

Construction of a Sentimental Word Dictionary

Eduard C. Dragut
Purdue University
edragut@purdue.edu

Clement Yu, Prasad
Sistla
Dept. of Computer Science
University of Illinois at Chicago
{yu, sistla}@cs.uic.edu

Weiyi Meng
Dept. of Computer Science
Binghamton University
meng@cs.binghamton.edu

ABSTRACT

The Web has plenty of reviews, comments and reports about products, services, government policies, institutions, etc. The opinions expressed in these reviews influence how people regard these entities. For example, a product with consistently good reviews is likely to sell well, while a product with numerous bad reviews is likely to sell poorly. Our aim is to build a sentimental word dictionary, which is larger than existing sentimental word dictionaries and has high accuracy. We introduce rules for deduction, which take words with known polarities as input and produce synsets (a set of synonyms with a definition) with polarities. The synsets with deduced polarities can then be used to further deduce the polarities of other words. Experimental results show that for a given sentimental word dictionary with D words, approximately an additional 50% of D words with polarities can be deduced. An experiment is conducted to find the accuracy of a random sample of the deduced words. It is found that the accuracy is about the same as that of comparing the judgment of one human with that of another.

Categories and Subject Descriptors

H.3 [INFORMATION STORAGE AND RETRIEVAL]:
Content Analysis and Indexing—*Dictionaries*

General Terms

Algorithm, Experimentation

Keywords

Sentimental word dictionary, Deduction, WordNet

1. INTRODUCTION

Opinions affect people from all walks of life. A person who wants to buy a product would choose one having excellent reviews and stay away from a product with poor reviews. Our objective is to build a dictionary of sentimental words, which facilitates the analysis and retrieval of opinionated

texts. The proposed dictionary has the following *majority sentiment property*: each word with a given part of speech has polarity p (positive or negative) if the majority sense of the word with that part of speech has polarity p . Existing sentimental word dictionaries do not possess this property. This property is significant, because it is rather natural for humans to classify the polarity of a word in such a manner. More importantly, it permits the deduction of other sentimental words with the same property.

Opinion mining (e.g., [15, 6, 10, 8]) and opinion retrieval (e.g., [16, 17, 9]) have significant practical interest. We aim to construct a sentimental word dictionary, which is of high quality and has a large coverage of sentimental words. Such a dictionary is essential, because it facilitates both opinion mining and opinion retrieval. Most existing sentimental word dictionaries (*General Inquirer* [11], *Appraisal Lexicon* [13], *Opinion Finder* [15]) are far from being exhaustive and do not possess the majority sentiment property. *SentiWordNet* [6] gives a score of positivity, a score of negativity and a score of objectivity to each synset, with the sum of scores = 1. Since these scores are obtained by classifiers and through several iterations of extending the seed sets of data, its accuracy is not high. For example, for the synset

{*iniquity, immorality, evil, wickedness*} (*morally objectionable behavior*),

SentiWordNet assigns the degrees $P = 0.75$, $N = 0$, $O = 0.25$, where P stands for positive polarity, N stands for negative polarity and O stands for neutral polarity. Most people would agree that it has a negative instead of a positive polarity. Quite a few such examples exist in *SentiWordNet*. Other researchers [1, 12] consider words to have degrees of positivity/negativity. We choose to have a single polarity associated with each word for the reason explained earlier.

The contributions of this paper are:

- By using a deduction approach, the resulting sentimental word dictionary contains approximately 50% more words than a given sentimental word dictionary;
- The accuracy of the deduced polarities is reasonably high; it is comparable to that of human judgment.

The rest of the paper is organized as follows. Section 2 defines the problem of constructing a sentimental word dictionary. Section 3 describes our approach of constructing the dictionary. Section 4 is the experimental results. Related works are given in section 5. Section 6 is the conclusion.

