7: Catchup I

Machine Learning and Real-world Data

Simone Teufel and Ann Copestake

Computer Laboratory
University of Cambridge

Lent 2017

Last session: uncertainty and human annotation

- In the last session, we used multiple human annotation and an appropriate agreement metric
- Can be appropriate in apparently “overly subjective” situations
- This way, we could define an defensible definition of “truth”
- This concludes the practical part about text classification.
- Today: catchup session 1

What happens in catchup sessions?

- Lecture and demonstrated session scheduled as in normal session.
- Lecture material for your information only, non-examinable.
- Time for you to catch-up in demonstrated sessions or attempt some starred ticks.
- Demonstrators help as per usual.
- Fridays are Ticking sessions, whether catchup or not.

Research on sentiment detection

- Unsupervised sentiment lexicon induction
  - Mutual information method
  - Coordination method
- Propagating sentiments from words to larger units
  - Negation treatment
  - Propagation by supervised ML
  - Symbolic-semantic propagation
- The function of text parts
  - plot description
  - recommendation
- Other
  - Aspect-based
  - Irony detection

Pointwise Mutual Information Method

- Due to Turney (2002)
- Estimate semantic orientation of any unseen phrase
- 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.
- Quantify tendency by pointwise mutual information and search engine hits.

PMI and SO

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

- 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:

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

Turney's second idea: context

- 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**

Examples:

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**
unethical 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.

Coordination Method

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

- Consider:
  1. The tax proposal was **simple** and **well-received** by the public.
  2. The tax proposal was **simplistic** but **well-received** by the public.
- **but** combines adjectives of opposite orientation; **and** adjectives of the same orientation
- This indirect information from pairs of coordinated adjectives can be exploited using a corpus.

Algorithm

- Extract all coordinated adjectives from 21 million word WSJ corpus
- 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
- Now cluster adjectives into two orientations, placing as many words of the same orientation as possible into the same subset

Classification Features

- number of modified noun
- type of coordination (**and**, **or**, **but**, **either-or**, **neither-nor**)
- syntactic context
  - **black or white** horse (attributive)
  - horse was **black or white** (predicative)
  - horse, **black or white**, gallopped away (appositive)
  - Bill laughed himself **hoarse and exhausted** (resultative)
- **and** is most reliable same-orientation predictor, particularly in predicative position (85%), this drops to 70% in appositive position.
- **but** has 31% same-orientation
- Morphological filter (un-, dis-) helps

Clustering adjectives with same orientation together

- When clustering, Interpret classifier's P(same-orientation) as similarity value.
- Perform non-hierarchical clustering via Exchange Method:
  - Start from random partition, locate the adjective which reduces the cost *c* most if moved.
  - Repeat until no movements can improve the cost; overall dissimilarity cost is now minimised.
- Call cluster with overall higher frequency "positive", the other one "negative"
- Results between 78% and 92% accuracy; main factor: frequency of adjective concerned
- Baseline: most frequent category (MFC) 51% negative

Examples

- **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.

Propagation of Polarity: Supervised ML

- Due to Wilson, Wiebe, Hoffman (2005)
- Learn propagation of word polarity into polarity of larger phrases
- Source of the sentiment lexicon we used in Task 1
- Whether words carry global polarity depends on the context (e.g., **Environmental Trust** versus **He has won the people's trust**)
- Cast task as supervised ML task
- **they have not succeeded, and will never succeed**, was marked as positive in the sentence, **They have not succeeded, and will never succeed, in breaking the will of this valiant people.**

And what are we going to do about negation?

