Lecture 5: Semantic Orientation

Lexical Semantics and Discourse Processing
MPhil in Advanced Computer Science

Simone Teufel

Natural Language and Information Processing (NLIP) Group
UNIVERSITY OF CAMBRIDGE
Simone.Teufel@cl.cam.ac.uk

February 4, 2011

Reading: Cruse (1986), chapters 9 and 11.3; Hatzivassiloglou and McKeown (1997); Turney (2002).

Oppositeness and Antonymy

- There are different kinds of opposites: complementaries and antonyms
- Semantic orientation: degree of positiveness/negativeness.
- Many antonyms have opposite semantic orientation. Exceptions: *verbose*—*terse*

Oppositeness and Antonymy

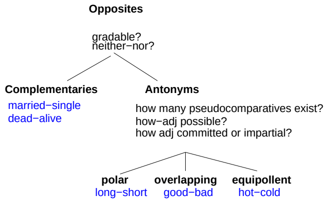

Complementaries

Complementaries between them exhaustively divide some conceptual domain into mutually exclusive compartments. Antonyms don't.
neither–nor test:

- *? Mary is neither married nor is she single.*
- *It's neither hot nor cold today.*

They are also not gradable:

- *? extremely true – extremely safe*
- *? more pregnant than most – longer than some*
- *? moderately female – moderately clean*

Antonyms 1: Pseudo comparatives and true comparatives

light–heavy and *hot–cold* do not behave in the same way:

- *This box is light, but it's heavier than that one.*
- *? Today it's cold, but hotter than yesterday.*

What is going on? *hot* seems to mean *hot, and to a larger degree*, whereas *heavier* seems to mean *of greater weight*.

- *hotter* is a true comparative of *hot*
- *heavier* is a pseudo-comparative of *heavy/1*, and a true comparative of *heavy/2*

Antonyms 2: How-adj questions

Are they possible for both antonyms?
Compare *long–short*:

- *How long is it?*
- *? How short is it?*

with *hot–cold*:

- *How cold is it?*
- *How hot is it?*

Antonyms 3: Impartiality of how-adj questions

Does one of the questions imply something about your presuppositions?
hot–cold:

- *How cold is it?* → committed
- *How hot is it?* → committed

clean–dirty:

- *How clean was the room?* → impartial
- *How dirty was the room?* → committed

Antonyms: Three types

- *good–bad* is an example of an **overlapping** antonym.
- Overlapping antonyms are evaluative, and thus carry semantic orientation in our sense.
- *hot–cold* is an example of an **equipollent** antonym.
- Equipollent antonyms are often correlated with sensory perceptions.
- *long–short* is an example of a **polar** antonym.
- Polar antonyms show the greatest level of abstraction, but are neutral/descriptive.

Linguistic polarity vs. natural polarity

- Can we predict from the linguistic form which one of the antonyms is more positive?
- Prediction: the more salient antonym often has a positive polarity.
- **Test 1**: The antonym that can be paraphrased as the other one plus a negative prefix is the less salient one.
- **Test 2**: The more salient antonym is associated with "more" properties:
  - *Something is dead when there is no life present.*
  - *? Something is alive when there is no deadness present.*
- *clean* and *safe* are exceptions in that
  - *Something is clean when there is no dirt present.*
  - *? Something is dirty when there is no cleanness present.*

Linguistic polarity vs. natural polarity, II

- **Test 3**: The more salient antonym yields the impartial interpretation in the how-adj question.
- In the case of verbs:
  - Antonymy in verbs often concerns directional actions, and reversive actions (Cruse, chapter 10)
  - The salient antonym is the one that results in "increased entropy" (*undress, dismount, disarrange, unscrew, unpack*. . .)

Automatic Detection of Semantic Orientation of Adjectives

Hatzivassiloglou and McKeown's (1997) algorithm classifies adjectives into those with positive or negative semantic orientation:

- **Semantic Polarity of an adjective:**
  - **Direction:** In which direction does the referent deviate from the norm in its semantic field?
  - **Evaluative:** Is this good or bad?
- If we know that two adjectives relate to the same property (e.g., *hot* and *cold*) but have different orientations they are usually antonyms

Idea

- In coordinations, these facts result in constraints on the semantic orientation:

(1)
  a. The tax proposal was **simple and well-received** by the public.
  b. The tax proposal was **simplistic but well-received** by the public.
  c. ?The tax proposal was **simplistic and well-received** by the public.

- *but* combines adjectives of opposite orientation; *and* adjectives of the same orientation
- This indirect information can be exploited using a corpus.

Algorithm

- Extract all coordinated adjectives from corpus
- Classify each extracted adjective pair as same or different orientation
- This results in graph with same or different links between adjectives
- Cluster into two orientations, placing as many words of the same orientation as possible into the same subset
- Cluster with higher overall frequency is labelled positive
- Evaluate against independently orientation-annotated gold standard set (1336 most frequent adjectives; 657 positive, 679 negative)

Coordinated adjectives

- Extract from POS tagged WSJ (21 million words) adjective pairs coordinated by *and*, *or*, *but*, *either-or*, *neither-nor*
- This results in 15048 adjective pairs (token); 9296 (type)
- Number of those where orientation of both partners is known (via gold standard): 4024 (token); 2748 (type)
  - *and* is most reliable same-orientation predictor, particularly in predicative position (85%), this drops to 70% in appositive position.
  - *but* has 31% same-orientation.

