Subjectivity Classification

Wilson, Wiebe and Hoffmann: Recognizing contextual polarity in phrase-level sentiment analysis

Wiltrud Kessler

Institut für Maschinelle Sprachverarbeitung
Universität Stuttgart

Sentiment Analysis
Summer 2012

Outline

Outline

Contextual Polarity

- Sentiment analysis is often done using word lists of positive and negative words (**subjectivity clues**).
- The polarity of a clue in the list is called **prior polarity** – what sentiment does it evoke out of context?
- The **contextual polarity** of a phrase in which a clue appears can be different from the clue's prior polarity.

Contextual Polarity – Example

Trust

well

reason

polluters

reasonable

Contextual Polarity – Example

Philip Clapp, president of the National Environment Trust , sums up well the general thrust of the reaction of environmental movements: "There is no reason at all to believe that the polluters are suddenly going to become reasonable ."

Motivation for Subjectivity Classification

- Simple application of subjectivity clues list: Always assume the prior polarity as contextual polarity of a clue.
- Accuracy: 48%

		Prior-Polarity Classifier				
		Neut	Pos	Neg	Both	Total
Gold	Neut	798	784	698	4	2284
	Pos	81	371	40	0	492
	Neg	149	181	622	0	952
	Both	4	11	13	5	33
	Total	1032	1347	1373	9	3761

Table 2: Confusion matrix for the prior-polarity classifier on the development set.

- 76% of errors result from words with non-neutral prior polarity appearing in phrases with neutral contextual polarity.
- First classifying if an expression is neutral or polar can significantly reduce errors.

Motivation for Subjectivity Classification

- Simple application of subjectivity clues list: Always assume the prior polarity as contextual polarity of a clue.
- Accuracy: 48%

		Prior-Polarity Classifier				
		Neut	Pos	Neg	Both	Total
Gold	Neut	798	784	698	4	2284
	Pos	81	371	40	0	492
	Neg	149	181	622	0	952
	Both	4	11	13	5	33
	Total	1032	1347	1373	9	3761

Table 2: Confusion matrix for the prior-polarity classifier on the development set.

- 76% of errors result from words with non-neutral prior polarity appearing in phrases with neutral contextual polarity.
- First classifying if an expression is neutral or polar can significantly reduce errors.

Overview

Goal:
Automatically identify contextual polarity of a subjectivity clue.

Approach:

- Identify all phrases containing subjectivity clues in the corpus.
- Classify each phrase containing a clue as *neutral* or *polar*.
- For all phrases marked as *polar*, disambiguate their contextual polarity (*positive*, *negative*, *both*, or *neutral*).

Overview

Goal:
Automatically identify contextual polarity of a subjectivity clue.

Approach:

- Identify all phrases containing subjectivity clues in the corpus.
- Classify each phrase containing a clue as *neutral* or *polar*.
- For all phrases marked as *polar*, disambiguate their contextual polarity (*positive*, *negative*, *both*, or *neutral*).

Overview

Goal:
Automatically identify contextual polarity of a subjectivity clue.

Approach:

- Identify all phrases containing subjectivity clues in the corpus.
- Classify each phrase containing a clue as *neutral* or *polar*.
- For all phrases marked as *polar*, disambiguate their contextual polarity (*positive*, *negative*, *both*, or *neutral*).

Outline

Data

- Subjectivity clues: Manually compiled list.
- A list of intensifiers (not specified).
- Passive verb patterns (not specified).
- Corpus: Multi-perspective Question Answering (MPQA) Opinion Corpus.

Subjectivity Clues

- Manually compiled list.
- Only single-word clues.
- Separate entries for different parts of speech.
- Each clue has a reliability class: *strongsubj* or *weaksubj*
- Each clue has a prior polarity: *positive*, *negative*, *both*, or *neutral*

Subjectivity Clues – Sources

- Subjectivity clues from [RW03] (manually collected from different sources):
  $\rightarrow$ Manually add prior polarity.
- Expansion from “dictionary and thesaurus”:
  $\rightarrow$ Manually add reliability and prior polarity.
- General Inquirer positive and negative word lists:
  $\rightarrow$ Manually add reliability, adjust prior polarity.
- Annotated adjectives from [HM97]:
  $\rightarrow$ Manually add reliability, adjust prior polarity.

