

Sentiment Composition Using a Parabolic Model

Baptiste Chardon^{1,2} Farah Benamara¹ Yannick Mathieu³
Vladimir Popescu¹ Nicholas Asher¹

¹ IRIT-CNRS, Toulouse University,

² Synapse Développement, Toulouse,

³ LLF-CNRS, Paris 7 University

{chardon, benamara, popescu, asher}@irit.fr
yannick.mathieu@linguist.jussieu.fr

Abstract

In this paper, we propose a computational model that accounts for the effects of negation and modality on opinion expressions. Based on linguistic experiments informed by native speakers, we distil these effects according to the type of modality and negation. The model relies on a parabolic representation where an opinion expression is represented as a point on a parabola. Negation is modelled as functions over this parabola whereas modality through a family of parabolas of different slopes; each slope corresponds to a different certainty degree. The model is evaluated using two experiments, one involving direct strength judgements on a 7-point scale and the other relying on a sentiment annotated corpus. The empirical evaluation of our model shows that it matches the way humans handle negation and modality in opinionated sentences.

1 Introduction

Sentiment composition is the process of computing the sentiment orientation of an expression or a sentence (in terms of polarity and / or strength) on the basis of the sentiment orientation of its constituents. This process, similar to the *principle of compositionality* (Dowty et al., 1989), aims to capture how opinion expressions interact with each other and with specific linguistic operators such as intensifiers, negations or modalities. For instance, the sentiment expressed in the sentence *This restaurant is good but expensive* is a combination of the prior sentiment orientation of the words *restaurant*, *good*, *but* and *expensive*. Similarly, in *My wife confirms that this restaurant is not good enough*, sentiment composition has to deal with the verb *confirm*, the adjective *good* and the adverbs *not* and *enough*.

Several computational models were proposed to account for sentiment composition. (Moilanen and Pulman, 2007) use a syntactic tree representation where nodes are associated to a set of specific hand-made composition rules that treat both negation and intensifier via three models: sentiment propagation, polarity conflict resolution and polarity reversal. (Shaikh et al., 2007) use verb frames representation for sentence-level classification and show that their compositional model outperforms a non-compositional rule-based system. (Yessenalina and Cardie, 2011) represent each word as a matrix and combine words using iterated matrix multiplication, which allows for modelling both additive (for negations) and multiplicative (for intensifiers) semantic effects. This matrix-space model is learned in order to assign ordinal sentiment scores to sentiment-bearing phrases. (Socher et al., 2011) model sentences in a vectorial representation and propose an approach based on semi-supervised recursive autoencoders in order to predict sentence-level sentiment distributions. (Wu et al., 2011) propose a graph-based method for computing a sentence-level sentiment representation. The vertices of the graph are the opinion targets, opinion expressions and modifiers of opinion and the edges represent relations among them (mainly, opinion restriction and opinion expansion). Finally (Socher et al., 2012) propose a matrix-vector representations with a recursive neural network. The model is build on a parse tree where the nodes are associated to

a vector. The matrix captures how each constituent modifies its neighbour. The model was applied to predict fine-grained sentiment distributions of adverb-adjective pairs.

Based on linguistic experiments informed by native speakers (Benamara et al., 2012), we propose a sentiment composition model based on a parabolic representation where an opinion expression is represented as a point on a parabola. Our model is designed to handle the interactions between opinion expressions and specific linguistic operators at the sub-sentential level. This paper focus particularly on modality and negation but our model can be used to treat intensifier as well. Within the model, negation are modelled as functions over this parabola whereas modality through a family of parabolas of different slopes; each slope corresponds to a different certainty degree. The model is applied for French but it can be easily instantiated for other languages like English. Its empirical evaluation shows that it has good agreement with the way humans handle negation and modality in opinionated sentences. Our approach is novel:

- it takes into account both negation and modality in a uniform framework. In our knowledge, our approach is the first study dealing with the semantic of modality for sentiment analysis,
- it distills the effect of these linguistic phenomena on opinion expressions depending on different types of negation and modality. We distinguish between three types of negation (Godard, 2013): *negative operators*, such as “not”, “without”, *negative quantifiers*, such as “ever”, “nobody” and *lexical negations*, such as “absence” and between three types of modality (Larreya, 2004) (Portner, 2009): *bouletic*, such as “hope”, “wish”, *epistemic* such as “definitely”, “probably” and *deontic*, such as “must”. (Benamara et al., 2012) empirically show that each type of negation and modality has a specific effect on the opinion expression in its scope: both on the polarity and the strength for negation and on the strength and/or the certainty degree for modality. These empirical results provide a basis for our computational model.
- it provides a lexicon independent representation of extra-propositional aspects of meaning.

The paper is organized as follow. We first give an overview of how existing sentiment analysis systems deal with negation and modality. We then give in section 3 the linguistic motivations behind our approach. The parabolic model and its evaluation are respectively described in section 4 and section 5.

