

RESEARCH

Open Access



Extracting Chinese polarity shifting patterns from massive text corpora

Ge Xu^{1,2,3*} and Chu-Ren Huang⁴

* Correspondence: xuge@pku.edu.cn

¹Department of Computer Science, Minjiang University, Fuzhou 350108, China

²Fujian Provincial Key Laboratory of Information Processing and Intelligent Control, Fuzhou, China
Full list of author information is available at the end of the article

Abstract

In sentiment analysis, polarity shifting means shifting the polarity of a sentiment clue that expresses emotion, evaluation, etc. Compared with other natural language processing (NLP) tasks, extracting polarity shifting patterns from corpora is a challenging one because the methods used to shift polarity are flexible, which often invalidates fully automatic approaches. In this study, which aimed to extract polarity shifting patterns that inverted, attenuated, or canceled polarity, we used a semi-automatic approach based on sequence mining. This approach greatly reduced the cost of human annotating, while covering as many frequent polarity shifting patterns as possible. We tested this approach on different domain corpora and in different settings. Three types of experiments were performed and the experimental results were analyzed, which will be reported in this paper.

Keywords: Sentiment analysis, Polarity shifting pattern, Sequence mining, Prior polarity

1 Background

Polarity shifting is a phenomenon in sentiment analysis that denotes the inversion, attenuation, or cancelation of the polarity of a sentiment clue (i.e., word or phrase) in a sentence. A polarity shifting pattern is a pattern that performs polarity shifting. There are many polarity shifting patterns that can shift the polarity of a sentiment clue in a sentence. The most common polarity shifting pattern in Chinese (and also in English) is the use of 不 *bu* “not” to negate the sentiment, as can be seen in Examples 1 and 2 below:

(1) 她现在不高兴

ta_xianzai_bu_gaoxing
she_now_NEG^a_happy
She is not happy now.

(2) 我不喜欢吃苹果

wo_bu_xihuan_chi_pingguo
I_NEG_like_eat_apple
I don't like eating apples.

In Examples 1 and 2, 不 *bu* “not” directly negates 高兴 *gaoxing* “happy” and 喜欢 *xihuan* “like,” and inverts the overall polarity of the sentences. Some polarity shifting is less obvious than 不 *bu* “not,” as shown in Examples 3 through 5 below:

(3) 如果她难过,她就会喝酒

ruguo__ta__nanguo, ta__jiuhui__hejiu

if__she__sad, she__will__drink

If she is sad, she will drink.

(4) 这家店以前很好

zhe__jia__dian__yiqian__hen__hao

this__CL__shop__previously__very__good

This shop used to be very good.

(5) 这道菜看起来很好吃

zhe__dao__cai__kanqilai__hen__haochi

this__CL__dish__seem__very__delicious

This dish seems very delicious.

如果 *ruguo* “if” in Example 3 converts a fact into a condition, so the polarity of 难过 *nanguo* “sad” is canceled. In Example 4, 以前 *yiqian* “used to be” is a time contrast that inverts the polarity of 好 *hao* “good” and implies that the shop is bad now. In Example 5, 看起来 *kanqilai* “seem” attenuates the polarity of 好吃 *haochi* “delicious”.

Finding polarity shifting patterns automatically in a corpus is not easy. The density of polarity shifting is normally low, and shifting patterns are subtle and flexible, making the task of extracting polarity shifting patterns difficult. Currently, human annotating is the main way to obtain polarity shifting patterns, which is time-consuming. This paper will present a semi-automatic approach to extracting polarity shifting patterns from specially designed corpora. This approach used sequence mining to generate frequent word sequences, which greatly reduced human labor in annotating polarity shifting patterns compared with annotating the original corpus directly.

2 Methods

Polanyi and Zaenen (2004), Quirk et al. (1985), and Kennedy and Inkpen (2005) categorized contextual valence shifters (i.e., polarity shifting patterns) into three classes: inverters, which invert the polarity of a polarized item; intensifiers, which intensify it; and attenuators, which diminish it. Recent work on polarity shifting has mainly pursued two of the three classes, namely, inverters and attenuators. In our study, we also limited our research to these two classes of polarity shifting patterns.

To find the polarity shifting patterns in a given corpus, our approach followed the steps below, which will be explained in detail in subsequent sections:

1. Select text lines from a corpus where polarity shifting occurs frequently.
2. Use sequence mining (see Section 2.1) to generate frequent polarity shifting pattern candidates.
3. Annotate polarity shifting patterns.
4. Repeat steps 2 and 3 until no more frequent patterns are generated.

2.1 Sequence mining

Our approach was based on sequence mining. A sequence mining algorithm can extract frequent sequences from a corpus. PrefixSpan (Pei et al. 2001) is one such algorithm. The term “frequent” here means that the frequency of an extracted sequence should be higher than a threshold, which is called “minimum support” in the industry of data mining. In our study, a sequence is considered a sequence of tokens (i.e., a word or TAG^b), and an interval is the space allowed between the tokens. An example of how such algorithms work is shown in Table 1:

In Table 1, (“a”, “c”), 4 means that the sequence (“a”, “c”) occurs four times in the lines and that an interval is allowed between “a” and “c”.

2.2 Criteria for annotating polarity

Commonly, polarity is classified as positive, negative, and neutral. These concepts are quite intuitive so there is no need to define them here. However, when researchers annotate the polarity of words or expressions from a corpus, a well-defined sentiment framework that defines the types of sentiments and other related issues can help them make a decision quickly and clearly.

A key part of our approach was to select high-quality polarity words, which helped to retrieve text lines with dense polarity shifting. Then, we extracted polarity shifting patterns from these text lines. In the following, we will introduce some sentiment frameworks briefly and offer a simple and practical criterion for Chinese sentiment annotation.

2.2.1 A simple and practical sentiment criterion

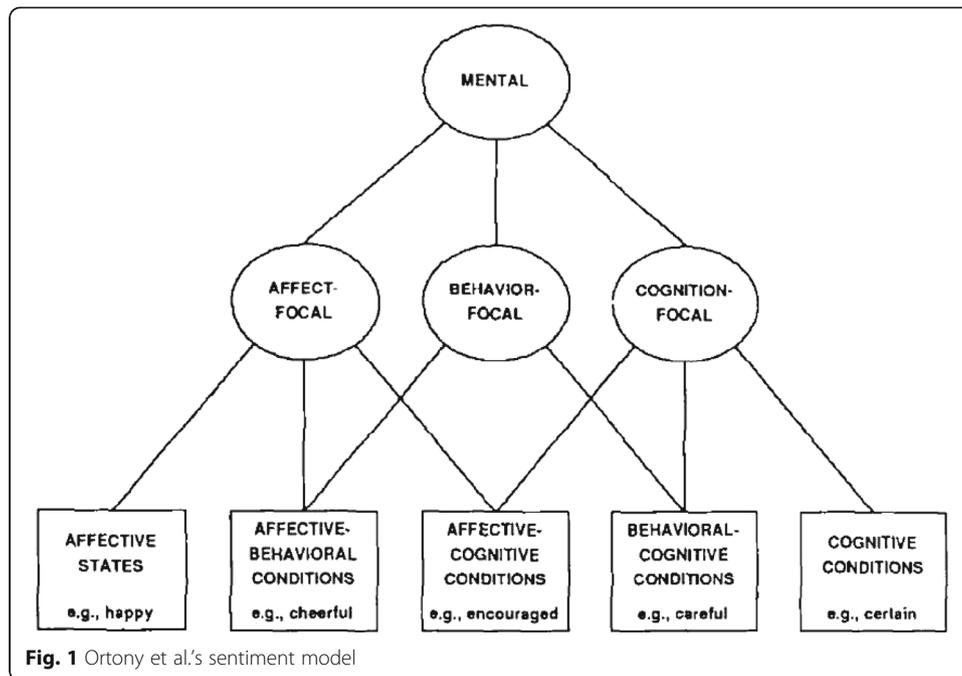
Ortony et al. (1987) presented a sentiment model that classified mental states into three categories: affect-focal, behavior-focal, and cognition-focal (see Fig. 1). Their model did not enumerate basic sentiments, but instead defined basic sentiments in a specific application. Namely, Ortony et al. (1987) offered only a general framework for sentiments, which greatly affected many works on sentiment analysis, such as WordNet-Affect (Strapparava and Valitutti 2004). The framework proposed in this paper was also heavily affected by Ortony et al.’s sentiment model.

