Steklov Mathematical Institute
St. Petersburg Department
27 Fontanka, St. Petersburg, Russia
https://ai.pdmi.ras.ru/

SAMSUNG

Improving unsupervised neural aspect extraction for online discussions using out-of-domain classification

Anton Alekseev, E. Tutubalina, V. Malykh, and S. Nikolenko

work was supported by Samsung Research

7th International Symposium on Language & Knowledge Engineering
Dublin, Ireland
29-31 Oct 2019

Aspect Extraction

"The **stew** was hot and delicious"

Objects of interest: aspects of the entities,
on which the opinions have been expressed

The tasks:

1) extract "**stew**" as an aspect
2) group other aspects of the similar kind into one **cluster**,
   "stew", "mole", "borscht", "goulash" … [~"**food**"?]

Different methods: rule-based, supervised learning, unsupervised learning

Unsupervised Aspect Extraction

- does not rely on labeled data
- allows to work with new domains by design

Dominant approaches until recently: BTM and LDA-based topic modeling variants; each aspect = topic

ACL2017, ABAE: pretrained word embeddings + self-attention

An Unsupervised Neural Attention Model for Aspect Extraction

Ruidan He[†‡], **Wee Sun Lee**[†], **Hwee Tou Ng**[†], and **Daniel Dahlmeier**[‡]
[†]Department of Computer Science, National University of Singapore
[‡]SAP Innovation Center Singapore
[†]{ruidanhe,leews,nght}@comp.nus.edu.sg
[‡]d.dahlmeier@sap.com

Inferred Aspects	**Representative Words**	**Gold Aspects**
Main Dishes	beef, duck, pork, mahi, filet, veal	Food
Dessert	gelato, banana, caramel, cheesecake, pudding, vanilla	
Drink	bottle, selection, cocktail, beverage, pinot, sangria	
Ingredient	cucumber, scallion, smothered, stewed, chilli, cheddar	
General	cooking, homestyle, traditional, cuisine, authentic, freshness	
Physical Ambience	wall, lighting, ceiling, wood, lounge, floor	Ambience
Adjectives	intimate, comfy, spacious, modern, relaxing, chic	
Staff	waitstaff, server, staff, waitress, bartender, waiter	Staff
Service	unprofessional, response, condescending, aggressive, behavior, rudeness	
Price	charge, paid, bill, reservation, came, dollar	Price
Anecdotes	celebrate, anniversary, wife, fiance, recently, wedding	Anecdotes
Location	park, street, village, avenue, manhattan, brooklyn	Misc.
General	excellent, great, enjoyed, best, wonderful, fantastic	
Other	aged, reward, white, maison, mediocrity, principle	

