

Detecting and Blocking False Sentiment Propagation

Hye-Jin Min

CS Department, KAIST
Daejeon, Republic of Korea
hjmin@nlp.kaist.ac.kr

Jong C. Park

CS Department, KAIST
Daejeon, Republic of Korea
park@cs.kaist.ac.kr

Abstract

Sentiment detection of a given expression involves interaction with its component constituents through rules such as polarity propagation, reversal or neutralization. Such compositionality-based sentiment detection usually performs better than a vote-based bag-of-words approach. However, in some contexts, the polarity of the adjectival modifier may not always be correctly determined by such rules, especially when the adjectival modifier characterizes the noun so that its denotation becomes a particular concept or an object in customer reviews. In this paper, we examine adjectival modifiers in customer review sentences whose polarity should either be propagated (SHIFT) or not (UNSHIFT). We refine polarity propagation rules in the literature by considering both syntactic and semantic clues of the modified nouns and the verbs that take such nouns as arguments. The resulting rules are shown to work particularly well in detecting cases of ‘UNSHIFT’ above, improving the performance of overall sentiment detection at the clause level, especially in ‘neutral’ sentences. We also show that even such polarity that is not propagated is still necessary for identifying implicit sentiment of the adjacent clauses.

1 Introduction

Detecting the sentiment of a given grammatical unit such as phrase, clause or sentence has received much attention in opinion mining and sentiment analysis. While earlier work simply detected the overall polarity by computing the majority of the polarity of words within the expression, researchers are now looking more into the composition of polarity of words within the expression (Wilson et al., 2005; Moilanen and Pulman, 2007; Choi and Cardie, 2008). They have utilized either word features (e.g., ‘Context

Valence Shifters’) or grammatical structures (e.g., ‘the Principle of Compositionality’). It is shown that a machine learning approach with these features performs better than a vote-based bag-of-words approach. While the importance of salient features such as ‘negation’ or ‘intensifier’ is fully recognized, it is not yet clearly understood when the polarity of a particular word is propagated or is sacrificed.

The polarity of an adjectival modifier is often propagated to the polarity of the modified noun or noun phrases with no inherent polarity. However, sometimes the polarity is not propagated to that of the enclosing clause or sentence at all. For example, the polarity of the word ‘real’ is not propagated to that of the sentence “Real 501’s are made of 14 oz canvas-like material.” in a customer review of the pants, even though there is no salient sentiment word except for the word ‘real’. Our observation shows that stopped propagation of this kind in customer reviews often appears because of the following reasons: 1) the word in question is mainly used to refer to the property or the identity of the product entity; 2) it is mainly used to describe certain processes about the author’s experiences or to provide a useful guide for potential customers.

It is important to make the correct decision about the polarity propagation in particular regarding no propagation of the polarity of such an adjectival modifier, in order to detect the sentiment over customer reviews at a deeper linguistic level. For example, the word ‘real’ above is chosen to refer to the other comparative entity, which is regarded as a ‘positive’ entity, as opposed to the present one that is not ‘real’. Hence, the ‘positive’ polarity should not be propagated to the polarity of the current reviewed product entity in the context. The benefit of this decision is that it will enhance the detection of the ‘neutral’ polarity of the sentences in a document. This decision can also be utilized in identifying

the underlying ‘negative’ sentiment of the given sentence. Although it is still hard to detect the case by just looking into the sentiment of the words at the surface level, this will still work as a good clue for the detection because such a word is in contrast to the phrase ‘these *Iconic Rigid* jeans’ as in “These *Iconic Rigid* jeans are made of some sleazy, much lighter material”, which is the sentence that follows. By considering these two sentences together, we can see that a ‘negative’ sentiment is conveyed. Previous work on sentiment detection from customer reviews mainly focuses on detecting sentiment of product features from the patternized sentences (Hu and Liu, 2004; Popescu and Etzioni, 2005; Titov and McDonald, 2008), so the sentences containing such implicit sentiment were not analyzed properly, despite of its importance.

In this paper, we examine adjectival modifiers in customer review sentences whose polarity should either be propagated (SHIFT) or not (UNSHIFT) when the modified noun has no inherent polarity. We refine the previous polarity propagation rules (Moilanen and Pulman, 2007) in order to enhance the performance of the propagation decision by considering both syntactic and semantic clues of the modified nouns and the verbs that take such modified nouns as arguments.

