

Word Network Topic Model: A Simple but General Solution for Short and Imbalanced Texts

Yuan Zuo · Jichang Zhao · Ke Xu

Received: Dec 18, 2014 / Revised: Jun 23, 2015 / Accepted: Aug 24, 2015

Abstract The short text has been the prevalent format for information of Internet, especially with the development of online social media. Although sophisticated signals delivered by the short text make it a promising source for topic modeling, its extreme sparsity and imbalance bring unprecedented challenges to conventional topic models like LDA and its variants. Aiming at presenting a simple but general solution for topic modeling in short texts, we present a word co-occurrence network based model named WNTM to tackle the sparsity and imbalance simultaneously. Different from previous approaches, WNTM models the distribution over topics for each word instead of learning topics for each document, which successfully enhance the semantic density of data space without importing too much time or space complexity. Meanwhile, the rich contextual information preserved in the word-word space also guarantees its sensitivity in identifying rare topics with convincing quality. Furthermore, employing the same Gibbs sampling as LDA makes WNTM easily to be extended to various application scenarios. Extensive validations on both short and normal texts testify the outperformance of WNTM as compared to baseline methods. And we also demonstrate its potential in precisely discovering newly emerging topics or unexpected events in Weibo at pretty early stages.

Keywords Word Co-occurrence Network · Topic Modeling · Short Texts · Imbalanced Texts

Yuan Zuo
State Key Lab of Software Development Environment, Beihang University
E-mail: skywatcher.buaa@gmail.com

Jichang Zhao
School of Economics and Management, Beihang University
E-mail: jichang@buaa.edu.cn

Ke Xu (✉)
State Key Lab of Software Development Environment, Beihang University
E-mail: kexu@nlsde.buaa.edu.cn

1 Introduction

With the rapid development of World Wide Web and the spur of various kinds of web applications, short texts have been becoming the dominating content of Internet, such as web search snippets, micro-blogs (tweets), forum messages, and news titles. Specifically, around 250 million active users in Twitter generate almost 500 million tweets everyday, which carry sophisticated signals reflecting the real world. Hence accurately mining topics behind these short texts are essential for a wide range of tasks, including content analysis (Ramage et al 2010; Tong et al 2014; Jiang et al 2014b), query suggestion (Jiang et al 2014a; Zhou et al 2015), document classification (Chen et al 2011) and text clustering (Jin et al 2011; Quan et al 2010). However, due to the severe sparse context information, revealing topics from short texts is still a challenging problem for traditional frameworks that initially designed to handle the normal text. Meanwhile, short texts also suffer from significantly imbalanced document distribution. For example, in social media like Weibo, the amount of entertainment tweets is much larger than the number of ones in other categories (Yu et al 2011, 2013; Fan et al 2014). Since the objective of most commonly used topic models is to maximize the probability of the observed data, they tend to sacrifice the performance on rare topics (Jagaramudi et al 2012). Consequently, those topic models may not perform well in extrinsic tasks (Chang et al 2009).

As a canonical form of existing topic models, LDA (Blei et al 2003) is a hierarchical parametric Bayesian approach for topic discovery in a large corpus. To be specific, LDA models documents as mixture of topics and each topic is a probability distribution over words in the vocabulary of the corpus. Statistical inference methods are then used to learn the probability distribution over words associated with each topic and the distribution over topics for each document. Generally speaking, LDA-like models group semantically related words into a single topic by utilizing document-level word co-occurrence information (Wang and McCallum 2006), which makes them extremely sensitive to the document length and the number of documents related to each topic. Since the short text contains low word counts, those models will fail to obtain the accurate picture of how words are related to each other. Moreover, if the distribution over documents for the topic is heavily skewed, LDA-like models tend to learn more general topics held by majority documents rather than rare topics contained only by fewer documents. A recent study from Tang et al (2014) suggests that if the distribution over documents for each topic is heavily skewed, identifying topics from a small number of documents will be extremely difficult for LDA. While in fact rare topics might be essential for newly emerging events discovery or real-time hot trends detection in online social media (Chen et al 2013a).

