Identifying Synonyms among Distributionally Similar Words

Dekang Lin and Shaojun Zhao
Department of Computing Science
University of Alberta
Edmonton, AB, Canada, T6G 2E8

Lijuan Qin and Ming Zhou
Microsoft Research Asia
No.49, Zhichun Road
Hai Dian District, Beijing, China, 100080

Abstract

There have been many proposals to compute similarities between words based on their distributions in contexts. However, these approaches do not distinguish between synonyms and antonyms. We present two methods for identifying synonyms among distributionally similar words.

1 Introduction

The distributional hypothesis states that words with similar meanings tend to appear in similar contexts [Harris, 1968]. Consider the words *adversary and foe.* Both of them are often used as the objects of the verbs:

batter, crush, defeat, demonize, deter, outsmart,...

and modified by the adjectives:

ardent, bitter, formidable, old, tough, worthy,...

There have been many proposals for computing distributional similarity of words [Hindle, 1990; Pereira *et al,* 1993; Lin, 1998]. The list (1) shows the top-20 distributionally similar words of *adversary,* obtained with Lin's method [Lin, 1998] on a 3GB newspaper corpus.

(1) *adversary:* enemy, foe, ally, antagonist, opponent, rival, detractor, neighbor, supporter, competitor, partner, trading partner, accuser, terrorist, critic, Republican, advocate, skeptic, challenger

Compared with manually compiled thesauri, distributionally similar words often offer much better coverage. Compare (1) with the entry for *adversary* in Webster Collegiate Thesaurus [Kay, 1988]:

(2) adversary: Synonyms: opponent, antagonist, anti, con, match, opposer, oppugnant; Related Words: assaulter, attacker Contrasted Words: backer, supporter, upholder; Antonyms: ally

The thesaurus entry missed many synonyms such as: *enemy, foe, rival, competitor* and *challenger.*

A problem with the distributionally similar words, however, is that many of them are antonyms, e.g., *ally* and *supporter* in (1). The problem gets worse if a word belongs to a semantic category with many members, since all of them tend to have similar distributions. This is demonstrated in the following list of top-20 distributionally similar words for *orange.*

(3) *orange:* yellow, lemon, peach, pink, lime, purple, tomato, onion, mango, lavender, avocado, red, pineapple, pear, blue, plum, cucumber, melon, turquoise, tangerine

In many applications, such as information retrieval and machine translation, the presence of antonyms or other types of semantically incompatible words (e.g., orange-pink) can be devastating. This paper presents two methods for identifying synonyms among distributionally similar words.

2 Methods

2.1 Patterns of Incompatibility

Consider the following phrasal patterns:

(4) a. from *X to Y*
b. either *X* or *Y*

If two words *X* and *Y* appear in one of these patterns, they arc very likely to be semantically incompatible. For example, the following table shows the queries and the hits (the number of return documents) from the search engine AltaVista:

(5)

Query	
adversary NEAR ally	
"from adversary to ally"	
"from ally to adversary"	19
"either adversary or ally"	1
"either ally or adversary"	2
adversary NEAR opponent	2797
"from adversary to opponent"	0
"from opponent to adversary"	0
"either adversary or opponent"	0
"either opponent or adversary"	0

Given a query *x NEARy,* AltaVista returns documents where the words *x andy* appear close to each other. When two words are unrelated, the hits for the NEAR query tend to be low.

Motivated by the above examples, we propose to identify semantically incompatible word pairs by searching on the Web for instantiations of the patterns in (4). We define a score:

$$(6)\quad score(x, y) = \frac{\text{hits}(x \; NEAR \; y)}{\sum_{pat \in P} \text{hits}(\text{pat}(x, y)) + \epsilon}$$

where *hits(query)* is the number of hits returned by AltaVista for the *query, P* is the set of patterns in (4) and *c* is a small constant to prevent the denominator of the above formula to be 0 (we set e=0.0001). The lower the score, the less likely that the words *x* and *y* are synonyms. To determine whether or not distributionally similar words *x* and *y* are synonyms, we compute *score(x, y).* If the value is higher than 0=2000, (x, *y)* is classified as a pair of synonyms.

2.2 Using Bilingual Dictionaries

The second method is based on the observation that translations of a word from another language are often synonyms of one another. For example, (7) contains the English translations of the French word *defenseur.* Many of them are synonyms.

