An Unsupervised Aspect-Sentiment Model for Online Reviews

Samuel Brody   Noèmie Elhadad

Department of Biomedical Informatics
Columbia University

NAACL-HLT 2010

Outline

Outline

Online Reviews

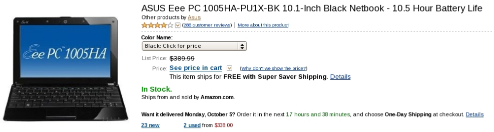

What we have:

- overall score
- details in free-form text

Online Reviews

What we have:

- overall score
- details in free-form text

What we need:

the relevant information in the text, for

- summary
- comparison
- pro/con lists

Why Unsupervised?

- manual annotation may not be feasible
- relevant information is unpredictable
- varying ways of expressing similar meaning
- spelling errors and typos

"The Pitch"

Aspect-Sentiment Model

- simple and elegant
- unsupervised - no labeled training data
- effective
- flexible across domains

Outline

Previous Approaches

Keyword Based

- manual annotation / ask users
- IE techniques (TF-IDF)
- use special lexicons

plus

- adapt across domains
- keyword expansion and clustering

Previous Approaches

Keyword Based

- manual annotation / ask users
- IE techniques (TF-IDF)
- use special lexicons

plus

- adapt across domains
- keyword expansion and clustering

Issues:

- manual annotation - expensive, may overlook important aspects
- keywords can't capture abstract/complex aspects

LDA for Aspects

Latent Dirichlet Allocation (LDA) - Blei et al. (2003)

- unsupervised - infers important topics from data
- can capture complex topics

LDA for Aspects

Latent Dirichlet Allocation (LDA) - Blei et al. (2003)

- unsupervised - infers important topics from data
- can capture complex topics

Issues:

- Problem: global vs. local topics
  (Titov and McDonald, 2008a)
- Problem: mapping topics to aspects
  (Titov and McDonald, 2008b)

LDA for Aspects

Latent Dirichlet Allocation (LDA) - Blei et al. (2003)

- unsupervised - infers important topics from data
- can capture complex topics

Issues:

- Problem: global vs. local topics
  (Titov and McDonald, 2008a)

  Solution: use sentences instead of documents

- Problem: mapping topics to aspects
  (Titov and McDonald, 2008b)

LDA for Aspects

Latent Dirichlet Allocation (LDA) - Blei et al. (2003)

- unsupervised - infers important topics from data
- can capture complex topics

Issues:

- Problem: global vs. local topics
  (Titov and McDonald, 2008a)

  Solution: use sentences instead of documents

- Problem: mapping topics to aspects
  (Titov and McDonald, 2008b)

  Solution: few topics correspond directly to aspects

Model Order - Number of Aspects

Validation Procedure (for a given k) :

1. For dataset D:
   - Run LDA with k topics on D
2. Sample random subset D^i of size $\delta|D|$
   - Run LDA on D^i
   - Calculate consistency between full and partial topics
3. Repeat 2nd step q times.
4. Return the average score over q iterations.

Data

- Restaurants - 50,000 restaurant reviews from Citysearch NY (`http://newyork.citysearch.com/`)
  - 3,400 sentences annotated by Ganu et al. (2009)
- Products - from Amazon (`http://www.amazon.com/`)
  - 1,086 reviews for four leading netbooks
  - 586 reviews for watches

Results in the Restaurant Domain

Inferred Aspect	Representative Words	Manual Aspect
Food - General Wine & Drinks Dishes Bakery	menu, fresh, sushi, fish, chef, cuisine wine, list, glass, drinks, beer, bottle chicken, sauce, rice, cheese, spicy, salad, hot, delicious, dessert, bagels, bread, chocolate	Food & Drink
Ambiance / Mood Physical Atmosphere	great, atmosphere, wonderful, music, experience bar, room, outside, seating, tables, cozy, loud	Atmosphere
Staff Service	service, staff, friendly, attentive, busy, slow table, order, wait, minutes, reservation, forgot	Staff
Value	portions, quality, worth, size, cheap	Price
Anec. - experience Anec. - location	dinner, night, group, friends, date, family out, back, definitely, around, walk, block	Anecdotes
Recommendation Location Misc. Description	best, top, favorite, city, NYC restaurant, found, Paris, (New) York, location place, eat, enjoy, big, often, stuff	Misc.

