

Singapore Management University

Institutional Knowledge at Singapore Management University

Research Collection School Of Computing and Information Systems

School of Computing and Information Systems

7-2016

Build emotion lexicon from the mood of crowd via topic-assisted joint non-negative matrix factorization

Kaisong SONG

Wei GAO

Singapore Management University, weigao@smu.edu.sg

Ling CHEN

Shi FENG

Daling WANG

See next page for additional authors

Follow this and additional works at: https://ink.library.smu.edu.sg/sis_research



Part of the [Databases and Information Systems Commons](#)

Citation

SONG, Kaisong; GAO, Wei; CHEN, Ling; FENG, Shi; WANG, Daling; and ZHANG, Chengqi. Build emotion lexicon from the mood of crowd via topic-assisted joint non-negative matrix factorization. (2016). *Proceedings of the 38th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval (SIGIR 2016)*. 773-776.

Available at: https://ink.library.smu.edu.sg/sis_research/4568

This Conference Proceeding Article is brought to you for free and open access by the School of Computing and Information Systems at Institutional Knowledge at Singapore Management University. It has been accepted for inclusion in Research Collection School Of Computing and Information Systems by an authorized administrator of Institutional Knowledge at Singapore Management University. For more information, please email cherylds@smu.edu.sg.

Author

Kaisong SONG, Wei GAO, Ling CHEN, Shi FENG, Daling WANG, and Chengqi ZHANG

Build Emotion Lexicon from the Mood of Crowd via Topic-Assisted Joint Non-negative Matrix Factorization

Kaisong Song¹, Wei Gao², Ling Chen³, Shi Feng^{1,4}, Daling Wang^{1,4}, Chengqi Zhang³

¹School of Computer Science and Engineering, Northeastern University, Shenyang, China

²Qatar Computing Research Institute, Hamad Bin Khalifa University, Doha, Qatar

³Centre for Quantum Computation and Intelligent Systems, University of Technology, Sydney, Australia

⁴Key Laboratory of Medical Image Computing (Northeastern University), Ministry of Education, China
kaisongsong@gmail.com, {fengshi, wangdaling}@cse.neu.edu.cn
wgao@qf.org.qa, {ling.chen, chengqi.zhang}@uts.edu.au

ABSTRACT

In the research of building emotion lexicons, we witness the exploitation of crowd-sourced affective annotation given by readers of online news articles. Such approach ignores the relationship between topics and emotion expressions which are often closely correlated. We build an emotion lexicon by developing a novel joint non-negative matrix factorization model which not only incorporates crowd-annotated emotion labels of articles but also generates the lexicon using the topic-specific matrices obtained from the factorization process. We evaluate our lexicon via emotion classification on both benchmark and built-in-house datasets. Results demonstrate the high-quality of our lexicon.

Keywords

emotion lexicon; joint NMF; emotion classification

1. INTRODUCTION

A basic task in sentiment analysis is classifying the sentiment polarity (positive or negative) of the given subjective text [8, 11, 13]. However, the binary scheme may be oversimplified. Recently, *emotion analysis* represents a natural evolution of *sentiment analysis* by modeling finer-grained emotions, e.g., happy, sad, angry, etc. [9, 14].

Emotion lexicons are the essential resources for emotion analysis. Compared to sentiment lexicons such as SentiWordNet¹ [1], where each entry is typically labeled with sentiment polarity, emotion lexicons are more complex in the sense that each entry may convey a mixture of multiple emotions which bear different emotion intensity² [14].

Most of the existing lexicon construction approaches [3, 12, 15] are based upon a set of hand-coded seed words. Consequently, the quality of lexicons is sensitive to the manual

¹<http://sentiwordnet.isti.cnr.it/>

²git.io/MqyoIg

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

SIGIR '16, July 17 - 21, 2016, Pisa, Italy

© 2016 Copyright held by the owner/author(s). Publication rights licensed to ACM. ISBN 978-1-4503-4069-4/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2911451.2914759>



Figure 1: The emotion distribution generated by “mood meter” on a news article

seed selection. Today, many news websites (e.g., rappler.com, corriere.it, etc.) allow users to express their feelings about an article with a simple click on a given set of emoticons. Figure 1 shows the emotion distribution based on such kind of votes from the crowd regarding an article on rappler.com via an GUI called “mood meter” which is embedded in each of its web page.