(7) advocate, attorney, counsel, fullback, intercessor, lawyer

When two such words are not synonyms, the reason is typically that the French word have multiple senses and the English words are translations of the French word in different senses. Under such circumstances, the distribution of the English words are usually quite different (e.g., *lawyer* and *fullback* appear in very different contexts). We can therefore identify synonyms of a word *w* by intersecting the set of words that share with *w* the same French (or any other language) translation and the set of distributionally similar words of *w.* For example, the top-20 distributionally similar words of *lawyer* are:

(8) *lawyer:* attorney, counsel, prosecutor, doctor, official, judge, executive, manager, investigator, consultant, aide, agent, physician, expert, banker, officer, politician, lobbyist, teacher, accountant

The intersection of (7) and (8) gives us the synonyms of *lawyer, attorney* and *counsel.*

Since this method generally has high precision and low recall (see the next section), we can use this method with multiple bilingual dictionaries separately and take the union of the results. In our experiments, we used 7 dictionaries from http://www.freewaresite.com/onldict/

English-Swedish, English-Spanish, English-Japanese, English-German, English-French and English-Esperanto.

3 Evaluation

Using the algorithm in [Lin, 1998] on a 3GB newspaper corpus, we computed the distributional similarity between about 45,000 words. We randomly selected 80 pairs of synonyms and 80 pairs of antonyms from Webster's Collegiate Thesaurus [Kay, 1988] that are also among the top-50 distributionally similar words of each other. We then used the methods presented in the previous section to determine which pairs are synonyms. Let *S* be the set of true synonym pairs and *S'* be the set of pairs classified as synonyms. The precision and recall measures are defined as follows:

$$precision = \frac{|S \cap S'|}{|S'|} \qquad recall = \frac{|S \cap S'|}{|S|}$$

The table in (9) shows the evaluation results:

Method	Precision	Recall
Pattern-based	86.4%	95.0%
Bilingual Dictionaries	93.9%	39.2%

4 Related Work

The problem we address here is related to semantic orientation. The semantic orientation of a word is positive (or negative) if it is generally associated with good (or bad) things. For example, The words *simple* and *simplistic* have similar meanings, but *simplistic* has a negative semantic orientation. The algorithm in [Hatzivassiloglou and McKeown, 1997] is based on the fact that conjoined adjectives generally have the same orientation. They use a small set of adjectives with known orientation to determine the orientations of other adjectives. [Turney, 2002] computed the degree of positive or negative semantic orientation of a word *w* with the hit counts from AltaVista for the queries *w NEAR excellent* and *w NEAR poor.* While semantic orientation is bipolar, the problem we are dealing with is multipolar. For example, Turney's method is not able to tell that the words *red, orange, yellow, green,...* have incompatible meanings.

5 Conclusion

Distributionally similar words include many antonyms and other semantically incompatible words, which minimizes their use in many applications. We have presented two methods for identifying synonyms among distributionally similar words. Our preliminary evaluation with known synonyms and antonyms extracted from Webster Collegiate Thesaurus has produced promising results.

References

[Harris, 1968] Zelig S. Harris. *Mathematical Structures of Language.* Wiley, New York, 1968.

[Hatzivassiloglou and McKeown, 1997] Vasileios Hatzivassiloglou and Kathleen R. McKeown. Predicting the semantic orientation of adjectives. In *Proceedings ofACL/EACL-97,* pages 174-181, Madrid, Spain, July 1997.

[Hindle, 1990] Donald Hindle. Noun classification from predicate-argument structures. In *Proceedings ofACL-90,* pages 268-275, Pittsburg, Pennsylvania, June 1990.

[Kay, 1988] Maire Weir Kay, editor. *Webster's Collegiate Thesaurus.* Merrian-Webster, 1988.

[Lin, 1998] Dekang Lin. Automatic retrieval and clustering of similar words. In *Proceedings of COLING/ACL-98,* pages 768-774, Montreal, 1998.

[Pereira *et al,* 1993] F. Pereira, N. Tishby, and L. Lee. Distributional Clustering of English Words. In *Proceedings ofACL-93,* pages 183-190, Ohio State University, Columbus, Ohio, 1993.

[Turney, 2002] Peter D. Turney. Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews. In *Proceedings of ACL-02,* pages 417-424, Philadelphia, July 2002.