- Negation may be local (e.g., **not good**)
- Negation may be less local (e.g., **does not really always look very good**)
- Negation may sit on the syntactic subject (e.g., **no one thinks that it's good**)
- Diminishers can act as negation (e.g., **little truth**)
- Negation may make a statement hypothetical (e.g., **no reason to believe**)
- Intensifiers can wrongly look as if they were negation (e.g., **not only good but amazing**)

Negation methods

- Fixed and syntactic windows
- Machine-learning of different syntactic constructions (Wilson et al. 2015)
- Treatment of affected words:
  - NEG-labelling of words (put **is_N not_N good_N** into NEG)
  - adding antonym in features for same class (add both **good_N** + **bad** into NEG)
  - adding negated word in a feature of **opposite** category (add **good** into POS)
- Very hard to show any effect with negation

Deep syntactic/semantic inference on sentiment

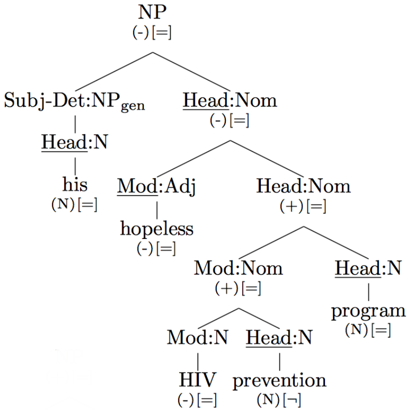

Moilanen and Pulman (2007)

Deep syntactic/semantic inference on sentiment

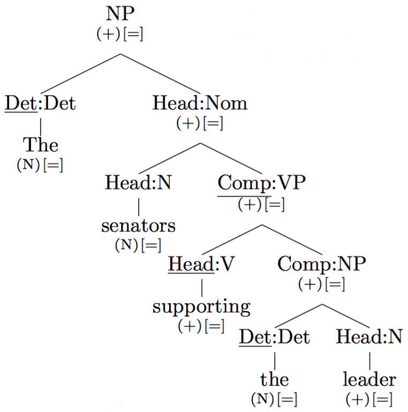

Deep syntactic/semantic inference on sentiment

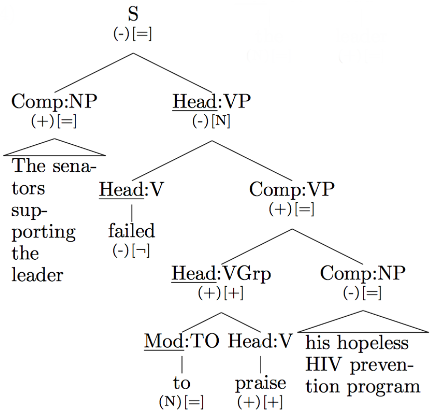

Spinout company: TheySay

Pang and Lee (2004)

- Idea: objective sentences should not be used for classification
- Plot descriptions are not evaluative
- Algorithm:
  - First classify each individual sentence as objective or subjective
  - Find clusters of similarly objective or subjective sentences inside the document (by Minimum Cut algorithm)
  - Exclude objective sentences; then perform normal BOW sentiment classification

Minimum Cut algorithm

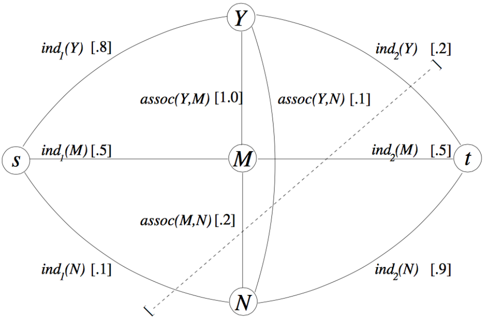

C_1	Individual penalties	Association penalties	Cost
{Y,M}	.2 + .5 + .1	.1 + .2	1.1
(none)	.8 + .5 + .1	0	1.4
{Y,M,N}	.2 + .5 + .9	0	1.6
{Y}	.2 + .5 + .1	1.0 + .1	1.9
{N}	.8 + .5 + .9	.1 + .2	2.5
{M}	.8 + .5 + .1	1.0 + .2	2.6
{Y,N}	.2 + .5 + .9	1.0 + .2	2.8
{M,N}	.8 + .5 + .9	1.0 + .1	3.3