What is ABAE, in brief: the model

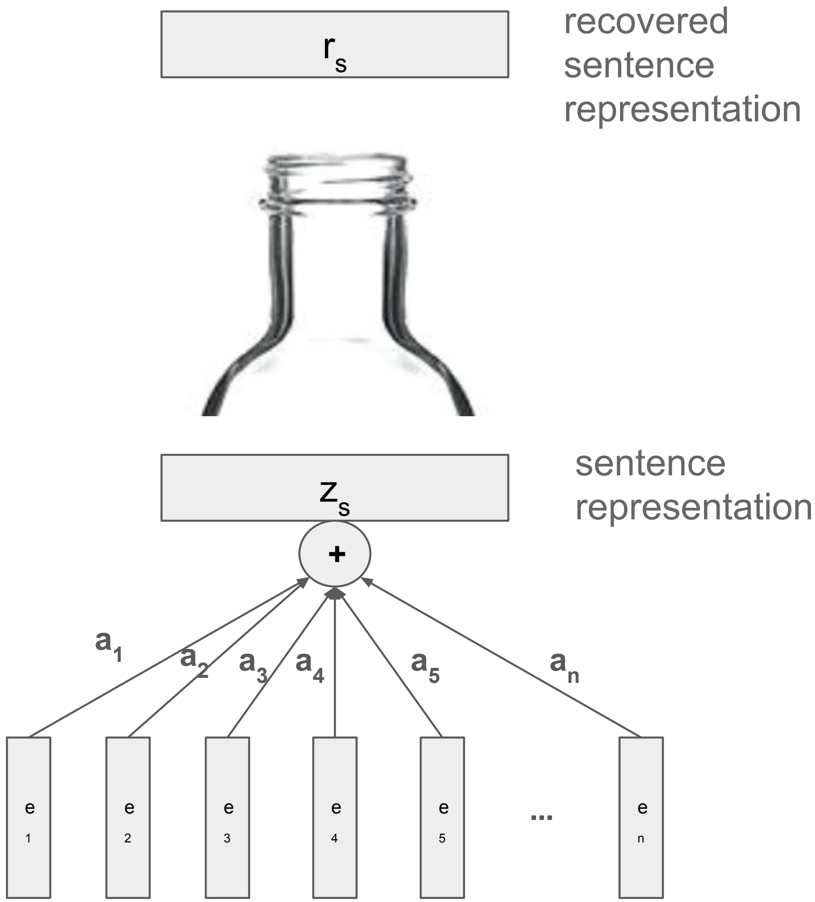

Two linear feedforward layers:

$$\mathbf{r}_s = \mathbf{T}^\top \cdot \mathbf{p}_t$$

$$\mathbf{p}_t = softmax(\mathbf{W} \cdot \mathbf{z}_s + \mathbf{b})$$

T -- aspects embedding matrix **K** x **d**
(K aspects representations, each with the same number of dimensions as word embeddings)

$\mathbf{p_t}$ -- a “share” of each aspect in a sentence in concern

Img src: http://mariegalerie.com/showcase/

What is ABAE, in brief: training

Negative sampling and max margin loss

$$J(\theta) = \sum_{s \in D} \sum_{i=1}^{m} \max(0, 1 - \mathbf{r}_s \mathbf{z}_s + \mathbf{r}_s \mathbf{n}_i)$$

+ the loss function that promotes topic diversity:
(**U** reaches maximum when **T** is orthogonal)

$$U(\theta) = \|\mathbf{T}_n \cdot \mathbf{T}_n^\top - \mathbf{I}\|$$

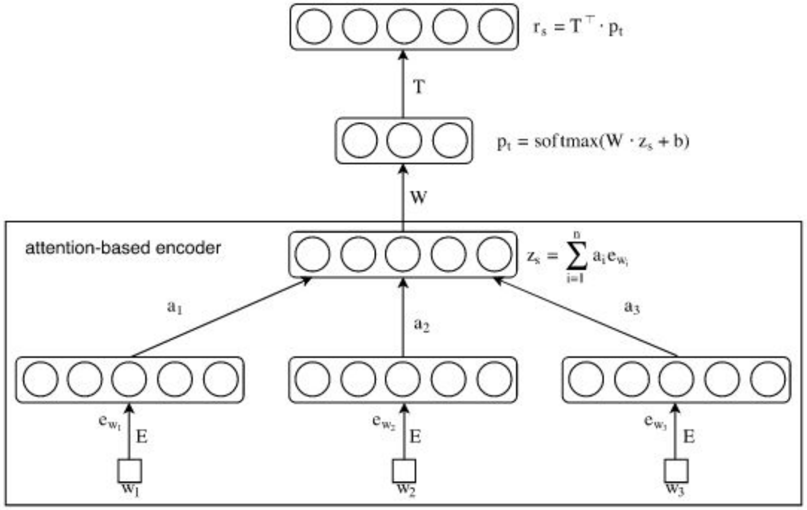

Figure 1: An example of the ABAE structure.

Why focusing on ABAE?