Staiano and Guerini [14] proposed a compositional semantics method that utilized crowd-based affective annotation, where they represented words and emotions in a high-dimensional space based on their occurrences in the document. A deficiency is that they ignored the versatility of affections among various contexts. It cannot distinguish accurately the emotion of words by disregarding different topics where words exist. Since documents and topics are of many-to-many correspondence in a collection, it would be more useful to consider emotions at topic level. Some researchers tried to model topic and sentiment simultaneously [4, 8] for joint sentiment-topic analysis. However, no work has considered topics when building emotion lexicons.

Intuitively, emotion expressions are pertinent to the topics in which they reside. For example, “predictable” suggests happiness for stock market, but for a movie it implies disappointment or even anger. We expect that topic-assisted workaround can produce finer-grained and more accurate entries for emotion lexicons. In this paper, we develop a novel joint non-negative matrix factorization model which associates words with emotions in a low-dimensional semantic space based on hidden topics. An emotion lexicon is built from word-topic and emotion-topic factor matrices, which result from the joint model, using matrix composition.

2. RELATED WORK

Emotion lexicons are typically built based on a set of seed words [3, 12, 14, 15]. Xu et al. [15] proposed a graph-based algorithm which ranked words according to a few manually selected seed words. Song et al. [12] and Feng et al. [3] improved this method by supplementing seed words with graphical emoticons or combined their effects. Differently, Staiano and Guerini [14] proposed a compositional semantics method using crowd-annotated articles crawled from the Internet. In this paper, we also resort to crowd-annotated articles while we incorporate topic-emotion relationship for lexicon construction which was not considered previously.

Non-negative matrix factorization (NMF) has been widely used in image or text representation. Lee and Seung [5] investigated the properties of the algorithm and emphasized the clustering aspect. Xu et al. [16] applied standard NMF to document clustering. In recent years, different extensions [7, 10] have been proposed for solving sentiment analysis and sentiment lexicon construction. Li et al. [7] creatively applied orthogonal NMF, proposed by Ding et al. [2], to sentiment classification by incorporating lexical prior knowledge. Peng and Park [10] proposed a constrained symmetric NMF method for sentiment lexicon construction, which considers synonyms and antonyms simultaneously.

Lee et al. [6] used a generic semi-supervised NMF (SSNMF) method which jointly incorporates the data matrix and the (partial) class label matrix into NMF. We base our model on SSNMF for lexicon construction by incorporating different factorization schemes for the supervision matrix, which naturally results in a lexicon from the estimated factor matrices. To our knowledge, this is the first attempt for building fine-grained emotion lexicon based on NMF models.

3. PRELIMINARIES

We first introduce a compositional semantics method [14] for building an emotion lexicon from crowd-annotated news. Then, we review semi-supervised NMF [6] which paves the way for developing our lexicon generation method.

3.1 Compositional Semantics Method (CS)

Let $D = \{d_1, \dots, d_{|D|}\}$ be a set of documents, and $W = \{w_1, \dots, w_{|W|}\}$ be the complete vocabulary set of the whole corpus. We define a word-document matrix \mathbf{M}_{WD} of size $|W| \times |D|$, using (w, d) as index and $M_{WD(w,d)}$ as entry value based on raw frequencies f , normalized frequencies nf or $tf-idf$. Given emotion set $E = \{e_1, \dots, e_{|E|}\}$, we can represent emotion labels from crowd-annotated resources as a document-emotion matrix \mathbf{M}_{DE} of size $|D| \times |E|$ whose entry values are based on crowd-sourced affective annotation (see Figure 1). Staiano and Guerini [14] built a word-emotion matrix \mathbf{M}_{WE} using the compositional semantics (CS) method by multiplying matrices \mathbf{M}_{WD} and \mathbf{M}_{DE} :

$$\mathbf{M}_{WE} = \mathbf{M}_{WD}\mathbf{M}_{DE} \quad (1)$$

An emotion lexicon can be obtained by first applying column-wise normalization to \mathbf{M}_{WE} and then scaling its row-wise data that sums up to one.

