High, Medium or Low? Detecting Intensity Variation Among polar synonyms in WordNet

Raksha Sharma and Pushpak Bhattacharyya
Dept. of Computer Science and Engineering
IIT Bombay, Mumbai, India
{raksha,pb}@cse.iitb.ac.in

Abstract

For fine-grained sentiment analysis, we need to go beyond zero-one polarity and find a way to compare adjectives (synonyms) that share the same sense. Choice of a word from a set of synonyms, provides a way to select the exact polarity-intensity. For example, choosing to describe a person as *benevolent* rather than *kind*[1] changes the intensity of the expression.

In this paper, we present a sense based lexical resource, where synonyms are assigned intensity levels, *viz., high, medium* and *low*. We show that the measure $P(s|w)$ (probability of a sense s given the word w) can derive the intensity of a word within the sense. We observe a statistically significant *positive correlation* between $P(s|w)$ and intensity of synonyms for three languages, *viz., English, Marathi* and *Hindi*. The average correlation scores are 0.47 for English, 0.56 for Marathi and 0.58 for Hindi.

1 Introduction

Sentiment analysis is a crucial task for various web and media outlets, such as, e-commerce websites, blogs and newspapers. The general approach of Sentiment Analysis is to summarize the semantic polarity (i.e., positive or negative) of sentences/documents (Riloff and Wiebe, 2003; Pang and Lee, 2004; Danescu-Niculescu-Mizil et al., 2009; Takamura et al., 2005; Baccianella et al., 2010a; Guerini et al., 2013). However, sentence intensity becomes crucial when we need to compare sentences having the same polarity orientation. In such scenarios, we can use intensity of words to judge the intensity of a sentence. Words that bear the same sense can be used interchangeably to upgrade or downgrade the intensity of the expression. The following example helps illustrate the problem we attempt to address.

- the synset (set of synonyms), {sound, levelheaded, intelligent, healthy} (Gloss: exercising or showing good judgment), are assigned a fixed positive polarity of *0.75* in SentiWordNet[2], while most people would agree that all the synonymous words are not equally positively intense. The use of *levelheaded* or *sound* makes a sentence more intensely positive in comparison to *healthy*, given that the sentence expresses the sense *exercising or showing good judgment*.

In addition to English, there exists polarity-intensity variation across synonyms in other languages also. Consider the following example from Hindi:

- The word सद्गुणी (Transliteration: Sadgunee, Translation: Virtuous) and लायक (Transliteration: Layak, Translation: Worthy) are synonymous words according to HindiWordNet for the sense *morally excellent*. Hindi native speakers confirm that the word सद्गुणी, is more intense than the word लायक in terms of polarity.

There are several manually or automatically created sense based lexical resources (Agerri and García-Serrano, 2010; Baccianella et al., 2010b) that assign the same positive or negative polarity to all synonymous words, making no distinction among them in terms of their intensity.

In this paper, we address the concept of polarity-intensity variation among synonyms and come up with a measure to predict the polarity intensity of a word for the given sense. We show that

---

[1]The words, *Benevolent* and *kind* are synonyms for the sense *well meaning and kindly* as per Oxford English dictionary.

there is a statistically significant *positive correlation* between $P(s|w)$ (probability of sense s given word w) and intensity of a word w within the sense s (Section 3). Hence, the measure $P(s|w)$ can be used to predict the intensity of a word within the sense. We extensively validate this *positive correlation* in three languages[3], *viz.*, English, Marathi and Hindi (Section 5). We observe a statistically significant *positive correlation* of 0.47 for English, 0.56 for Marathi and 0.58 for Hindi (Section 7).

Our Contribution: Our work contributes an automatically generated sense based lexical resource where words which belong to the same sense are assigned three intensity levels, *viz., high, medium* and *low*. This resource can be used to derive intensity information of a subjective sentence or document, which essentially empowers existing sentiment analysis systems. In addition to this, intensity information of words can be used to reduce or enhance an over-expressed or under- expressed text respectively.