Our rules incorporating these clues into the previous rules have an important role in detecting the ‘UNSHIFT’ case. We found that our rules help the overall sentiment detection at the clause level especially regarding the ‘neutral’ cases but found also that even such polarity with no propagation is also necessary identifying the implicit sentiment of the adjacent clauses.

The rest of the paper is organized as follows. Section 2 introduces previous work analyzing the sentiment in customer reviews focusing on the detection of the polarity. Section 3 summarizes compositionality-based polarity detection in this paper. Sections 4 and 5 describe basic and refined polarity decision rules for adjectival modifiers. Section 6 analyzes our experimental results and Section 7 discusses its importance and limitation. Section 8 concludes the paper with future work.

2 Related Work

Previous work on the detection of the opinions and sentiments to a given product can be divided into three groups: graph-based method with polarity words, rule-based and machine learning-

based methods focusing on sentiment detection in a compositional way. Hu and Liu (2004) identified the sentiment by counting the relevant adjectives that belong to each polarity class with a graph-based polarity lexicon. Popescu and Etzioni (2005) determined the polarity of opinion-containing phrases by identifying the polarity of the words based on relaxation labeling.

The rule-based sentiment identification methods are based on the Principle of Compositionality (Moilanen and Pulman, 2007; Neviarouskaya et al., 2009). Such methods determine the polarity of a given unit basically by composing the inherent polarity of its component lexical units or other linguistic constituents. In addition, a certain type of unit called ‘valence shifters’ works to contextually neutralize or intensify the polarity of the given phrase or sentence (Polanyi and Zaenen, 2004; Kennedy and Inkpen, 2006). Our work is also based on the polarity decision rules proposed by the previous work, and we modified some of them for our purpose. The benefit of rule-based approach is that it is easy to incorporate the additional rules into a rule-based framework for further detailed classification with additional categories.

Some researchers incorporated rule-based sentiment identification into machine learning techniques (Wilson et al., 2005; Choi and Cardie, 2008). Wilson and colleagues (2005) developed the classifier using AdaBoost.HM based on the idea of contextual valence shifters in order to identify contextual polarity at the phrase level. One of the features they considered is modification feature, *modifies* (parent with polarity), and *modified* (children with polarity), though they did not examine the context in which these types of feature may or may not contribute to the overall polarity. Choi and Cardie (2008) developed a machine-learning based sentiment detection method by adopting the Principle of Compositionality (Moilanen and Pulman, 2007) in order to examine whether such computational semantic-based idea can be made empirically effective. Their results show that the method incorporating compositionality performed best among all the methods. Our work is similar to their work in that we followed the idea of the Principle of Compositionality. However, our focus is on examining the characteristics of context surrounding a given adjectival modifier when its polarity is either propagated or not propagated and seeing how this propagation result affects the overall polarity of the clause.

3 Sentiment Detection based on Compositionality

Previous work on deciding the overall polarity of the given expression based on the Principle of Compositionality (Moilanen and Pulman, 2007; Neviarouskaya et al., 2009) takes into account how component lexical units are syntactically combined and develops rules to handle contextual polarity propagation, reversal, conflict or neutralization when combining the inherent polarities of the component lexical units.


We follow the polarity decision rules from previous work (Moilanen and Pulman, 2007; Shaikh et al., 2007; Neviarouskaya et al., 2009) as shown below. We apply the rules to each sentence with its dependency tree structure acquired from the Stanford parser (Klein and Manning, 2003; Marneffe et al., 2006).