Results in the Restaurant Domain

Inferred Aspect	Representative Words	Manual Aspect
Food - General Wine & Drinks Dishes Bakery	menu, fresh, sushi, fish, chef, cuisine wine, list, glass, drinks, beer, bottle chicken, sauce, rice, cheese, spicy, salad, hot, delicious, dessert, bagels, bread, chocolate	Food & Drink
Ambiance / Mood Physical Atmosphere	great, atmosphere, wonderful, music, experience bar, room, outside, seating, tables, cozy, loud	Atmosphere
Staff Service	service, staff, friendly, attentive, busy, slow table, order, wait, minutes, reservation, forgot	Staff
Value	portions, quality, worth, size, cheap	Price
Anec. - experience Anec. - location	dinner, night, group, friends, date, family out, back, definitely, around, walk, block	Anecdotes
Recommendation Location Misc. Description	best, top, favorite, city, NYC restaurant, found, Paris, (New) York, location place, eat, enjoy, big, often, stuff	Misc.

Comparison with MAS - Titov and McDonald (2008b)

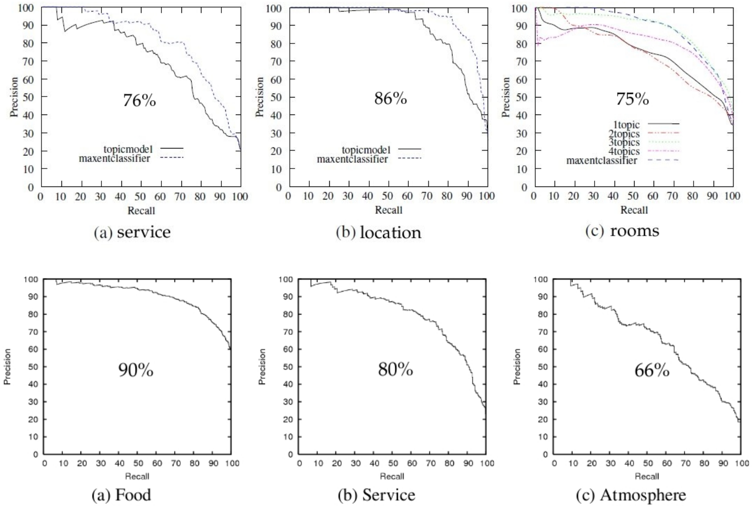

(a) service (b) location (c) rooms

(a) Food (b) Service (c) Atmosphere

Products Domain

Netbooks

Aspect	Representative Words
Performance	power, performance, mode, fan, quiet
Hardware	drive, wireless, bluetooth, usb, speakers, webcam
Memory	ram, 2GB, upgrade, extra, 1GB, speed
Software	using, office, software, installed, works, programs
Usability	internet, video, web, movies, music, email, play
Battery	battery, life, hours, time, cell, last
Size	screen, keyboard, size, small, enough, big

..., Portability, Comparison, Mouse, General, Purchase, Looks, OS

Watches

Aspect	Representative Words
General	buy, perfect, husband, gift, beautiful, deal
Appearance	looks, looking, look, nice, titanium, quality
Display	read, little, date, display, digital, set
Performance	atomic, day, accurate, battery, solar, adjust

Summary

- automatically infer relevant aspects

Summary

- automatically infer relevant aspects
- aspects missed by annotators
  - Food & Drink, Staff & Service, Wireless

Summary

- automatically infer relevant aspects
- aspects missed by annotators
  - Food & Drink, Staff & Service, Wireless
- flexible with regard to domain
  - services, products