Subjectivity Clues – Statistics

label	percentage	# words
positive	33.1%	2718
negative	59.7%	4911
neutral	6.9%	570
both	0.3%	21
total	100%	8221

Subjectivity Clues – Examples

positive	great, good, easy, like, just, will, sound, better, even, nice, want, light, excellent, best, pretty, easily, large, free, clear, love, clarity, ...
negative	too, little, long, hard, need, poor, bad, problem, down, although, low, difficult, less, expensive, cheap, lack, extremely, alarm, limited, annoying, flimsy, ...
neutral	absolute, activist, anyhow, apparent, belief, contemplate, consider, deeply, exact, eyebrows, feel, firm, giant, high, imperative, ...
both	brag, covet, demand, fawn, gloat, implore, infatuated, lust, plead, ...

MPQA Opinion Corpus

- MPQA corpus from [WWC05].
- This corpus contains annotations for "subjective expressions": Phrases used to express an opinion, emotion, evaluation, stance, speculation, ...
- Existing annotations for "subjective expressions" have been annotated with polarity ($\rightarrow$ MPQA Opinion Corpus).
- Annotators were asked to judge the contextual polarity, i.e. the polarity in context of the whole sentence.
- Possible labels: *positive*, *negative*, *both*, or *neutral*

MPQA Opinion Corpus – Annotation Examples

(5) Thousands of coup supporters celebrated (**positive**) overnight, waving flags, blowing whistles …

(6) The criteria set by Rice are the following: the three countries in question are repressive (**negative**) and grave human rights violators (**negative**) …

(7) Besides, politicians refer to good and evil (**both**) only for purposes of intimidation and exaggeration.

(8) Jerome says the hospital feels (**neutral**) no different than a hospital in the states.

MPQA Opinion Corpus – Agreement

- Study on 10 documents (447 subjective expressions).
- Two annotators.

	Neutral	Positive	Negative	Both	Total
Neutral	123	14	24	0	161
Positive	16	73	5	2	96
Negative	14	2	167	1	184
Both	0	3	0	3	6
Total	153	92	196	6	447

Table 1: Agreement for Subjective Expressions (Agreement: 82%, κ: 0.72)

MPQA Opinion Corpus – Corpus Statistics

part	# documents	# sentence	# expressions
total	425	8 984	15 991
development	66	1 373	2 808
train/test	359	7 611	13 183

Sentences with	percentage
0 subjective expression	28%
1 subjective expression	25%
$\geq$ 2 subjective expressions	47%

MPQA Opinion Corpus – Training Data for Classifier

- Items to be classified are subjectivity clues.
- Let x be a subjectivity clue, se a subjective expression with label $l(se)$ (which can be positive, negative, neutral or both).
- Gold standard label for x ...
  - x not in a $se \Rightarrow neutral$
  - x in exactly one $se \Rightarrow l(se)$
  - x in se_i and se_j
    - $l(se_i) = l(se_j) \Rightarrow l(se_i)$
    - $l(se_i) = positive$ and $l(se_j) = negative \Rightarrow both$
    - $l(se_i) = positive|negative$ and $l(se_j) = neutral \Rightarrow l(se_i)$

Subjectivity Clues in the MPQA Opinion Corpus

label	# instances
total development set	3761
total train/test set	19 506
positive (dev set)	492
negative (dev set)	952
neutral (dev set)	2284
both (dev set)	33

Outline

AdaBoost (Adaptive Boosting)

- Main idea: Combine many simple, inaccurate classifier into a single, very accurate classifier.
- Each simple classifier (weak learner) is trained on the examples that have been misclassified by the previously learned classifiers.
- AdaBoost can be used with any classifier as the weak lerner.
- The weak learner does not have to be a good classifier.
- BoosTexter is a version of AdaBoost applied to text classification.