2. SENTIMENTAL WORD DICTIONARY

Our approach of constructing a sentimental word dictionary is as follows. We start from a *seed dictionary* D , i.e., a small dictionary of sentimental words where each word in

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

CIKM'10, October 26–30, 2010, Toronto, Ontario, Canada.
Copyright 2010 ACM 978-1-4503-0099-5/10/10 ...\$10.00.

the dictionary has the *majority sentiment property*. Second, we attempt to deduce the polarities of as many additional sentimental words as possible, based on D .

In practice, we construct the sentimental dictionary on top of the electronic dictionary WordNet [7]. The goal is to annotate each word (and its synsets) in WordNet with their corresponding polarities. The idea is to start from a small number of words whose polarities are known and to use the relationships between words in the dictionary to deduce the polarities of the rest of the words. In this work two words are related if they share at least a synset. The deduction process is the main contributions of this paper.

WordNet distinguishes between four part of speeches: noun, verb, adjective and adverb. It groups words into sets of synonyms called *synsets*. Each word with a part of speech has a set of senses, with each sense defined by a synset. A sense has a short, general definition and some examples. An example synset with a definition is:

(1)S: (adj) {bland, flat, flavorless, flavourless, insipid, savourless, savourless, vapid} (lacking taste or flavor or tang) “a bland diet”; “insipid hospital food”; “flavorless supermarket tomatoes”; “vapid beer”; “vapid tea”

This is one of the three senses of the word **bland**. “1” is the frequency count of the synset w.r.t. **bland**. The part delimited by { and } is the synset, that delimited by (and) is the definition and the rest are the examples.

Note that the synset in the example is shared by all the words present in the synset. Hence, the words **flat**, **flavorless**, etc. have the synset among their synsets. The synset may however have different frequency counts with respect to each word. For instance, its frequency count is 0 with respect to the word **flat**. Whenever a synset has a 0 frequency count we replace it with 0.1. Usually the senses with rare usages of a word have a 0 count. Increasing their frequency counts by a small constant is a standard smoothing technique. If two senses of a word have the same synset, then they can be differentiated by their definitions. To ease the presentation, we will not mention the part of speech and the definition of a word, and assume that the sense of each word is characterized by its synset.

We assume that each synset has a unique polarity. For example, the word **bland** has three senses, with the first two senses being negative and the last sense being positive. A synset of a word has a frequency of use, indicating the “probability” that the word is used in the sense conveyed by the synset. Suppose the frequencies of use of the word **bland** in the three senses are f_1 , f_2 and f_3 respectively. For instance, the frequency of the sense of the word **bland** described above is 1. Then, the *relative frequency* of the synset in the first sense, which denotes the probability that the word is used in the first sense is $f_1/(f_1 + f_2 + f_3)$. The probability that the word expresses a negative sentiment is $(f_1 + f_2)/(f_1 + f_2 + f_3)$. If this probability is greater than .5, then the majority sense of the word is negative. This property allows the polarities of other words to be automatically deduced. The following definition formally states this property.

DEFINITION 1. (Polarity) *Let w be a word and S its set of synsets. Each synset in S has an associated polarity and a relative frequency with respect to w . The word w has polarity p , $p \in \{\text{positive, negative}\}$ if there is a subset of synsets $S' \subseteq S$ such that each synset $s \in S'$ has polarity p and the sum of the frequencies of the synsets in S' is larger than 0.5. S' is called a **polarity dominant subset**. If there does not exist any such subset then the word has a neutral polarity.*

3. POLARITY INFERENCE FRAMEWORK

Given a seed dictionary D , the inference process can be classified into two categories: (1) determine the polarities of as many synsets as possible and (2) determine the polarities of as many additional words not in D as possible. Once the polarities of synsets are determined we employ Definition 1 (and other techniques) to determine the polarities of additional words. We first describe how the polarities of the synsets are inferred from the polarities of the words and then cover the deduction of polarities for words.