89 citations* and numerous diverse applications, including:

- **Extractive summaries** from multiple reviews
  [Angelidis et al. 2018] 1808.08858

- Summary extraction, **user profiling**
  [Micheltree et al. 2018] 1804.08666

- Text-based **recommender systems**
  [Nikolenko et al. 2019] 1901.07829

* Google Scholar, 29.10.2019

Idea: what if we apply to non-review data?

...to [possibly] enhance other tasks
[as topic modeling already did]

Yields aspects of challengeable quality!

No surprise: trained on sentences ⇨
we implicitly assume there are aspects in **each** sentence

Possible cause: in non-review texts,
authors are *less focused* on the topic/object of discussion ⇨
not every sentence is on the topic

RQ: can we improve the aspects coherence with doing data preprocessing in a slightly more sophisticated manner?

Improving coherence with text preprocessing

- **tweet pooling** by hashtags in order to improve LDA topics [Mehrotra et al.’16]

- term-weighting approach for the LDA input in order to **promote named entities** [Krasnashchok et al.’18]

- thesaurus **relations-based** LDA weights modifications improve coherence [Loukachevitch et al.’18]

Proposed approach: removing out-of-domain sentences

Given:

- ID (in-domain) target text collection we are to extract aspects from
- OOD (out-of-domain): unrelated, out-of-domain texts (collected)

Method:

1) split all texts into sentences
2) train a probabilistic classifier separating ID sentences from OOD sentences
3) compute the trained classifier scores for all the sentences in ID
4) remove sentences with scores lower than certain threshold
5) train the aspect extraction model on the remaining sentences

Out-of-domain classifier's scores for sentences from the *sci.electronics* newsgroup

Score	**Sentence after preprocessing with NLTK**
0.844	paul simundza writes probably tell dc blocking capacitor series one chip single ended audio amp spe
0.836	open look power amp ic
0.047	fairly obvious
0.466	replace one connected dead output
0.668	well one thing poke around terminal power amp chip

Experimental settings

Data: selected diverse topics from 20Newsgroups

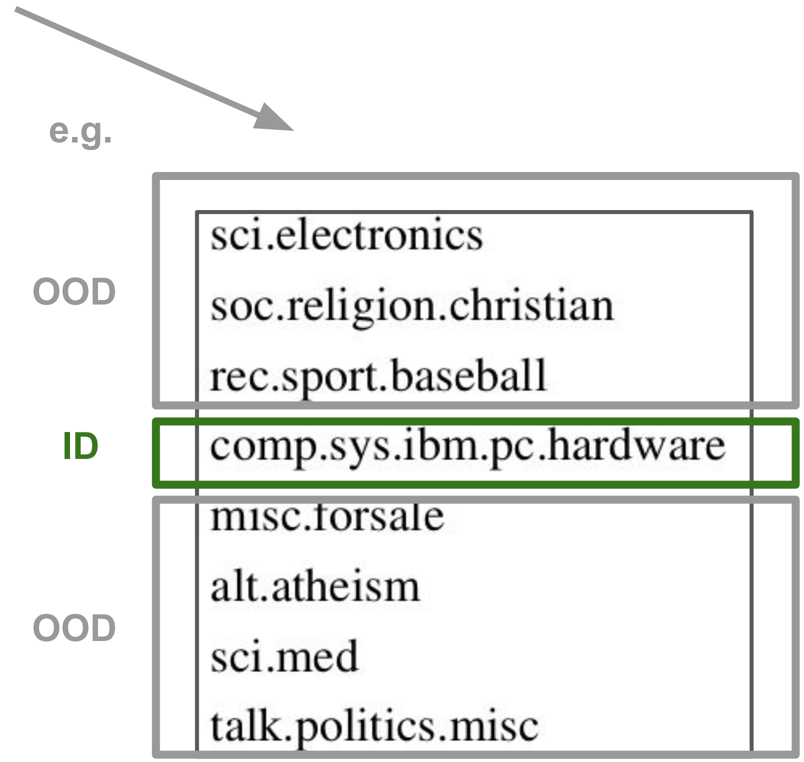

Model (ABAE):
- 15 aspects (topics)
- 20 negative samples
- 10 epochs
- batch size of 256 on one GPU

word2vec: SGNS vectors, trained on the corresponding domain (newsgroup)
dimension is 200,
window size equals 10,
5 negative samples