AdaBoost Algorithm

Given: $(x_1, y_1), \ldots, (x_m, y_m)$ where $x_i \in X$, $y_i \in Y = \{-1, +1\}$
Initialize $D_1(i) = 1/m$.
For $t = 1, \ldots, T$:

- Train weak learner using distribution D_t.
- Get weak hypothesis $h_t : X \to \{-1, +1\}$ with error

$$\epsilon_t = \Pr_{i \sim D_t}\left[h_t(x_i) \neq y_i\right].$$

- Choose $\alpha_t = \frac{1}{2} \ln\left(\frac{1 - \epsilon_t}{\epsilon_t}\right)$.
- Update:

$$\begin{aligned} D_{t+1}(i) &= \frac{D_t(i)}{Z_t} \times \begin{cases} e^{-\alpha_t} & \text{if } h_t(x_i) = y_i \\ e^{\alpha_t} & \text{if } h_t(x_i) \neq y_i \end{cases} \\ &= \frac{D_t(i) \exp(-\alpha_t y_i h_t(x_i))}{Z_t} \end{aligned}$$

where Z_t is a normalization factor (chosen so that D_{t+1} will be a distribution).

Output the final hypothesis:

$$H(x) = \text{sign}\left(\sum_{t=1}^{T} \alpha_t h_t(x)\right).$$

AdaBoost Example (1)

Toy Example

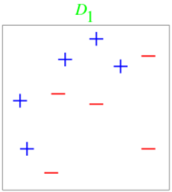

weak classifiers = vertical or horizontal half-planes

[Schapire: Theory and Applications of Boosting, Tutorial @ NIPS 2007]

AdaBoost Example (2)

Round 1

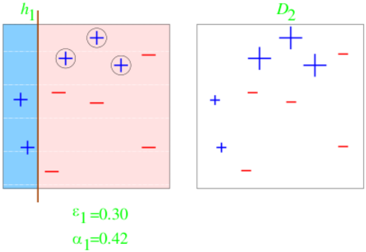

[Schapire: Theory and Applications of Boosting, Tutorial @ NIPS 2007]

AdaBoost Example (3)

Round 2

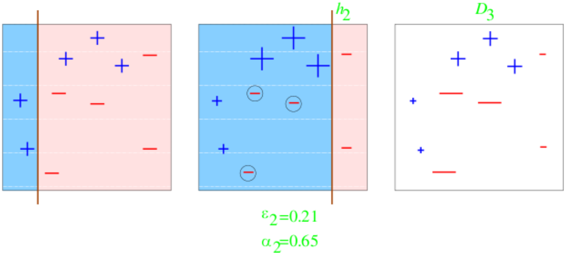

[Schapire: Theory and Applications of Boosting, Tutorial @ NIPS 2007]

AdaBoost Example (4)

Round 3

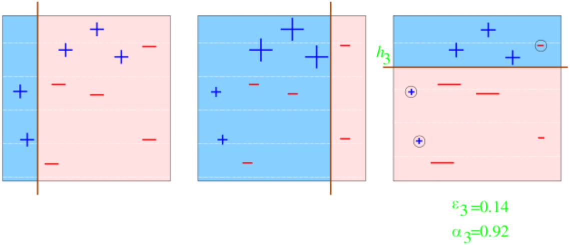

AdaBoost Example (5)

Final Classifier

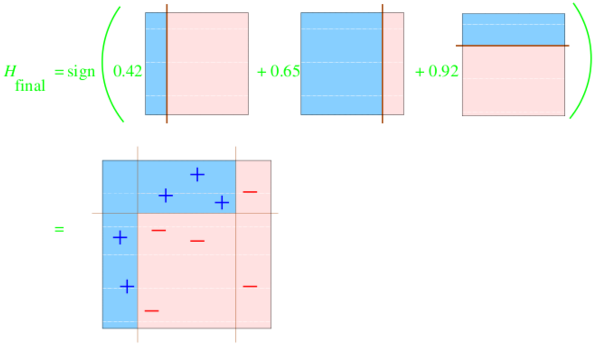

[Schapire: Theory and Applications of Boosting, Tutorial @ NIPS 2007]

The Weak Learner in BoosTexter

- We need to output a hypothethis for every training example x and class l.
- Our prediction $h(x, l)$ for will be c_{0l} if $w \notin x$ and c_{1l} if $w \in x$, where w is a feature.
- We calculate all values for all possible features and define a score for the resulting prediction.
- In the end, the prediction with the lowest score is selected and returned by the weak learner.