3.2 Semi-Supervised NMF (SSNMF)

Non-negative Matrix Factorization (NMF) [5] is an unsupervised algorithm widely used in image or text representation. A generic semi-supervised NMF (SSNMF) algorithm [6] was proposed to incorporate the data matrix

$\mathbf{X} = [x_1, \dots, x_n] \in \mathbb{R}_+^{m \times n}$, where m is the dimension of data vectors, and the class label matrix $\mathbf{Y} = [y_1, \dots, y_n] \in \mathbb{R}_+^{c \times n}$, where c is the number of classes. The objective function, which involves the non-negative two-factor decomposition of \mathbf{X} and \mathbf{Y} sharing a common factor matrix $\mathbf{S} \in \mathbb{R}_+^{k \times n}$, is to minimize the following:

$$\mathcal{J} = \|\mathbf{U} \odot (\mathbf{X} - \mathbf{AS})\|^2 + \alpha \|\mathbf{V} \odot (\mathbf{Y} - \mathbf{BS})\|^2 \quad (2)$$

where α is a tradeoff parameter adjusting the importance of the supervised term, $\mathbf{A} \in \mathbb{R}_+^{m \times k}$ and $\mathbf{B} \in \mathbb{R}_+^{c \times k}$ are basis matrices for \mathbf{X} and \mathbf{Y} , \mathbf{U} and \mathbf{V} are weights matrices, both of which can be fixed as *all-ones matrices* to make an NMF that fully uses labeled data [6]. Notation $\|\cdot\|^2$ denotes the squared sum of all the elements in the matrix and \odot represents element-wise product.

4. OUR JOINT NMF MODEL

Inspired from SSNMF, we can jointly model the hidden topics and the explicit crowd-based emotions of articles by customizing the factorization process. Let $T = \{t_1, \dots, t_{|T|}\}$ be a set of topics in low-dimensional space with $|T| \ll \min\{|D|, |W|\}$. Given the word-document matrix \mathbf{M}_{WD} and document-emotion matrix \mathbf{M}_{DE} , we decompose them based on equation 2 by minimizing:

$$\mathcal{J} = \|\mathbf{M}_{WD} - \mathbf{M}_{WT}\mathbf{M}_{DT}^\top\|^2 + \alpha \|\mathbf{M}_{DE}^\top - \mathbf{M}_{ET}\mathbf{M}_{DT}^\top\|^2 \quad (3)$$

to learn the three topic-specific factor matrices \mathbf{M}_{WT} , \mathbf{M}_{DT} and \mathbf{M}_{ET} , where \mathbf{M}_{ET} represents the strength that emotions are associated with topics. Then, we can get the word emotion distributions by using a *variant* of compositional semantics approach (see equation 1) as below:

$$\mathbf{M}'_{WE} = \mathbf{M}_{WT}\mathbf{M}_{ET}^\top \quad (4)$$

A deficiency of directly applying SSNMF is that the emotion modeling is still coarse-grained for lexicon construction which is concerned about word-level emotion. We enhance the model by modeling the emotions of subjective texts using a weighted linear combination of emotion words, which will result in an additional term of the 3-factor decomposition of \mathbf{M}_{DE} based on formula 3 as follow:

$$\mathcal{J}' = \mathcal{J} + \beta \|\mathbf{M}_{DE} - \mathbf{M}'_{WD}\mathbf{M}_{WT}\mathbf{M}_{ET}^\top\|^2 \quad (5)$$

where β is a tradeoff parameter, \mathbf{M}_{WD} is fixed and considered as word weights, and \mathbf{M}_{WT} and \mathbf{M}_{ET} are variables whose product happens to be \mathbf{M}'_{WE} . With the last term, the joint model aims to improve the estimation of the topic-specific factor matrices by approximating the document-level emotions \mathbf{M}_{DE} based on the word-level emotions \mathbf{M}'_{WE} .