Baseline: OnlineLDA model [Hoffman et al.' 10] trained with gensim [Rehurek et al.' 10] with default parameters (same vocabulary, same number of aspects)

Experiment results: aspects

ABAE trained on all sentences in a post (less coherent)	ABAE trained on selected sentences (more coherent)
<num> <pad> raffle anyone copy	wiring green cable box gfci grounded case
time frequency chip source take much	voltage input supply output signal power circuit
<num>greggo <unk>mc68882rc33 <pad> raffle	edu university uk mail fax email internet
<unk>raffle <pad>greggo mc68882rc33 <num> ca input	dc digital wave drive per data decimal state
dtmedin b30 catbyte ingr uunet com uucp look	com dtmedin b30 catbyte uunet ingr uucp al everywhere
mail edu university writes com email uk	radar detector number someone radio law shack
copy anyone know could would help get	ca mb bison baden inqmind de sys6626 mind bb bari
edu university uk henry toronto mail	best year around machine least seems band
input pin output data latch voltage	phone neoucom departmentedu oh usa computer uhura
input output data voltage pin high	pin input latch output data voltage supply
ca mb bison baden inqmind de sys6626 bb	ground wire neutral conductor box outlet grounding
connected outlet hot wire grounding neutral	uk mail university email com edu fax internet
ground wire conductor neutral outlet connected	would anyone know copy get could want
mc68882rc33 <pad>greggo input raffle voltage	pin input neutral voltage connected wire current
phone neoucom edu department oh usa computer service	copy anyone know could would help

Evaluation: PMI and NPMI

Standard PMI-coherence: averaged per-topic PMI value for every pair of top **N** tokens computed either on the training set or the heldout data

$$C_{\text{PMI}} = \frac{2}{N(N-1)} \sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \text{PMI}(w_i, w_j),$$

where

$$\text{PMI}(w_i, w_j) = \log \frac{P(w_i, w_j) + \epsilon}{P(w_i)P(w_j)};$$

And its normalized modification

$$\text{NPMI}(w_i, w_j) = \left(\frac{\text{PMI}(w_i, w_j)}{-\log P(w_i, w_j) + \epsilon} \right)$$

Evaluation

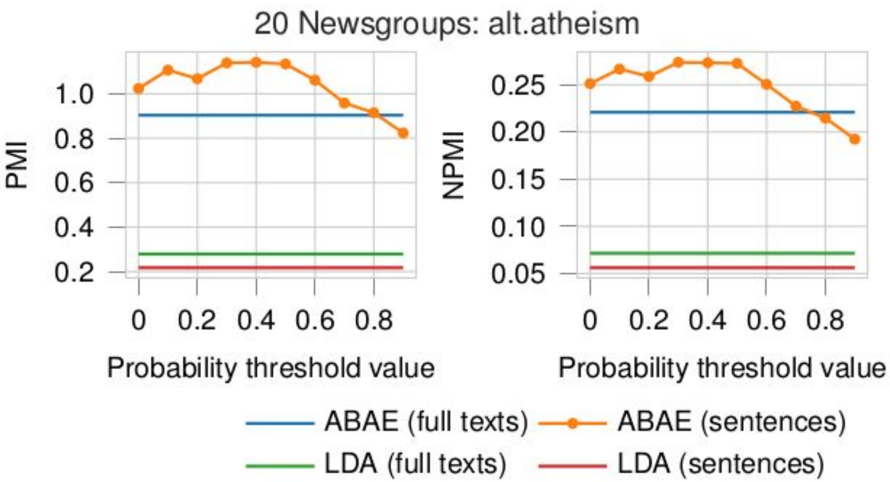

20 Newsgroups: alt.atheism
PMI
1.0
0.8
0.6
0.4
0.2
NPMI
0.25
0.20
0.15
0.10
0.05
0 0.2 0.4 0.6 0.8
Probability threshold value
ABAE (full texts)
ABAE (sentences)
LDA (full texts)
LDA (sentences)