Aspect-based sentiment detection challenge 2016

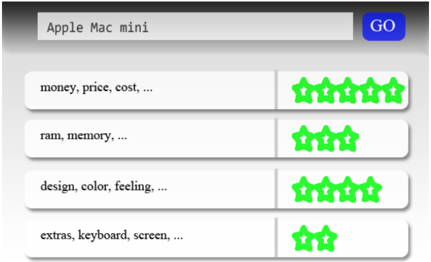

- 8 languages, 39 large datasets

Aspect-based sentiment detection challenge 2016

1. Service was slow, but the people were friendly. → {trg: "Service", pol: "negative"}, {trg: "people", pol: "positive"}
2. Snelle bediening en vriendelijk personeel moet ook gemeld worden!! → {trg: "bediening", pol: "positive"}, {trg: "personeel", pol: "positive"}
3. Le service est impeccable, personnel agréable. → {trg: "service" , pol: "positive"}, {trg: "personnel", pol: "positive"}
4. Про сервис ничего негативного не скажешь-быстро подходят, все улябаются, подходят спрашивают, всё ли нравится. → {trg: "сервис", pol: "neutral" }
5. También la rapidez en el servicio. → {trg: "servicio", pol: "positive" }
6. Servisi hızlı valesi var. → {trg: "Servisi", pol: "positive"}
7. الخدمة جيدة جدا و سريعة .. → {trg: "الخدمة" , pol: "positive"}

1. It is extremely portable and easily connects to WIFI at the library and elsewhere. → {cat: "LAPTOP#PORTABILITY", pol: "positive"} , {cat: "LAPTOP#CONNECTIVITY", pol: "positive"}
2. Apps starten snel op en werken vlot, internet gaat prima. → {cat: "SOFTWARE#OPERATION_PERFORMANCE", pol: "positive"}, {cat: "PHONE#CONNECTIVITY", pol: "positive"}
3. 当然屏幕这么好 →{cat: "DISPLAY#QUALITY", pol: "positive"}
4. 更 轻 便 的 机 身 也 便 于 携 带。→ {cat: "CAMERA# PORTABILITY", pol: "positive"}

Irony-detection in Twitter

- Gonzalez-Ibanez et al. (2011)

@UserName That must suck.
I can't express how much I love shopping on black Friday.
@UserName that's what I love about Miami. Attention to detail in preserving historic landmarks of the past.
@UserName im just loving the positive vibes out of that!

Irony-detection: features

S-P-N	S-NS	S-N	S-P
Negemo(PP)	Posemo(PP)	Posemo(PP)	Question
Posemo(PP)	Present(LP)	Negemo(PP)	Present(LP)
Smiley(Pr)	Question	Joy(WNA)	ToUser(Pr)
Question	ToUser(Pr)	Affect(PP)	Smiley(Pr)
Negate(LP)	Affect(PP)	Anger(PP)	AuxVb(LP)
Anger(PP)	Verbs(LP)	Sad(PP)	Ipron(LP)
Present(LP)	AuxVb(LP)	Swear(PP)	Negate(LP)
Joy(WNA)	Quotation	Smiley(Pr)	Verbs(LP)
Swear(PP)	Social(PP)	Body(PP)	Time(PP)
AuxVb(LP)	Ingest(PP)	Frown(Pr)	Negemo(PP)

Ticking today

- Task 5 – Crossvalidation
- Task 6 – Kappa implementation

Literature

- **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 the ACL.
- **Pang, Lee** (2004): A sentimental education: sentiment analysis using subjectivity summarisation based on minimum cuts. Proceedings of the ACL.
- **Pang, Lee, Vaithyanathan** (2002): Thumbs up? Sentiment Classification Using Machine Learning Techniques. Proceedings of EMNLP.
- **Wilson, Wiebe, Hoffmann** (2005): Recognising contextual Polarity in phrase-level sentiment analysis, Proceedings of HLT.
- **Gonzalez-Ibanez, Muresan, Wacholder** (2011). Identifying Sarcasm in Twitter: A Closer Look. Proceedings of the ACL.
- **Moilanen, Pullman** (2007): Sentiment Composition. Proceedings of RANLP.
- **Pontiki et al.** (2016): SemEval-2016 Task 5: Aspect Based Sentiment Analysis. Proceeding of SemEval.