Although researchers have contributed much work in this field, it is still a bit difficult to apply it directly to Chinese sentiment annotation. For example, Ortony et al.’s model was proposed before the sentiment analysis of NLP, so polarity was not explicitly considered. In addition, some redundancy exists in the WordNet-Affect model, such as the overlapping of the attitude and cognitive states, where attitude is a complicated concept that may contain cognition, emotion, desire, etc.

Based on the analysis of these models, and with an aim to improve Chinese sentiment annotation, we propose a simple and practical criterion (see Table 2) for Chinese sentiment annotation that will describe concepts such as emotion, evaluation, contextual polarity, prior polarity, and neutrality. The model first separates the basic sentiment elements (e.g., cognition, desire, emotion) as much as possible. For a complex concept,

Table 1 Sequence mining

Input	lines = (“caabc”, “abcb”, “cab”, “abbca”), minSup = 3
Output	(“a”, 4; (“a”, “b”), 4; (“a”, “b”, “c”), 4; (“a”, “c”), 4; (“b”), 4; (“b”, “c”), 4; (“c”), 4; (“c”, “a”), 3; (“c”, “b”), 3



a combination of basic sentiment elements is used to describe it. Moreover, the criterion is used for classifying concepts, not words. If a word has different senses, these senses can be seen as different sentiment elements.

The following explains the information in Table 2:

- Polarity is normally divided into three categories—positive, negative, and neutral. In most cases, positive polarity and negative polarity are considered.
- Although objective entities or behaviors are not subjective, they can arouse polarity, which strongly relies on human experiences and social norms; thus, some entities

Table 2 Basic subjectivity

Subclass	Positive	Negative	Neutral
Emotion	高兴 <i>gaoxin</i> "happy", 轻松 <i>qingsong</i> "relaxed"	愤怒 <i>fennu</i> "angry", 悲伤 <i>beishang</i> "sad"	
Desire			思念 <i>sinian</i> "miss", 想 <i>x-iang</i> "want to", 希望 <i>xi-wang</i> "hope", 期待 <i>qidai</i> "expect"
Cognition	相信 <i>xiangxin</i> "believe", 明白 <i>mingbai</i> "know it clearly", 理解 <i>lijie</i> "understand"	怀疑 <i>huaiyi</i> "doubt", 困惑 <i>kunhuo</i> "confuse", 不解 <i>bujie</i> "wonder"	吃惊 <i>chijing</i> "surprise", 认为 <i>renwei</i> "think", 看上去 <i>kanshangqu</i> "look like", 相同 <i>xiangtong</i> "same", 常见 <i>changjian</i> "common", 空前 <i>kongqian</i> "unprecedented"
Evaluation	聪明 <i>congming</i> "smart", 结实 <i>jieshi</i> "strong", 热情 <i>reqing</i> "passionate", 和睦 <i>hemu</i> "concord", 正义 <i>zhengyi</i> "justice"	愚蠢 <i>yuchun</i> "stupid", 垃圾 <i>laji</i> "garbage", 吵杂 <i>chaoza</i> "noise", 残酷 <i>canku</i> "cruel"	一般 <i>yiban</i> "just so so", 普通 <i>putong</i> "just so so"
Physiology	饱 <i>bao</i> "full", 暖 <i>nuan</i> "warm", 柔软 <i>rouruan</i> "soft", 悦耳 <i>yue'er</i> "euphonic"	饿 <i>e</i> "hungry", 渴 <i>ke</i> "thirsty", 困 <i>kun</i> "sleepy", 累 <i>lei</i> "tired", 晕眩 <i>yunxuan</i> "faint"	

or behaviors bear polarity. As the classification of objectivity is incomplete, we left this to specific domains.

- Contextual polarity, such as the words 大 *da* “big,” 小 *xiao* “small,” 多 *duo* “many,” 少 *shao* “few,” and 大胆 *dadan* “fearless,” is not included in Table 2. These words are polar, but it is hard to distinguish if no enough context is provided. Therefore, we did not regard these words as neutral, but instead as potential polar words.
- The emotion subclass can be positive or negative, but not neutral. As such, there is no need to describe this subclass without polarity.
- The desire subclass has special subjectivity and is only neutral.
- The cognition subclass mainly contains some concepts related to thought, but not to emotion, evaluation, desire, or physiology. Although most concepts of cognition are neutral, some may bear negative polarity, for example, 怀疑 *huaiyi* “doubt”.
- The physiology subclass is similar to the emotion subclass and thus cannot be neutral. The main difference between them is that emotion is cognition-triggered, whereas physiology is not.

Furthermore, many sentiment concepts are a combination of basic sentiment elements. Since we only considered prior polarity in our experiments, and the categories listed in Table 2 are sufficient for annotating prior polarity, we will not discuss these concepts in this paper.

2.3 Selecting polarity clues

2.3.1 Selecting polar clue candidates from corpora

For any given corpus, researchers need to obtain a domain-oriented sentiment lexicon to cover all the sentiment clues. It has been proven in many experiments that general sentiment lexicons are insufficient for domain-related tasks. Normally, general sentiment lexicons have low coverage, and different standards of word segmentation can affect the usage of general sentiment lexicons. As a result, researchers have to construct ways to mine data from a domain corpus to obtain a high-quality sentiment lexicon. In Chinese, degree adverbs can be used to find possible sentiment clues (i.e., words or phrases) for annotation. Although this process takes some time to complete, such processing guarantees both coverage and precision.

2.3.2 Chinese degree adverbs

In English, the use of degree adverbs is flexible, for example:

- I am very happy. (The degree adverb “very” appears before the modified word “happy”.)
- I am the happiest person. (The degree adverb is replaced by the superlative degree.)
- I support you very much. (The degree adverb “very much” appears after the modified word “support”.)
- I support this project you proposed very much. (The degree adverb “very much” is far from the modified word “support”.)

In Chinese, this task seems much easier due to how the degree adverbs are used with sentiments.

Degree adverbs are frequent adverbs in Chinese, and they have a strong relationship with sentiments. Compared with English, Chinese degree adverbs have close, simple, and frequent combinations with modified words (mainly sentiment expressions). Except for some cases in which degree adverbs occur after modified adjectives; in most cases, the degree adverbs occur immediately before the modified words. This relationship between degree adverbs and modified words is suitable for large-scale automatic text analysis. For example, degree adverbs can be used to extract multi-word sentiment expressions, as shown in Examples 6 through 9 below:

(6) 非常适合老年人使用

feichang__shihe__laonian__ren__shiyong
 very__suitable__older__people__use
very suitable for older people

(7) 非常有创意

feichang__you__chuangyi
 very__have__creative
very creative

(8) 真是个垃圾

zhen__shi__ge__laji
 really__is__a__garbage
really a piece of garbage

(9) 很男人

hen__nanren
 very__man
very manly

In our experiments, we scanned through each corpus to extract all word sequences with less than four words, and then let the annotators decide whether a word sequence was a sentiment expression; if it was, they annotated its polarity.

2.3.3 Rules for selecting polarity clues

Several issues were considered when we selected polarity clues:

1. A polarity clue was any adjective or word sequence (i.e., word phrase) that appeared after degree adverbs.
2. Contextual polarity words, such as 大 *da* "big," 高 *gao* "high," 长 *chang* "long," etc., were not selected. If these words were included, the set of text lines that was inputted into the sequence mining algorithm would be much larger, and the density of polarity shifting in text lines would be smaller, thus annotating would be more inefficient. Therefore, only prior polarity clues were selected, whose polarity is independent of context.
3. Words denoting desire were not selected because desire is one type of polarity shifting, such as 希望 *xiwang* "hope," 建议 *jianyi* "recommend," etc.
4. Words that have lost their positive semantics through language development or in a specified domain were not regarded as polarity clues. For example, in chatting

about online sales, 朋友 *pengyou* “friend” and 美女 *meiniu* “beauty” could be understood as “buyer” and “girl,” respectively.

5. Words that contain opposite polarities, such as 悲喜 *beixi* “sorrow and happiness,” were not chosen as polarity clues.

For more details on sentiment definitions, please refer to Section 2.2.