- **Basic Propagation** (Moilanen and Pulman, 2007; Neviarouskaya et al., 2009): The polarity of the lexical unit at the upper level in the dependency structure of the text unit has a higher priority. If the word at the upper level has no inherent polarity, the polarity of its dependent word (at the lower level) is propagated to the polarity of the text unit.
- **OBJ/Comp domination for the case of transfer verbs** (Moilanen and Pulman, 2007): The polarity of the constituent as an object or a complement of the transfer verbs that “transmit sentiments among the arguments” (Neviarouskaya et al., 2009) is dominant when there is a polarity conflict among arguments of such verbs. (e.g., “My good old size started showing up too big or wouldn’t shrink right.”)
- **Reversal** (Moilanen and Pulman, 2007; Neviarouskaya et al., 2009): The negation words (e.g., ‘not’, ‘no’, ‘hardly’, ‘reduce’) reverse the polarity. We added more verb entries containing the meaning of ‘reversal’ from other existing review corpora.
- **Reversal for Conflict** (Both *negative* adverbs and *negative* verbs are combined from (Shaikh et al., 2007)): When two lexical units with ‘negative’ polarity are combined, the polarity of the unit covering both units is reversed. (e.g., “They are 501, it’s hard to go wrong with these.”: positive)

- **Neutralizing operators** (Condition operators (e.g., if) from Neviarouskaya et al., 2009): The polarity of the main clause or the sentence is neutralized when there are adverbial conditional clauses. We added the markers ‘wh-words’ or ‘how’ as well as the conditional marker ‘if/unless’. (e.g., “How can I go wrong with the classic 501?”)

4 Basic Rules for Adjectival Modifier

By following the compositionality-based polarity decision rules, the polarity of a noun or a noun phrase that has no inherent polarity is determined by its modifier’s polarity. In other words, the polarity of the modifier is propagated to the upper level node in the dependency tree structure. For example, the noun phrase ‘a perfect dryer’ becomes to have ‘positive’ polarity by the result of polarity propagation. And such polarity may or may not be propagated depending on its syntactic role of the noun phrase at the clause level. If the phrase is a subject, it gets lower priority than the one that works as an object or a complement, but if the word at the upper level or the word with higher priority at the same level (i.e., object or complement) has no inherent polarity, its polarity can be propagated up to the root level by the ‘Basic Propagation’. Figure 1 illustrates this case.



Aux = {aux, neg, intensifier, ...}

R = {nsobj, dobj, prep, complement}

Sibling Xn: Children of the Verb with dependency relation R_2

Sibling Yn: Children of the Root with dependency relation R_3

Figure 1. Basic propagation

Nonetheless, we found that the polarity should not be propagated in some cases as shown in Example (1a).

- (1) a. Real 501’s are made of 14 oz canvas-like material.
- b. It’s a real 501’s.

There is no word with inherent polarity except ‘real’ in (1a), so the overall polarity could be de-

cided as ‘positive’ by the rules just like (1b), but it is actually closer to ‘neutral’ sentence. The reason is that the adjective is utilized to refer to another product entity, which is ‘the original Levis 501 Jean’ in this context.

Interestingly, we see that such phrases often appear in the customer reviews of a product which is a steady seller and whose quality is already well known. To detect whether the polarity of the adjectival modifier is propagated or not is crucial especially when there are no other salient polarity words except for the adjective. It is mainly used to refer to the other product entity for contrastive purposes.

In this paper, we examine the types of clues that affect the propagation of the adjectival modifier’s polarity at the clause level. We also refine the previous polarity decision rules by incorporating additional clues. With the refined rules, we define our problem as follows. For a given adjectival modifier modifying a noun or a noun phrase with no inherent polarity, we label it with ‘SHIFT/UNSHIFT’ tags depending on the nature of propagation. If it is propagated, we label it with the ‘SHIFT’ tag, and if not, we do it with the ‘UNSHIFT’ tag.

The basic rules for labeling by considering only syntactic clues from the previous polarity decision rules are as follows.

- SHIFT: 1) if the syntactic role of the noun phrase is complement; 2) if the syntactic role of the noun phrase is object of verbs or prepositions
- UNSHIFT: 1) if the syntactic role of the noun phrase is subject (e.g., (1a)); 2) if the syntactic role of the noun phrase is object of the verb whose syntactic type is either ‘gerund (*Ving*)’ or ‘infinitive (*to V*)’

We also regarded the case as ‘UNSHIFT’ where the noun phrase has lower priority than its sibling phrases in the dependency tree; for example, if there is an object with non-neutral polarity and the syntactic role of the given noun phrase is subject, the labeling is done with ‘UNSHIFT’. Example (2) shows ‘SHIFT’ and ‘UNSHIFT’ for the adjectival modifier ‘good’ and ‘great’, respectively.

