of questions ties in with the analysis of lexicalized sentence stems, as many questions have stock beginnings:

- (11) But why....
 Couldn't PRONOUN (Couldn't they... Couldn't you...)
 How come...
 How ADJECTIVE is PRONOUN (How perfect is that? How pathetic is it?)
 So what... (So what went wrong? So what's good about it?)
 Why BARE-INFINITIVE (Why buy junk?)
 Why should/would... (Why should the American people trust Hillary Clinton?)
 Why won't...

4.3 Rhetorical figures

In our corpus, there are two salient categories of rhetorical figures—tropes and schemes⁴. Tropes represent an “artful deviation” from the ordinary or principal signification of a word, while schemes present an artful deviation from the ordinary arrangement of words. Among the most prominent tropes in the newspaper comments we find the use of metaphors, in most cases comparing a person, a group or an idea to an undesirable source.

- (12) This is not well-known because the fundamentalists have the upper-hand (control the message) and have the ear of the media. They (the fundamentalists) have become the tree which hides the forest.
- (13) Under the trickled down upon theory, a rising tide lifts all yachts. We are all the guinea pigs of medical science.

Another salient category is hyperbole, a rhetorical exaggeration that is often accomplished via comparisons, similes, and metaphors. In (14), we see a hyperbolic description of ‘losers’ in sports as those who finish in 64th place (and yet receive funding). In (15), the hyperbole is accentuated by the series of descriptors (*intolerant, organized, structured, financed*), and the repetition of intensifiers (*most, best*).

- (14) Here's yet another leftist bleeding-heart that would rather we go back to the years of giving money equally to 64th place losers for whom getting a “personal best” is seen as an accomplishment.
- (15) I only got to know what fundamentalist Islam was when I came to Canada. It is here that I got to know the most intolerant, the best organized, the most structured and the best financed groups, with means and worldwide connections. It was quite a shock.

Negative evaluations to a considerable level are expressed with indirect figures of speech. Among those are figures that substitute a more favourable expression for a pejorative or socially delicate term. These are referred to as euphemisms and litotes, that is, deliberate understatements. Example (16) illustrates the use of litotes:

- (16) Bernard-Henri Lévy has influence in most democracies of Europe. He helped French government during Libya revolution anti Kadafy. Now he would be more fairplay if he recognized right to live to Palestinians of Gaza strip and Cisjordania.

Sarcasm and irony are tropes that involve changes in meaning, by extending, turning or alluding to a meaning that needs to be re-interpreted in the new context. There is typically a disconnect between what is said and how it is meant to be interpreted. Sarcasm involves four different components (Goddard, 2016): (i) counter-expectation, (ii) exaggeration, (iii) intention to be funny, (iv) aggressiveness.

⁴ All the definitions of rhetorical figures are cited from <http://rhetoric.byu.edu>.

The definitions of sarcasm and irony are very close, with irony being defined as “speaking in such a way as to imply the contrary of what one says, often for the purpose of derision, mockery, or jest”, while sarcasm implies “the use of mockery, verbal taunts, or bitter irony”. For that reason, we choose to group them together—although there are clear differences, in particular in the more contextual nature of verbal irony (Alba-Juez & Attardo, 2014). Example (17) demonstrates the use of sarcasm in the online comments:

(17) Christians don't silence their nutbar extremists either. It would serve our own best interests - trust me. Like always, trees are not the problem, the people are.

Less prominent but equally creative are wordplays and puns that are occasionally found in our data, such as *trumped*, *trump cards*, (referring to then presidential candidate Donald Trump), as well as oxymoron, in which two opposite ideas are joined to create an effect (*open secret*, *seriously funny*, *original copies*).

The group of schemes is predominantly expressed with antithesis, juxtaposition of contrasting words or ideas (*often*, *although not always*, found in parallel structures) in online comments, as in Examples (18) and (19).

(18) We are caught in war, wanting peace.

(19) it's a medal or nothing.

Repetitions do not occur only in antithesis. Different types of linguistic units may be repeated, from sounds (alliteration), to words and grammatical structures. In parallelism, a frequent repetition is of the same beginning for a sentence, with different endings, as in Example (20). Note that the example also showcases the use of sarcasm in the expression *surge on*, a word play with the concept of surge in war, and probably echoing the expression *game on*. In this case, as in many others, effective negativity usually combines several features.