Many efforts actually have been devoted to tackle the incompetence of LDA in modeling short texts. For example, related short texts can be aggregated into lengthy pseudo-documents before training the topic model (Weng et al 2010) or models trained from external data (e.g. Wikipedia) can be used to help the topic inference in short texts (Phan et al 2008). Besides, many arbitrary manipulations of LDA have also been introduced to satisfy the demands of specific short texts analyses (Zhao et al 2011; Chua and Asur 2013; Chen et al 2013a). Different from aforementioned approaches that are highly data-dependent or task-dependent, topic models focusing on the general-domain short text is emerging recently. A typical example is the biterm topic model (Yan et al 2013), which works well on short texts. However, biterm topic model is a special form of the mixture of unigram and not based on LDA. Therefore it does not overcome the shortcomings of LDA-like approaches on short texts and its flexibility is also greatly constrained. Another example is the dual-sparse topic

model (Lin et al 2014), which modifies LDA to learn focused topics for each short document and focused terms for each topic.

On the other hand, with respect to the topic imbalance, performance improvement of LDA is mainly obtained by adding prior information to guide the topic learning progress (Jagrlamudi et al 2012; Andrzejewski et al 2009) or using asymmetric Dirichlet prior over the document-topic distribution (McCallum et al 2009). While note that in practice, the knowledge about underlying structure of a given corpus is often undiscovered, so the prior information is not easily acquirable. As to asymmetric Dirichlet priors, how to determine proper parameter estimations is sophisticated and scenario-dependent for different applications and assuming symmetric Dirichlet priors help most variants of LDA keep the flexibility. Therefore, alleviating topic imbalance of LDA with symmetric Dirichlet priors is actually quite desirable.

To sum up, approaches mentioned above are neither scenario independent nor easy to be extended. In order to handle the sparsity and imbalance of short texts simultaneously through a general framework, we propose *Word Network Topic Model* (WNTM) based on the word co-occurrence network. The main idea of WNTM comes from the following observations. 1) When texts are short, word-by-document space is extremely sparse, while word-word space is still rather dense. Since the topic quality can be guaranteed in the dense word-word space (Arora et al 2013), we conjecture that learning topic components from word co-occurrence network rather than document collection is more reliable. 2) Intuitively, the number of words connected to rare topics often exceeds the amount of documents related to those topics. So the distribution over topics for words is greatly less skewed than the distribution over topics for documents. 3) Since the distribution over topics for each document can not be learned accurately in short or imbalanced texts, we should learn the distribution over topics for each word instead. 4) Different from the existing solutions, a new framework should be simple enough to guarantee its scalability in different application scenarios. Hence WNTM employs the standard Gibbs sampling for LDA to discover latent word groups (i.e. topics) (Henderson and Eliassi-Rad 2009) and learns distribution over topics for words rather than topics for documents. Learning words' topics rather than documents' topics makes WNTM less sensitive to the document length or heterogeneity of the topic distribution. In addition, the word co-occurrence network can be constructed with any type of given texts, which makes WNTM further simple and general in real-world applications.

Extensive experiments are conducted on various data sets to compare WNTM and baseline methods in three aspects, including topic quality, word semantic similarity and document semantic classification. And results suggest that WNTM can discover the most coherent topics in short texts. Meanwhile, WNTM outperforms all baseline methods on word similarity and document categorization in both short and normal texts. Particularly, WNTM shows much better capability than LDA in rare topic detection in extremely imbalanced texts. Major contributions of this paper are 1) WNTM is a generative model for a word network rather than a collection of documents, 2) and it learns topics for each word rather than topics for each document, therefore 3) it is less sensitive to document length or document distribution over topics. 4) Since WNTM uses the standard Gibbs sampling for LDA, it's general and very easy to be applied in different scenarios.

The rest of the paper is organized as follows. We first give a short review of related works in section 2. This is followed by detail introductions of our model and re-weighting method in section 3 and section 4. Experimental results are illustrated and explained in section 5. Finally, we conclude the present work briefly in section 6 and several possible directions in future are also pointed out.