2.4 Annotating polarity shifting patterns

After using a sequence mining algorithm to generate word sequence patterns, we asked the annotators to decide whether a word sequence displayed a polarity shifting pattern. To annotate polarity shifting pattern candidates (i.e., the output of the sequence mining algorithm), the annotators had three options:

1. Yes—the pattern was a polarity shifting pattern or it contained a polarity shifting pattern.
2. No—the pattern had nothing to do with polarity shifting; none of the words in the pattern is a part of a polarity shifting pattern.
3. Not sure—as the current pattern may be part of a polarity shifting pattern, longer patterns containing the current pattern would have to be found in the following rounds before making a decision.

Two issues warrant attention. First, since we focused on words denoting cognition (e.g., surprised), desire (e.g., hope, wish, should), possibility (e.g., probably, maybe), limitation (e.g., only), and comparison (e.g., had thought), it is very possible that these words would display a polarity shifting pattern. Second, a special type of polarity shifting was simply to use a negative clue to negate the positive clues, such as in Examples 10 and 11 below:

(10) 傻瓜喜欢
shagua__xihuan
stupid__like
the stupid like it

(11) 鬼相信
gui__xiangxin
ghost__believe
nobody believes

In our experiments, we did not regard such negative clues as polarity shifting patterns. After each round of annotating, we compressed a set of text lines for the next round of sequence mining:

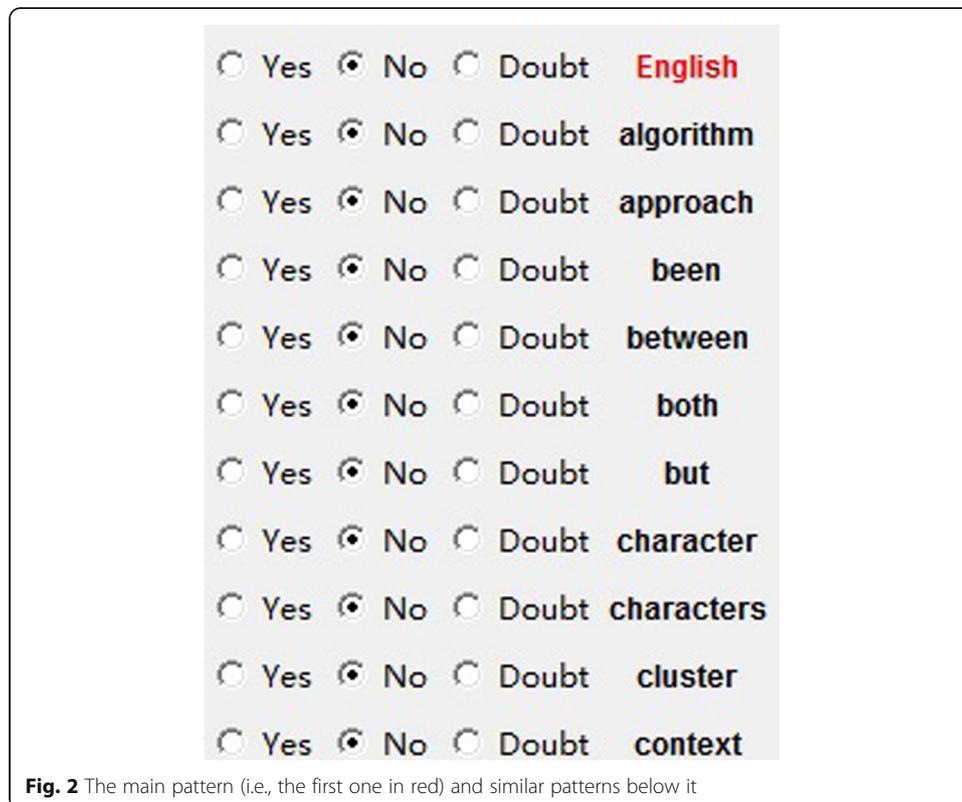
1. If a pattern contained a polarity shifting pattern, any line containing the pattern was removed from the set of text lines.
2. If a pattern had nothing to do with polarity shifting, or if any word did not contribute to detecting polarity shifting, it was removed from the text. If a line of text was too short after removing such words, it was removed from the set of text lines.

In both cases, the set of text lines was compressed, making the task easier. For example, the pattern 不X *bu X* “not X” was a polarity shifting pattern, any text line that contains the pattern is deleted from the set of text lines. In another case, the pattern 桌子的X *zhuozi de X* “X of table” had nothing to do with polarity shifting, so 桌子 *zhuozi* “table” and 的 *de* “of” were deleted from the set of text lines. In both cases, we compress the set of text lines.

For the “not sure” cases, the decision was made based on the next round (or rounds) of sequence mining. For example, for the pattern 以为X *yiwei X* “think X,” it could not be determined whether it displayed a polarity shifting pattern or not, so it was annotated as “not sure”. In the next round of sequence mining, let us suppose that we meet a less frequent pattern 原本以为X *yuanben yiwei X* “had thought X,” which was a polarity shifting pattern, and it was annotated as “yes”.

2.5 Annotating more patterns at the same time

To further facilitate annotation, we used the Word2Vec^c model to help find patterns similar to a pattern. Here, we call the pattern to be annotated as main pattern, and its similar patterns can be found by using the Word2Vec model. Normally, most of similar patterns have the same annotation tag as the main pattern, so one click or keystroke can annotate many patterns (see Fig. 2). We used the default function of sentence similarity in the Word2Vec model. If the similarity between a pattern and the main pattern was larger than a given threshold, the pattern was displayed below the main pattern. In our experiments, we found that listing similar patterns when annotating the main pattern greatly improved annotation efficiency.



3 Experimental setting

We chose to use PrefixSpan (Pei et al. 2001) as the sequence mining algorithm. We set the MinSupRatio to 0.05 and the MinSup to 3. First, we extracted frequent patterns by multiplying the MinSupRatio by the number of lines and then gradually decreased the MinSupRatio to find less frequent patterns in the following rounds. When the algorithm was unable to mine further patterns by MinSup, we stopped the algorithm. Among the remaining text lines, there were still some polarity shifting patterns of low frequency that could be fully annotated manually if so desired.

3.1 Corpora

We used two domain-specific corpora—food reviews and product reviews—as the text sources for our experiments for the following reasons:

- Compared with other styles of texts, such as news, and technical reports, reviews contain more sentiment expressions, thus potentially more polarity shifting patterns.
- Due to their commercial value, reviews arouse interest in companies and organizations. Producers and sellers want to know how people evaluate a product and what they want as products and then improve products or services accordingly.

We used the Jieba package^d to perform word segmentation and parts-of-speech tagging on the corpora.

3.1.1 Food reviews

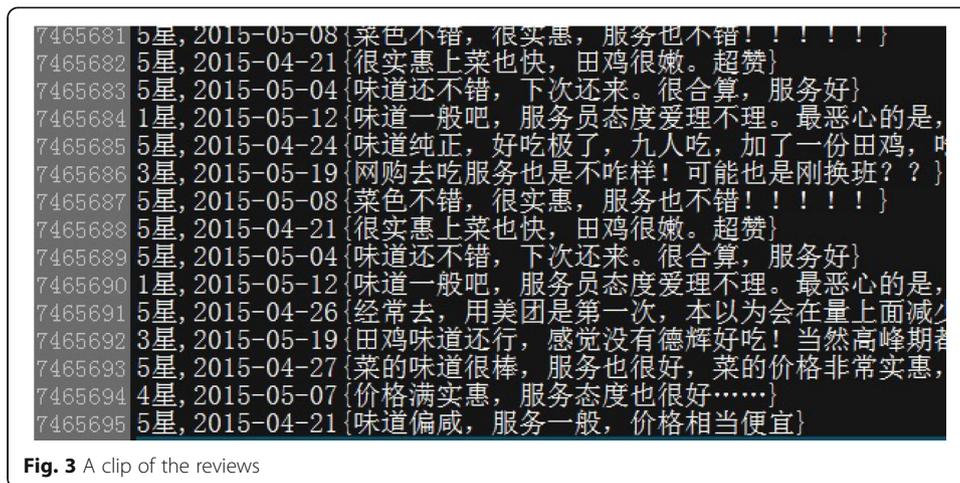
We collected food reviews from 美团 *meituan* (<https://github.com/fxsjy/jieba>), a famous online food shopping platform. In total, we collected 7,465,696 lines of reviews (400+ MB). An example of a line from the food reviews can be seen in Example 12 below:

(12)味道还不错,下次还来,很合算,服务好。
 weidao__hai__bucuo, xia__ci__hai__lai, hen__hesuan, fuwu__hao
 taste__yet__good, next__time__still__come, very__affordable, service__good
Good taste, will come again, worthwhile, good service.