Computation: Factor matrices \mathbf{M}_{WT} , \mathbf{M}_{ET} and \mathbf{M}_{DT} are first randomly initialized, and then updated iteratively by the following updating formulas:

$$\mathbf{M}_{WT} \leftarrow \mathbf{M}_{WT} \odot \frac{\mathbf{M}_{WD}\mathbf{M}_{DT} + \beta\mathbf{M}_{WE}\mathbf{M}_{ET}}{\mathbf{M}_{WT}\mathbf{M}_{DT}^\top\mathbf{M}_{DT} + \beta\mathbf{M}_{WW}\mathbf{M}_{WT}\mathbf{M}_{ET}^\top\mathbf{M}_{ET}}$$

$$\mathbf{M}_{ET} \leftarrow \mathbf{M}_{ET} \odot \frac{\alpha\mathbf{M}_{DE}^\top\mathbf{M}_{DT} + \beta\mathbf{M}'_{WE}\mathbf{M}_{WT}}{\alpha\mathbf{M}_{ET}\mathbf{M}_{DT}^\top\mathbf{M}_{DT} + \beta\mathbf{M}_{ET}\mathbf{M}'_{WT}\mathbf{M}_{WW}\mathbf{M}_{WT}}$$

$$\mathbf{M}_{DT} \leftarrow \mathbf{M}_{DT} \odot \frac{\mathbf{M}'_{WD}\mathbf{M}_{WT} + \alpha\mathbf{M}_{DE}\mathbf{M}_{ET}}{\mathbf{M}_{DT}\mathbf{M}'_{WT}\mathbf{M}_{WT} + \alpha\mathbf{M}_{DT}\mathbf{M}_{ET}^\top\mathbf{M}_{ET}}$$

where $\mathbf{M}_{WW} = \mathbf{M}_{WD}\mathbf{M}_{WD}^\top$, $\mathbf{M}_{WE} = \mathbf{M}_{WD}\mathbf{M}_{DE}$, and the divisions are all element-wise. The update formulas can be easily induced based on the derivatives in [6].

Example entries	afraid	amused	angry	annoyed	dont_care	happy	inspired	sad
crime#n	.000 (.119)	.000 (.088)	.994 (.272)	.000 (.104)	.000 (.085)	.006 (.096)	.000 (.094)	.000 (.142)
dead#a	.000 (.218)	.000 (.059)	.000 (.173)	.000 (.080)	.000 (.075)	.000 (.055)	.000 (.057)	1.00 (.283)
criminal#a	.001 (.145)	.000 (.092)	1.00 (.233)	.000 (.137)	.000 (.086)	.000 (.117)	.000 (.059)	.000 (.131)
interesting#a	.000 (.034)	1.00 (.252)	.000 (.046)	.000 (.098)	.000 (.142)	.000 (.181)	.000 (.206)	.000 (.041)
monitor#v	.511 (.238)	.000 (.098)	.485 (.157)	.000 (.124)	.000 (.078)	.004 (.109)	.000 (.097)	.000 (.099)
funny#a	.000 (.055)	.977 (.278)	.000 (.065)	.000 (.125)	.001 (.203)	.022 (.110)	.000 (.093)	.043 (.071)
sad#a	.000 (.050)	.000 (.081)	.000 (.112)	.000 (.104)	.000 (.172)	.000 (.071)	.000 (.109)	1.00 (.301)

Table 1: Example entries in our constructed emotion lexicon. Emotion scores higher than 20% are highlighted for readability purposes. The brackets enclose the scores given by the CS method as a baseline

Emotion	afraid	amused	angry	annoyed
votes _{avg.}	7.8%	10.6%	10.9%	5.9%
Emotion	dont_care	happy	inspired	sad
votes _{avg.}	5.9%	34.1%	10.3%	14.5%