2 Related Work and Discussion

Several researchers have made successful attempts for finding opinion words (Wiebe, 2000; Taboada and Grieve, 2004; Takamura et al., 2005; Wilson et al., 2005; Kanayama and Nasukawa, 2006; Liu, 2010; Dragut et al., 2010; Ohana and Tierney, 2009; Agerri and García-Serrano, 2010; Sharma and Bhattacharyya, 2013); however, finding intensity of words still considered as a challenging task.

There have been some works on scaling adjectives by their strength, independent of the sense they express. The first work in the direction of adjectival scale was done by Hatzivassiloglou and McKeown (1993). They exploited linguistic knowledge available in the corpora to compute similarity between adjectives. However, their approach did not consider polarity orientation of adjectives, they provided ordering among non-polar adjectives like, *cold, lukewarm, warm, hot*. Kim et al. (2013) demonstrated that vector off-set can be used to derive scalar relationship amongst adjectives. De Melo and Bansal (2013) used a pat-

Language	Variables	Cor-value
English	Annotator-1~TFC	-0.09
English	Annotator-2~TFC	-0.09
Hindi	Annotator-1~TFC	-0.04
Hindi	Annotator-2~TFC	-0.09
Marathi	Annotator-1~TFC	-0.11
Marathi	Annotator-2~TFC	-0.10

Table 1: Correlation between Total Frequency Count (TFC) and intensity score assigned by two annotators in each language.

tern based approach to identify intensity relation among adjectives, but their approach had a severe coverage problem. Ruppenhofer et al. (2014) provided ordering among polar adjectives that bear the same semantic property.

None of the existing works address intensity variation among synonyms. However, choice of a word from a set of synonyms provides a way to intensify the expression. Our approach pin-points the polarity-intensity variation across synonyms.

3 Polarity-intensity Variation And Synonymous Words

The classical *semantic bleaching theory*[4] states that a word which has *high* frequency of use tends to have *low* intensity in comparison to a word having less frequency of use. For example, the frequent use of the word *good* makes it less intense, while rare use of the word *great* makes it more intense (Kim and de Marneffe, 2013). However, *good* and *great* are not synonyms according to SentiWordNet. The *semantic bleaching phenomenon* throws light on the *positive* association between frequency and intensity regardless of any semantic relation (for example, *synonymy*). But, when we computed correlation between total-frequency (Section 5) and polarity-intensity within a sense (Section 6), we observed a negative correlation. Table 1 shows the correlation values obtained for three languages, *viz.,* English, Hindi and Marathi. The negative correlation shown in table 1, substantiates that total-frequency of a word cannot predict the polarity-intensity of a word within a particular sense. The *semantic bleaching* phenomenon compares total-frequency of words (sum

---

[3]A person who is a linguist as well as a native speaker of the language can annotate words with more accuracy. The availability of the linguists, who are also native speakers of Hindi and Marathi made us to choose these two languages other than English. Hindi and Marathi are two of the 23 official languages of India, which have approximately 258 and 73 million speakers respectively. English is chosen, because most of the lexical resources which we had pointed out in our work are in English only.

[4]The *semantic bleaching phenomenon* in words was reported in US edition of *New York Times*:http://www.nytimes.com/2010/07/18/magazine/18onlanguage-anniversary.html?_r=0

of the word's count in all its senses), while count of a word in the sense is the potential clue for intensity of the word in the sense.

In this paper, we define polarity-intensity variation among synonyms (words having the same sense), specifically for the polar words. Words having the same sense cannot be directly compared on the basis of their frequency count in the sense, because their total frequency of usage (total count of the words in all its senses) are different. We need a relative count of the synonymous words, that is, (*Count of the word with the sense / Total count of the word*).

The cause of overuse of a word is its use in multiple senses (Durkin, 2009). Therefore, use of a word in multiple senses increases the total frequency of use, but the word loses its frequency count with a particular sense relative to the total frequency count of the word. Considering this frequency distribution as a base, we hypothesize polarity-intensity variation among words belonging to a particular sense.

{*A word which has high relative frequency for a sense is high intense in comparison to a word which has low relative frequency for the sense.*}

Consider the following derivation that validates our proposed hypothesis for polarity-intensity variation across synonyms. According to the *semantic bleaching* phenomenon:

$$TC(w_1) < TC(w_2) \Rightarrow I(w_1) > I(w_2) \quad (1)$$

Where, TC is a function that gives total-count of a word and I is a function that gives intensity of a word.