2 Related Works

The computational treatment of negation and modality has recently become an emerging research area. These complex linguistic phenomena have been shown to be relevant in several NLP applications such as sentiment analysis (Wiegand et al., 2010), information retrieval (Jia and Meng, 2009), recognizing contrasts and contradictions (de Marneffe and Manning, 2008) and biomedical text processing (Szarvas, 2008). Due to the emergence of this field, several workshops and conferences have been organized such as the Negation and Speculation in Natural Language Processing (NeSp-NLP 2010) workshop, the Extra-Propositional Aspects of Meaning in Computational Linguistics (ExPRom 2012) workshop, and the publication of a special issue of the journal Computational Linguistics. A number of resources annotated with factuality information are also available. Among them, we can cite the BioScope corpus (Vincze et al., 2008) and FactBank (Saurí and Pustejovsky, 2009).

In sentiment analysis, the presence of modalities is generally used as a feature in a supervised learning setting for sentence-level opinion classification (Kobayakawa et al., 2009). However, to our knowledge, no work has investigated how modality impacts on opinions. There are two ways of treating negation when computing the contextual polarity an opinion expression at the sentence-level: (a) *polarity reversal* (Polanyi and Zaenen, 2006; Moilanen and Pulman, 2007; Choi and Cardie, 2008) that flips the prior polarity of the expression to its opposite value. For instance, if the score of the adjective “excellent” is +3, then the opinion in “this student is not excellent” is -3 ; (b) *polarity shift* (Taboada et al., 2011) that assumes that negation affects both the polarity and the strength. For instance, the opinion in “this student is not excellent” cannot be -3 ; it rather means that the student is not good enough. Two main types of

negation were taken into account in these models: negators such as “not” and / or *content word negators* (Choi and Cardie, 2008) that can be positive polarity shifters (like *abate*) or negative polarity shifters (like *lack*). Few studies take into account other types of negation. (Taboada et al., 2011) treat negative polarity items (NPIs) (as well as modalities) as “irrealis blockers” by ignoring the semantic orientation of the word under their scope. For example, the opinion word “good” will just be ignored in “any good movie in this theater”. We think that ignoring NPIs is not suitable and a more accurate analysis is needed. In addition, no work has investigated the effect of multiple negatives on opinions.

All the previous studies have focused on English. In French, as far as we know, main existing research in sentiment analysis treat negation as polarity reversal and do not take into account modality (Vernier et al., 2007). Thus, there is little existing work for us to compare ourselves to.

3 Linguistic motivations

Our analysis of negation is based on the lexical-syntactic classification of (Godard, 2013) as part of the “Grande Grammaire du Français” project (Abeillé and Godard, 2010). We distinguish between four types of negation in French¹.

- *Negative operators*, denoted by `NEG`: they are the adverbs “pas” (“not”), “plus” (“no more”), “non” (“no one”), the preposition “sans” (“without”) and the conjunction “ni” (“neither”). These operators always appear alone in the sentence and they cannot be combined with each other. The semantic of negative operators are similar to the negation used in logic since they can be paraphrased by “it is not true”.
- *Negative quantifiers*, denoted by `NEG_quant`, express both a negation and a quantification. They are, for example, the nouns and pronouns “aucun” (“none”), “nul” (“no”), “personne” (“nobody”), “rien” (“nothing”), or the adverbs “jamais” (“never”) and “aucunement”/“nullement” (“in no way”)². `Neg_quant` have three main properties: (i) they can occur in positive sentences (that is not negated), particularly in interrogatives, when they are employed as indefinite (as in *Jean travaille toute la semaine mais jamais le dimanche* (*Jean works all the week but never on Sunday*) or when they appear after the relative pronoun “que” (“that”) (as in *Il a réussi sans qu’il ait jamais fait d’efforts* (*He was successful without doing any efforts*), (ii) in negative contexts, they are always associated to the adverb “ne” (“not”) and (iii) they can be combined with each other as well as with negative operators. Here are some examples of this type of negation extracted from our corpus of French movie reviews: “on ne s’ennuie jamais” (“you will never be bored”), “je ne recommande cette série à personne” (“I do recommend this movie to nobody”)
- *Lexical negations* denoted by `NEG_lex` which are implicit negative words, such as “manque de” (“lack of”), “absence de” (“absence of”), “carence” (“deficiency”), “manquer de” (“to lack”), “dénué de” (“deprived of”). `NEG_lex` can be combined with each other as well as with the two previous types of negation.
- *Multiple negatives*. In some languages, double negatives cancel the effect of negation, while in negative-concord languages like French, double negations usually intensify the effect of negation³. In French, multiple negatives that preserve negation concern two cases: the combinations composed of negative quantifiers and the combination of a negative quantifier and a negative operator. Note that the combination of a lexical negation with a lexical quantifier or a lexical negation with a negative operator cancel the effect of `NEG_lex`. Here is an example of a positive opinion

¹This classification does not cover words such as *few* or *only*, since we consider them as weak intensifiers (strength diminishers) rather than negations.

²In this paper, all examples are in French along with their direct translation in English. Note however that there are substantial semantic differences between the two languages.