Table 2: Emotion distribution over Rappler dataset

Lexicon Construction: Given \mathbf{M}_{WT} and \mathbf{M}_{ET} represented in the $|T|$ -dimensional topic space, we can build a word-emotion matrix \mathbf{M}'_{WE} based on equation 4. After normalizing its entries, we obtain a $|E|$ -dimensional vector for each word w :

$$\overline{\mathbf{M}'_{WE(w, \cdot)}} = \frac{1}{Z_w} \langle \mathbf{M}'_{WE(w, 1)}, \dots, \mathbf{M}'_{WE(w, |E|)} \rangle \quad (6)$$

where each element $\mathbf{M}'_{WE(w, e)}$ indexed by (w, e) represents the emotion score of word w belonging to emotion category $e \in E$, and $Z_w = \sum_{e \in E} \frac{\mathbf{M}'_{WE(w, e)}}{\sqrt{\sum_{w \in W} \mathbf{M}'_{WE(w, e)}^2}}$ is the normalization term for w (column-wise normalization ensures that different columns for all the emotions are comparable).

Our created emotion lexicon contains 31,806 entries in total. Table 1 presents several example entries. Similar to [14], we lemmatize and Part-of-Speech (PoS) tag all the documents (PoS we considered are adjective, noun, verb and adverb), and we only keep those *lemma#PoS* entries in the lexicon which also appear in WordNet for eliminating noise words. We can see that each entry has at least one main emotion (e.g., *monitor#v* has two main emotions of *afraid* and *angry*), and our lexicon differentiates the emotions better by assigning discriminative weighting scores as compared to the CS baseline.

5. EXPERIMENTAL EVALUATION

5.1 Data Resources

To build our lexicon, we crawled 31,107 English news articles published before 2015-11-06 from *rappler.com*. We used Stanford CoreNLP³, an integrated suite of natural language processing tools, to tokenize, PoS tag and lemmatize all text data. In Table 2, we report the average percentage of votes for each emotion over all the documents in the corpus. From Table 2, we find that the emotion of *happy* has a lot more votes than the others, which reflects that readers' emotion preference is consistent to the general observation – positive sentiment dominates in real-world data. The crawled resources and generated lexicons have been made publicly available⁴.

To evaluate the lexicon, we applied it for emotion classification on news headlines as [14] did. We used two datasets:

³<http://nlp.stanford.edu/software/corenlp.shtml>

⁴<https://sites.google.com/site/emolexdata/>

Rappler	angry	sad	afraid	happy	inspired	other
SemEval	anger	sadness	fear	joy	surprise	–

Table 3: Emotion label mapping over two test sets

(1) A benchmark dataset from SemEval-2007⁵ on identifying ‘‘Affective Text’’, which contains 1k annotated headlines. As SemEval-2007 test set consists of only six emotions, we adopted an emotion mapping method as displayed in Table 3 to map them to our pre-defined emotions in the lexicon. (2) A built-in-house dataset with total 31k headlines of the crawled *Rappler* articles.

We implement the algorithms using Matlab and run them on a high performance Linux cluster.

5.2 Experiments and Results

We evaluate the quality of the emotion lexicons in two ways: (1) we examine the quality of lexicons created by our method and other competitive methods using the crawled *Rappler* news articles via an emotion classification task; (2) we compare our created lexicons with publicly available state-of-the-art lexicons in similar size.

5.2.1 Parameter Setting

We tune the number of topics $|T|$ by performing a grid search over all values of $10 * x$ with $x \in \{1, 2, \dots, 30\}$. The tradeoff parameters α and β are tuned over all values in $\{0.1, 1, 10, 100\}$. The tuning is based on the performance of emotion classification on the headlines in SemEval’s trial set and a held-out set which consists of 20% headlines randomly selected from the *Rappler* articles. Finally, we set $|T| = 250$, $\alpha = 1$ and $\beta = 10$. We set the number of iterations as 300, which is large enough for ensuring convergence according to our observation on the drop of \mathcal{J} and \mathcal{J}' values.