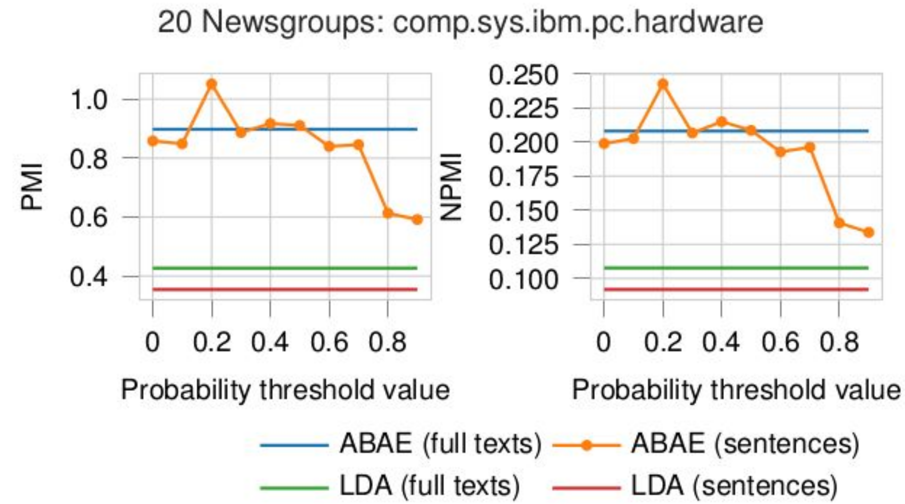

20 Newsgroups: comp.sys.ibm.pc.hardware
PMI
1.0
0.8
0.6
0.4
NPMI
0.250
0.225
0.200
0.175
0.150
0.125
0.100
0 0.2 0.4 0.6 0.8
Probability threshold value
ABAE (full texts)
ABAE (sentences)
LDA (full texts)
LDA (sentences)

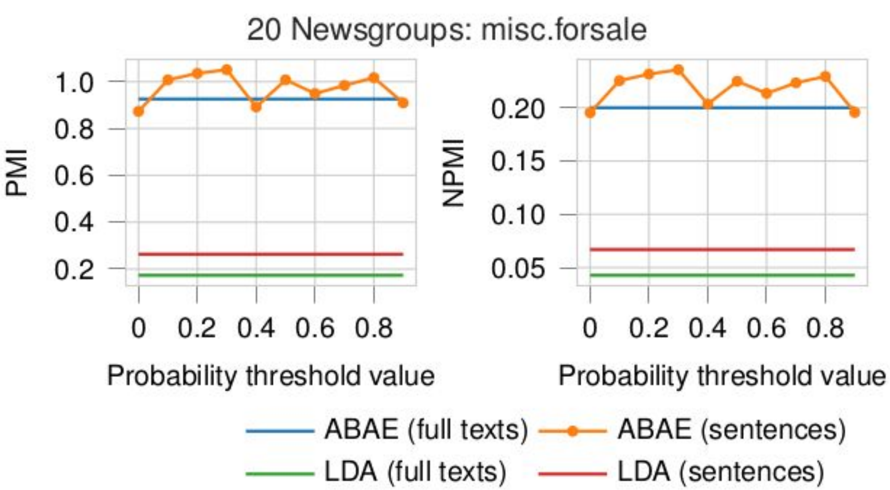

20 Newsgroups: misc.forsale
PMI
1.0
0.8
0.6
0.4
0.2
NPMI
0.20
0.15
0.10
0.05
0 0.2 0.4 0.6 0.8
Probability threshold value
ABAE (full texts)
ABAE (sentences)
LDA (full texts)
LDA (sentences)