3.1 Synset Polarity Inference

We classify the synset inference rules into single word and multi-word inference rules. We first introduce three single word rules.

A synset of a word is called *dominant* if its relative frequency with respect to the word is larger than 0.5. Synsets enjoying this property are important because, if a synset is of polarity p and is dominant for the word w , then the word has polarity p . Conversely, if the word w is known to have polarity p and it has a dominant synset s , and possibly some other synsets, then s must have polarity p . This rule is one of the basic means of deducing the polarities of the synsets given the polarities of the words.

INFERENCE RULE 1. *Let w be a word having a dominant synset s . If w has polarity $p \in \{\text{positive, negative}\}$ then s has polarity p , too.*

INFERENCE RULE 2. *Let w be a word with polarity p and exactly two synsets. The synsets have the same relative frequency with respect to w (i.e., 0.5). If p is either positive or negative polarity then both synsets have polarity p .*

In inference rule 2, if either synset does not have polarity p , then the word w cannot have polarity p . Thus, both synsets must have polarity p .

As an illustration, the adjective **advance** has positive polarity in **General Inquirer** and it has two senses with identical relative frequencies in WordNet. Hence, we deduce that both its synsets have positive polarities.

The next rule is for words with an arbitrary number of synsets. It partitions the set of synsets of a word into those whose polarities are known and those whose polarities are unknown. If the two sets satisfy certain constraints the polarities of all those synsets whose polarities are unknown can be determined.

DEFINITION 2. *Let w be a word and S be its set of synsets. A subset S' of S is a minimally dominant subset of synsets if the sum of the relative frequencies of the synsets in S' is larger than 0.5 and the removal of any synset s from S' will make the sum of the relative frequencies of the synsets in $S' - \{s\}$ smaller than 0.5.*

For example, the word **consummate** has three synsets in WordNet, each of them with the same relative frequency 1/3. It is easy to check that any two of its synsets form a minimally dominant subset (their sum of frequencies is 2/3 and $2/3 > 1/2$).

INFERENCE RULE 3. *Let w be a word with polarity p , negative or positive, and S be its set of synsets. Suppose S can be partitioned into two sets, S_1 and S_2 . S_1 is the set of the synsets whose polarities are known and are all different from p . S_2 is the set of the synsets whose polarities are either unknown or known but equal to p . If S_2 is a minimally dominant subset then all its synsets must have polarity p .*

The word **consummate** has positive polarity in Opinion Finder dictionary. Suppose one of its synsets has negative polarity. Since the other two synsets form a minimally dominant subset, they must each have positive polarity.

We give an example of an inference rule that involves two words sharing a number of synsets. We have defined and implemented 11 multi-word rules.

INFERENCE RULE 4. *Let w and v be two words, with polarities $\sim p$ (\sim is the negation symbol) and p , respectively, where $p \in \{positive, negative\}$. Neither w nor v have a dominant synset. w has $n \geq 3$ synsets $\{s_1, s_2, \dots, s_n\}$. v has 3 synsets $\{s_0, s_1, s_2\}$ (i.e., the two words share $\{s_1, s_2\}$). If the sum of the relative frequencies of s_1 and s_2 with respect to w is greater than 0.5, then the polarity of synset s_0 is p .*

Most of our rules involve words with up to four synsets. The explanation lies in the fact that more than 95% of the words in WordNet have up to four synsets. Additionally, all the words with five or more synsets share many of their synsets with words having up to four synsets.

3.2 Word Polarity Inference

The inference of polarities for words is as follows. The polarity of a word $w \notin D$ may be determined (i) from its underlying synsets, using Definition 1 and (ii) by comparing the subsets of synsets of w and those of a word with known polarity. The latter case is useful when there are not enough synsets with known polarities so that Definition 1 can be applied. This section covers this case.