³In French, there are at most three negative words in a multiple negative. However, this case is relatively rare in opinion text and we only deal with two negatives

extracted from our corpus of French movie reviews: *Cette série télé n'a jamais manqué de me surprendre* (*This TV series never fails to amaze me*) where we have two negatives: the negative quantifier *jamais* (*never*) and the lexical negation *manqué* (*fail*).

Drawing partly on (Portner, 2009) and on (Larreya, 2004) for French, we have chosen to split modality in three categories:

- *Bouletic*, denoted by Mod_B . It indicates the speaker's desires/wishes. This type of modality is expressed via a closed set of verbs denoting hope e.g. "I *wish* he were kind".
- *Epistemic*, denoted by Mod_E . It indicates the speaker's belief in the propositional content he asserts. They are expressed via adverbs expressing doubt, possibility or necessity such as "perhaps", "definitely", "certainly", etc., and via the French verbs "devoir" ("have to"), "falloir" ("need to/must") and "pouvoir" ("may/can"), e.g. "The movie *might* be good",
- *Deontic*, denoted by Mod_D . It indicates a possibility or an obligation (with their contrapositives, impossibility and permission, respectively). They are only expressed via the same modal verbs as for epistemic modality, but with a deontic reading, e.g., "You *must* go see the movie".

(Benamara et al., 2012) consider that effect of each modal category on opinion expression is on their *strength* – for instance, the strength of the recommendation "You must go see the movie, it's a blast" is greater than for "Go see the movie, it's a blast", and *certainty degree* – for instance, "This movie is *definitely* good" has a greater certainty than "This movie is good". The certainty degree has three possible values, in line with standard literature (Saurí and Pustejovsky, 2009): *possible*, *probable* and *certain*. However, as in (Benamara et al., 2012), we consider that, in an opinion analysis context, the frontier between the first two values is rather vague, hence we conflate them into a value that we denote by *uncertain*. We thus obtain two certainty degrees, from which we build a three-level scale, by inserting between these values a "default" certainty degree for all expressions which are not modalities or in the scope of a modality.

(Benamara et al., 2012) structure the effects of each negation type as a set of hypotheses *PolNeg*, *StrNeg*, *QuantNeg*, *LexNeg* and *MultiNeg* that have been empirically validated by volunteer native French speakers through two protocols: one for *PolNeg* and *StrNeg*, with 81 subjects and one for the three other hypotheses with 96 subjects. Similarly, the effects of modality are structured as a set of six hypotheses that have been empirically validated via a set of three evaluation protocols. Respectively 78, 111 and 78 subjects participated in these studies. The table 1 gives an overview of our set of hypotheses, as well as the results (as the average agreement and disagreement between the subjects' answers and the hypotheses). Regarding these results, only valid hypotheses (i.e that obtain more that 50% agreement) are plugged in our parabolic model. We leave lexical negations for future work since their effect is closely related to the semantic of the word used to express negation.

4 Parabolic Model

Let T be an explicitly subjective phrase that contains one opinion expression exp about one topic. exp can be an adjective, a noun or a verb, and can be modified by a set of linguistic operators (e.g., intensifier, negation, modality) that we denote by OP_i for $i = 1 \dots n$. Their cumulative effect on exp is represented by the nesting $OP_1(OP_2 \dots (OP_n((exp))))$, where the order of operators reflects their scope over exp . Here are some examples of T , along with their corresponding semantic representations, operators are in bold font:

- (1) *Cet étudiant est brillant* (*this student is brilliant*), $T = \text{brilliant}$
- (2) *Cet étudiant n'est pas brillant* (*this student is **not** brilliant*), $T = \text{NEG}(\text{brilliant})$
- (3) *Personne n'est brillant* (***nobody** is brilliant*), $T = \text{NEG}_{\text{quant}}(\text{brilliant})$