Classifier

- Features:
  - Type of coordination
  - Type of modification (attributive, predicative, appositive, resultative ("*Bill laughed himself hoarse*")
  - Number of modified noun (singular or plural)
- Simple derivational morphological analysis suggests additional different orientations: Out of the labelled adjectives, 97% of morphologically related pairs (102) have different orientation
- Log-linear regression model with linear predictor; best classifier achieves 82%
- Baseline: always predict same-orientation: 79%
- But-rule: different if seen with *but*, same-orientation otherwise: 82%

Clustering adjectives with same orientation

- Interpret classifier's P(same-orientation) as dissimilarity value.
- Perform non-hierarchical clustering via Exchange Method
- Start from random partition, locate the adjective which reduces the cost c most if moved.

$$c = \sum_{i=1}^{2} \left(\frac{1}{|C_i|} \sum_{x,y \in C_i, x \neq y} d(x,y) \right)$$

- Repeat until no movements can improve the cost; overall dissimilarity cost is now minimised.
- At final iteration, move any adjective which violates the following constraint:

$$\frac{1}{|C|-1} \sum_{y \in C, x \neq y} d(x,y) < \frac{1}{|\bar{C}|} \sum_{y \in \bar{C}} d(x,y)$$

Labelling Clusters as Positive or Negative

- In antonym pairs, the one which is semantically unmarked is also in most cases the positive one.
- Semantically unmarked ones should occur overall more frequently → group with overall higher frequency count gets labelled as positive.

Results

- Dependent on how sparse the test set is, results between 78% and 92% correct
- Baselines: MFC 51% negative
- Classified as positive: *bold, decisive, disturbing, generous, good, honest, important, large, mature, patient, peaceful, positive, proud, sound, stimulating, straightforward, strange, talented, vigorous, witty.*
- Classified as negative: *ambiguous, cautious, cynical, evasive, harmful, hypocritical, inefficient, insecure, irrational, irresponsible, minor, outspoken, pleasant, reckless, risky, selfish, tedious, unsupported, vulnerable, wasteful.*

Discussion

Strengths:

- Fully unsupervised, algorithm starts from nothing
- Convincing results

Weaknesses:

- Analysis of isolated adjectives, not phrases
- Needs large corpus in order to contain enough coordinated adjectives
- Clustering algorithm is not optimal (problem is NP-hard); it is a steepest-descending hill climbing method, which is at least guaranteed to converge (but might run algorithm repeatedly with different start partitions)

Turney's 2002 method

- Determine semantic orientation of phrases, not just single adjectives
- Single adjectives do not always carry full orientation; context is needed. *unpredictable plot* vs. *unpredictable steering*
- Unsupervised method based on distributional semantics
- Assign a numerical ranking indicating strength of orientation
- Use search engine hits to estimate semantic orientation of a phrase

Idea

- If an adjectival phrase has a positive semantic orientation, it will appear more frequently in the intermediate vicinity of known positive adjectives, and vice versa.
- Measure an adjective's tendency to appear in positive or negative vicinity via PMI-IR
  - Pointwise mutual information determines similarity of a pair of phrases
  - Use IR to quantify effect
- Measure success indirectly via classification of entire reviews

PMI and SO

$$PMI(word_1, word_2) = log_2\left(\frac{P(word_1, word_2)}{P(word_1)P(word_2)}\right)$$

- Semantic Orientation:
  SO(phrase) = PMI(phrase, *excellent*) - PMI (phrase, *poor*)
- Counts are calculated via search engine hits
- Altavista's NEAR operator – window of 10 words

Therefore:

$$SO(phrase) = log_2\left(\frac{hits(phrase\ NEAR\ excellent)hits(poor)}{hits(phrase\ NEAR\ poor)hits(excellent)}\right)$$

Results: indirectly via classification of documents

- 74% accuracy on classifying 410 reviews from Epinions
- 66% accuracy on movie reviews

An example:

little difference	-1.615	virtual monopoly	-2.050
clever tricks	-0.040	other bank	-0.850
programs such	**0.117**	extra day	-0.286
possible moment	-0.668	**direct deposits**	**5.771**
unetical practices	-8.484	**online web**	**1.936**
old man	-2.566	**cool thing**	**0.395**
other problems	-2.748	**very handy**	**1.349**
probably wondering	-1.830	lesser evil	-2.288

Total: -1.218. Rating: Not recommended.

Discussion

Strengths:

- Fully unsupervised
- Nominal context makes adjective semantics more interpretable

Weaknesses:

- No direct evaluation of SO provided
- Very simple model
- Requires many searches (too many without API)
- NEAR no longer supported
- Results depend substantially on lexical items chosen, but choice largely unmotivated

References

Hatzivassiloglou and McKeown (1997): Predicting the Semantic Orientation of Adjectives. Proceedings of the ACL.

Turney (2002): Thumbs up or down? Semantic Orientation Applied to Unsupervised Classification of Reviews. Proceedings of ACL.

Homework

- Familiarize yourself with the organisation of WordNet (e.g., `www.wordnet-online.com`)
- Explore lexical neighbourhood in WordNet of some of the examples given in either of the two papers discussed today
- Are they part of an antonym pair, and if so, which type of antonymy is it? Support your answer with linguistic tests.