- (2)a. It’s a good buy.
- b. A great shave takes a little more commitment than just breaking out a can of foam and a disposable razor.

We decided not to use machine learning techniques with the following reasons. First, our goal is not to enhance the overall performance of sentiment detection in general but to examine what kinds of additional clues are called especially for the decision of the polarity propagation of the adjectives modifying the noun with no polarity. Following Kennedy and Inkpen (2006)’s work for measuring the impact of valence shifters on sentiment classification, we believe it is not straightforward to identify major factors for the improvement with a machine learning algorithm. Second, our work focuses on relatively small cases among all the cases in the whole review sentences (See Table 5), so it is reasonable to directly apply refined rules to each case without an unnecessary training process handling other cases. We believe that the rules of this kind by taking a closer look at the focused cases could be extended regarding scalability with the help of the machine leaning techniques in the future.

5 Rule Refinement

We refined the rules with additional clues as follows because the basic rules do not work properly in some context.

5.1 Phrase-level Clues

The basic rules mainly consider the syntactic types of the noun phrase and the verb taking the noun phrase as the argument. However, the following clues at the phrase level may also affect the propagation.

- Quoted / Capital Letters (UNSHIFT)
- Types of noun in the product reviews

Example (3) shows quoted adjectival modifiers. The quotes indicate that its inherent polarity is not effective. We can see that the author of the review intentionally used them to indicate such neutralization.

- (3)a. The fit on these “relaxed” jeans is just that--relaxed but not loose.
- b. If you love “Happy Hippy” shower gel, this fun bath product will impress you.

Examples (4) and (5) show the polarity propagation depending on the types of modified noun by the adjectives. While the polarities in Example (4) are propagated, those in Example (5) are not propagated.

- (4)a. This product also arrived in good condition and in good time. (Delivery)
 b. I will never order anything from Levi again until they come back to the original levi material. (Product feature)
- (5)a. They were also not much cheaper than I would have paid from a real retailer. (Location & time)
 b. People with healthy, not-so-damaged, and normal hair do not need to use this stuff. (User Information)
 c. All it takes for a great shave is a good safety razor, a decent brush, shaving soap and a mug. (Purpose)
 d. After a quick rinse my skin looked less dull, and helped clear blotches (Process).
 e. Use with Chanel loose powder and Chanel concealer for a perfectly flawless finish. (Product name)

The types of noun in Example (4) are related to explicit sentiment of the product. On the other hand, the types of noun in Example (5) are related to implicit sentiment or additional background information provided for the potential customers. In order to distinguish SHIFT cases from UNSHIFT cases resulting from such different types of noun we built a lexicon for the noun type as shown in Table 1. We collected the words belonging to each type by utilizing three different methods. We first manually collected words belonging to each noun type from the sample review texts and extended the entries by including synonyms of the seed words in WordNet (Synonyms; Miller, 1995). Some synsets of WordNet such as ‘body_part’/‘illness’ and ‘shop’ are appropriate for ‘User information’ and ‘Location’. We combined several synsets for such type (WordNet Synsets). The noun types such as ‘Product name’, ‘Product feature’, ‘Purpose’, and ‘Process’ are domain dependent, we collected the words by utilizing Point-wise Mutual Information (PMI).

Propagation	Noun Type	Method
SHIFT	Delivery	Synonyms
	Deal	
	Product feature	PMI
UNSHIFT	Time	Synonyms
	Location (shop, store)	WordNet Synsets
	User info (body part/illness)	
	Product name	PMI
	Process/Purpose	

Table 1. Types of noun

5.2 Clause-level Clues

The main reason for the second rule of the ‘UNSHIFT’ basic rules in Section 4 is that we assumed the given phrase/clause could be regarded as a secondary concept or topic for the main concept or topic as shown in Example (6).

- (6)a. Anyone who is that determined to make the best product on the market, obviously will do whatever it takes to make it happen.
 b. Getting an outstanding shave from this razor should be a cinch.

However, the given phrase/clause should be regarded as the main concept or an independent concept as shown in Example (7).