Since w_2 has higher total-frequency (overused) than w_1, we can deduce that w_2 has more senses in comparison with w_1. Let us assume that w_1 and w_2 are synonyms for i^{th} sense and w_1 has only one sense, that is, i^{th} sense and w_2 has n ($n > 1$) senses. Now, we rewrite equation 1 in terms of count of words in their senses in equation 2. Here, $SC_{W_j}^k$ represents sense-count, that is, count of the word 'w_j' with the sense 'k'.

$$SC_{w_1}^1 < \sum_{i=1}^n SC_{w_2}^i \Rightarrow I(w_1) > I(w_2) \quad (2)$$

Now, to compare the synonymous words w_1 and w_2 in the i^{th} sense, we need their relative counts in the sense (Equation 3). Relative count is the count of the word with the sense divided by total-count of the word. Since w_1 has only one sense, so its

sense-count and total-count will remain the same. Hence, the fraction $\frac{SC_{w_1}^i}{SC_{w_1}^1}$ will be 1, which is the maximum possible value for the fraction. While, w_2 has more than one senses, so its sense-count will always be less than total-count. Hence, the fraction $\frac{SC_{w_2}^i}{\sum_{i=1}^n SC_{w_2}^i}$ will always be less than 1. On the other hand, w_1 has only one sense, so the intensity relation between w_1 and w_2, given by the *semantic bleaching phenomenon* will remain the same.

$$\frac{SC_{w_1}^i}{SC_{w_1}^1} > \frac{SC_{w_2}^i}{\sum_{i=1}^n SC_{w_2}^i} \quad (3)$$
$$\Rightarrow I^1(w_1) > I^i(w_2)$$

We observe a reversal of the sign $<$ to $>$ in case of relative frequency comparison of w_1 and w_2, but the intensity relation remains intact. Essentially, a word that shows its majority occurrence with the sense or has a higher relative frequency count, is more intense for the sense than the other synonymous words.

A few such instances of polarity-intensity variation in a sense are shown in table 2. We asked two linguists in each language to compare the polarity-intensity of the exemplified synonymous words for the given sense. They mutually agreed on the fact that the first word is more intense than the second word for the considered sense. The same intensity relation between the synonyms can be inferred from the relative frequency counts (sense-count/total-count) of words. The relative frequency count of the first word is higher than the second word for all the senses given in table 2. The total-count and sense-count values are obtained from English and Hindi sense annotated corpus (section 5).

4 Probability of Sense Given Word

Statistically, a relative frequency count of a word is nothing but the *probability of sense given word* ($P(s|w)$). The function $C(w_i, s_j)$ gives count of w_i with the sense s_j, while the function $C(w_i)$ gives total-count (aggregation of count in all senses). The measure $P(s_j|w_i)$ is defined as follows:

$$P(s_j|w_i) = P(s_j, w_i)/P(w_i)$$
$$= C(w_i, s_j)/C(w_i), \quad (4)$$
$$\text{Where, } C(w_i) = \sum_K C(w_i, S_k)$$

Synonymous-Words(w_1,w_2)	Sense-Definition	Total-Count(w_1)	Total-Count(w_2)	Sense-Count(w_1)	Sense-Count(w_2)
Awful, Painful	exceptionally bad or displeasing	10	14	1	12
Proficient, Good	Having or showing knowledge and skill and attitude	4	263	3	2
लाभदायक, उपयोगी (Translation: Beneficial, Useful)	Giving an advantage	22	35	22	15
उत्तम, अच्छा (Translation: Exquisite, Substantial)	Having or marked by unusual and impressive intelligence	181	270	181	1

Table 2: Examples of polarity-intensity variation from English and Hindi. In all cases, first word is more intense than the second word for the given sense.