Calculation the Prediction

- Let $X_0 = \{x : w \notin x\}$ and $X_1 = \{x : w \in x\}$. For each label l, and each $j \in \{0, 1\}$ and $b \in \{-1, +1\}$:

$$W_b^{j\ell} = \sum_{i=1}^{m} D_t(i, \ell) [\![x_i \in X_j \wedge Y_i[\ell] = b]\!] .$$

- c_{il} is then calculated as:

$$c_{j\ell} = \frac{1}{2} \ln \left(\frac{W_+^{j\ell}}{W_-^{j\ell}} \right) ,$$

For more details see [SS00].

BoosTexter Example Terms [SS00]

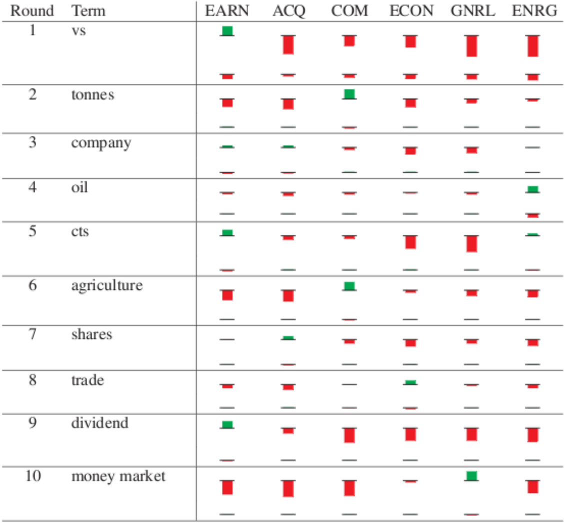

Round	Term	EARN	ACQ	COM	ECON	GNRL	ENRG
1	vs						
2	tonnes						
3	company						
4	oil						
5	cts						
6	agriculture						
7	shares						
8	trade						
9	dividend						
10	money market						

Outline

Features for a Subjectivity Classifier

- The classifier uses 28 different features.
- The classifier is trained on / applied to occurrences of subjectivity clues from the list that are found in the data.
- Define

w_i	the current subjectivity clue to be classified
w_{i-1}	the word immediately preceding w_i
w_{i+1}	the word immediately following w_i.

Word-level Features

word token: bag of words for w_i?

word part-of-speech: (bag of?) POS of w_i?

word context: bag of words from w_{i-1}, w_i, w_{i+1}.

prior polarity: *positive*, *negative*, *both*, or *neutral* as indicated in subjectivity clue list.

reliability class: *strongsubj* or *weaksubj* as indicated in subjectivity clue list.

Necessary information: Subjectivity clues list (polarity, reliability), POS.

Modification Features (1)

preceeded by adjective True if w_{i-1} is an adjective.

preceeded by adverb True if w_{i-1} is an adverb (other than *not*).

preceeded by intensifier True if w_{i-1} is in the list of intensifiers and w_i has the appropriate POS.

is intensifier True if w_i is in the list of intensifiers and w_{i+1} has the correct POS.

Necessary information: Intensifiers list, POS.

Modification Features (2)

Binary features extracted from the dependency parse of the sentence. Only applied if parent and child are in a *adj*, *mod* or *vmod* relationship.

modifies strongsubj True if parent's reliability is *strongsubj*.

modifies weaksubj True if parent's reliability is *weaksubj*.

modified by strongsubj True if a child's reliability is *strongsubj*.

modified by weaksubj True a child's reliability is *weaksubj*.

Necessary information: Dependency parse, subjectivity clues list (reliability).

Structure Features

Binary features that look at the path from the clue instance to the root in the dependency tree of the sentence.

in subject	True if a *subj* relationship is found.
in copular	True if in subject is false and if a node along the path is both a main verb and a copular verb.
in passive	True if a passive verb pattern is found.

Necessary information: Dependency parse, POS, passive verb patterns.

Sentence-level Features (1)

Counts the number of subjectivity clues of a certain reliability in the previous, current and next sentence. These have been used in previous work for sentence subjectivity classification.

strongsubj clues in current sentence
strongsubj clues in previous sentence
strongsubj clues in next sentence
weaksubj clues in current sentence
weaksubj clues in previous sentence
weaksubj clues in next sentence

Necessary information: Subjectivity clues list (reliability).