Hypothesis	Description	Results
<i>PolNeg</i>	The negation always reverses the polarity of an opinion expression. Exp. <i>exceptionnel (exceptional)</i> and <i>pas exceptionnel (not exceptional)</i> .	90.7 %
<i>StrNeg</i>	The strength of an opinion expression in the scope of a negation is not stronger than of the opinion expression alone.	100 %
<i>QuantNeg</i>	The strength of an expression when in the scope of a <i>NEG_quant</i> is greater than when in the scope of a <i>NEG</i> . Exp. <i>jamais exceptionnel (never exceptional)</i> is stronger than <i>pas exceptionnel (not exceptional)</i> .	67 %
<i>LexNeg</i>	<i>NEG_lex</i> has the same effect as <i>NEG</i> . Exp. <i>lack of taste</i> and <i>no taste</i>	43 %
<i>MultiNeg</i>	The strength of an expression when in the scope of multiple negatives is greater than when in the scope of each negation alone. Exp. <i>plus jamais bon (no longer ever good)</i> is stronger than <i>plus bon (no longer good)</i>	64 %
<i>BoulMod</i>	<i>Mod_B</i> alters the certainty degree of opinion expressions in their scope and is weaker than the certainty degree of the opinion expression itself. Exp. <i>I hope this movie is funny</i> there is less certainty than in <i>This movie is funny</i>	86.5 %
<i>EpisMod1</i>	<i>Mod_E</i> alters the certainty degree of opinion expressions in their scope. For adverbial <i>Mod_E</i> , this degree is altered according to the certainty of the respective adverb: if the latter is uncertain, then the certainty of the opinion in the scope of the adverb is reduced; otherwise, the certainty is augmented	72 %
<i>EpisMod2</i>	The certainty of opinion expressions in the scope of a verbal <i>Mod_E</i> is always lower than when not in the scope of such a modality and varies according to the certainty of the respective verb, from <i>pouvoir</i> – lowest certainty, as in “the film <i>might</i> be good”, to <i>devoir</i> and <i>falloir</i> – greater certainty, as in “the film <i>must</i> be good”.	79 %
<i>EpisMod3</i>	The certainty degrees of opinion expressions in the scope of epistemic <i>devoir</i> and <i>falloir</i> are the same.	57 %
<i>DeonMod1</i>	<i>Mod_D</i> alters the strength of opinion expressions in their scope. Hence, strength varies according to the verb: <i>pouvoir</i> reduces the strength of the opinion, whereas <i>devoir</i> and <i>falloir</i> boost it.	54 %
<i>DeonMod2</i>	The strengths of opinion expressions in the scope of deontic <i>devoir</i> and <i>falloir</i> are the same.	60 %

Table 1: An overview of our set of hypotheses and their associated results

- (4) *Cet étudiant n’apportera jamais rien de bon (This student will never bring anything good)*, $T = NEG_quant(NEG_quant(bon))$
- (5) *Cet étudiant n’est définitivement pas brillant (this student is definitely not brilliant)*, $T = Mod_E(NEG(brillant))$

We assume that *exp* is characterized by a prior score $s = pol \cdot str$ encoded in a lexicon, where $pol \in \{-1, +1\}$ is the polarity of *exp* and $str \in (0, MAX]$ is its strength. For example, if we have a three-value scale to encode opinion strength, we can put $s(brillant) = +3$. The key question is: how can we compute the contextual score of *exp*? i.e. what is the value of $s(T)$? Knowing contextual score of opinion expressions at the sub-sentential level is a necessary step in a sentiment analysis system since the $s(T)$ scores have to be aggregated in order to determine the overall polarity orientation and/or the overall rating at the document level.

To compute the contextual polarity of *exp*, we propose a parabolic model where an opinion expression *exp* is represented by a point *E* of the parabola of focus *F* and summit *O*, such that $E \neq O^4$. This parabola belongs to a family of three parabolas of the same focus and different slopes. The slopes correspond to certainty degrees. By convention, we set a reference value p_0 for “default” certainty degrees, $p_1 > p_0$ for “certain” and $p_2 < p_0$ for “uncertain”. The certainty degree of *exp* being “default”, we place it on the parabola of slope p_0 . The polarity and strength of *exp* on this parabola are then characterized by the angle θ between the lines *EF* and *OF* (see Figure 1).

Our model is parametrised by *pol*, *str* and *MAX*. Hence, θ is obtained as a mapping $\phi : \{pol\} \times \{str\} \rightarrow (0; \pi)$, such that: $\phi = \varphi_2 \circ \varphi_1$ where $\varphi_1 : \{str\} \rightarrow (0; 1)$ and $\varphi_2 : \{pol\} \times (0; 1) \rightarrow (0; \pi)$. To compute φ_1 , we rely on a “pivot” word *exp₀*, such that when in the scope of a negative operator (see Section 3), its polarity is reversed, while its strength, denoted by *str₀*, is preserved. This generally corresponds to words with relatively weak strengths like “good” or “bad” in English. We set $\varphi_1(str_0)$ to $\frac{1}{2}$. This parameter is set to this value in order to be consistent with our elementary operation for negation operators Σ_{neg} (cf. description below). Then, for any expression *exp*, its new strength is computed as follows:

⁴*E* cannot be on the summit of the parabola, since this would correspond to a non-opinionated expression, and our model does not apply to such expressions.

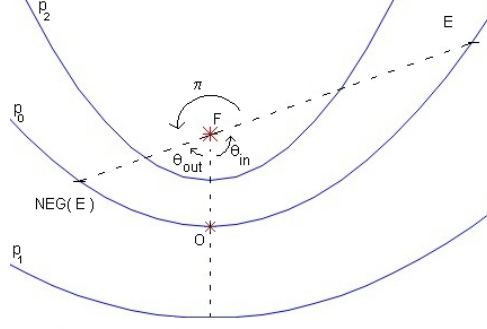


Figure 1: Parabolic model, with negation and modality.

$$\varphi_1(str) = \begin{cases} \frac{str}{str_0} \cdot \frac{1}{2} & , \text{ if } str \leq str_0 \\ \frac{1}{2} + \left(\frac{str - str_0}{MAX + 1 - str_0} \cdot \frac{1}{2} \right) & , \text{ else} \end{cases}$$