Summary

- automatically infer relevant aspects
- aspects missed by annotators
  - Food & Drink, Staff & Service, Wireless
- flexible with regard to domain
  - services, products
- handle non-keyword aspects
  - Food, Atmosphere, Portability

Summary

- automatically infer relevant aspects
- aspects missed by annotators
  - Food & Drink, Staff & Service, Wireless
- flexible with regard to domain
  - services, products
- handle non-keyword aspects
  - Food, Atmosphere, Portability
- handle misspellings and domain specific words
  - *desert, decour/decore, anti-pasta, creme-brule, sandwhich, omlette*
  - six common misspellings of *restaurant*
  - *Korma, Edamame, Dosa, Pho*

Outline

Methodology

Design Principles:

- graph based approach

  (following Hatzivassiloglou and McKeown (1997))

- start with a seed and propagate
- allow for supervised or unsupervised seeds
- keep it aspect-specific

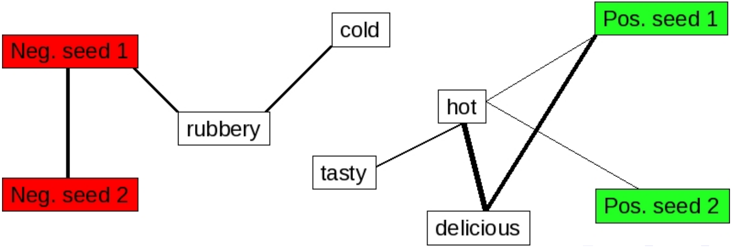

Graph Representation - Edges

"The food was tasty and hot, but our waiter was not friendly."

Graph Representation - Edges

“The food was tasty and hot, but our waiter was not friendly.”

(tasty, food), (hot, food), (not-friendly, waiter)

Graph Representation - Edges

"The food was tasty and hot, but our waiter was not friendly."

⇩

(tasty, food), (hot, food), (not-friendly, waiter)

⇩

```
weight{(tasty,hot)}++;
```

Polarity Seed

Negation Based Seed:

- explicit negation:
  - *friendly* vs. *not friendly*
- morphological indicators:
  - *(dis)courteous, (un)interesting, (in)expensive*

Polarity Seed

Negation Based Seed:

- explicit negation:
  - *friendly* vs. *not friendly*
- morphological indicators:
  - *(dis)courteous*, *(un)interesting*, *(in)expensive*

Discarded:

- disjunctions
  - *“convenient but noisy location”*

Polarity Seed

Negation Based Seed:

- explicit negation:
  - *friendly* vs. *not friendly*
- morphological indicators:
  - *(dis)courteous, (un)interesting, (in)expensive*

Discarded:

- disjunctions
  - *"convenient but noisy location"*
  - *"cheap but tasty food", "dainty but strong necklace"*

Polarity Seed

Negation Based Seed:

- explicit negation:
  - *friendly* vs. *not friendly*
- morphological indicators:
  - *(dis)courteous, (un)interesting, (in)expensive*

Discarded:

- disjunctions
  - *"convenient but noisy location"*
  - *"cheap but tasty food", "dainty but strong necklace"*
- antonyms from dictionary (WordNet)
  - *round* vs. *square*, *cool* vs. *warm*

Human Gold Standard

2. Ambiance (page 1 of 8)

These adjectives describe things related to ambiance in a restaurant (such as decor, music, environment, mood etc.).

*** 1. Please rate each of the adjectives below.**

	Strongly Negative	Weakly Negative	Neutral	Weakly Positive	Strongly Positive	Can't Tell / Unclear
pretty	○	○	○	○	○	○
diverse	○	○	○	○	○	○
artsy	○	○	○	○	○	○
smart	○	○	○	○	○	○

