L114 Lexical Semantics

Session 7: Antonymy and Sentiment Detection

Simone Teufel

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

2013/2014

1. Semantic Orientation of Adjectives
   - Antonymy
   - Linguistic tests for complementaries and antonymy type
   - Linguistic vs. natural polarity

2. Automatic Detection of Sem. Orientation
   - Hatzivassiloglou and McKeown
   - Turney (PMI Method)

Opposites, Antonyms and Semantic Orientation

- There are different kinds of opposites: complementaries and antonyms
- Antonyms are closely related to semantic orientation (degree positiveness/negativeness).
- If we know that two adjectives relate to the same property (e.g., *hot* and *cold*) but have different semantic orientations they are typically antonyms.
- Exceptions: *verbose*—*terse* (same semantic orientation)

Oppositeness and Antonymy

Opposites

gradable?
neither−nor?

- **Complementaries**
  - married−single
  - dead−alive
- **Antonyms**

  how many pseudocomparatives exist?
  how−adj possible?
  how adj committed or impartial?

  - **polar**: long−short
  - **overlapping**: good−bad
  - **equipollent**: hot−cold

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.* ($\rightarrow$ complementaries)
- *It's neither hot nor cold today.* ($\rightarrow$ antonyms)

Gradability test:

- *? 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?

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? *Heavier* seems to mean *of greater weight* here (relative property), whereas *hot* seems to express a more absolute property.

- *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 how-adj questions 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? (this is about neutrality)

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

Excursion: presuppositions

- A presupposition is an implicit assumption about the world or background belief relating to an utterance whose truth is taken for granted in discourse.
- Examples:
    - Jane no longer writes fiction.
      Presupposition: Jane once wrote fiction.
    - Have you stopped eating meat?
      Presupposition: you had once eaten meat.
    - Have you talked to Hans?
      Presupposition: Hans exists.
    - If the notice had only said 'mine-field' in Welsh as well as in English, we would never have lost poor Llewellyn.
      Presupposition: The notice didn't say 'mine-field' in Welsh.

Presupposition triggers

- Many words and constructions are presupposition triggers
- *regret, realise, manage, forget, try, again, since X happened, Carol is a better linguist than Mary...*

Presuppositions

- Negation of utterance does not cancel its presuppositions.
- This distinguishes it from entailment.
- A presupposition of a sentence must normally be part of the common ground of the utterance context (the shared knowledge of the interlocutors) in order for the sentence to be felicitous.
- If not, presupposition accommodation takes place unless this leads to inconsistency. ("My wife is a dentist", said to somebody who does not know that you have a wife.)

Three types of antonyms

- *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 which one of the antonyms is more "salient" (that is typically the one that is 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.*
- **Test 3**: The more salient antonym yields the impartial interpretation in the how-adj question.

That picture again

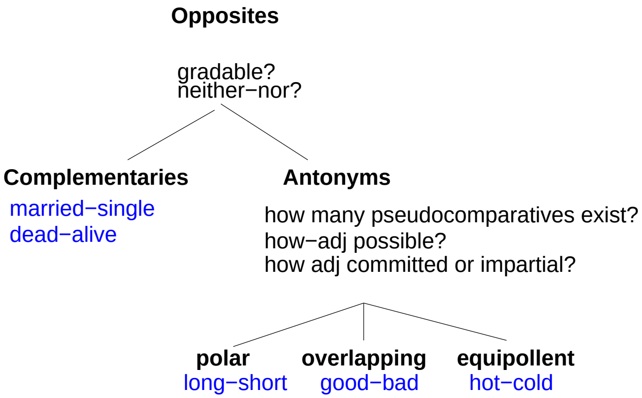

Automatic Detection of Semantic Orientation of Adjectives

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

- In coordinations, antonymy results 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 21 million word WSJ corpus (*and*, *or*, *but*, *either-or*, *neither-nor*)
- 15048 adj pairs (token), 9296 (type)
- Classify each extracted adjective pair as same or different orientation (82% accuracy)
- This results in graph with same or different links between adjectives

Classification

- features used: number of modified noun; type of coordination; type of modification (attributive, predicative, appositive, resultative (“*Bill laughed himself hoarse*”)
- *and* is most reliable same-orientation predictor, particularly in predicative position (85%), this drops to 70% in appositive position.
- *but* has 31% same-orientation.
- This information comes from an independently annotated gold standard (1336 most frequent adjectives; 657 positive, 679 negative)
- Additional different orientations comes from simple morphological analysis: Out of the labelled adjectives, 97% of morphologically related pairs (102) have different orientation

Clustering adjectives with same orientation

- Now cluster adjectives into two orientations, placing as many words of the same orientation as possible into the same subset
- 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.

Exchange method; final step

- 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

- Hatzivassiloglou empirically find that the cluster with overall higher frequency tends to be the positive one; so this is the final step in their algorithm
- Possible reason: In overlapping antonym pairs, the positive adjective tends to be semantically unmarked (as we heard earlier today).
- Semantically unmarked adjectives should occur more frequently in language (if only because of neutral questions etc)

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:

- Algorithm only needs gold standard list
- 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

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

- 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:

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

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

Reading

- Cruse (1986), chapters 9 and 11.3;
- **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.