2 Related works

Probabilistic topic models such as PLSA (Hofmann 1999) and LDA (Blei et al 2003) have been extensively applied in exploring text corpora. Particularly, LDA is a more complete generative model since it extends PLSA by adding Dirichlet priors on topic distributions. Due to their extensibility, many complicated variants of LDA and PLSA have been proposed in the last decade, such as the dynamic topic model (Blei and Lafferty 2006), social topic model (Cha and Cho 2012), author-topic model (Rosen-Zvi et al 2004) and author-topic-community model (Li et al 2015) etc. While most of them are designed to handle normal texts with special additional properties, such as time, social relationship and authorship.

The sparse short texts have also attracted much research interest in the previous literature and most early studies mainly focus on increasing data density through utilizing auxiliary information. For example, Hong and Davison (2010) train topic models on aggregated tweets that share the same word, and find those models work better than those being directly trained on original tweets. Sahami and Heilman (2006) propose a search-snippet-based similarity measure for short texts. Jin et al (2011) learn topics on short texts via transfer learning from auxiliary long text data. Another way to deal with data sparsity in short texts is to apply special topic models. For example, Zhao et al (2011) assume each tweet only covers a single topic. Yan et al (2013) propose a special form of mixture of unigrams (Nigam et al 2000), which is called biterm topic model to improve topic modeling on short texts. Lin et al (2014) introduce sparse constraints for document-topic distributions and topic-term distributions in order to perform topic modeling on short texts.

While regarding to the topic modeling on imbalanced texts, the prior knowledge has been widely used to alleviate skewed distributions over documents of different topics. Andrzejewski et al (2009) propose Dirichlet forest priors to incorporate must-links and cannot-links constraints into topic models. Chen et al (2013b) use general lexical knowledge to help discovering coherent topics. It is worthy noting that the document level knowledge can also be utilized. For example, Ramage et al (2009) and Rubin et al (2012) take labels into account for the generative process of the corpora and Blei and McAuliffe (2007) propose a supervised topic model to predict the category labels for input labeled documents, which successfully improve performance of topic modeling on short texts.

However, different from above approaches, we try to figure out a simple but general solution to take care of sparsity and imbalance in texts simultaneously. To the best of our knowledge, little attention has been paid to this issue and our proposed topic model is the first one handling short and imbalanced texts in a general way without exploiting any external knowledge.

3 Word network topic model

Commonly used topic models implicitly take advantage of rich word co-occurrence patterns in documents. However, short texts naturally lack of enough contextual information. Furthermore, the goal of traditional topic models is to maximize the probability of generating observed documents, and rare topics that reflected by fewer documents are tended to be ignored. As a result, directly applying conventional topic models on short or imbalanced texts can not perform as well as that on normal balanced texts. In order to solve the problem mentioned above through a simple but general method, we propose a new framework, which applies the same Gibbs sampling (Heinrich 2005) as LDA to discover latent word groups in a word co-occurrence network. Here latent word groups of the network are taken as topic




Fig. 1 An illustration of the sliding window and word pair weighting schema.

components of a corpus. In addition, the distribution over latent groups for each word is also learned by our model. The details of the new framework are presented in the following subsections.

3.1 Word co-occurrence network

In a word co-occurrence network (which may also be denoted as *word network* in the following text if there is no conflict), nodes are words occurring in the corpus and an edge between two words indicates that the connected two words have co-occurred in the same context at least once. Here the context can refer to a document or a sliding window with fixed size. To limit the size of word network and reserve only the local context for each word, we take a sliding window of fixed size as the context in the present work. According to the work of Peirsman et al (2008), a sliding window with size of 10 or larger can sufficiently capture topical similarities among words. However, larger window size results in higher computational complexity as shown in Section 4.1. Therefore, we choose to set the size of sliding window to 10 for both normal and short texts. For short texts, one can also directly take each short document as an individual sliding window. And we recommend to use sliding window for short texts, especially when its average document length is larger than window size. For instance, the average document length of micro-blogs we used in experiments is 17.2. Degree of a node is defined as the sum of weights over its adjacent links, and activity of a node is defined as the averaged weight of its adjacent links.