(20) One more surge and Afghanistan will be pacified.

And then one more surge and Iraq will be pacified.

And then one more surge and Syria will be pacified.

Surge on.

Negative words and expressions tend to be rare, as we discussed earlier. They can, however, be repeated for rhetorical effect. Potts (2011) comments on a use of negative items whereby the repetition seems to be used for intensification (*I do not not not like...*).

An additional phenomenon that characterizes online comments is ellipsis, the omission from speech or writing of a word or words that are superfluous or able to be understood from contextual clues. This happens especially when the ellipsis follows a long sentence, as illustrated with the phrase *rare, indeed* in (21).

- (21) A King, who shows care and genuine concern for his people but also for those outside his kingdom.
A King, who believes he has a moral obligation to speak out. Rare, indeed.

5. Conclusions

The goal of this paper was to investigate the linguistic expression of negative evaluation, and in particular to sketch a preliminary typology of the expression of negative evaluation. The extensive analyses of evaluative language in our previous work led us to study the expression of negativity and negative evaluation, following the claim that humans have an inherent negativity bias and perceive negative events more salient. We first discussed the linguistic manifestations of that negativity bias, and why it is so important to pay attention to negative language: Negative language permeates and colours our perception of what is said around it. Then, using corpus-assisted analysis, we classified some of the resources that play a role in the expression of negative evaluation into phenomena in the lexicogrammar, and phenomena that belong in discourse semantics. This general classification of evaluative devices overlaps with the planes of expression in Systemic Functional linguistics. The third component of the paper is a corpus-based analysis of three specific phenomena that play a role in the expression of negativity: adverb-adjective combinations, questions, and rhetorical figures. Through our data analysis, we show how apparently innocuous expressions are deployed to subtly convey a negative evaluation of objects, persons, or situations. The corpus analysis is limited, and includes only two very specific types of data (online reviews and online comments to news articles). We believe, however, that the specific phenomena that we discuss are present across different types of texts. We also plan to investigate other types of data.

Although the emphasis in this paper is on corpus methodology to study a pragmatic phenomenon, this is not reproduceable research, and represents an initial step in our larger project. The objective of our future research is to make it quantitative, in extracting negativity patterns and resources in an automated or semi-automated manner. Future research will involve a detailed analysis of the phenomena in our taxonomy, first manually through human annotations. Once we understand the phenomena well, and are able to find patterns in the language that convey negativity, we plan to develop automatic methods to extract those, through similar methods to those used in sentiment analysis (Benamara et al., to appear; Liu, 2015; Taboada et al., 2011).

Acknowledgements

We would like to thank Emilie Francis for assistance in data collection and preparation.

Funding

This research was funded by a Collaborative Travel Grant (Simon Fraser University and Griffith University), and by a grant from the Social Sciences and Humanities Research Council of Canada (435-2014-0171, "Opinion, subjectivity and point of view in language").