Then, we determine the angle corresponding to *exp* from its polarity and new strength as follows:

$$\theta \equiv \varphi_2(pol; \varphi_1(str)) = pol \cdot \varphi_1(str) \cdot \pi$$

The table 2 below shows normalized values in case of a three-points discrete strength such that the “pivot” word *good* is associated to the score +1:

Opinion score s	Normalized angular score θ	Example
+1	$\pi/2$	good
+2	$2\pi/3$	brilliant
+3	$5\pi/6$	excellent
-1	$-\pi/2$	bad
-2	$-2\pi/3$	disgusting
-3	$-5\pi/6$	outrageous

Table 2: Normalization on a 3-points scale

The next step is to compute the score of T , given that T contains one single phrase of the type $OP_1(OP_2\dots(OP_n((exp))))$. Negations and modalities are modeled as functions Σ over the angle θ and the slope p of the parabola where the expressions are placed: $\Sigma : (\theta_{in}; p_{in}) \mapsto (\theta_{out}; p_{out})$. Σ is customized with respect to the operator type: we have both “primitive” and “composition” functions. We have four “primitive” functions:

- Σ_{neg} for negative operators NEG. It consists in adding/subtracting π to/from θ , which ensures that negating of a high-strength opinion expression yields a low-strength one, which is in line with observed behaviour in Hypotheses *PolNeg* and *StrNeg* (cf. table 1):

$$\theta_{out} = \begin{cases} \theta_{out} = \theta_{in} + \pi & , \text{ if } \theta_{in} < 0 \\ \theta_{out} = \theta_{in} - \pi & , \text{ if } \theta_{in} > 0 \end{cases} ; p_{out} = p_{in}$$

Table 3 shows how this formula can be applied in case of a three-points strength scale for positive values. As expected, “not good” has a stronger score than “not excellent”.

- Σ_{int} for intensity modifiers, i.e. deontic modalities (MOD_D) or intensity adverbs. This operation consists in an angle adjustment: it can either increase or decrease the value of θ . We denote these effects by the two sub-functions Σ_{int+} and Σ_{int-} , respectively:

$$\Sigma_{int+}(\theta) = \begin{cases} 2 \cdot \frac{|\theta|}{\pi} \cdot |\theta| & , \text{ if } |\theta| \leq \frac{\pi}{3} \\ \frac{|\theta|}{\pi} \cdot \left(\frac{\pi}{2} + \frac{|\theta|}{2} \right) & , \text{ else;} \end{cases}$$

θ_{in}	$\Sigma_{neg}(\theta_{in})$	Example
$\pi/2$	$-\pi/2$	good / not good
$2\pi/3$	$-\pi/3$	brilliant / not brilliant
$5\pi/6$	$-\pi/6$	excellent / not excellent

Table 3: Negation primitive function on a 3-points scale

$$\Sigma_{int-}(\theta) = \pi - \Sigma_{neg}(\pi - \theta)$$

Table 4 shows an example of these functions in case of a three-points strength scale for positive values.

θ_{in}	$\Sigma_{int+}(\theta_{in})$	Example	$\Sigma_{int-}(\theta_{in})$	Example
$\pi/2$	$3\pi/4$	definitely good	$\pi/4$	possibly good
$2\pi/3$	$5\pi/6$	definitely brilliant	$\pi/3$	possibly brilliant

Table 4: Modality primitive functions on a 3-points scale

- Σ_{cert} for modalities that alter the certainty degree of the expressions in their scope (epistemic MOD_E), according to Hypotheses *BoulMod*, *EpisMod1* to *EpisMod3*. It consists in altering the slope of the parabola, according to the certainty degree c of the modality:

$$\theta_{out} = \theta_{in}; p_{out} = \begin{cases} 2 & , \text{if } c = \text{“certain”} \\ 0.5 & , \text{if } c = \text{“uncertain”}, \end{cases}$$

- Σ_{cert0} for buletic modalities. This operation consists in cancelling the opinion by setting the parameter p to 0.

We have two “composition” functions, Σ_{neg_quant} and Σ_{neg_m} , that account for negative quantifiers and multiple negations, respectively. These functions adjust the output angle yielded by Σ_{neg} and ϕ according to Hypotheses *QuantNeg*, *DeonMod1*, *DeonMod2* and *MultiNeg*. These “composition” functions are defined as follows.

$$\Sigma_{neg_quant} : \theta_{out} = \Sigma_{int+}(\Sigma_{neg}(\theta_{in}; p_{in})), p_{out} = p_{in}.$$

$$\Sigma_{neg_m} : \theta_{out} = \Sigma_{int+}(\Sigma_{int+}(\Sigma_{neg}(\theta_{in}; p_{in}))); p_{out} = p_{in}.$$

Table 5 illustrates these functions.