In order to convert the given document collection into a word network, we first filter out low frequency and stopping words, and then a sliding window is moved to scan each document. As the window scanning word by word through the document, any two distinct words that appear in the same window would be regarded as co-occurred with each other. Times that two words co-occurred are accumulated and defined as the weight of the corresponding edge between them. Note that, a pair of words might be counted multiple times, which we refer to as word pair weighting schema as illustrated in Fig. 1. Words that co-occur near each other will be counted more times than those that co-occur far from each other. For instance, word W_5 and word W_4 (or W_6) are counted 4 times, meanwhile, W_5 and W_1 (or W_9) are counted only once. In this we can encourage the model to put words that co-occur near each other into the same topic, which can greatly benefit the coherence of learned topics, since cohesive words generally are strongly semantically related.

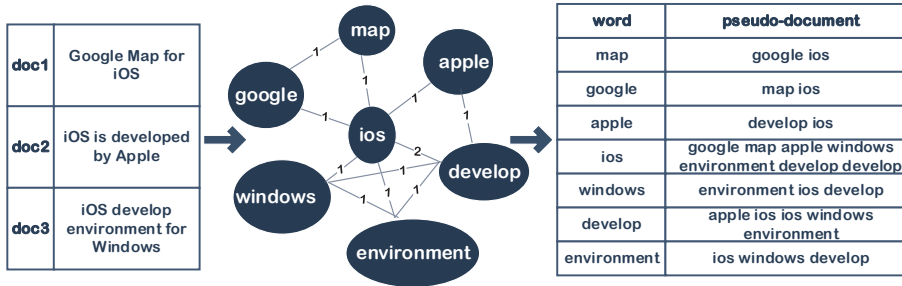


Fig. 2 An example of how to generate pseudo-documents from original input corpus in WNTM.

Note that in topic models, a topic can be viewed as a bag of words co-occurred frequently in the same document, which is very similar to latent word groups (or communities) in the word network, since words co-occurred frequently in the same sliding windows are closely connected in the semantic space and they could appear in the same document with high probabilities. Therefore, we can take latent word groups in our word network based model as the topics in LDA. At the same time, learning topics from word co-occurrence network, a special form of word-word space, has a theoretical guarantee for topic coherence according to the work of Arora et al (2013). What's more, rare topics may form compact latent word groups in word network, therefore a topic model based on the word network could effectively find word groups that correspond to rare topics. Based on the considerations above, we propose our *word network topic model* (WNTM). In order to keep the new model simple and general to employ at different scenarios, we take a similar approach as Henderson and Eliassi-Rad (2009) did in to discover latent word groups in word network.

3.2 Word network topic model (WNTM)

The standard Gibbs sampling for LDA can be used to discover latent word groups in large word network. While in order to reserve the standard Gibbs sampling, we first have to represent the word co-occurrence network back to a pseudo-document set. We assume the word network as undirected and weighted. As illustrated in Fig. 2, each word in the network can be treated as a pseudo-document with content constituted by the list of its adjacent words. Obviously since the word network is weighted, the adjacent words may occur multiple times in the text of the pseudo-document.

Although WNTM uses the same Gibbs Sampling as LDA, the rationalities underlying the generative process of them are different. LDA learns to generate a collection of documents by using topics and words under those topics. However, WNTM learns to generate each word's adjacent word-list in the network by using latent word groups and words belonging to those groups. More specifically, WNTM learns the statistical relations between words, latent word groups and words' adjacent word-lists by assuming that each word's adjacent word-list is generated semantically by a particular probabilistic model. It first supposes that there is a fixed set of latent word groups in the word network, and each latent word group z is associated with a multinomial distribution over the vocabulary Φ_z , which is drawn from a Dirichlet prior $Dir(\beta)$. The generative process of the whole pseudo-document collection converted from the word network can be interpreted as follows:

1. For each latent word group z , draw $\Phi_z \sim Dir(\beta)$, a multinomial distribution over words for z
2. Draw $\Theta_i \sim Dir(\alpha)$, a latent word group distribution for the adjacent word-list L_i of the word w_i
3. For each word $w_j \in L_i$:
 - (a) select a latent word group $z_j \sim \Theta_i$
 - (b) select the adjacent word $w_j \sim \Phi_{z_j}$

In WNTM, the Θ distributions represent the probability of latent word groups appearing in each word’s adjacent word-list and the Φ distributions stand for the probability of words belonging to each latent word group. Given the observed corpus, WNTM first converts it to a word network, then generate the pseudo-document set and finally the same Gibbs sampling implementation that developed for conventional LDA is employed to infer values of the latent variable in both Φ and Θ . Because each word’s adjacent word-list actually represents its global context information, so different from previous LDA-like approaches, WNTM models the distribution over latent word groups for each word instead of the distribution over topics for each document.