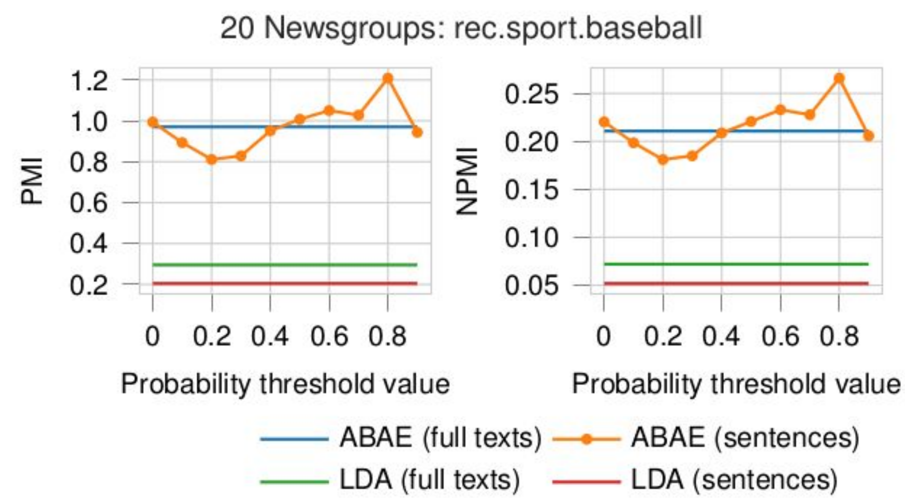

20 Newsgroups: rec.sport.baseball
PMI
1.2
1.0
0.8
0.6
0.4
0.2
NPMI
0.25
0.20
0.15
0.10
0.05
0 0.2 0.4 0.6 0.8
Probability threshold value
ABAE (full texts)
ABAE (sentences)
LDA (full texts)
LDA (sentences)