θ_{in}	$\Sigma_{neg_quant}(\theta_{in})$	Example	$\Sigma_{neg_m}(\theta_{in})$	Example
$\pi/2$	$-3\pi/4$	good / never good	$-7\pi/8$	good / no longer ever good
$2\pi/3$	$-2\pi/3$	brilliant / never brilliant	$-5\pi/6$	brilliant / no longer ever brilliant

Table 5: Composition functions on a 3-points scale

5 Empirical validation

In order to validate empirically our model, we conducted two complementary evaluations. The first one relies on a set of linguistic protocols that aims at evaluating at what extent our model matches the way humans handle negation and modality in opinionated sentences. The second one relies on manually annotated review product corpus and aims at comparing the score that annotators give to elementary discourse segments to the score computed by our model. In both evaluation settings, we compare our model with some baselines and with the (Taboada et al., 2011)’s system which is the state-of-the art model that is the most closer to our. Indeed, (Taboada et al., 2011)’s model shifts the score of an expression to the

opposite polarity by a fixed amount. Thus a +2 adjective is negated to a -2, but the negation of a -3 adjective (for instance, *sleazy*) is only slightly positive.

5.1 Assessing the parabolic model via linguistic protocols

We designed three protocols: *P_NegOp1* and *P_NegOp2* to assess our model with respect to negative operators and one protocol, namely *P_NegQuantMulti*, to evaluate our model with respect to negative quantifiers and to multiple negatives. Since the function Σ_{cert} simply alters the slope of the parabola following the already validated hypothesis *BoulMod* and *EpisMod1* to *EpisMod3* (cf. Table 1), we do not give its evaluation here (see (Benamara et al., 2012) for more details).

In our framework, the strength of the opinion is discretized on a three-level scale, going from 1 (minimal strength) to 3 (maximal strength). Several types of scales have been used in sentiment analysis research, going from continuous scales to discrete ones. Since our negation hypotheses have to be evaluated against human subjects, the chosen length of the scale has to ensure a trade-off between a fine-grained categorisation of subjective words and the reliability of this categorisation with respect to human judgments. We thus use in our framework a discrete 7-point scale, going from -3 (which corresponds to “extremely negative” opinions) to +3 (for “extremely positive” ones) to quantify the strength of an opinion expression. Note that 0 corresponds to cases where in the absence of any context, the opinion expression can be neither positive nor negative.

5.1.1 The experimental setup

The first protocol *P_NegOp1* was already used for evaluating Hypothesis *PolNeg*. It is needed to check whether the scores yielded by the parabolic model match those elicited from human subjects. A set of six questions are shown to subjects. In each question, an opinionated sentence is presented, along with its negation using negative operators, as in “This student is brilliant” and “This student is *not* brilliant”. The strengths of the opinions vary from one question to another on a discrete scale. A set of 81 native French speakers were asked to indicate the strength of each sentence in a question on the same 7-point scale. In the second protocol *P_NegOp2*, the same subjects are given 6 couples of sentences with negative operators, where we vary the strength of the opinion expression in the scope of the negation, while keeping their polarity, e.g. “This student is not brilliant” and “This student is not exceptional”. We ask them to compare, within each couple, the strengths of its members. A set of 96 native French speakers participated in this study. *P_NegOp2* is needed in order to discriminate between our model and different, baseline or state-of-the-art ones (see below), in case of equal performance according to the first protocol. In the third and last protocol, named *P_NegQuantMulti*, we give subjects a set of sentences where each contains an opinion expression of a distinct strength. Each sentence is presented with three forms: one with a negative operator, one with a negative quantifier and one with multiple negation. We then ask subjects to rank each sentence on our 7-point scale. 96 volunteers participate in this protocol.

Given that negation alters only the polarity and strength of an expression (and hence its angle in the model), we first perform a mapping between the angle obtained by applying Σ_{neg} , Σ_{neg_quant} and Σ_{neg_m} , and the 7-point scale, used by human subjects. This mapping is based on the fact that, φ_1 and φ_2 being bijections, their composition $\phi = \varphi_2 \circ \varphi_1$ is a bijection as well. Hence, the inverse mapping $\phi^{-1} = \varphi_1^{-1} \circ \varphi_2^{-1}$ is also a bijective function. Thus, for any angle θ , we get a real-numbered score σ_θ in $[-3, 3]$, which is further discretized via the nearest integer function, yielding the integer $\lfloor \sigma_\theta \rfloor$ on the 7-point scale. The evaluation is performed in two steps: (i) verifying, via *P_NegOp1*, and *P_NegQuantMulti* that, for a given expression, its $\lfloor \sigma_\theta \rfloor$ corresponds to the score given by the subjects; (ii) verifying, via *P_NegOp2*, that, for a set of expressions, the ordering of their $\lfloor \sigma_\theta \rfloor$ s is identical to the ordering of the scores given by subjects. The assessments are quantified as subjects-model agreements.