References

- Alba-Juez, Laura, & Attardo, Salvatore. (2014). The evaluative palette of verbal irony. In Geoff Thompson & Laura Alba-Juez (Eds.), *Evaluation in Context* (pp. 93-116). Amsterdam: John Benjamins.
- Benamara, Farah, Taboada, Maite, & Mathieu, Yannick. (to appear). Evaluative language beyond bags of words: Linguistic insights and computational applications. *Computational Linguistics*.
- Biber, Douglas, & Barbieri, Federica. (2007). Lexical bundles in university spoken and written registers. *English for Specific Purposes*, 26, 263-286.
- Biber, Douglas, & Finegan, Edward. (1989). Styles of stance in English: Lexical and grammatical marking of evidentiality and affect. *Text*, 9(1), 93-124.
- Biber, Douglas, Johansson, Stig, Leech, Geoffrey, Conrad, Susan, & Finegan, Edward. (1999). *Longman Grammar of Spoken and Written English*. Harlow, Essex: Pearson Education.
- Borchers, Timothy. (2006). *Rhetorical Theory: An Introduction*. Long Grove, IL: Waveland.
- Boucher, Jerry D., & Osgood, Charles E. (1969). The Pollyanna hypothesis. *Journal of Verbal Learning and Verbal Behaviour*, 8, 1-8.
- Burke, Kenneth. (1966). Definition of man. In Kenneth Burke (Ed.), *Language as symbolic action: Essays on life, literature, and method* (pp. 3-24). Berkeley: University of California Press.
- Curtis, Valerie. (2013). *Don't Look, Don't Touch, Don't Eat*. Chicago: University of Chicago Press.
- Davies, Mark. (2008-). The Corpus of Contemporary American English: 1990-present. Available online at <http://corpus.byu.edu/coca/>. Retrieved from <http://corpus.byu.edu/coca/>
- Dodds, Peter Sheridan, Clark, Eric M., Desu, Suma, Frank, Morgan R., Reagan, Andrew J., Williams, Jake Ryland, . . . Danforth, Christopher M. (2015). Human language reveals a positivity bias. *Proceedings of the National Academy of Sciences*, 112(8), 2389-2394.
- Eggs, Suzanne. (1994). *An Introduction to Systemic Functional Linguistics*. London: Pinter.
- Garcia, David, Garas, Antonios, & Schweitzer, Frank. (2012). Positive words carry less information than negative words. *EPJ Data Science*, 1(3).
- Gardiner, Becky, Mansfield, Mahana, Anderson, Ian, Holder, Josh, Louter, Daan, & Ulmanu, Monica. (2016). The dark side of Guardian comments. *The Guardian*. Retrieved from <https://www.theguardian.com/technology/2016/apr/12/the-dark-side-of-guardian-comments>
- Ginsberg, Benjamin. (2011). *The Fall of the Faculty: The rise of the all-administrative university and why it matters*. Oxford: Oxford University Press.
- Goddard, Cliff. (2016). *The semantics of "sarcasm"*. Paper presented at the Australasian Humour Studies Network Conference, Sydney, Australia.
- Goddard, Cliff, Taboada, Maite, & Trnavac, Radoslava. (2016a). Semantic Descriptions of 24 Evaluational Adjectives, for Application in Sentiment Analysis: Simon Fraser University.
- Goddard, Cliff, Taboada, Maite, & Trnavac, Radoslava. (2016b). The semantics of evaluational adjectives: Perspectives from Natural Semantic Metalanguage and Appraisal. *Functions of Language*, Under review.
- Goldberg, Adele E. (2006). *Constructions at Work: The Nature of Generalization in Language*. Oxford: Oxford University Press.

- Greenberg, Joseph H. (1966). *Language Universals, with Special Reference to Feature Hierarchies*. The Hague: Mouton.
- Halliday, Michael A. K. (1985). *An Introduction to Functional Grammar* (1st ed.). London: Arnold.
- Halliday, Michael A. K., & James, Z. L. (1993). A quantitative study of polarity and primary tense in the English finite clause. In John Sinclair, Michael Hoey & Gwyneth Fox (Eds.), *Techniques of description: Spoken and written discourse* (pp. 32-66). New York: Routledge.
- Halliday, Michael A. K., & Matthiessen, Christian M.I.M. (2014). *An Introduction to Functional Grammar* (4th ed.). London: Arnold.
- Horn, Laurence R. (1989). *A Natural History of Negation*. Chicago: University of Chicago Press.
- Hunston, Susan. (2011). *Corpus Approaches to Evaluation: Phraseology and Evaluative Language*. New York: Routledge.
- Hunston, Susan, & Francis, Gill. (2000). *Pattern Grammar: A Corpus-Driven Approach to the Lexical Grammar of English*. Amsterdam and Philadelphia: John Benjamins.
- Israel, Michael. (2004). The pragmatics of polarity. In Laurence Horn & Gregory Ward (Eds.), *The Handbook of Pragmatics* (pp. 701-723). Malden, MA: Blackwell.
- Izard, Carroll. (1971). *The Face of Emotions*. New York: Appleton-Century-Crofts.
- Jing-Schmidt, Zhuo. (2007). Negativity bias in language: A cognitive-affective model of emotive intensifiers. *Cognitive Linguistics*, 18(3), 417-443.
- Kahneman, Daniel, & Tversky, Amos. (1979). Prospect Theory: An analysis of decision under risk. *Econometrica*, 47(2), 263-291.
- Keith, William M., & Lundberg, Christian O. (2008). *The Essential Guide to Rhetoric*. Boston, MA: Bedford/St. Martin's.
- Liebrecht, Christine, Hustinx, Letticia, & Van Mulken, Margot. (2017). A disastrously bad or an extraordinarily good hotel? The perceived strength of intensified evaluations. *67th Annual ICA Conference*. San Diego, CA.
- Liu, Bing. (2015). *Sentiment Analysis: Mining Opinions, Sentiments, and Emotions*. Cambridge: Cambridge University Press.
- Louw, Bill. (1993). Irony in the text or insincerity in the writer? The diagnostic potential of semantic prosodies. In Mona Baker, Gill Francis & Elena Tognini-Bonelli (Eds.), *Text and Technology: In Honour of John Sinclair* (pp. 157-176). Amsterdam: Benjamins.
- Mann, William C., & Thompson, Sandra A. (1988). Rhetorical Structure Theory: Toward a functional theory of text organization. *Text*, 8(3), 243-281.
- Martin, James R. (1992). *English Text: System and Structure*. Amsterdam and Philadelphia: John Benjamins.
- Martin, James R. (2014). Evolving systemic functional linguistics: Beyond the clause. *Functional Linguistics*, 1(3), 1-24.
- Martin, James R., & White, Peter R. R. (2005). *The Language of Evaluation*. New York: Palgrave.
- Matlin, Margaret W., & Stang, David J. (1978). *The Pollyanna Principle: Selectivity in language, memory, and thought*. Cambridge, MA: Schenkman.