Each review had three parts: the number of stars given by the reviewer (star score), the date of publication, and a comment. Of these three parts, the star scores and comments were of primary importance in our experiments. The number of reviews (7,356,695) was smaller than the number of lines (7,465,696) because some reviews took up more than one line. A short clip from the reviews is shown in Fig. 3, while the distribution of star scores is shown in Table 3 below:

As Table 3 shows, positive comments were dominant, which can be explained by the following:

1. Customers tended to choose restaurants according to their taste, so they were satisfied in most cases.



2. Many of the meals were purchased through Groupon or a Groupon-type application whose discounts are high, which might have led costumers to offer more favorable reviews on their experience.
3. Customers chose to give high ratings even if the meal was not quite perfect, which is consistent with the Pollyanna principle. Many of the comments stating that the dining experience was “just so-so” habitually granted five stars to the restaurant.

3.1.2 Product reviews

To compare the polarity shifting patterns extracted from different domains, we performed an experiment that shifted positive clues in bad comments using both food reviews and product reviews.

To construct the product review corpus, we collected bad comments on several best-selling products from 京东 *jingdong* (<http://www.jd.com>), a popular online shopping platform. We broke the comments down into sentences (by 。 !?), then phrases (by,;,:), leaving us with a total of 440,000+ lines of phrases.

3.2 Schemes for selecting polarity clues

There are generally three schemes used to select polarity clues (i.e., words or phrases) for annotation:

Table 3 Distribution of star scores

Stars	Frequency	Percentage (%)
5 stars	4,765,510	64.8
4 stars	1,642,601	22.3
3 stars	636,393	8.7
2 stars	163,289	2.2
1 star	148,902	2.0
Total	7,356,695	100

1. Select the N most frequent clues in the corpus to annotate, where N depends on how many hours a person is willing to spend on annotation. This scheme is the most time-consuming, but it also results in the highest coverage.
2. Intersect words in the corpus with polar words in a general sentiment lexicon.
In our experience, many domain-related sentiment words are not covered in general lexicons, which are low quality from the perspective of a specific domain. Normally, this scheme alone cannot satisfy experimental requirements.
3. Use aspects from both the first scheme and the second one, manually annotating the intersection set of the second scheme.

As the first scheme is very time-consuming and the second scheme cannot provide high-quality annotation, we chose the third scheme for our experiments, using existing Chinese sentiment lexicons to narrow the range of annotation and save on labor:

1. Hownet sentiment lexicon^e
 2. NTUSD (NTU Sentiment Dictionary)^f
 3. Emotion Ontology published by the Dalian University of Technology^g
- These three sentiment lexicons contain both positive and negative terms. We merged all the positive/negative terms as one set of positive/negative terms to help annotate positive/negative clues in our corpora.

4 Results and Discussion

4.1 Experiment I: shifting positive clues in bad comments

4.1.1 Basic steps

Given a corpus, this experiment:

1. Initialize PS (the set of polarity shifting patterns) to be empty.
2. Perform segmentation and parts-of-speech tagging on the corpus, and the result is C .
3. Annotate whether polarity clues (please see Section 3.2 for selecting polarity clues) are prior positive, and the result is POSLIST.
4. Break C into phrases, and the result is PHRASE.
5. Retrieve all lines that contain at least one element in POSLIST from PHRASE, and all positive clues in these lines are replaced by POSTAG. The result is POSLINE.
6. Given a minimum support, run the sequence mining algorithm on POSLINE, and the output of the algorithm is PATLIST.
7. If PATLIST is empty or minimum support is lower than a threshold, terminate the experiment.
8. Prune PATLIST by keeping the patterns that contain at least one positive word in POSLIST.
9. Annotate PATLIST for polarity shifting patterns, which are added to PS; compress PHRASE.
10. Reduce the minimum support, and go to step 6.

We used the method introduced in Section 2.3 to generate polar clue candidates and then used the criterion introduced in Section 2.2 to select the prior positive clues that expressed positive sentiments independent of context; if such a positive clue occurred

in a bad comment, it was highly possible that there is a polarity shifting pattern. For example, 好吃 *haochi* “delicious” was a prior positive word, so if 好吃 *haochi* “delicious” occurred in a bad comment, such as 不好吃 *bu haochi* “not delicious,” it was possible that 不*X* *bu X* “not *X*” was a potential polarity shifting pattern, where *X* denoted a positive word.

However, a contextual word made it harder to find a polarity shifting pattern. For example, 大 *da* “big” is a contextual polar word. If it is classified as a positive word, then it should fail in helping to find a polarity shifting pattern. For example, 噪声很大 *zaosheng hen da* “the noise is very big” is a bad comment, but 很*X* *hen X* “very *X*” is not a polarity shifting pattern. For this reason, we did not include contextual polar words in the three experiments in this study.

4.1.2 Why were bad comments considered?

According to the Pollyanna theory, people are inclined to use positive words, so in free texts, there are more “like” than “hate” and more “good” than “bad”. Furthermore, when expressing negative concepts, people are inclined to use the negation of a positive expression to alleviate sentiments, instead of using negative expressions directly. Having witnessed such a phenomenon (see Table 4), we hypothesized that polarity shifting would show a higher correlation with positive words than with negative words. In experiment I, we only considered polarity shifting on positive words.

4.1.3 Retrieving text lines with positive clues

After obtaining a set of prior positive clues, we used it to find text lines that contained at least one positive clue. The reason for performing this step was to obtain texts in which polarity shifting patterns were dense. Of course, one could skip the steps described in Section 2.3 and in this section and directly annotate polarity shifting in any given corpus. However, according to our experience, the possibility that annotators will detect polarity shifting patterns in an arbitrary unlabeled corpus is quite low. Therefore, to obtain a specified corpus where polarity shifting patterns occur frequently, this experiment used bad comments with at least one positive clue, assuming that polarity shifting inverted or attenuated the polarity of the positive clue.

4.1.4 Product reviews

We collected about 19,000 lines of product reviews that contained at least one positive clue. The main statistics of the annotating process are shown in Table 5 below:

Table 5 shows that only 1065 (i.e., the sum of the last column) patterns needed to be annotated, of which some were automatically annotated according to their nesting relationship. The whole process of annotation took about one and half hours. Without using this

Table 4 Pollyanna phenomenon

Phrase	Google hits
不喜欢 <i>bu xihuan</i> “not like”	16,000,000
不讨厌 <i>bu taoyan</i> “not hate”	587,000
不好 <i>bu hao</i> “not good”	30,300,000
不坏 <i>bu huai</i> “not bad”	914,000

Table 5 Process of Annotating

Round	MinSupRatio	Number of phrases	Number of patterns
1	0.05	18,936	16
2	0.05	10,818	2
3	0.025	9,849	9
4	0.0125	9,602	29
5	0.0125	8,688	7
6	0.0125	8,480	2
7	0.00625	8,374	48
8	0.00625	7,652	8
9	0.00625	7,510	1
10	0.003125	7,498	92
11	0.003125	6,713	14
12	0.0015625	6,610	159
13	0.0015625	5,135	48
14	0.0015625	4,879	30
15	0.00078125	4,766	392
16	0.00078125	3,695	208

approach, we would have had to annotate approximately 15,000+ (18,936–3695) lines of phrases, which would have taken much longer to accomplish.