Hence, we deduce that if a word possess higher value for the measure $P(s|w)$, then it is more intense than other synonymous words. Equation 5 generalizes the proposed hypothesis.

$$P(s|w_1) > P(s|w_2) \Rightarrow I^s(w_1) > I^s(w_2)$$

Where, w_1 and w_2 belong to the same sense s. (5)

In summary, when we compare words within a sense, we need to account for the participation of these words in other senses also. The proposed probabilistic measure, *probability of sense given word* considers the participation of a word in other senses also in the form of its total-count. We observe a statistically significant *positive correlation* between polarity-intensity levels assigned by linguists and the value of $P(s|w)$ (relative frequency of a word w in a sense s) (Section 7).

A high value of $P(s|w)$ is possible in the following scenarios.

- If w is rarely found with the sense s, then it should be rare in all.

- If w is very frequent, then the majority part of its total occurrences should be with the sense s only.

5 Dataset

We validate our hypothesis using three languages, *viz., English, Hindi*, and *Marathi*.

English: For English, we extracted all the adjective synsets whose polarity (positive or negative) value is greater than 0.5 as per SentiWordNet,

except the synsets that have only one word. We ignored the synsets having polarity values less than or equal to 0.5, considering them a weak candidate for polarity-intensity variation phenomenon. With the threshold value of 0.5, we extracted a total of 1116 synsets. However, SentiWordNet is an automatically compiled lexical resource, which assigns polarity values based on corpus dependent probabilistic measures. To make our English dataset potentially conclusive, we asked two linguists in English to manually inspect the polarity orientation of synsets (senses). Table 3 is a confusion matrix, that summarizes the results of manual inspection of English dataset extracted from SentiWordNet (SWN).

		Polarity Orientation in SWN		
		Negative	Positive	Objective
Actual	Negative	599	37	0
	Positive	77	311	0
	Objective	84	8	–

Table 3: Confusion matrix

A few examples of wrong polarity orientation by SentiWordNet are given in table 4. We considered the correct synsets for our experiment. Consequently, intensity ordering is demonstrated for 1024 $(1116 - 92)$ English synsets.

Hindi and Marathi: For Hindi and Marathi, we asked two linguists in each language to extract polar synsets (senses) from HindiWordNet and MarathiWordNet[5]. Manual extraction of

---

[5]Available at: http://www.cfilt.iitb.ac.in/wordnet/

Synonymous-Words	Sense-Definition	Polarity by SWN	Actual Polarity
Murderous, homicidal	Having a tendency towards killing another human beings	Positive(0.625)	Negative
Enthralled, entranced	Filled with wonder and delight	Negative(0.75)	Positive
Unmarried, Single	Not married	Negative(0.75)	Objective

Table 4: Examples of synsets (Senses), which are assigned wrong polarity by SentiWordNet.

senses were required, because HindiWordNet and MarathiWordNet do not have polarity information for synsets. The total number of observed senses and words in each language are specified in table 5.

Language	Senses	Words
English	1024	3397
Hindi	172	2614
Marathi	325	1346

Table 5: Observed synset statistics

C(w_i,s_j): For English words, the value of the function C is obtained from the English WordNet database file, that is, 'cntlist'[6]. For Hindi and Marathi, we used a sense marked corpus in tourism and health domain[7] (Khapra et al., 2010). The total number of sense marked words in each domain are depicted in table 6. If a word shows zero frequency of use for any particular sense, we replace it with 0.1 according to a standard smoothing technique (Han et al., 2006).

POS category	Tourism	Health
Noun	72932	52230
Verb	26086	24291
Adjective	32499	22699
Adverb	9820	855

Table 6: Hindi/Marathi sense marked corpus statistics

6 Gold Standard Data Preparation

We asked two linguists in each language to assign words to different intensity levels, *viz., high* (3), *medium* (2), and *low* (1) within a synset. A discrete scale with only three intensity levels is chosen to reduce the subjectivity issue in annotation, consequently complexity of annotation.

Consider the following example of synonymous words, where intensity levels are assigned by English linguists.

- Grievous (Intensity: 3) > dangerous (Intensity: 2) > serious (Intensity: 1) for the sense, *causing fear or anxiety by threatening great harm.*

Table 7 shows the inter annotator agreement for each language, computed using *weighted Cohen's kappa* measure.