3.3 Inferring topics in a document

As discussed in the previous section, WNTM does not model the document generation process. Therefore, we cannot directly obtain topics in a document from the result of Gibbs sampling. Since WNTM models the generation process of each word’s adjacent word-list which stands for the word’s global contextual information, we can take topic proportions of word w_i ’s adjacent word-list Θ_i as topic proportions in w_i . Given topic proportions for all words, topics of each document can be obtained accordingly. Specifically, to infer topics in a document, we assume that the expectation of the topic proportions of words generated by a document equals to the topic proportions of the document, i.e.,

$$P(z|d) = \sum_{w_i} P(z|w_i)P(w_i|d), \quad (1)$$

where $P(z|w_i)$ equals to $\Theta_{i,z}$, which has been learned in WNTM. As to $P(w_i|d)$, we simply take the empirical distribution of words in the document as a estimation, i.e.,

$$P(w_i|d) = \frac{n_d(w_i)}{Len(d)}, \quad (2)$$

where $n_d(w_i)$ is the word frequency of w_i in document d and $Len(d)$ is the length of d . It is worthy noting the above strategy is straight-forward and easy to implement, which guarantees WNTM’s simplicity further.

To sum up, when texts are short and sparse, learning topics in word-by-document space will suffer from the severe sparsity problem, while learning topics in a word-word space has a theoretical guarantee for topic coherence, which has been proved in (Arora et al 2013). Meanwhile, as the distribution over documents for each topic is imbalanced, rare topics tend to be ignored by LDA-like models. However, we conjecture that words related to rare topics would still form a semantically compact latent group in the word co-occurrence network. So latent groups standing for rare topics will also be detected by WNTM. Therefore, the rich contextual information in word-word space facilitates WNTM to discover topics in word co-occurrence network other than directly reveal topics from document collection.

4 Complexity analysis and word network re-weighting

Although our inference on WNTM uses the same Gibbs sampling as LDA, the running time and space complexities of them are different. We will compare the time complexity in detail first, and then give a brief discussion about the space complexity of the two models. Finally, in order to reduce the time and space complexity of WNTM, we further propose a method to perform word network re-weighting.

4.1 Complexity analyses

The time complexity of LDA is $O(N_d K_z L_d)$, where N_d is the number of documents, K_z is the number of topics and L_d is the average document length. Similarly, the time complexity for WNTM is $O(N_p K_g L_p)$, where N_p is the number of pseudo-documents, i.e., the size of the vocabulary, K_g is the number of latent word groups (topics) and L_p is the average pseudo-document length. Since the maximum number of sliding windows in a corpus is $N_d L_d$, and each sliding window can generate $\binom{c}{2}$ edges, where c is the size of the sliding window. Thus, approximately $N_p L_p$ can be rewritten as

$$N_p L_p \approx N_d L_d c(c-1). \quad (3)$$

Supposing K_z equals to K_g , the time complexity of WNTM is $o(c^2)$ times larger than LDA's cost. In practice, for short texts, the average document length $\langle l \rangle$ is often close to c . So when applied to short texts, the time complexity of WNTM is acceptable. However, with respect to normal texts, it becomes unacceptable since c can be set to a large number.

The space complexity of LDA is $O(N_d K_z + N_d L_d)$ and the space complexity of WNTM is $O(N_p K_g + N_p L_p)$. Similar to the time complexity, if we assume N_d is equal to N_p and K_z is equal to K_g , then the memory WNTM consumes is $o(c^2)$ times of the size that LDA needs. In order to reduce the time (or space) complexity, for instance, to decrease the time (or space) complexity to linear times the cost of LDA, we would propose a word network re-weighting method in the next subsection, which will help boost the learning process of WNTM effectively.