After the 16th round of annotation, the minimum support was 3, no frequent patterns were extracted, and the last 3695 phrases mainly contained low-frequency words. By our observation, most of the phrases were composed of domain-dependent nouns and verbs. Furthermore, we assumed that adjectives should not be part of the polarity shifting pattern, so we removed nouns, verbs, and adjectives from the phrases and only kept phrases longer than two words; thus, we obtained a set of 616 phrases. Next, we applied fully manual annotating. After we broke down the bad comments from sentences to phrases, we found that not every phrase was a bad comment. Of the 616 phrases, many did not display polarity shifting. We will not report the detailed annotations of these 616 phrases in this paper because the detected polarity shifting patterns in these phrases were both low frequency and long. Finally, we categorized the polarity shifting patterns extracted from the above iterative annotating process into nine classes, which are shown in Table 6 below:

4.1.5 Food reviews

We chose the same amount of lines from the food reviews that contained a positive clue and had only one star^h. The details of the annotation process for the food reviews will not be reported in this paper because they were quite similar to the annotation process for the product reviews. For a comparison, the extracted polarity shifting patterns were also categorized into nine classes, as shown in Table 7 below:

Compared with the results from the product reviews, we found that the patterns in Table 7 were fewer and simpler than those in Table 6. One reason for this may be that the costumers who wrote the food reviews were less emotional than those who wrote the product reviews, as explained from the following three perspectives:

Table 6 Results of shifting positive clues to bad comments in the product reviews (X = positive)

Negation	真不X <i>zhenbu</i> X “really not X”, 再也不能X <i>zai ye buneng</i> X “cannot X any more”, 不够X <i>bugou</i> X “not X enough”, 再也不会X <i>zai ye buhui</i> X “never X again”, 不该X <i>bugai</i> X “should not X”, 没什么X <i>mei shenme</i> X “not X”, 不敢X <i>bugan</i> X “dare not X”, 不X <i>bu</i> X “not X”, 不X! <i>bu!</i> “not X!”, 非X <i>fei</i> X “not X”, 没法X <i>meifa</i> X “cannot X”, 不是X <i>bushi</i> X “be not X”, 别X <i>bie</i> X “do not X”, 没有X <i>meiyou</i> X “there is no X”, 不能X <i>buneng</i> X “cannot X”, 不要X <i>buyao</i> X “do not X”, 没X <i>mei</i> X “there is no X”, 不会X <i>buhui</i> X “will not X”, 不 怎么X <i>bu zenme</i> X “not X very much”
Doubt	怎么X啊 <i>zenme</i> X a “why X”, 搞不懂X <i>gaobudong</i> X “cannot understand X”, 奇怪X <i>qiguai</i> X “be weird that X”, X? X? “X?”, 为什么X <i>weishenme</i> X “why X”, 怀疑X <i>huaiyi</i> X “doubt that X”, 是不是X <i>shibushi</i> X “is X or not”, 怎么会X! <i>zenme shi!</i> “how could it be X!”, 是否X <i>shifou</i> X “is X or not”, 难道X <i>nan-dao</i> X “don’t tell me X”, 如何X <i>ruhe</i> X “how to X”, 是否是 <i>shifoushi</i> X “is X or not”
Desire	望X <i>wang</i> X “hope X”, 不想 <i>buxiang</i> X “do not want X”, 希望X <i>xiwang</i> X “hope X”, 辜负X <i>gufu</i> X “fail to live up to X”, 枉费X <i>wangfei</i> X “waste X”, 但愿X <i>danyuan</i> X “wish X”, 枉X <i>wang</i> X “waste X”
Comparison	原先X <i>yuanyuan</i> X “X in the past”, 原以为X <i>yuanyiywei</i> X “had thought X”, 以前X啊 <i>yiqian</i> X a “X in the past”, 之前X 才 <i>zhiqian</i> X <i>cai</i> “only X in the past”, 本来以为X <i>benlai yiywei</i> X “had thought X”, 本来应该X <i>benlai yinggai</i> X “should have been X”, 本以为 <i>benyiywei</i> X “had thought X”
Transition	才X <i>cai</i> X “X unless”, X就是 <i>jiushi</i> “X but”, 虽然X <i>suiran</i> X “although X”, X但是 <i>danshi</i> “X but”, X 但 <i>dan</i> “X but”, X 可是 <i>keshi</i> “X but”, 就算X <i>jiusuan</i> X “even if X”, X不过 <i>wuguo</i> “X but”
Condition	如果X <i>ruguo</i> X “if X”, 若X <i>ruo</i> X “if X”, 要是X <i>yaoshi</i> X “if X”
Limitation	只有X <i>zhiyou</i> X “only X”, 勉强X <i>mianqiang</i> X “manage X with difficulty”, 才能X <i>caineng</i> X “X only if”, 唯一X <i>weiyi</i> X “only X”, 仅仅X <i>jinjin</i> X “only X”, 要不是X <i>yaobushi</i> X “if not X”, X而已 <i>eryi</i> “only X”
Uncertainty	也许X <i>yexu</i> X “maybe X”, 像是X <i>xiangshi</i> X “seems X”
Others	失去X <i>shiqu</i> X “lose X”, 过于X <i>guoyu</i> X “too much X”, 怪我太X <i>guaiwo tai</i> X “blame me X too much”

- The money spent for a meal in China per person was about 30~100 RMBs on average, but was hundreds of RMBs or more for a product. If buyers were unsatisfied with the food or the product, those who spent more money became angrier and spent more time on writing a bad review.
- The purchasing and reviewing of food was mainly accomplished via mobile phones, while purchasing and reviewing products was mainly done on personal computers, which made writing longer and versatile comments more convenient.
- Food sellers and buyers were in the same city, and it was easier for buyers to know about a food seller through recommendations by friends and relatives. So, although

Table 7 Results of shifting positive clues to bad comments in the food reviews (X = positive)

Negation	不X <i>bu</i> X “not X”, 不够X <i>bugou</i> X “not X enough”, X不够 <i>X bugou</i> “not X enough”, 不能X <i>buneng</i> X “cannot X”, 不如X <i>buru</i> X “worse than X”, 不是X <i>bushi</i> X “be not X”, 不算X <i>busuan</i> X “not X”, 不想X <i>buxiang</i> X “not want X”, 没X <i>mei</i> X “there is no X”, 不会X <i>buhui</i> X “will not X”, 不 怎么X <i>bu zenme</i> X “not X very much”, 没什么X <i>mei shenme</i> X “not X”, 没啥X <i>mei sha</i> X “not X”, 没有X <i>meiyou</i> X “there is no X”, 不愿X <i>buyuan</i> X “wish not X”, 无X <i>wu</i> X “there is no X”
Doubt	是不是X <i>shibushi</i> X “is X or not”, X? X? “X?”, 为什么X <i>weishenme</i> X “why X”, 什么X啊 <i>shenme</i> X a “not X”, X 吗 <i>X ma</i> “what X”
Desire	应该X <i>yinggai</i> X “should X”
Comparison	以前X <i>yiqian</i> X “X in the past”, 从前X <i>congqian</i> X “X in the past”, X现在 <i>X xianzai</i> “X but now”, 之前X <i>zhiqian</i> X “X in the past”
Transition	X 但 <i>X dan</i> “X but”
Condition	如果X <i>ruguo</i> X “if X”
Limitation	很少X <i>henshao</i> X “rarely X”, X而已 <i>X eryi</i> “only X”
Uncertainty	
Others	

unsatisfied, the gap between expectations and reality was not very big. In contrast, product sellers and buyers were normally in different cities, and many good comments to buyers were faked and can be purchased. In this case, the gap between expectations and reality was sometimes very big.

These comparisons suggest that polarity shifting in different domains varies in terms of quantity and type. These results confirm that the best way to construct polarity shifting patterns is from the corpora themselves.

4.2 Experiment II: shifting negative clues in good comments

Although we hypothesized that polarity shifting would occur more frequently when negating or canceling a positive clue, we needed to find out how polarity shifting behaved in a good comment when customers shifted negative clues. We used the corpus of food reviews to perform this experiment.

4.2.1 Basic steps

Given a corpus, this experiment:

1. Initialize PS (the set of polarity shifting patterns) to be empty.
2. Perform segmentation and parts-of-speech tagging on the corpus, and the result is C.
3. Annotate prior negative clues (please see Section 3.2 for selecting polarity clues), and the result is NEGLIST.
4. Break C into phrases, and the result is PHRASE.
5. Retrieve all lines that contain at least one element in NEGLIST from PHRASE, and all the negative clues in the lines are replaced by NEGTAG. The result is NEGLINE.
6. Given a minimum support, run the sequence mining algorithm on NEGLINE, the output of the algorithm is PATLIST.
7. If PATLIST is empty or minimum support is lower than a threshold, terminate the experiment.
8. Prune PATLIST by keeping the patterns that contain at least one element in NEGLIST.
9. Annotate PATLIST for polarity shifting patterns, which are added to PS; compress PHRASE.
10. Reduce the minimum support, and go to step 6.