PROPOSITION 1. *Let w and v be two words. Let S_w and S_v be the sets of synsets of w and v , respectively. Suppose that $S_w \subseteq S_v$. If every minimal dominant subset S_w is also a dominant subset in S_v then, if word w has polarity p , $p \in \{positive, negative\}$, then word v has polarity p , too.*

We noticed that this statement is satisfied in many instances by words with multiple spellings (e.g., British vs. American English): e.g., **brutalise** vs. **brutalize**, the two words have identical sets of synsets, i.e., $S_w = S_v$. Besides the multi-lingual spelling examples, there are also other motivating examples, e.g. the verbs **dehydrate** and **desiccate**. These two words have identical sets of synsets.

PROPOSITION 2. *Let w and v be two words. Let S_w and S_v be the sets of synsets of w and v , respectively. Suppose that $S = S_w \cap S_v \neq \emptyset$ and $S' = S_w - S_v \neq \emptyset$. The polarity of w is p , $p \in \{positive, negative\}$, and the polarities of all the synsets in S' are known and all are different from p . S must contain some dominant subset of w . If every minimal dominant subset of w in S is also a dominant subset w.r.t. v , then the polarity of word v is p , too.*

The results can be used to derive clusters of words with the same polarity. Suppose that there are pairs of words $(w, v_1), (w, v_2), \dots, (w, v_n)$ such that each pair satisfies the conditions of Proposition 1. If w has polarity p , then v_1, v_2, \dots, v_n have polarity p . The same principle applies to word pairs satisfying the conditions of Proposition 2. This increases the number of words whose polarities can be determined.

4. EXPERIMENTAL RESULTS

4.1 Data sets

The data sets used in our experiments consist of: WordNet [7] and the sentimental dictionaries **General Inquirer** [11], **Appraisal Lexicon** [13] and **Opinion Finder** [15]. We use

POS	Words	Synsets	Op. Finder	Gen. Inq.	App. Lex.
Noun	117,798	82,115	1,907	1,444	2
Verb	11,529	13,767	1,501	1,041	0
Adjective	21,479	18,156	2,608	1,188	1,440
Adverb	4,481	3,621	775	51	317
Total	155,287	117,659	6,791	3,724	1,759

Table 1: Distribution of words and synsets

POS	Input Words	Inferred Words	Inferred Synsets
Noun	2,315	1,460	1,683
Verb	1,617	844	1,079
Adjective	2,937	1,407	1,907
Adverb	925	364	430
Total	7,794	4,075	5,099

Table 2: Inference result using the union dictionary.

WordNet 3.0, whose statistics are given in Table 1. The table shows the distribution of the words and synsets per part of speech. Columns 2 and 3 pertain to WordNet. For example, there are 21,479 adjective words, which have 18,156 synsets. In total WordNet 3.0 has 155,287 words and 117,659 synsets.

Data cleaning

The three sentimental dictionaries are organized in triplets of the form $\langle word, pos, polarity \rangle$ (where pos stands for part of speech), e.g., $\langle bland, Adjective, negative \rangle$. **General Inquirer** has 3,994 distinct triplets, **Appraisal Lexicon** has 1,888 distinct triplets and **Opinion Finder** has 8,223 distinct triplets. But not all their original entries are useable. The main reason is that a triplet $\langle w, pos, p \rangle$ is present in a sentimental word dictionary, but (1) the word w does not appear in WordNet (regardless of pos) or (2) w with pos is not present in WordNet.

After cleaning, as shown in Table 1, there are 3,724 entries in **General Inquirer**, 1,759 entries in **Appraisal Lexicon** and 6,791 entries in **Opinion Finder** which appear in WordNet. Herein, whenever we refer to these dictionaries we refer to their cleaned versions.