P_NegOp1 and *P_NegOp2* aim, in addition, to assess our model, along with three other negation models: (i) a “switch” model, which only changes the polarity of the prior score of an expression, while keeping the strength unchanged; (ii) a “flat” model, where the strengths of expressions in the scope of negations are either +1 for negative expressions or -1 for positive ones; (iii) “Tab et al.” model,

standing for (Taboada et al., 2011)’s model. In this model negation boils down to a ± 4 shift of the scores of the opinion expressions on a scale of $\{-5, -4, \dots, 4, 5\}$; hence, polarity is not preserved (Hypothesis *PolNeg* not validated). The assessment according to *P_NegOp1* allows us to indirectly compare these three models to our model. To this end, we first need to perform a scale adjustment for prior scores in (Taboada et al., 2011)’s model: first, our prior scores are linearly mapped to Taboada et al.’s scale, then their model is applied and finally the results are re-mapped to our scale.

5.1.2 Results

In Table 6 we evaluate the subjects-model agreement measure of the four models. We thus assess their ability to provide scores that reflect subjects’ intuition (protocol *P_NegOp1*). In case of equal performance according to this measure, the models are further assessed with respect to their ability to provide the same score orderings as the subjects (protocol *P_NegOp2*). Concerning the correspondence between subject and model scores (*P_NegOp1*), we observe that the “flat” and parabolic models perform best. The “switch” and “Tab et al.” models reflect to a lesser extent subjects’ assessments. The “Tab et al.” model exhibits lower performance figures because, unlike the “flat” and parabolic models, it does not systematically reverse polarity, whereas subjects do so. The parabolic and flat models show the same performance because in both models negation boils down to assigning ± 1 strengths to negated expressions and, in fact, discretizing the output of the parabolic model on the $\{-3, \dots, 3\}$ scale boils down to applying the same formula as for the “flat” model. Hence, in order to further distinguish between the “flat” and parabolic models, we performed the second evaluation, with respect to score orderings (*P_NegOp2*). In this setting, we remark that (Taboada et al., 2011)’s and our parabolic model perform best, which shows that the “switch” and “flat” models fail to provide a score ranking in agreement with subjects’ intuitions. Our model has the same performance as (Taboada et al., 2011)’s model because both are order-preserving shifting models and hence yield the same score ordering for the negated expressions, starting from the same prior score ordering for the expressions.

Model	<i>P_NegOp1</i>	<i>P_NegOp2</i>
Switch	27.03 %	5.80 %
Flat	61.43 %	21.16 %
Tab et al.	47.77 %	73.04 %
Parabolic	61.43 %	73.04 %

Table 6: Empirical validation of the parabolic model

Finally, using *P_NegQuantMulti*, the agreement between the parabolic model and subjects that are in concordance with Hypothesis *PolNeg* is 85.96% for negative quantifiers and 78 % for multiple negatives. Our results show that the adjustment function Σ_{int+} performs well.

5.2 Assessing the parabolic model on manually annotated data

In order to validate our model as a whole, we conducted an experiment on manually annotated data. The data consists in a set of 133 reviews on various subjects: films, TV series, books, and video games. The annotation includes opinion information both at the expression level, with polarity and strength on a three-point scale for opinion words, and with the operators associated to them, and at the discourse segment level, with polarity and strength after application of the operators. While annotating, annotators are not asked to determine the semantic category of negation and modality. For our evaluation, we first automatically determine the type of each operator (i.e negative operator, negative quantifier, multiple negative, epistemic modality, boulic modality as well as intensifiers where we distinguish between adverbs that increase (vs. decrease) the opinion strength) using a dedicated lexicon. Then, we compare the score of discourse segments with those given by annotators. The corpus used for the evaluation contains 393 segments. Table 7 shows the results obtained in terms of accuracy.

We observe that the three models obtain good results, especially in case of intensifiers. Indeed, this kind of operation is usually well supported by each model. Concerning negation, switch model loses an

Model	Accuracy
Switch	59.5 %
Tab et al.	64.7 %
Parabolic	68.8 %

Table 7: Empirical validation of the parabolic model

important part of discourse segments when dealing with high strength opinions; Tab et al. model performs better on most negation, but loses some segments especially when high intensity opinion expressions are concerned: Tab et al model doesn't forecast a polarity switch, and we showed with hypothesis *PolNeg* that this is not the best behaviour for French. On the contrary our model deals correctly with in these cases. In addition, our model performs well on multiple negative and negative quantifiers, which are not taken into account neither in the switch nor in the Tab et al. model. Finally, we also observe that our results for modality are very good, with a F-measure of 88%. However, these results need to be assessed on a larger corpus (we had few instances of epistemic and deontic modalities in our corpus).

6 Conclusion

In this paper, we propose a way to compute the opinion orientation at the sub-sentential level using a parabolic model. Our approach takes into account both negation and modality in a uniform framework and distils the effect of these linguistic phenomena on opinion expressions depending on different types of negation and modality. The empirical evaluation of our model shows that it has good agreement with the way humans handle negation and modality in opinionated sentences. In further work, we plan to study the effect of cumulative modalities, as in “you definitely must see this movie” and of co-occurring negation and modality, as in *you should not go see this movie*, on opinion expressions. At the moment, our model is based on the assumption that a subjective text span contains a single opinion expression. This assumption is far from being verified. Hence, we plan to extend our parabolic model so that it can compute the overall opinion of a text containing several opinion expressions. The focus of the family of three parabolas can correspond to a couple (topic, holder), hence we have as many families of parabolas as opinions expressed towards different topics and/or by different holders. Sentiment composition can then be parametrized by the topic or the holder of the opinion. Finally, we plan to instantiate our model in other languages in order to compare its prediction on standard datasets available in the literature.