Step 5 was similar to the fifth step in Section 4.1, so please refer to it for more details. We used the method introduced in Section 2.3 to generate polar clue candidates, and then used the criterion introduced in Section 2.2 to select the prior negative clues, which expressed negative sentiments independent of context.

Once a negative clue occurred in a good comment, it was highly possible that a polarity shifting pattern was met, as shown in Example 13 below:

- (13) 谁说难吃?
 shui__shuo__nanchi
 who__say__unpalatable
Who said that it has a bad taste?

In Example 13, 难吃 *nanchi* “bad taste” is a prior negative word, so if 难吃 *nanchi* “bad taste” occurred in a good comment, such as 谁说X? *shuishuo X?* “who said that it is X?,” then this indicates a possible polarity shifting pattern, where X denotes any polar word. Similarly, we did not include contextual words in this experiment.

4.2.2 Experimental results

Some of the polarity shifting patterns are shown in Table 8. Compared with Table 7, we found that the transition class was used more frequently when shifting negative clues in good comments than when shifting positive clues in bad comments. In particular, 就是 *jiushi* “only that” as a kind of concession seemed to occur only in Table 8, which suggests that in a positive conversation circumstance, shifting to the negative needs to be done more carefully and euphemistically. In contrast, when shifting positive clues in bad comments, the speaker has no desire to maintain the conversation and will use more straightforward ways (other than a transition) to shift from good to bad. In addition, we found many attenuating patterns in this experiment, which confirms that in a positive conversation circumstance, shifting to the negative needs to be done more carefully and euphemistically. This phenomenon is consistent with the Pollyanna principle.

4.3 Experiment III: shifting between positive clues and negative clues

4.3.1 Basic steps

In this experiment, we wanted to see what the polarity shifting looked like when a positive clue and a negative clue occurred in the same line. Given a corpus, to find the polarity shifting patterns in it, our approach:

1. Initialize PS (the set of polarity shifting patterns) to be empty.
2. Perform segmentation and parts-of-speech tagging on the corpus, and the result is C.
3. Annotate prior positive clues and prior negative clues (please see Section 3.2 for selecting polarity clues), and the results are POSLIST and NEGLIST respectively.
4. Retrieve all lines that contain at least one element in POSLIST and one element in NEGLIST, then all the positive/negative clues are replaced by X/Y. The result is POSNEGLINE.

Table 8 Results of shifting negative clues to good comments (X = negative)

Negation	不X <i>bu X</i> “not X”, 没有X <i>meiyou X</i> “there is no X”, 不能X <i>buneng X</i> “cannot X”, 不是X <i>bushi X</i> “be not X”, 不算X <i>busuan X</i> “not X”, 没X <i>mei X</i> “there is no X”, 不会X <i>buhui X</i> “will not X”, 没啥X <i>mei sha X</i> “not X”, 不用X <i>buyong X</i> “need not X”, 无X <i>wu X</i> “there is no X”
Doubt	为什么X <i>weishenme X</i> “why X”, 是不是X <i>shibushi X</i> “is X or not”
Desire	可惜X <i>kexi X</i> “it is a pity that X”, 希望X <i>xiwang X</i> “hope X”
Comparison	以前X <i>yiqian X</i> “X in the past”, 从前X <i>congqian X</i> “X in the past”
Transition	就是X <i>jiushi X</i> “only that X”, 虽然X <i>suiran X</i> “although X”, X 但是X <i>danshi X</i> “X but”, X但X <i>dan X</i> “X but”, 但X <i>dan X</i> “but X”, X不过X <i>buguo X</i> “X but”, 不过X <i>buguo X</i> “but X”
Condition	
Limitation	只X <i>zhi X</i> “only X”, 除了X <i>chule X</i> “except X”, 唯一X <i>weiyi X</i> “only X”
Attenuator	有点X <i>youdian X</i> “somewhat X”, 有些X <i>youxie X</i> “somewhat X”, 稍微X <i>shaowei X</i> “a bit X”
Uncertainty	可能X <i>keneng X</i> “maybe X”, 好像X <i>haoxiang X</i> “seems X”
Others	

5. Given a minimum support, run the sequence mining algorithm on POSNEGLINE; the output of the algorithm is PATLIST.
6. If PATLIST is empty or minimum support is lower than a threshold, terminate the experiment.
7. Prune PATLIST by keeping the patterns that contain at least one element in POSLIST and one element in NEGLIST.
8. Annotate PATLIST for polarity shifting patterns, which are added to PS; compress POSNEGLINE.
9. Reduce the minimum support, and go to step 5.

As we already obtained the POSLIST and the NEGLIST from experiments I and II, we did not need to annotate them again. A positive clue and a negative clue often occurred between phrases, so we did not break down the sentences into phrases in this experiment. Furthermore, there were too many lines that contained at least a negative word and a positive word, so we limited our corpora to the first 1,000,000 lines, resulting in a 27-MB file that contained 318,000 lines. Then, we used sequence mining to extract the polarity shifting pattern candidates and annotate them. Some of the extracted polarity shifting patterns are shown in Table 9 below:

Table 9 Results of shifting between positive clues and negative clues (X= positive; Y= negative)

