

Learning Attitudes and Attributes from Online Reviews

Julian McAuley, Jure Leskovec, Dan Jurafsky

Leland Stanford Junior University

November 16, 2012



'Partridge in a Pear Tree', brewed by 'The Bruery'

Dark brown with a light tan head, minimal lace and low retention. Excellent aroma of dark fruit, plum, raisin and red grape with light vanilla, oak, caramel and toffee. Medium thick body with low carbonation. Flavor has strong brown sugar and molasses from the start over bready yeast and a dark fruit and plum finish. Minimal alcohol presence. Actually, this is a nice quad.

Feel: 4.5 Look: 4 Smell: 4.5 Taste: 4 Overall: 4

'Partridge in a Pear Tree', brewed by 'The Bruery'

Dark brown with a light tan head, minimal lace and low retention. Excellent aroma of dark fruit, plum, raisin and red grape with light vanilla, oak, caramel and toffee. Medium thick body with low carbonation. Flavor has strong brown sugar and molasses from the start over bready yeast and a dark fruit and plum finish. Minimal alcohol presence. Actually, this is a nice quad.

Feel: 4.5 Look: 4 Smell: 4.5 Taste: 4 Overall: 4

Tasks

1. Learn language models for each aspect [RR09]
2. Summarize reviews and review corpora [BE10, GEM09, TM08]
3. Recover 'missing' ratings from users reviews [GDFH10, LOCT11]

DATASET	ASPECTS	#REVIEWS
Beer (beeradvocate)	feel, look, smell, taste, overall	1,586,259
Beer (ratebeer)	feel, look, smell, taste, overall	2,924,127
Pubs (beeradvocate)	food, price, quality, selection, service, vibe	18,350
Toys & Games (amazon)	durability, educational, fun, overall	373,974
Audio Books (audible)	author, narrator, overall	10,989
CitySearch [GEM09]	food, ambiance, price, staff	652

- ▶ 1,000 labeled sentences per dataset (labeled by me)
- ▶ 10,000 labeled beer sentences (labeled by oDesk)

Preference and **A**tribute **L**earning from **L**abeled **G**roundtruth and **E**xplicit **R**atings.

Preference and **A**tttribute **L**earning from **L**abeled **G**roundtruth and
Explicit **R**atings.

PALE LAGER

The Pale Lager Model


$P(\text{aspect}(s) = k \mid \text{sentence } s, \text{rating } v) \propto$

$$\exp \sum_{w \in s} \left\{ \underbrace{\theta_{kw}}_{\text{aspect weights}} + \underbrace{\phi_{kv_k w}}_{\text{sentiment weights}} \right\}$$

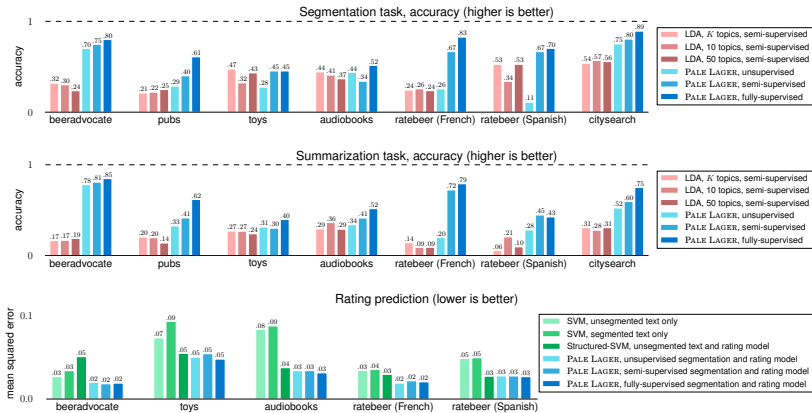
Learning proceeds by choosing the *aspect labels*, and the *model parameters* that maximize the likelihood of the corpus.

The Pale Lager Model

- s_1 : Clear copper colored brew, medium cream colored head.
- s_2 : Floral hop nose, caramel malt.
- s_3 : Caramel malt front dominated by a nice floral hop background.
- s_4 : Grapefruit tones.
- s_5 : Very tasty hops run the show with this brew.
- s_6 : Thin to medium mouth.
- s_7 : Not a bad choice if you're looking for a nice hop treat.



Results



Aspect k

Aspect words θ_k

2-star sentiment words

5-star sentiment words

$\phi_{k,2}$

$\phi_{k,5}$

Feel



Look



Smell



Taste



Overall



Bibliography



S. Brody and N. Elhadad.

An unsupervised aspect-sentiment model for online reviews.

In *ACL, 2010*.



N. Gupta, G. Di Fabbrizio, and P. Haffner.

Capturing the stars: predicting ratings for service and product reviews.

In *HLT Workshops, 2010*.



G. Ganu, N. Elhadad, and A. Marian.

Beyond the stars: Improving rating predictions using review text content.

In *WebDB, 2009*.



B. Lu, M. Ott, C. Cardie, and B. Tsou.

Multi-aspect sentiment analysis with topic models.

In *Workshop on SENTIRE, 2011*.



D. Rao and D. Ravichandran.

Semi-supervised polarity lexicon induction.

In *ACL, 2009*.



I. Titov and R. McDonald.

A joint model of text and aspect ratings for sentiment summarization.

In *ACL, 2008*.