4.2 Word network re-weighting

The above analysis shows that the time and space complexity is unaffordable for Gibbs sampling over the pseudo-document collection, which is directly generated by the weighted word network. To reduce both the time and space consumption, we need to decrease $N_p L_p$. While N_p is fixed as it is determined by the size of the given corpus's vocabulary, so L_p is the only tunable parameter. Representing the length of a pseudo-document, L_p also equals to the degree of the node corresponding to the pseudo-document. Therefore, decreasing the weights of edges in the network can shorten the length of pseudo-documents, and then reduce the time and space complexity of WNTM accordingly. In order to reserve the relative closeness of different words in the process of tuning weights, a re-weighting method is illustrated in Algorithm 1.

The weight of each edge is divided by the activity of its end with lower degree, denoted as w_e . Along this line, we can decrease the whole weighted degree of the entire word network, which can then decrease the time and space complexity of WNTM. Since the weighted degree of a node must be larger than $c-1$, then the averaged weighted degree

Algorithm 1 Word network re-weighting algorithm

Require: the original word network $G = (V, E, W)$, where V is the set of words, E is the set of edges and W is the set of weights for edges.

Ensure: the re-weighted word network $G' = (V, E, W')$.

- 1: compute degree $D(n)$ and activity $A(n)$ of each node $n \in V$
- 2: **for all** $e = (n_1, n_2) \in E$ **do**
- 3: set $w_e = \left\lceil \frac{w_e}{A(n_i)} \right\rceil, \text{ argmin}\{D(n_i), i = 1, 2\}$
- 4: **end for**

of the word network is actually much larger than $c - 1$. Therefore, the averaged length of pseudo-documents should be smaller than $\frac{L_p}{c-1}$, where L_p is the averaged length before re-weighting. Then from Eq. 3 we can easily get that the time and space cost of re-weighted WNTM is $O(c)$ times the complexity of LDA, which is just a linear scale-up. Hence the re-weighting algorithm successfully reduces the cost of WNTM and guarantees its feasibility in both short and normal texts.

5 Experiments

In this paper, we evaluate our approach in three measures, including topic quality, word similarity and document classification. For each measure, extensive experiments are performed on real-world short texts and normal texts respectively. For short texts, We take LDA, biterm topic model (BTM) (Yan et al 2013) and dual-sparse topic model (DSTM) (Lin et al 2014) as baseline methods. The reason to compare with BTM and DSTM is that they are state-of-the-art topic models for short texts. We also give comparison between LDA and WNTM on normal texts, since we want to demonstrate that WNTM can also be applied to normal texts. Note that the comparison between WNTM and LDA can also demonstrate the strengths and weaknesses of learning topics from document collection and word co-occurrence network. As to normal texts, we omit the comparison with BTM or DSTM because of their intense time complexity when applied on normal texts, besides, both approaches are designed specifically for short texts.

Most experiments in this section are carried out on a Windows Server with an Intel Xeon 2.40GHz CPU and 12G memory except for the experiments using Wikipedia data set. Due to the large volume of the Wikipedia data set, corresponding experiments are conducted on a Linux cluster with 13 nodes. Each node contains 2 Intel Xeon 2.27Hz CPUs and 12 GB memory. For both LDA and WNTM, we use a Java open-source implementation JGibbLDA¹ on short texts and an MPI open-source implementation titled PLDA² on normal texts. For BTM, we use the source code opened by the authors³. For DSTM, we implement the method by ourself. For JGibbLDA and BTM, we set $\alpha = 50/k$ and $\beta = 0.01$, where k is the number of topics. For PLDA, we set $\alpha = 0.1$ and $\beta = 0.01$. For DSTM, we set parameters the same way as the original paper dose. In all experiments, number of topics is set to 100, length of the sliding window for WNTM is set to 10 and each model’s Gibbs sampling is run for 2,000 iterations. Except for document classification on news content, the results reported here are the average over 10 rounds.

¹ <http://jgibbllda.sourceforge.net/>

² <http://code.google.com/p/plda/>

³ <http://code.google.com/p/btm/>

5.1 Evaluation of the topic quality

It is a typical way for evaluating topic models through comparing the perplexity on a held-out test set. However, WNTM does not model the generation process of documents. Hence, the perplexity is not suitable in this paper. Furthermore, recent research shows that the perplexity does not always correlate with semantically interpretable topics (Chang et al 2009). Therefore, here we utilize the topic coherence (Mimno et al 2011) as an evaluation metric for the topic discovery, which has been found to correlate well with human judgments of the topic quality.