Negation	X没什么Y <i>X meishenme</i> Y "X not Y", X没有Y <i>X meiyou</i> Y "X not Y", Y不如X <i>Y buru</i> X "Y worse than X", Y不算X <i>Y busuan</i> X "Y not X", Y没X <i>Y mei</i> X "Y not X", X不Y <i>X bu</i> Y "X not Y", X不会Y <i>X buhui</i> Y "X will not Y", X不是Y <i>X bushi</i> Y "X is not Y", Y不是X <i>Y bushi</i> X "Y is not X", X没Y <i>X mei</i> Y "X not Y", 没XY <i>mei X Y</i> "not X Y", Y不够X <i>Y bugou</i> X "Y not X enough", X不能Y <i>X buneng</i> Y "X cannot Y", X不够Y <i>X bugou</i> Y "X not Y enough", X不用Y <i>X buyong</i> Y "X not Y", X不算Y <i>X busuan</i> Y "X not Y", X没Y <i>X mei</i> Y "X not Y", Y没什么X <i>Y meishenme</i> X "Y not X", Y没X <i>Y mei</i> X "Y not X", Y不怎么X <i>Y buzhenme</i> X "Y not very X", 不能X <i>Y buneng</i> X Y "cannot X Y", 没什么X <i>Y meishenme</i> X Y "not X Y", 不够X <i>Y bugou</i> X Y "not X enough Y", 不怎么X <i>Y buzhenme</i> X Y "not X enough Y", 不算X <i>Y busuan</i> Y X "not Y X", X不要Y <i>X buyao</i> Y "X not want Y"
Doubt	
Desire	Y希望X <i>Y xiwang</i> X "Y hope X", X可惜Y <i>X kexi</i> Y "X it is a pity Y", Y建议X <i>Y jianyi</i> X "Y suggest X", Y要是X <i>Y yaoshi</i> X "Y if X"
Comparison	Y其他X <i>Y qita</i> X "Y others X", X唯一Y <i>X weiyi</i> Y "X only Y", 以前XY <i>y-iqian</i> X Y "X in the past Y", X以前Y <i>X yiqian</i> Y "X Y in the past", X其他Y <i>X qita</i> Y "X others Y", Y以前X <i>Y yiqian</i> X "Y X in the past", Y其它X <i>Y qita</i> X "Y others X", 不如XY <i>buru</i> X Y "worse than X Y", 以前YX <i>y-iqian</i> Y X "Y in the past X", 之前XY <i>zhiqian</i> X Y "X in the past Y", X后来Y <i>X houlai</i> Y "X Y in the later", X之前Y <i>X zhiqian</i> Y "X Y in the past", X平时Y <i>X pingshi</i> Y "X Y usually", X本来Y <i>X benlai</i> Y "X Y originally", Y之前X <i>Y zhiqian</i> X "Y X in the past", Y别的X <i>Y biede</i> X "Y other than X", 上次XY <i>shangci</i> X Y "last time X Y", 平时XY <i>pingshi</i> X Y "usually X Y", 平时Y <i>X ping-shi</i> Y X "usually Y X", 本来XY <i>benlai</i> X Y "originally X Y"
Transition	X就是Y <i>X jiushi</i> Y "X only that Y", X不过Y <i>X buguo</i> Y "X but Y", Y不过X <i>Y buguo</i> X "Y but X", X但是Y <i>X danshi</i> Y "X but Y", Y还是X <i>Y haishi</i> X "Y still X", Y但X <i>Y dan</i> X "Y but X", Y但是X <i>Y danshi</i> X "Y but X", X虽然Y <i>X suiran</i> Y "X although Y", X可是Y <i>X keshi</i> Y "X but Y", X但Y <i>X dan</i> Y "X but Y", 虽然XY <i>suiran</i> X Y "although X Y", X可Y <i>X ke</i> Y "X but Y", X却Y <i>X que</i> Y "X but Y", Y可是X <i>Y keshi</i> X "Y but X", Y虽然X <i>Y suiran</i> X "Y although X", 虽然YX <i>suiran</i> Y X "although Y X"
Condition	X如果Y <i>X ruguo</i> Y "X if Y", Y如果X <i>Y ruguo</i> X "Y if X", 如果XY <i>ruguo</i> X Y "if X Y"
Limitation	X只是Y <i>X zhishi</i> Y "X only Y", X只Y <i>X zhi</i> Y "X only Y", 只是XY <i>zhishi</i> X Y "only X Y", 只是YX <i>zhishi</i> Y X "only Y X", 唯一YX <i>weiyi</i> Y X "only Y X", X除了Y <i>X chule</i> Y "X except Y", Y只X <i>Y zhi</i> X "Y only X", Y唯一X <i>Y weiyi</i> X "Y only X", 除了XY <i>chule</i> X Y "except X Y", 除了YX <i>chule</i> Y X "except Y X"
Attenuator	X有点Y <i>X youdian</i> Y "X somewhat Y", X稍微Y <i>X shaowei</i> Y "X somewhat Y", X略Y <i>X lue</i> Y "X a bit Y", 略YX <i>lue</i> Y X "a bit Y X", 稍微YX <i>shaowei</i> Y X "a bit Y X", X有些Y <i>X youxie</i> Y "X a bit Y", X稍Y <i>X shao</i> Y "X a bit Y"
Uncertainty	X好像Y <i>X haoxiang</i> Y "X seems Y", 看起来XY <i>kanqilai</i> X Y "seems X Y"
Others	

4.3.2 Experimental results

As can be seen in Table 9, there are no polarity shifting patterns in the doubt class¹. Suppose a negative clue is shifted by a “doubt” shifting pattern and a positive clue follows, the positive clue seems an answer to the doubt, which is like answering one’s own question. Such baffling expressions are rare in the literature. In contrast, there are many patterns in the comparison and transition classes. Since this experiment required that sentences include both a positive clue and a negative clue to qualify as having polarity shifting patterns, it is not difficult to understand why there were so many instances found in these two classes. This is a similar case for the negation class, as speakers can easily negate a polar clue and put it with a clue of opposite polarity, which is quite common in informal texts.

In the attenuator class in Table 9, only Y (i.e., negative) was attenuated, which is consistent with the results in Tables 6 through 8. Normally, people try to attenuate bad things instead of good things. Furthermore, there were some shifting patterns in the limitation class. Normally, a speaker selects the only positive/negative thing from the overall negative/positive comments.

The most important finding in experiment III is that the experiment did not need star rating information, which is sometimes hard to obtain. Therefore, we can apply our approach to *almost unlimited raw corpora* as long as a polarity lexicon is constructed. Although some classes (especially doubt) of polarity shifting patterns were not abundant in this experiment, it did provide most of the polarity shifting patterns.

5 Related work

In Wiegand et al. (2010), the authors presented a survey on the role of negation in sentiment analysis and analyzed it based on negation modeling, the scope of negation, and the detection of negation. The negation discussed in Wiegand et al. (2010) included modal, possibility, diminisher, etc. The work also offered an interesting observation—“not bad” is not the same as “good”—and showed that sentiment analysis is a very subtle and hard task.

In Polanyi and Zaenen (2004), the different types of negations were modeled on contextual valence shifting. The model assigned scores to polar expressions (i.e., positive scores for positive polar expressions and negative scores for negative polar expressions, respectively). If a polar expression was negated, its polarity score was simply inverted. Moreover, Kennedy and Inkpen (2005) evaluated a negation model that was similar to the one proposed by Polanyi and Zaenen (2004).

Some research has reported performance improvements when considering polarity shifting and its tasks. Pang et al. (2002) reported improvements by adding artificial features to plain words in which negation was not considered. Ikeda et al. (2008) presented a model to detect polarity shifting in sentences, which improved sentiment classification. To improve the performance of sentiment classification, Morsy and Rafea (2012) considered different categories of contextual valence shifters (i.e., intensifiers, negators, and polarity shifters) and frequency information. The results of all these experiments showed a significant improvement from the baselines in terms of accuracy, precision, and recall, which indicates that the proposed feature sets were effective in document-level sentiment classification. In Li et al. (2013), the authors performed experiments on a multi-domain sentiment classification corpus and the Cornell movie

review dataset. After all the trigger words were used, the improvement was found to be significant (5.3 higher on average). Moreover, improvements from the baseline were all significant (at the 0.01 level).

In Li et al. (2010), the authors manually checked 100 sentences that were explicitly polarity-shifted and offered the following eight types of polarity shifting structures:

- Explicit negation (not, no, without)
- Contrast transition (but, however, unfortunately)
- Implicit negation (avoid, hardly)
- False impression (look, seem)
- Likelihood (probably, perhaps)
- Counterfactual (should, would)
- Exception (the only)
- Until (until)

This work offered a rough distribution of polarity shifting in the corpora and some typical polarity shifting patterns. However, for a practical system, the patterns manually extracted from only 100 sentences are not enough. Our work aimed to extract more polarity shifting patterns from massive text corpora.

Boubel et al. (2013) presented an automatic approach to extracting contextual valence shifters (i.e., polarity shifting patterns). The system depended on two resources in French, a corpus of reviews, and a lexicon of valence terms to build a list of French contextual valence shifters. The work had a target similar to our approach, which aimed to extract Chinese polarity shifting patterns. However, our work did not use any syntactic parsers, and the extracted patterns were word sequences other than a term in a syntactic tree, which makes our approach more general.

6 Conclusions

Polarity shifting is a phenomenon in sentiment analysis that denotes the inversion, attenuation, or cancelation of the sentiment of a polarity clue (i.e., word or phrase) in a sentence. A polarity shifting pattern is a pattern that performs polarity shifting. Previous research has shown that the detection of polarity shifting can improve the performance of sentiment analysis. Currently, researchers often manually check a corpus line by line to find sentences containing polarity shifting and then extract the polarity shifting pattern. However, this process is time-consuming and may miss many polarity shifting patterns.

This study used an iterative annotating approach to find polarity shifting patterns from massive text corpora, which greatly reduced the labor used for annotating. To find the polarity shifting patterns, our approach followed the steps below:

1. Select text lines from a corpus where shifting occurs frequently. For example, if a positive word occurs in a bad comment, normally, polarity shifting should occur.
2. Use sequence mining to generate frequent polarity shifting pattern candidates. In the candidates, all the polar words (positive/negative) were generalized to variables (X/Y) so we could focus on the patterns of polarity shifting instead of the polar words.

3. Annotate polarity shifting patterns. Sentences were annotated as “yes” for polarity shifting patterns, “no” for sequences that had nothing to do with polarity shifting, and “not sure” for uncertainty.
4. Repeat steps 2 and 3 until no more frequent patterns are generated.

In choosing polar words for our approach, we avoided contextual polarity words, such as 大 *da* “big” and 高 *gao* “high,” and words from the desire class because it is difficult for an algorithm to decide the polarity (positive/negative) of desire words or contextual polarity words without contexts, which would make it hard to decide whether polarity shifting existed in a sentence.