5.2.2 Comparison of Lexicon Building Methods

For emotion classification, we use a straight-forward voting-based algorithm [12, 14] to assign emotion labels to a test headline h . We conduct element-wise sum over the emotion words in the headline by looking up the lexicon and then average the sums by word counts, i.e., $\overline{V}_h = \frac{1}{Z_h} < \sum_w \overline{\mathbf{M}'_{WE(w, 1)}}, \dots, \sum_w \overline{\mathbf{M}'_{WE(w, |E|)}} >$ where Z_h is the number of emotion words in h . We then normalize \overline{V}_h with the min-max normalization and map each emotion element into a binary decision with fixed thresholds. We set threshold at 0.5 for SemEval-2007 test set and 0.35 for *Rappler* test set, empirically⁶. We use F-1 measure to assess the classification performance on each emotion. Tables 4 and 5 show the re-

⁵<http://nlp.cs.swarthmore.edu/semeval/tasks/>

⁶We set a lower threshold for *Rappler* test set since only a single emotion in each news receives more than 50% votes

Method		fear	anger	joy	sadness	surprise
CS	<i>f</i>	.301	.080	.292	.349	.096
	<i>nf</i>	.329	.090	.291	.386	.081
	<i>tf-idf</i>	.338	.090	.289	.354	.094
Joint (\mathcal{J})	<i>f</i>	.288	.098	.231	.383	.099
	<i>nf</i>	.306	.087	.256	.387	.095
	<i>tf-idf</i>	.372	.091	.283	.357	.132
Joint (\mathcal{J}')	<i>f</i>	.309	.101	.252	.350	.082
	<i>nf</i>	.349	.105	.222	.393	.058
	<i>tf-idf</i>	.361	.082	.270	.359	.133

Table 4: SemEval-2007 emotion classification (F1)

Method		afraid	amused	angry	annoyed
CS	<i>f</i>	.304	.293	.361	.160
	<i>nf</i>	.324	.292	.375	.162
	<i>tf-idf</i>	.341	.277	.377	.145
Joint (\mathcal{J})	<i>f</i>	.331	.289	.370	.208
	<i>nf</i>	.325	.302	.360	.191
	<i>tf-idf</i>	.338	.267	.366	.138
Joint (\mathcal{J}')	<i>f</i>	.333	.279	.371	.198
	<i>nf</i>	.351	.311	.386	.198
	<i>tf-idf</i>	.328	.274	.368	.148
Method		dont_care	happy	inspired	sad
CS	<i>f</i>	.142	.653	.297	.429
	<i>nf</i>	.143	.652	.291	.445
	<i>tf-idf</i>	.133	.654	.284	.427
Joint (\mathcal{J})	<i>f</i>	.177	.596	.270	.430
	<i>nf</i>	.172	.613	.300	.434
	<i>tf-idf</i>	.129	.624	.266	.416
Joint (\mathcal{J}')	<i>f</i>	.191	.608	.275	.400
	<i>nf</i>	.199	.596	.298	.461
	<i>tf-idf</i>	.152	.613	.268	.419

Table 5: Rappler emotion classification (F1)

sults by averaging 20 independent runs (with random initial matrices) on *SemEval* and *Rappler* test sets, respectively.

Our joint models perform better than CS for most emotions especially under *nf* configuration. This indicates that normalized frequency can prevent the bias towards long documents and our method considering topic is effective. Moreover, joint model \mathcal{J}' also performs \mathcal{J} at most cases, implying the usefulness of considering word-level emotion in decomposition. Surprisingly, the results under *tf-idf* configuration are unstable, which suggests that introducing *idf* is sub-optimal. This is because frequent emotion words, e.g., “good”, receive low *tf * idf*, thus are not learned well.

5.2.3 Comparison with Available Emotion Lexicons

We compare our lexicons with the original lexicons released by Staiano and Guerini [14]. We assess them via emotion classification on the larger built-in-house *Rappler* test set. Figure 2 demonstrates that our lexicon configured as *nf* achieves the best results on nearly all emotions, which suggests the high usability of our created lexicon.

6. CONCLUSIONS

We present a joint NMF method which incorporates crowd-based emotion labels on articles and generates topic-specific factor matrices for building emotion lexicons via compositional semantics. Experiments conducted on the benchmark and built-in-house datasets demonstrate our method outperforms the competitive methods on emotion classification. Moreover, our created lexicon outperforms the competitive counterpart on emotion classification task. Our future work will study emotion-specific word embeddings for lexicon construction using deep learning.