Acknowledgement

This work was supported by a DGA-RAPID project under grant number 0102906143.

References

- Abeillé, A. and D. Godard (2010). The grande grammaire du français project. In *Proceedings of the 7th LREC*.
- Benamara, F., B. Chardon, Y. Mathieu, V. Popescu, and N. Asher (2012). How do negation and modality impact on opinions? In *Proceedings of the Workshop on Extra-Propositional Aspects of Meaning in Computational Linguistics*, pp. 10–18. Association for Computational Linguistics.
- Choi, Y. and C. Cardie (2008). Learning with compositional semantics as structural inference for sub-sentential sentiment analysis. In *Proceedings of EMNLP*, pp. 793–801.
- de Marneffe, Marie-Catherine, A. N. R. and C. D. Manning (2008). Finding contradictions in text. In *In Proceedings of ACL 2008*, pp. 10391047.

- Dowty, D. R., R. E. Wall, and S. Peters (1989). *Introduction to Montague Semantics*, Volume 11. D. Reidel.
- Godard, D. (2013). Les négateurs. In *La Grande Grammaire du français*. Godard Danièle, Anne Abeillé and Annie Delaveau. Actes Sud.
- Jia, L. Yu, C. T. and W. Meng (2009). The effect of negation on sentiment analysis and retrieval effectiveness. In *CIKM*, pp. 1827–1830.
- Kobayakawa, T., T. Kumano, H. Tanaka, N. Okazaki, J. Kim, and J. Tsujii (2009). Opinion classification with tree kernel svm using linguistic modality analysis. In *Proceedings of the 18th ACM conference on Information and knowledge management*, pp. 1791–1794.
- Larreya, P. (2004). L’expression de la modalité en français et en anglais (domaine verbal). *Revue belge de philologie et d’histoire* 82(3), 733–762.
- Moilanen, K. and S. Pulman (2007, September 27-29). Sentiment composition. In *Proceedings of RANLP*, pp. 378–382.
- Polanyi, L. and A. Zaenen (2006). Contextual valence shifters. In *Computing Attitude and Affect in Text: Theory and Applications*, Volume 20 of *The Information Retrieval Series*, pp. 1–10.
- Portner, P. (2009). *Modality*, Volume 1. Oxford University Press, USA.
- Saurí, R. and J. Pustejovsky (2009). FactBank: a corpus annotated with event factuality. *Language Resources and Evaluation* 43(3), 227–268.
- Shaikh, M. A. M., H. Prendinger, and M. Ishizuka (2007). Assessing sentiment of text by semantic dependency and contextual valence analysis. In *Affective Computing and Intelligent Interaction, Second International Conference, ACII 2007, Lisbon, Portugal, September 12-14, 2007, Proceedings*, Lecture Notes in Computer Science, pp. 191–202. Springer.
- Socher, R., B. Huval, C. D. Manning, and A. Y. Ng (2012). Semantic Compositionality Through Recursive Matrix-Vector Spaces. In *Proceedings of EMNLP*.
- Socher, R., J. Pennington, E. H. Huang, A. Y. Ng, and C. D. Manning (2011, July). Semi-supervised recursive autoencoders for predicting sentiment distributions. In *Proceedings of EMNLP*, pp. 151–161.
- Szarvas, G. (2008). Hedge classification in biomedical texts with a weakly supervised selection of keywords. In *In Proceedings of ACL 2008*, pp. 281289.
- Taboada, M., J. Brooke, M. Tofiloski, K. Voll, and M. Stede (2011). Lexicon-based methods for sentiment analysis. *Computational Linguistics* 37, 267–307.
- Vernier, M., Y. Mathet, F. Rioult, T. Charnois, S. Ferrari, and D. Legallois (2007). Classification de textes d’opinions: une approche mixte n-grammes et sémantique. *DEFT’07*.
- Vincze, V., G. Szarvas, R. Farkas, G. Móra, and J. Csirik (2008). The bioscope corpus: biomedical texts annotated for uncertainty, negation and their scopes. *BMC bioinformatics* 9(Suppl 11), S9.
- Wiegand, M., A. Balahur, B. Roth, D. Klakow, and A. Montoyo (2010). A survey on the role of negation in sentiment analysis. In *Proceedings of the Workshop on Negation and Speculation in Natural Language Processing, NeSp-NLP ’10*, pp. 60–68.
- Wu, Y., Q. Zhang, X. Huang, and L. Wu (2011). Structural opinion mining for graph-based sentiment representation. In *Proceedings of EMNLP*, pp. 1332–1341.
- Yessenalina, A. and C. Cardie (2011). Compositional matrix-space models for sentiment analysis. In *Proceedings of EMNLP*, pp. 172–182.