Kendall's tau coefficient (τ_k) and Kendall's distance (D_k)

$$G = \{(a,b) : a, b \in Gold \wedge a \prec b\}$$

$$\tau_k = \frac{|Same| - |Reverse|}{|G|} \qquad D_k = \frac{|Reverse| + p \cdot |Tied|}{|G|}$$

$$-1 \leq \tau_k \leq +1 \qquad 0 \leq D_k \leq 1$$

correlation: more is better

distance: less is better

Evaluation Results

	Auto.		Lexicon	
Aspect	$\tau_k \uparrow$	$D_k \downarrow$	$\tau_k \uparrow$	$D_k \downarrow$
Mood	0.53	0.23	0.56	0.22
Staff	0.57	0.22	0.60	0.20
Main Dishes	0.19	0.40	0.38	0.31
Physical Atmo.	0.34	0.33	0.25	0.37
Bakery	0.33	0.33	0.35	0.33
Food - General	0.19	0.41	0.41	0.30
Wine & Drinks	0.32	0.34	0.52	0.24
Service	0.41	0.30	0.54	0.23
Average	0.36	0.32	0.45	0.27

Summary

- requires no manual annotation

Summary

- requires no manual annotation
- sentiment indicators missed by annotators
  - “neutral” adjectives convey sentiment in a certain domain

Summary

- requires no manual annotation
- sentiment indicators missed by annotators
  - "neutral" adjectives convey sentiment in a certain domain
- sentiment depends on aspect
  - *warm, cheap, busy*

Summary

- requires no manual annotation
- sentiment indicators missed by annotators
  - “neutral” adjectives convey sentiment in a certain domain
- sentiment depends on aspect
  - *warm, cheap, busy*
- handle misspellings and domain specific words
  - *exelent, tastey*
  - *New-Yorky, orgasmic*

Outline

Summary

Aspect-Sentiment Model

- infers relevant information
- flexible across domains
- no annotation / supervision
- robust to noise and error

Future Directions

- aspect:
  - closer integration of aspect and sentiment
  - cross-sentence interaction
- sentiment:
  - other sentiment indicators
  - other graph methods

Thank You!

Bibliography

Blei, David M., Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.

Fahrni, Angela and Manfred Klenner. 2008. Old Wine or Warm Beer: Target-Specific Sentiment Analysis of Adjectives. In *Proc.of the Symposium on Affective Language in Human and Machine, AISB 2008 Convention*. pages 60 – 63.

Ganu, Gayatree, Noemie Elhadad, and Amelie Marian. 2009. Beyond the stars: Improving rating predictions using review text content. In *WebDB*.

Griffiths, Thomas L. and Mark Steyvers. 2004. Finding scientific topics. *Proceedings of the National Academy of Sciences of the United States of America* 101(Suppl 1):5228–5235.

Hatzivassiloglou, Vasileios and Kathleen R. McKeown. 1997. Predicting the semantic orientation of adjectives. In *Proc. of the 35th Annual Meeting of the Association for Computational Linguistics*. ACL, Madrid, Spain, pages 174–181.

Jijkoun, Valentin and Katja Hofmann. 2009. Generating a non-english subjectivity lexicon: Relations that matter. In *Proc. of the 12th Conference of the European Chapter of the ACL (EACL 2009)*. ACL, Athens, Greece, pages 398–405.

Popescu, Ana-Maria and Oren Etzioni. 2005. Extracting product features and opinions from reviews. In *HLT '05: Proc. of the conference on Human Language Technology and Empirical Methods in Natural Language Processing*. ACL, Morristown, NJ, USA, pages 339–346.

Titov, Ivan and Ryan McDonald. 2008a. Modeling online reviews with multi-grain topic models. In *WWW '08: Proc. of the 17th international conference on World Wide Web*. ACM, New York, NY, pages 111–120.

Titov, Ivan and Ryan McDonald. 2008b. A joint model of text and aspect ratings for sentiment summarization. In *Proc. of ACL-08: HLT*. ACL, Columbus, Ohio, pages 308–316.

Turney, Peter. 2002. Thumbs up or thumbs down? semantic orientation applied to unsupervised classification of reviews. In *Proc. of 40th Annual Meeting of the Association for Computational Linguistics*. ACL, Philadelphia, Pennsylvania, USA, pages 417–424.