4.2 Automatic Discovery of Sentimental Words

One of the main experimental results is the deduction of polarities for additional words using the polarities of the words in the *disagreement-free* union of the three dictionaries. A triplet $\langle w, pos, p \rangle$ is included into the *disagreement-free* union dictionary if (1) $\langle w, pos, p \rangle$ is in one of the three dictionaries and there does not exist another triplet $\langle w, pos, p' \rangle$ in another dictionary such that $p \neq p'$ or (2) $\langle w, pos, x \rangle$ is in one of the three dictionaries, but in a different dictionary there is $\langle w, pos, y \rangle$ such that $x \neq y$; in which case, we determine the polarity of $\langle w, pos \rangle$ to be p by consulting WordNet and applying Definition 1.

The result of inferring the polarities of new words and synsets using the union dictionary is summarized in Table 2. The table breaks down the outcome by part of speeches. For example, we are given the polarities of 2,315 nouns (out of 7,794 words) and we deduce the polarities of 1,460 new nouns. In the process we deduce the polarities of 1,683 synsets. Overall, we are given the polarities of 7,794 words and we are able to deduce the polarities of 4,075 additional words and the polarities of 5,099 synsets. In other words, an additional 52.2% of the number of input words with polarities are deduced. This percentage seems to be a fair estimate of the inference power of our environment as confirmed by running the inference rules on the individual dictionaries. Specifically, we obtain 55% for **Opinion Finder**, 68% for **General Inquirer** and 54% for **Appraisal Lexicon**.

While from a practical point of view this number seems to be a fair characterization of our inference framework, analytically it is not quite fair. The ability to infer the polarities of new words is dependent on two aspects: (1) the structure of WordNet and (2) the distribution of the words in an input dictionary. We briefly discuss these issues next.

There are 40,077 *isolate words* in WordNet whose polarities cannot be automatically determined. These words have the property that their synsets are not shared with any other word in WordNet. In addition, the input words are *badly* distributed over WordNet, both per part of speech as well as overall. For instance, in WordNet there are 3,554 *non-trivial connected components* for the adjective part of speech and the input words fall into only 563 of them, which is less than 16%. A non-trivial connected component has at least two words and any two words (synsets) are connected to each other by paths. Overall, the distribution is far poorer. The words fall into less than 5% (1,443 of 33,015) of the non-trivial connected components. This has a negative implication: even if more rules are developed a large portion of WordNet still remains inaccessible to polarity inferencing.

4.3 Polarity Inference Accuracy

In Table 2, 4075 additional words with polarities were deduced. 100 words are randomly chosen from these 4075 words according to their distributions among nouns (22.2%), adjectives (37.5%), adverbs (11.5%) and verbs (28.8%). These 100 words are shown to 3 humans. They are asked to determine their polarities. If they are familiar with a word, they can determine its polarity directly using their experience; otherwise, they can utilize WordNet and Definition 1 to determine its polarity. The Kappa statistic [3] is used to measure the agreement, denoted agreement A-H (Automatic versus Human), between the deduced polarities of the words and those given by the humans. It is also used to measure the agreement, denoted H-H (Human versus Human), between the polarities of the words given by different humans. The agreement A-H is 63.3%, which is slightly larger than that of H-H, which is 62%. It should be noted that the deduction process does not give rise to errors. The “errors” only illustrate the discrepancies among the dictionaries.

If the Kappa statistics is over 75% the agreement is assumed to be excellent. The reason why the agreement accuracy among humans is relatively low, 62%, is that people have preconceived notions of the polarities of words/phrases. As an example, a phrase to be judged by humans is the verb *eat at*. A human may associate the phrase to mean “eating at a restaurant” and assigns either a neutral or a positive polarity. However, this phrase means “become deteriorate”, according to WordNet. Thus a human who consults WordNet assigns a negative polarity to it.