- (7)a. It seemed to have a rich sophistication which goes with horseback riding or polo.
 b. It’s wonderful doing everything I need, including making my hair nice and shiny, without the heaviness.

In order to capture these differences, we refined the rules as shown in Table 2.

Infinitive (<i>to V</i>)	Gerund (<i>Ving</i>)
IF the head of the infinitive has auxiliary characteristics such as ‘seem’ and ‘need’ THEN label it with SHIFT. Otherwise, label it with ‘UNSHIFT’.	IF the phrase/clause including the gerund is clausal subject THEN label it with ‘UNSHIFT’. Otherwise label it with ‘SHIFT’.

Table 2. Refined rules for ‘UNSHIFT’

The rule for the object of prepositions should also be refined. As we mentioned in Section 5.1, the reason for mentioning some particular types of object in the review is to explain additional background information as a guide for the potential customer as well as showing the sentiment about the product. The types of noun at the phrase level cannot always solely determine the polarity propagation because such decision is still affected by the presence of other constituents in the context at the clause level. For example, by comparing (5c) with the sentence “it provides a very close, smooth shave”, the polarity of ‘smooth’ is propagated while that of ‘great’ is not propagated.

To handle this case properly, we consider ‘Clause-level Semantic Label’ at a shallow level by taking into account both some preposition

types and the noun types together that frequently appear as shown in Table 3. We named the labels by referring to ‘Frame Index’ from FrameNet data (Baker et al., 1998). This list of the pairs filters further ‘UNSHIFT’ cases from the ‘SHIFT’ labeled cases by the basic rules.

Semantic Label (FrameNet)	Prep. & N.Type	Ex.
Purpose (related to Shopping)	‘for/with’ & Purpose (Use)	(5c)
Body mark	‘for/with’ & User info (Body part)	(5b)
Process (related to Using)	‘after’ & Process	(5d)
Place (related to Shopping)	‘on/from’ & Time/ Location	(5a)

Table 3. Semantic Label for filtering ‘UNSHIFT’

The last clue is about the sentence type. Even if the polarity of the adjectival modifier is propagated to the top node word at the clause level, the type of the sentence may block it for the overall polarity of the whole sentence. We consider three types of sentences that turn the ‘SHIFT’ label into the ‘UNSHIFT’ label as shown in Table 4.

Sentence Type	Examples
Condition	You will not be sorry <i>if</i> you are looking for the <u>perfect</u> brush to go with your <u>perfect</u> dryer (the featherweight)!
Experience (Perfect Tense & verb class)	<i>I have been searching</i> for what seems like forever for a <u>nice</u> cologne or perfume that is not all flowery and overpowering
Guide (Types of nominal subject)	The <u>best</u> <i>solution</i> for this is ... A personal <i>tip</i> I would like to suggest ...
Guide (Imperative sentence)	Make sure you shake the bottle before using for <u>best</u> color results (as mentioned on the <u>packaging</u>).

Table 4. Types of sentences for ‘UNSHIFT’

As a number of previous researches also considered, we canceled the detected sentiment at the conditional clause. In addition, we considered two domain specific types of sentences, namely, *experiences sentences* and *guide sentences* as the clues for ‘UNSHIFT’ cases, because these types of sentences also give background information rather than explicitly mentioning the sentiment

so that the polarity of the adjective tends not to be propagated. We defined *experience sentence* whose main subject is the author and which has present or past perfect tenses with purchase related verbs (e.g., buy, search, try or return). We also defined *guide sentence* that is an imperative sentence with no main subject or with the subject referring to the potential customer such as ‘you’ or ‘people’.

The preprocessing steps before applying the rules above are as follows. First, we get the dependency relation pairs for each input sentence acquired from the Stanford parser (Klein and Manning, 2003; Marneffe et al., 2006), and constructed the dependency tree structure for tree traversal in order to process polarity propagation. Then we assigned each word to its inherent polarity (if it has one) by looking up the sentiment lexicon, ‘Subjectivity Lexicon’ (Riloff and Wiebe, 2003). We adapted the lexicon to product reviews by modifying the inherent polarity of 36 lexical entries (e.g., white, positive to neutral) and adding 105 additional words frequently used (e.g., small with neutral). In order to apply rules to particular types of adjective and verb such as transfer verbs or contextually polarized adjectives, we added an additional field such as ‘type’ into each lexical entry to show their identities (The original types of 22 entries in ‘Subjective Lexicon’ are modified). As for extracting clues, we utilized dependency relations for syntactic types of nouns and verbs. For semantic types of nouns and verbs, we utilized the semi-automatically constructed lexicon as mentioned in Section 5.1. In addition, in order to identify ‘*experience sentences*’ and ‘*guide sentences*’, we extracted tense information and noun subject by utilizing dependency parse tree.