Evaluation

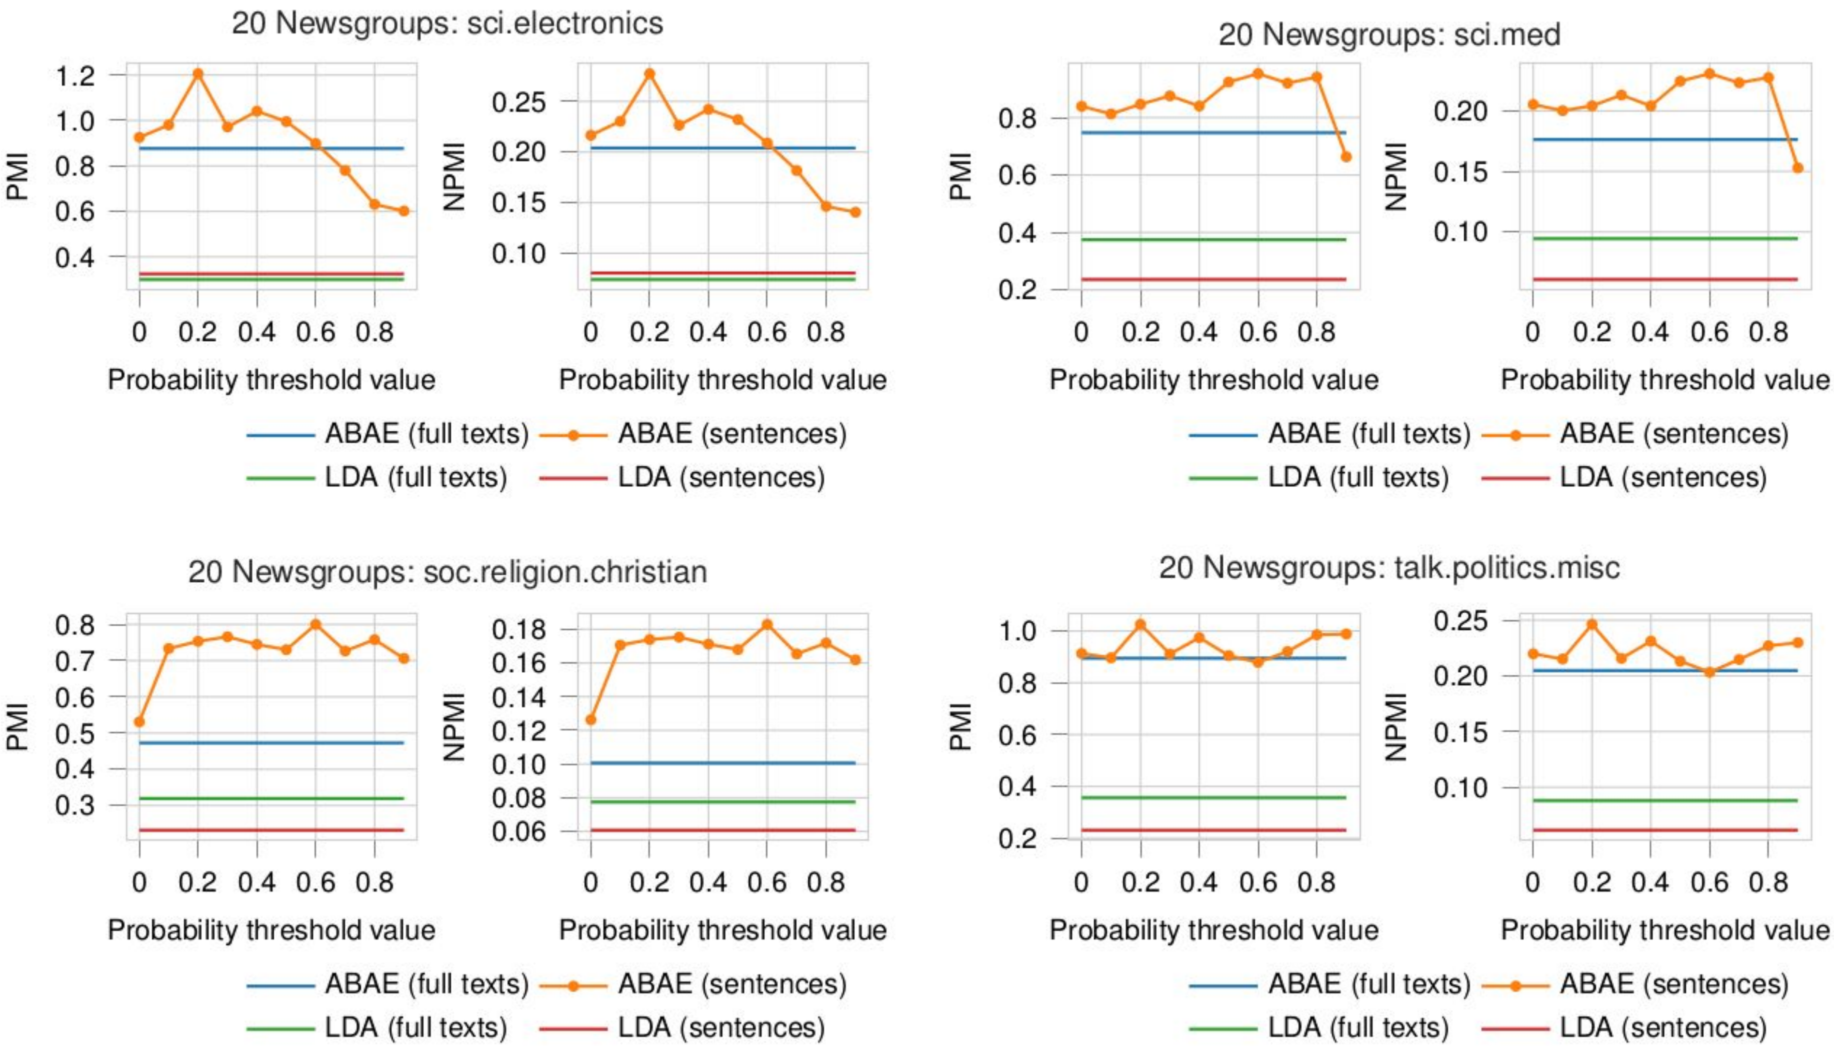
20 Newsgroups: sci.electronics
20 Newsgroups: sci.med
20 Newsgroups: soc.religion.christian
20 Newsgroups: talk.politics.misc
PMI
NPMI
Probability threshold value
ABAE (full texts)
ABAE (sentences)
LDA (full texts)
LDA (sentences)