Language	Inter-annotator Agreement
English	53%
Hindi	69%
Marathi	64%

Table 7: Weighted Cohen's Kappa in percent

7 Empirical Validation

We validate the hypothesized relation between polarity-intensity and probabilistic measure: $P(s|w)$ by finding Pearson product-moment correlation. To test the significance of correlation value, we perform a directional test, that is, *t-test* using *cor.test* function of R[8]. We obtain a statistically significant *positive correlation* between gold standard intensity levels and $P(s|w)$ for all the three languages. Table 8 shows the correlation values, t-values, p-values and confidence intervals. The statistically significant positive correlation parameter allows us to conclude that the polarity-intensity of a word in a sense can be inferred by the relative frequency ($P(s|w)$) of the word in the sense.

8 Error Analysis

The observed scenarios that affect the proposed hypothesis negatively are as follows.

---

[6]Detail available at:
http://wordnet.princeton.edu/wordnet/man/cntlist.5WN.html
[7]Available at: www.cfilt.iitb.ac.in.

[8]R is a language and environment for statistical computing and graphics. Detail available at: http://www.r-project.org/

Lang.	Variable	Cor-value	t-Value	p-value	95% Confidence-interval
English	$P(s\|w) \sim$ Linguist1	0.48	32.36	$< .0001$	0.46 to 0.51
	$P(s\|w) \sim$ Linguist2	0.44	28.96	$< .0001$	0.42 to 0.47
Marathi	$P(s\|w) \sim$ Linguist1	0.58	26.10	$< .0001$	0.54 to 0.62
	$P(s\|w) \sim$ Linguist2	0.53	22.91	$< .0001$	0.49 to 0.58
Hindi	$P(s\|w) \sim$ Linguist1	0.60	38.33	$< .0001$	0.56 to 0.63
	$P(s\|w) \sim$ Linguist2	0.55	33.66	$< .0001$	0.53 to 0.58

Table 8: Statistically significant correlation values with the results of t-test

1 There are words which do not have their all senses in WordNets. For instance, the word *bastard* as an adjective has only one sense, that is, *fraudulent; having a misleading appearance* as per WordNet, but according to Oxford dictionary, it has one more sense, that is, *born of parents not married to each other; illegitimate*. The exclusion of such senses leads to wrong total-count of the word in the English WordNet database file.

2 There are words which are not found in the corpus with any sense of them. In this case, besides the frequency count of the word with the sense, we fail to get evidence for total-count of the word.

In such cases, the probabilistic measure, that is, $P(s|w)$ fails to result in a strong value that can insinuate the correct polarity-intensity of a word, which leads to fall in the correlation estimate.

9 Conclusion

In this paper, we addressed the concept of polarity-intensity variation among synonyms. We show that the relative frequency of a word w in a sense s, that is, $P(s|w)$ is a predictor of polarity-intensity of the word in the sense. We present a sense based lexical resource in three languages, where polar synonyms are annotated with the intensity levels, *viz., high, medium* and *low*.

Manual checking of sentiment WordNets for intensity variation is a difficult endeavor. Therefore, a by-product of our polarity-intensity analysis is that sentiment WordNets can become more informative resource for sentiment analysis. In addition to this, intensity information of words can be used to reduce or enhance an over-expressed or under-expressed text respectively.

10 Acknowledgment

References

Rodrigo Agerri and Ana García-Serrano. 2010. Q-wordnet: Extracting polarity from wordnet senses. In *LREC*.

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010a. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *LREC*, volume 10, pages 2200–2204.

Stefano Baccianella, Andrea Esuli, and Fabrizio Sebastiani. 2010b. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *LREC*, volume 10, pages 2200–2204.

Cristian Danescu-Niculescu-Mizil, Gueorgi Kossinets, Jon Kleinberg, and Lillian Lee. 2009. How opinions are received by online communities: A case study on amazon.com helpfulness votes. In *Proceedings of the 18th International Conference on World Wide Web*, WWW '09, pages 141–150, New York, NY, USA. ACM.

Gerard De Melo and Mohit Bansal. 2013. Good, great, excellent: Global inference of semantic intensities. *Transactions of the Association for Computational Linguistics*, 1.