6 Experimental Results

We performed two types of experiment in order to examine the performance of our refined polarity propagation rules and the contribution of the propagation results to the sentiment detection at the clause level.

Table 5 shows the data sets of customer reviews we used for the experiments. We first tested our rules with Set 1 (Beauty positive), a corpus utilized in (Blitzer et al., 2007) because all the reviews are classified as ‘positive’, so we assume that there are many adjectival modifiers with ‘positive’ polarity. We then performed both propagation decision and sentiment classification experiments with Set 2 (Levi’s Jean), which is

crawled review data from Amazon.com by ourselves. The reasons why we chose this product are as follows. First, it is a steady-selling product so that most of the reviews are regarded as positive, which makes it more important to identify negative or neutral opinions than other kinds of reviews. Hence, it is crucial to consider correct decision of propagation of the adjectival modifiers with ‘positive’ polarity that is mostly not propagated. Second, after the initial observation, we found that a particular expression about ‘changes in quality’ frequently appears in such reviews (about 20%) and the adjectival modifiers with ‘positive’ polarity in such expression are mostly not propagated because it would refer to other particular entities or be used to describe a certain process.

Data Sets	# for exp.	Total	%
Beauty Positive set 1	444	6,126	7.2%
Levi’s Jeans set 2	147	1,655	8.8%

Table 5. Data sets

Table 6 shows the numbers of propagation rules and Table 7 shows the propagation decision results. Compared to the results by the basic rules, the performance is enhanced in general. However, we notice that the rules related to VerbType are effective on recall but not on precision for ‘SHIFT’. On the other hand, as for ‘UNSHIFT’ the rules are effective on precision but not on recall. Rules taking into account both noun types and prepositions slightly enhance the overall performance. The overall rules that include sentence type score the best precision and recall figures, which are both effective for ‘SHIFT’ and ‘UNSHIFT’.

Next, we apply these rules to our data set 2. Table 8 shows the propagation decision results. The accuracy for the overall test clauses is almost similar to that for set 1. While precision for ‘UNSHIFT’ and recall for ‘SHIFT’ rose, precision for ‘SHIFT’ and recall for ‘UNSHIFT’ dropped. We analyzed False Negative errors of ‘UNSHIFT’ cases. Most of them are unknown cases for each rule except due to parsing errors. This also led to the drop of the precision for ‘SHIFT’. The strong restriction for ‘UNSHIFT’ also affected the result of recall for ‘SHIFT’.

Table 9 shows the sentiment detection results at the clause level for set 2. The performance of ‘positive’ label is not much enhanced but that of ‘neutral’ label is enhanced. We believe that this is because if the polarity of the top node word is

explicitly ‘positive’ because of its inherent polarity the overall polarity of the clause is obviously ‘positive’ regardless of the result of the polarity propagation decision. On the other hand, in the case of ‘neutral’ clause, the correct polarity propagation decision for ‘UNSHIFT’ is critical for detecting the overall polarity. This confirms that our rules have a critical role in detecting the sentiment of ‘neutral’ sentences.