Conclusion and future work

Proposed technique can improve aspects coherence -- even with a simple discriminative BoW classifier without proper tuning

Future work:

- try more advanced classifications methods
- develop a reliable technique for the filtering threshold selection and make it a tool

References

M. Hoffman, F.R. Bach and D.M. Blei, Online learning for latent dirichlet allocation, in: advances in neural information processing systems, 2010, pp. 856–864.

K. Krasnashchok and S. Jouili, *Improving Topic Quality by Promoting Named Entities in Topic Modeling*, in: Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers), 2018, pp. 247–253.

N. Loukachevitch, K. Ivanov and B. Dobrov, *Thesaurus-Based Topic Models and Their Evaluation*, in: Proceedings of the 8th International Conference on Web Intelligence, Mining and Semantics, ACM, 2018, p. 11.

R. Mehrotra, S. Sanner, W. Buntine and L. Xie, *Improving lda topic models for microblogs via tweet pooling and automatic labeling*, in: Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval, ACM, 2013, pp. 889–892.

C. Mitcheltree, V. Wharton and A. Saluja, *Using Aspect Extraction Approaches to Generate Review Summaries and UserProfiles*, in: Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 3 (Industry Papers), 2018, pp. 68–75.

S.I. Nikolenko, E. Tutubalina, V. Malykh, I. Shenbin and A. Alekseev, *AspeRa: Aspect-Based Rating Prediction Model*,in: Advances in Information Retrieval, L. Azzopardi, B. Stein,N. Fuhr, P. Mayr, C. Hauff and D. Hiemstra, eds, Springer International Publishing, Cham, 2019, pp. 163–171. ISBN ISBN978-3-030-15719-7.

R. Rehurek and P. Sojka, Software Framework for Topic Mod-elling with Large Corpora, in: Proceedings of the LREC 2010Workshop on New Challenges for NLP Frameworks, ELRA,Valletta, Malta, 2010, pp. 45–50, http://is.muni.cz/publication/884893/en.

On topic modeling evaluation:

G. Bouma, Normalized (pointwise) mutual information in col-location extraction, Proceedings of GSCL (2009), 31–40.
D. Newman, S. Karimi and L. Cavedon, *External evaluation of topic models*, in: in Australasian Doc. Comp. Symp., 2009, Citeseer, 2009
J.H. Lau, D. Newman and T. Baldwin, Machine reading tealeaves: Automatically evaluating topic coherence and topicmodel quality, in: Proceedings of the 14th Conference of theEuropean Chapter of the Association for Computational Lin-guistics, 2014, pp. 530–539.

Steklov Mathematical Institute
St. Petersburg Department
27 Fontanka, St. Petersburg, Russia
https://ai.pdmi.ras.ru/

SAMSUNG

Thank you for your attention!
Q&A

Improving unsupervised neural aspect extraction for online discussions using out-of-domain classification

Anton Alekseev, E. Tutubalina, V. Malykh and S. Nikolenko
anton.m.alexeyev@gmail.com

7th International Symposium on Language & Knowledge Engineering
Dublin, Ireland
29-31 Oct 2019