Eduard C Dragut, Clement Yu, Prasad Sistla, and Weiyi Meng. 2010. Construction of a sentimental word dictionary. In *Proceedings of the 19th ACM international conference on Information and knowledge management*, pages 1761–1764. ACM.

Philip Durkin. 2009. *The Oxford guide to etymology*. Oxford University Press.

Marco Guerini, Lorenzo Gatti, and Marco Turchi. 2013. Sentiment analysis: How to derive prior polarities from sentiwordnet. *arXiv preprint arXiv:1309.5843*.

Jiawei Han, Micheline Kamber, and Jian Pei. 2006. *Data mining: concepts and techniques*. Morgan kaufmann.

Vasileios Hatzivassiloglou and Kathleen R McKeown. 1993. Towards the automatic identification of adjectival scales: Clustering adjectives according to meaning. In *Proceedings of the 31st annual meeting on Association for Computational Linguistics*. Association for Computational Linguistics.

Hiroshi Kanayama and Tetsuya Nasukawa. 2006. Fully automatic lexicon expansion for domain-oriented sentiment analysis. In *Proceedings of the 2006 Conference on Empirical Methods in Natural Language Processing*, pages 355–363. Association for Computational Linguistics.

Mitesh M. Khapra, Anup Kulkarni, Saurabh Sohoney, and Pushpak Bhattacharyya. 2010. All words domain adapted wsd: Finding a middle ground between supervision and unsupervision. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, ACL '10, pages 1532–1541, Stroudsburg, PA, USA. Association for Computational Linguistics.

Joo-Kyung Kim and Marie-Catherine de Marneffe. 2013. Deriving adjectival scales from continuous space word representations. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1625–1630, Seattle, Washington, USA, October. Association for Computational Linguistics.

Bing Liu. 2010. Sentiment analysis and subjectivity. *Handbook of natural language processing*, 2:568.

Bruno Ohana and Brendan Tierney. 2009. Sentiment classification of reviews using sentiwordnet. In *9th. IT & T Conference*, page 13.

Bo Pang and Lillian Lee. 2004. A sentimental education: Sentiment analysis using subjectivity summarization based on minimum cuts. In *Proceedings of the 42Nd Annual Meeting on Association for Computational Linguistics*, ACL '04, Stroudsburg, PA, USA. Association for Computational Linguistics.

Ellen Riloff and Janyce Wiebe. 2003. Learning extraction patterns for subjective expressions. In *Proceedings of the 2003 conference on Empirical methods in natural language processing*, EMNLP '03, pages 105–112, Stroudsburg, PA, USA. Association for Computational Linguistics.

Josef Ruppenhofer, Michael Wiegand, and Jasper Brandes. 2014. Comparing methods for deriving intensity scores for adjectives. *EACL 2014*, 117.

Raksha Sharma and Pushpak Bhattacharyya. 2013. Detecting domain dedicated polar words. In *Proceedings of the Sixth International Joint Conference on Natural Language Processing*, pages 661–666, Nagoya, Japan, October. Asian Federation of Natural Language Processing.

Maite Taboada and Jack Grieve. 2004. Analyzing appraisal automatically. In *Proceedings of AAAI Spring Symposium on Exploring Attitude and Affect in Text (AAAI Technical Re# port SS# 04# 07), Stanford University, CA, pp. 158q161. AAAI Press*.

Hiroya Takamura, Takashi Inui, and Manabu Okumura. 2005. Extracting semantic orientations of words using spin model. In *Proceedings of the 43rd Annual Meeting on Association for Computational Linguistics*, ACL '05, pages 133–140, Stroudsburg, PA, USA. Association for Computational Linguistics.

Janyce Wiebe. 2000. Learning subjective adjectives from corpora. In *AAAI/IAAI*, pages 735–740.

Theresa Wilson, Paul Hoffmann, Swapna Somasundaran, Jason Kessler, Janyce Wiebe, Yejin Choi, Claire Cardie, Ellen Riloff, and Siddharth Patwardhan. 2005. Opinionfinder: A system for subjectivity analysis. In *Proceedings of hlt/emnlp on interactive demonstrations*, pages 34–35. Association for Computational Linguistics.