Sentence-level Features (2)

Features from previous work on sentence subjectivity classification.

adjectives in sentence Count of adjectives in current sentence.

adverbs in sentence Count of adverbs (other than not) in current sentence.

cardinal number in sentence True if the current sentences contains a cardinal number.

pronoun in sentence True if the current sentences contains a pronoun.

modal in sentence True if the current sentences contains a modal (other than will).

Necessary information: POS.

Document-level Features

Feature representing the topic of the document. Unclear how this topic is assigned.

topic One out of a list of 15 topics.

Necessary information: ??

All Features

Word Features	Sentence Features	Structure Features
word token	strongsubj clues in current sentence: count	in subject: binary
word part-of-speech	strongsubj clues in previous sentence: count	in copular: binary
word context	strongsubj clues in next sentence: count	in passive: binary
prior polarity: positive, negative, both, neutral	weaksubj clues in current sentence: count	
reliability class: strongsubj or weaksubj	weaksubj clues in previous sentence: count	
Modification Features	weaksubj clues in next sentence: count	Document Feature
preceeded by adjective: binary	adjectives in sentence: count	document topic
preceeded by adverb (other than not): binary	adverbs in sentence (other than not): count	
preceeded by intensifier: binary	cardinal number in sentence: binary	
is intensifier: binary	pronoun in sentence: binary	
modifies strongsubj: binary	modal in sentence (other than will): binary	
modifies weaksubj: binary		
modified by strongsubj: binary		
modified by weaksubj: binary		

Table 3: Features for neutral-polar classification

Outline

Experiments for Subjectivity Classification

Compare:

word token Use only token as feature.

word+priorpol Use token and prior polarity as features.

28 features Use all 28 features.

Setup:

- 10-fold cross-validation on train/test data.
- Use BoosTexter AdaBoost.HM (?MH?) classifier
- 5000 rounds of boosting.

Experiments for Subjectivity Classification

Compare:

word token Use only token as feature.

word+priorpol Use token and prior polarity as features.

28 features Use all 28 features.

Setup:

- 10-fold cross-validation on train/test data.
- Use BoosTexter AdaBoost.HM (?MH?) classifier
- 5000 rounds of boosting.

Results

	Acc	Polar Rec	Polar Prec	Polar F	Neut Rec	Neut Prec	Neut F
word token	73.6	45.3	72.2	55.7	89.9	74.0	81.2
word+priorpol	74.2	54.3	68.6	60.6	85.7	76.4	80.7
28 features	75.9	56.8	71.6	63.4	87.0	77.7	82.1

Table 4: Results for Step 1 Neutral-Polar Classification

[The difference in accuracy between the 28-feature classifier and the other two classifiers is significant]

Discussion

- Using only the word token gives a higher precision than the 28-feature classifier, but lower recall.
- Polar recall is still relatively low.
- 5 671 instances are classified as polar (out of 19 506 instances).

Outline

Summary

- Sentiment words have a prior polarity (out of context) and a contextual polarity.
- A considerable amount of mistakes when considering only prior polarity for classifying the polarity of words in context comes from words with non-neutral prior polarity appearing in phrases with neutral contextual polarity..
- Two steps to contextual polarity classification:
  1. classify subjectivity,
  2. for subjective (polar) words, classify polarity.
- A subjectivity classifer with 28 features has been presented.
- Reported classification accuracy for the 28-feature classifier is 75.9%.

References

[HM97] Vasileios Hatzivassiloglou and Kathleen McKeown. Predicting the semantic orientation of adjectives. In Proceedings of the Joint ACL/EACL Conference, pages 174-181, 1997.

[SS00] Robert E. Schapire and Yoram Singer. 2000. BoosTexter: A boosting-based system for text categorization. Machine Learning, 39(2/3):135–168.

[RW03] Ellen Riloff and Janyce Wiebe. 2003. Learning extraction patterns for subjective expressions. In EMNLP-2003.

[WWC05] Janyce Wiebe, Theresa Wilson, and Claire Cardie. 2005. Annotating expressions of opinions and emotions in language. Language Resources and Evalution (formerly Computers and the Humanities), 1(2).

[WWH05] Theresa Wilson, Janyce Wiebe and Paul Hoffmann. 2005. Recognizing contextual polarity in phrase-level sentiment analysis. In Proceedings of HLT '05, pages 347-354.