- Matthiessen, Christian M.I.M. (2006). Frequency profiles of some basic grammatical systems: An interim report. In Geoff Thompson & Susan Hunston (Eds.), *System and Corpus: Exploring Connections* (pp. 10-142). London: Equinox.
- Mochales, Raquel, & Moens, Marie-Francine. (2011). Argumentation mining. *Artificial Intelligence and Law*, 19(1), 1-22.
- Pang, Bo, & Lee, Lillian. (2008). Opinion Mining and Sentiment Analysis. *Foundations and Trends in Information Retrieval*, 2(1-2), 1-135.
- Perelman, Chaïm, & Olbrechts-Tyteca, Lucie. (1969). *The New Rhetoric: A Treatise on Argumentation*. Notre Dame: University of Notre Dame Press.
- Potts, Christopher. (2011). On the negativity of negation *Proceedings of SALT 20: Semantics and Linguistic Theory* (pp. 636-659). Vancouver.
- Ribeiro, Filipe N., Araújo, Matheus, Gonçalves, Pollyana, Gonçalves, Marcos André, & Benevenuto, Fabrício. (2016). SentiBench: A benchmark comparison of state-of-the-practice sentiment analysis methods. *EPJ Data Science*, 5(23), 1-29.
- Rozin, Paul, Berman, Loren, & Royzman, Edward B. (2010). Biases in use of positive and negative words across twenty natural languages. *Cognition and Emotion*, 24(3), 536-548.
- Rozin, Paul, & Royzman, Edward B. (2001). Negativity bias, negativity dominance, and contagion. *Personality and Social Psychology Review*, 5(4), 296-320.
- Taboada, Maite. (2008). SFU Review Corpus [Corpus]. Vancouver: Simon Fraser University, http://www.sfu.ca/~mtaboada/research/SFU_Review_Corpus.html.
- Taboada, Maite, Brooke, Julian, Tofiloski, Milan, Voll, Kimberly, & Stede, Manfred. (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics*, 37(2), 267-307.
- Taboada, Maite, & Trnavac, Radoslava (Eds.). (2013). *Nonveridicality and evaluation: Theoretical, computational and corpus approaches*. Leiden: Brill.
- Toulmin, Stephen. (1958). *The Uses of Argument*. Cambridge: Cambridge University Press.
- Trnavac, Radoslava, & Taboada, Maite. (2012). The contribution of nonveridical rhetorical relations to evaluation in discourse. *Language Sciences*, 34(3), 301-318.
- Walton, Douglas, Reed, Chris, & Macagno, Fabrizio. (2008). *Argumentation Schemes*. Cambridge: Cambridge University Press.
- Wiebe, Janyce, Wilson, Theresa, Bruce, Rebecca, Bell, Matthew, & Martin, Melanie. (2004). Learning subjective language. *Computational Linguistics*, 30(3), 277-308.