5. RELATED WORK

The problem of determining the polarities of words has been studied by many researchers (e.g., [5, 8, 9, 10, 15]). The general theme is that there are sets of seed words that have known polarities; e.g., the seed word “good” is positive and the seed word “bad” is negative. These sets of seed terms are grown by various means. For example, synonyms and antonyms [5] are used to expand the sets of seeds. Another technique to expand the seed sets is to apply the *pointwise mutual information measure* [2] to each target word t w.r.t. each seed word t_i . The pointwise mutual information measure is obtained from the co-occurrence frequency of t and t_i in documents retrieved by a search engine and the frequencies of the individual words in those documents. Machine learning algorithms [4] can be employed to classify words

into different polarities. According to [5], the performance of [4] is comparable or better than those in [10, 14].

Unlike *SentiWordNet*, our view is that each synset does not have a degree associated with each polarity. Instead, each synset is 100% positive, 100% negative or 100% neutral.

6. CONCLUSIONS

In this paper, we introduce the concept of deducing the polarities of words based on the polarities of other words. Experimental results show that the number of new words with polarities deduced is approximately 50% of the size of the original sentimental word dictionary. There are quite a few operators such as *hyponym*, *antonym* and *similar-to* in WordNet, which can be used for deduction, as pointed out in [5]. Intuitively, a hyponym synset of a synset with polarity p also has polarity p and an antonym of a word with polarity p has polarity $\sim p$. However, we found numerous exceptions to the intuition. We believe that we can find remedies to these situations. Operators such as antonyms and hyponyms behave differently from the inference rules described in Section 3. Specifically, the former type of operations do not require the words and synsets to be within a connected component for deductions to operate properly.

7. REFERENCES

- [1] A. Andreevskaia and S. Bergler. Mining wordnet for fuzzy sentiment: Sentiment tag extraction from wordnet glosses. In *EACL*, 2006.
- [2] K. W. Church and P. Hanks. Word association norms, mutual information, and lexicography. *Comput. Linguist.*, 16(1), 1990.
- [3] J. Cohen. A coefficient of agreement for nominal scales. *Educational and Psychological Measurement*, 20(1):37–46, April 1960.
- [4] A. Esuli and F. Sebastiani. Determining the semantic orientation of terms through gloss classification. In *CIKM*, pages 617–624, 2005.
- [5] A. Esuli and F. Sebastiani. Determining term subjectivity and term orientation for opinion mining. In *EACL*, 2006.
- [6] A. Esuli and F. Sebastiani. Sentiwordnet: A publicly available lexical resource for opinion mining. In *LREC*, 2006.
- [7] C. Fellbaum. Wordnet: An on-line lexical database and some of its applications. 1998.
- [8] V. Hatzivassiloglou and J. Wiebe. Effects of adjective orientation and gradability on sentence subjectivity. In *COLING*, 2000.
- [9] X. Huang and W. Croft. A unified relevance model for opinion retrieval. In *CIKM*, 2009.
- [10] J. Kamps, M. Marx, R. Mokken, and M. de Rijke. Using wordnet to measure semantic orientation of adjectives. In *LREC*, 2004.
- [11] P. Stone, D. Dunphy, M. Smith, and J. Ogilvie. The general inquirer: A computer approach to content analysis. In *MIT Press*, 1996.
- [12] P. Subasic and A. Huettnner. Affect analysis of text using fuzzy semantic typing. *TFS*, 29(4), 2001.
- [13] M. Taboada and J. Grieve. Analyzing appraisal automatically. In *AAAI Spring Symposium on Exploring Attitude and Affect in Text*, 2004.
- [14] P. D. Turney and M. L. Littman. Measuring praise and criticism: Inference of semantic orientation from association. *ACM Trans. Inf. Syst.*, 2003.
- [15] T. Wilson, J. Wiebe, and P. Hoffmann. Recognizing contextual polarity in phrase-level sentiment analysis. In *HLT/EMNLP*, 2005.
- [16] K. Yang, N. Yu, A. Valerio, and H. Zhang. Wudit in trec 2006 blog track. In *TREC*, 2006.
- [17] W. Zhang, C. Yu, and W. Meng. Opinion retrieval from blogs. In *CIKM*, pages 831–840, 2007.