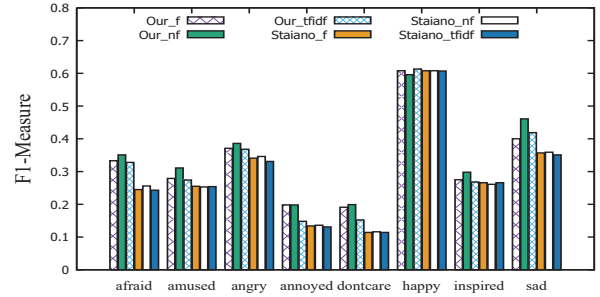


Figure 2: Comparison among different lexicons

Acknowledgments

This work is supported by the National Natural Science Foundation of China (61370074, 61402091), the Fundamental Research Funds for the Central Universities (N130604002, N140404012) and the ARC Discovery Project (DP140100545).

7. REFERENCES

- [1] S. Baccianella, A. Esuli, and F. Sebastiani. Sentiwordnet 3.0: An enhanced lexical resource for sentiment analysis and opinion mining. In *LREC*, pages 2200–2204, 2010.
- [2] C. H. Q. Ding, T. Li, W. Peng, and H. Park. Orthogonal nonnegative matrix t-factorizations for clustering. In *SIGKDD*, pages 126–135, 2006.
- [3] S. Feng, K. Song, D. Wang, and G. Yu. A word-emoticon mutual reinforcement ranking model for building sentiment lexicon from massive collection of microblogs. *World Wide Web*, 18(4):949–967, 2015.
- [4] Y. He, C. Lin, W. Gao, and K.-F. Wong. Dynamic joint sentiment-topic model. *ACM Trans. Intell. Syst. Technol.*, 5(1):6:1–6:21, 2013.
- [5] D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401, 1999.
- [6] H. Lee, J. Yoo, and S. Choi. Semi-supervised nonnegative matrix factorization. *IEEE Signal Process. Lett.*, 17(1):4–7, 2010.
- [7] T. Li, Y. Zhang, and V. Sindhwani. A non-negative matrix tri-factorization approach to sentiment classification with lexical prior knowledge. In *ACL*, pages 244–252, 2009.
- [8] Q. Mei, X. Ling, M. Wondra, H. Su, and C. Zhai. Topic sentiment mixture: modeling facets and opinions in weblogs. In *WWW*, pages 171–180, 2007.
- [9] S. M. Mohammad and P. D. Turney. Crowdsourcing a word-emotion association lexicon. *Computational Intelligence*, 29(3):436–465, 2013.
- [10] W. Peng and D. H. Park. Generate adjective sentiment dictionary for social media sentiment analysis using constrained nonnegative matrix factorization. In *ICWSM*, pages 273–280, 2011.
- [11] K. Song, L. Chen, W. Gao, S. Feng, D. Wang, and C. Zhang. Persentiment: A personalized sentiment classification system for microblog users. In *WWW (Companion Volume)*, pages 255–258, 2016.
- [12] K. Song, S. Feng, W. Gao, D. Wang, L. Chen, and C. Zhang. Build emotion lexicon from microblogs by combining effects of seed words and emoticons in a heterogeneous graph. In *Hypertext*, pages 283–292, 2015.
- [13] K. Song, S. Feng, W. Gao, D. Wang, G. Yu, and K.-F. Wong. Personalized sentiment classification based on latent individuality of microblog users. In *IJCAI*, pages 2277–2283, 2015.
- [14] J. Staiano and M. Guerini. Depeche mood: a lexicon for emotion analysis from crowd annotated news. In *ACL*, pages 427–433, 2014.
- [15] G. Xu, X. Meng, and H. Wang. Build chinese emotion lexicons using a graph-based algorithm and multiple resources. In *COLING*, pages 1209–1217, 2010.
- [16] W. Xu, X. Liu, and Y. Gong. Document clustering based on non-negative matrix factorization. In *SIGIR*, pages 267–273, 2003.