Rules	SHIFT		UNSHIFT	
	#	Σ	#	Σ
Basic rules (B)	3	3	2	2
B + PhClues	0	3	4	6
B + PhClues + VerbType	2	5	2	8
B + PhClues + VerbType + SemLabel	0	5	3	11
All	0	5	3	14

Table 6. Numbers of propagation rules

Rules	SHIFT		UNSHIFT		All Acc.
	P	R	P	R	
Basic rules (B)	0.84	0.83	0.72	0.73	0.79
B + PhClues	0.84	0.83	0.72	0.74	0.80
B + PhClues + VerbType	0.83	0.89	0.79	0.70	0.82
B + PhClues + VerbType + SemLabel	0.86	0.89	0.80	0.76	0.84
All	0.90	0.88	0.80	0.84	0.86

Table 7. Propagation decision results (set 1)

	SHIFT		UNSHIFT		All Acc.
	P	R	P	R	
Basic Rules	0.71	0.89	0.82	0.59	0.74
Refined Rules	0.78	0.90	0.85	0.69	0.80

Table 8. Propagation decision results (set 2)

	Basic Rules		Refined Rules	
	P	R	P	R
Positive	0.82	0.86	0.87	0.87
Negative	0.85	0.67	0.90	0.70
Neutral	0.45	0.56	0.51	0.70
All (Acc.)	0.74		0.79	

Table 9. Sentiment detection results

By the importance of ‘neutral’ polarity, we conducted an error analysis on 18 False Positive cases for ‘neutral’ polarity as shown in Table 10.

7 Discussion

Types	Description	#
Implicit sentiment	The overall sentiment should be detected by further deep linguistic analysis.	8
Incorrect 'UNSHIFT'	The 'UNSHIFT' result is incorrect.	3
Incorrect polarity	The polarity result is incorrect due to other lexical entries	3
Parsing errors	The propagation is made incorrectly due to incorrect dependency relation.	2
Others	Comparison without 'Positive/Negative' sentiment	2

Table 10. Error distribution

We note that the reason for considering specific sentence types as addressed in this paper is that we assume that these sentences are better suited to demonstrate the need for blocking the propagation of the polarity of the given adjectival modifier.

Although we considered certain types in a limited way, we haven't fully observed what types of sentence are actually involved in propagation. In addition, we found that some sentences in the data set we considered initially as having the sentence type that blocks the propagation of polarity of the adjectival modifier do not convey 'neutral' but convey 'positive' or 'negative' polarity implicitly as shown in Example (8).

- (8)a. If there is a more perfect shampoo, I haven't found it.
 b. *Previously*, I *had to* visit my favorite store more than once to get my size.
 c. I've had it for a year and the *elastic* is totally *stretched out* with normal wear.

The main clause in (8a) conveys 'positive' polarity implicitly even though there is no polarity-bearing word. Further processing is necessary including a proper account of negation. The phrases in (8b) and (8c) are about product entities contrastive to the currently reviewed product so that the inherently assigned polarity of 'favorite' and 'normal' is not applicable to the currently reviewed product. In order to detect the implicit intention of this kind, we should also detect the clues for contrast such as 'previously' or the relation between the phrases 'elastic' and 'be stretched out'.

Although the propagation decision for 'UNSHIFT' itself is correct, such inherent polarity of the adjectival modifier may help to identify the

implicit sentiment of the adjacent clause as shown in Example (9).

- (9)a. I washed them repeatedly in my very *efficient* and eco-friendly Asko washer, but the smell remained.
 b. I have paid much more for *inferior* brand jeans and I can say that I won't be doing that anymore.

The implicit polarity of the underlined clause in (9a) may be both 'negative' and 'positive' depending on the context. By utilizing both the inherent polarity of 'efficient' and the role of the conjunction 'but', the conventional polarity detection rule along with conjuncts (Meena and Prabhakar, 2007) can correctly detect its polarity as 'negative'. As for (9b), by the inherent polarity 'negative' of 'inferior' and negation on the underlined clause we can detect the 'positive' polarity of the underlined clause. However, the possibility of the correctness of the detection is still chancy, and a further analysis of the underlying meaning of the clause or the sentence is called for. For example, if we label the clause containing 'inferior' in (9b) as 'action for goal achievement', we can detect the polarity of the underlined clause as 'negative' by the rule taking such label and another label related to its continuity.

8 Conclusion

In this paper, we refined the previous polarity propagation rules in order to better decide whether the polarity of the adjectival modifier should be propagated or not. We considered both phrase-level and clause-level clues by considering syntactic and semantic types of nouns and verbs. Our rules incorporating these clues into the basic rules detected the 'UNSHIFT' case particularly well. The detection results of the overall sentiment at the clause level are meaningfully enhanced as compared to those based on the previous polarity propagation rules regarding especially 'neutral' sentences. However, despite the correct decision for 'UNSHIFT', we found that such polarity of the modifiers may also help to identify the implicit sentiment without further deeper linguistic analysis.