The highlights of this approach are as follows:

- We generated frequent word sequences for annotating, since we knew that polarity shifting patterns are short word sequences normally containing 1~3 words. For a frequent polarity shifting pattern, we only needed to annotate it once, instead of many times when annotating line by line, thus saving a great deal of human labor.
- We did not need to modify the criterion of annotation that was used for annotating the corpora in a fully manual way (i.e., checking each corpus line by line). Researchers using our approach can design their own annotation criterion.

We performed three experiments to extract polarity shifting patterns:

- I: Shifting positive clues in bad comments
- II: Shifting negative clues in good comments
- III: Shifting between positive clues and negative clues

To perform our experiments, we collected food comments and product comments from the Internet, where polar words are abundant. Most important of all, each comment had a star score (from 1~5), which means that customers had already annotated bad comments or good comments for us, so we could filter the comments we needed for various experiments directly.

Wiegand et al. (2010) reported that “not bad” is different from “good,” although people think that “not good” is almost the same as “bad”. We wanted to see the difference between these two different kinds of polarity shifting, so we performed the first and the second experiments. Furthermore, after seeing how the transition between “good” and “bad” behaved, we performed the third experiment. We did not perform experiments in a machine-learning scenario because we focused on the linguistic analysis of polarity shifting patterns in this study. Furthermore, experiments on performance improvement via polarity shifting have been shown in works such as Pang et al. (2002), Kennedy and Inkpen (2005), and Li et al. (2013).

The results were analyzed and compared, and some interesting phenomena were reported. For example, 就是 *jiushi* “only that” as a kind of concession occurred only in the experiment that shifted negative clues in good comments. Furthermore, in all the experiments only *Y* (i.e., negative) was attenuated. Both cases showed that representing the “negative” needs to be done more carefully and

euphemistically. Compared with the fully manual method, our approach detected subtler polarity shifting as we extracted far more polarity shifting patterns. For more details, see Sections 4.1, 4.2, and 4.3. In the future, we would like to generalize this approach to cover similar annotation tasks.

7 Endnotes

^aFor ease of reading, 不 *bu* is translated as “NEG” in the separate examples, and as “not” in the running text and tables

^bIn this paper, *X* and *Y* are used to denote positive and negative clues, respectively, as noted.

^c<https://pypi.python.org/pypi/gensim>

^d<https://github.com/fxsjy/jieba>

^e<http://www.keenage.com/download/sentiment.rar>

^f<https://docs.google.com/forms/d/e/1FAIpQLSe20EyOE3bp9cKT0gF6R4DodTHOm-riIGegkGYa03oHYejhi9g/viewform?c=0&w=1>

^g<https://docs.google.com/forms/d/e/1FAIpQLSe20EyOE3bp9cKT0gF6R4DodTHOm-riIGegkGYa03oHYejhi9g/viewform?c=0&w=1>

^hWe considered a review with one star to be a bad review.

ⁱWe did not annotate all the text lines but only the frequent ones, which could be the reason for the missing of polarity shifting patterns in the doubt class.

Acknowledgements

This research is supported by the National Natural Science Foundation of China (No. 61300156 and No. 61300152) and the Education Office Foundation of Fujian Province in China (No. JA13257).

Authors' contributions

GX designed and implemented the study, while C-RH provided linguistic modeling and interpretation of the result. All authors read and approved the final manuscript.

Competing interests

The authors declare that they have no competing interests.

Author details

¹Department of Computer Science, Minjiang University, Fuzhou 350108, China. ²Fujian Provincial Key Laboratory of Information Processing and Intelligent Control, Fuzhou, China. ³Internet Innovation Research Center of Humanities and Social Sciences Base of Colleges and Universities in Fujian, Fuzhou, China. ⁴Department of Chinese and Bilingual Studies, The Hong Kong Polytechnic University, Kowloon, Hong Kong.

Received: 16 August 2015 Accepted: 2 June 2016

Published online: 15 September 2016

References

- Boubel, Noémi, Thomas François, and Hubert Naets. 2013. Automatic extraction of contextual valence shifters. In Proceedings of Recent Advances in Natural Language Processing (RANLP), ed. Galia Angelova, Kalina Bontcheva, and Ruslan Mitkov, 98–104. Shoumen, Bulgaria: INCOMA Ltd.
- Ikeda, Daisuke, Hiroya Takamura, Lev-Arie Ratinov, and Manabu Okumura. 2008. Learning to shift the polarity of words for sentiment classification. In Proceedings of the Third International Joint Conference on Natural Language Processing (IJCNLP), 296–303. Hyderabad, India: International Institute of Information Technology, India.
- Kennedy, Alistair, and Diana Inkpen. 2005. Sentiment classification of movie and product reviews using contextual valence shifters. In Proceedings of FINEXIN-05, Workshop on the Analysis of Informal and Formal Information Exchange during Negotiations, 11–22. Ottawa, Canada.
- Li, Shoushan, Sophia Yat Mei Lee, Ying Chen, Chu-Ren Huang, and Guodong Zhou. 2010. Sentiment classification and polarity shifting. In 23rd International Conference on Computational Linguistics (COLING), ed. Chu-Ren Huang and Dan Jurafsky, 635–643. Beijing, China.
- Li, Shoushan, Zhongqing Wang, Sophia Yat Mei Lee, and Chu-Ren Huang. 2013. Sentiment classification with polarity shifting detection. In Proceedings of the 2013 International Conference on Asian Language Processing (IALP), ed. Guohong Fu, Haoliang Qi, Minghui Dong, Min Zhang, Yusufu Aibaidula, and Weimin Pan, 129–132. IEEE.
- Morsy, Sara A, and Ahmed Rafea. 2012. Improving document-level sentiment classification using contextual valence shifters. In Natural Language Processing and Information Systems, ed. Gosse Bouma, Ashwin Ittoo, Elisabeth Métais, and Hans Wortmann, 253–258. Springer.

- Ortony, Andrew, Gerald L. Clore, and Mark A. Foss. 1987. The referential structure of the affective lexicon. *Cognitive Science* 11: 341–364.
- Pang, Bo, Lillian Lee, and Shivakumar Vaithyanathan. 2002. Thumbs up? Sentiment classification using machine learning techniques. In *Proceedings of the 2002 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 79–86. Stroudsburg, PA: Association for Computational Linguistics.
- Pei, Jian, Jiawei Han, Behzad Mortazavi-Asi, Helen Pinto, Qiming Chen, Umeshwar Dayal, and Mei-Chun Hsu. 2001. PrefixSpan: Mining sequential patterns efficiently by prefix-projected pattern growth. In *Proceeding of the 17th International Conference on Data Engineering (ICDE01)*, 215–224. IEEE.
- Polanyi, Livia, and Annie Zaenen. 2004. Contextual lexical valence shifters. In *Proceedings of the AAAI Spring Symposium on Exploring Attitude and Affect in Text: Theories and Applications*, 106–111. AAAI.
- Quirk, Randolph, Sidney Greenbaum, Geoffrey Leech, and Jan Svartvik. 1985. *A comprehensive grammar of the English language*. Longman.
- Strapparava, Carlo, and Alessandro Valitutti. 2004. Wordnet-affect: An affective extension of Wordnet. In *Proceedings of the 4th International Conference on Language Resources and Evaluation*, ed. Maria Teresa Lino, Maria Francisca Xavier, Fátima Ferreira, Rute Costa, Raquel Silva, Carla Pereira, Filipa Carvalho, Milene Lopes, Mónica Catarino, and Sérgio Barros, 1083–1086. Paris: European Language Resources Association.
- Wiegand, Michael, Alexandra Balahur, Benjamin Roth, Dietrich Klakow, and Andre's 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*, ed. Roser Morante and Carolinne Sporleder, 60–68. Stroudsburg, PA: Association for Computational Linguistics.

Submit your manuscript to a SpringerOpen[®] journal and benefit from:

- ▶ Convenient online submission
- ▶ Rigorous peer review
- ▶ Immediate publication on acceptance
- ▶ Open access: articles freely available online
- ▶ High visibility within the field
- ▶ Retaining the copyright to your article

Submit your next manuscript at ▶ springeropen.com