In order to detect implicit sentiment, we will examine the clues for detecting contrast among product entities or product features for the future work. We will also classify the roles of the clause at a fine-grained level that is related to the detection of the implicit sentiment.

Acknowledgements

This work was supported in part by the National Research Foundation of Korea (NRF) grant funded by the Korea government (MEST) (No. 2011-0018262), and in part by the Intelligent Robotics Development Program, one of the 21st Century Frontier R&D Programs funded by the Ministry of Knowledge Economy of Korea. We thank the three anonymous reviewers for helpful comments and insightful suggestions.

References

- Collin F. Baker, Charles J. Fillmore, and John B. Lowe. 1998. The Berkeley FrameNet Project. *In Proceedings of COLING*, Volume 1, pp. 86–90, Morristown, NJ, USA, Association for Computational Linguistics.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, Bollywood, Boom-boxes and Blenders: Domain Adaptation for Sentiment Classification. *In Proceedings of ACL*, pp. 440–447, Prague, Association for Computational Linguistics.
- Yejin Choi and Claire Cardie. 2008. Learning with Compositional Semantics as Structural Inference for Subsentential Sentiment Analysis. *In Proceedings of HLT/EMNLP*, pp. 793–801, Honolulu, Association for Computational Linguistics.
- Minqing Hu and Bing Liu. 2004. Mining and summarizing customer reviews. *In Proceedings of ACM SIGKDD*, pp. 168–177, ACM Press.
- Alistair Kennedy and Diana Inkpen. 2006. Sentiment Classification of Movie and Product Reviews Using Contextual Valence Shifters. *Computational Intelligence* 22(2):110–125.
- Dan Klein and Christopher D. Manning. 2003. Accurate Unlexicalized Parsing. *In Proceedings of ACL*, pp. 423–430, Sapporo, Japan, Association for Computational Linguistics.
- Arun Meena and T.V. Prabhakar. 2007. Sentence Level Sentiment Analysis in the Presence of Conjunctions Using Linguistic Analysis. *In Proceedings of ECIR 2007, LNCS 4425*, pp. 573–580.
- Marie-Catherine de Marneffe, Bill MacCartney and Christopher D. Manning. 2006. Generating Typed Dependency Parses from Phrase Structure Parses. *In Proceedings of LREC 2006*, pp. 449–454.
- George A. Miller. 1995. WORDNET: A Lexical Database for English. *Communications of ACM* (11): pp. 39–41.
- Karo Moilanen and Stephen Pulman. 2007. Sentiment Composition. *In Proceedings of RANLP-2007*, pp. 378–382, Borovets, Bulgaria.
- Alena Neviarouskaya, Helmut Prendinger, and Mitsuru Ishizuka. 2009. Semantically distinct verb classes involved in sentiment analysis. *In Proceedings of IADIS International Conference Applied Computing*, pp. 27–34.
- Livia Polanyi and Annie Zaenen. 2004. Contextual valence shifters. *In Working Notes of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*, pp. 106–111, Menlo Park, CA, The AAAI Press.
- Ana-Maria Popescu and Oren Etzioni. 2005. Extracting product features and opinions from reviews. *In Proceedings of HLT/EMNLP*, pp. 339–346, Vancouver, Association for Computational Linguistics.
- Ellen Riloff and Janyce Wiebe. 2003. Learning extraction patterns for subjective expressions. *In Proceedings of EMNLP*, pp. 105–112, Sapporo, Japan, Association for Computational Linguistics.
- Mostafa Al Masum Shaikh, Helmut Prendinger, and Mitsuru Ishizuka. 2007. Assessing sentiment of text by semantic dependency and contextual valence analysis. *In Proceedings of ACII 2007, LNCS 4738*, pp. 191–202.
- Theresa Wilson, Janyce Wiebe, and Paul Hoffman. 2005. Recognizing Contextual Polarity in Phrase-Level Sentiment Analysis. *In Proceedings of HLT/EMNLP*, pp. 347–354, Vancouver, Association for Computational Linguistics.