5.1.1 Topic coherence

Topic coherence (also called UMass measure) is a comprehensive and automated evaluation measure for topic models, which measures the score of a single topic by computing the semantic similarity degree between high probability words in the topic. Higher topic coherence often indicates better topic quality, i.e., better topic interpretability. The topic coherence is defined as

$$C(z; M^{(z)}) = \sum_{t=2}^T \sum_{l=1}^{t-1} \log \frac{D(m_t^{(z)}, m_l^{(z)}) + \epsilon}{D(m_t^{(z)})}, \quad (4)$$

where $M^{(z)} = (m_1^{(z)}, \dots, m_T^{(z)})$ is the list of the T most probable words in topic z , $D(m)$ counts the number of documents containing the word m , $D(m, m')$ counts the number of documents containing both m and m' , and $\epsilon = 10^{-12}$ is used to avoid taking the log of zero for words that never co-occur and to smooth the score for completely unrelated words. We use the average coherence score of all topics as the evaluation metric for topic quality of different topic models.

5.1.2 Topic coherence on short texts

To investigate WNTM’s ability of learning high quality topics from real-world short texts, we carry out experiments on one day’s micro-blogs⁴ sampled from Weibo. As a Twitter-like service in China, it also imposes a limited length for each tweet, i.e., no more than 140 Chinese characters. Since the textual content of micro-blogs is not formal, careful preprocessing is quite necessary. In the preprocessing, we take the following steps to wash the collected corpus: (a) using NLPiR⁵ to do tokenization; (b) removing stopping words; (c) removing words with frequency less than 20; (d) filtering out URLs, emoticons, Hashtags and words with non-Chinese characters; (e) removing micro-blogs with length less than 10. Note that abbreviations, acronyms and slangs are remained for latter topic modeling. Finally, 189,223 micro-blogs retained with 20,942 distinct words in total. The average number of tokens in documents is 17.2.

We compare WNTM with LDA, BTM and DSTM on this micro-blog collection. For all models, we set the number of topics to 100. The average topic coherence results of four models are listed in Table 1, where the size of top words set in each topic, denoted as T , ranges from 5 to 20. We find that the average topic coherence of WNTM is obviously higher than other three models, which indicates that WNTM outperforms LDA, BTM and DSTM in learning high quality topics from short texts. And the improvement made by WNTM is

⁴ Publicly available at <http://ipv6.nlsde.buaa.edu.cn/zhaojichang/paper/wntm.rar>

⁵ <http://ictclas.nlpir.org/downloads>

Table 1 The average topic coherence results on the micro-blog collection. A larger value stands for more coherent topics.

T	5	10	20
LDA	-36.6±1.8	-221.4±4.0	-1484.6±39.4
DSTM	-36.5±1.9	-199.6±7.7	-1260.5±29.7
BTM	-37.4±1.9	-207.5±8.2	-1235.9±30.7
WNTM	-32.5±1.3	-181.6±5.5	-1056.6±22.8

statistically significant (p -value < 0.001 by Mann-Whitney U test). The outperformance of WNTM as compared to LDA is in accordance with our understanding that learning topics from dense word-word space can guarantee the topic quality even in short texts. Remarkably, WNTM outperforms BTM and DSTM significantly, which indicates WNTM can learn more coherence topics than state-of-the-art models. Although WNTM and BTM both induce topics from word-word space instead of from documents, WNTM can learn more coherence topics. The reason might be that WNTM can give higher weights for word pairs that co-occur more near each other while BTM does not, furthermore, WNTM is based on LDA while BTM is based on the mixture of unigram, which might further limits the performance of BTM. The comparison between WNTM and DSTM indicates that applying sparse constraints on topic models is less efficient than inducing topics from word-word space according to above topic coherence results in short texts scenario.

5.1.3 Topic coherence on normal texts

We conduct experiments on Wikipedia data provided by Phan et al (2008) to investigate WNTM’s ability of learning high quality topics from real-world normal texts. The Wikipedia data set contains 71,986 documents with 60,649 distinct words. The average number of tokens in documents is 423.5. The number of topics is set to 100 for each model and the average topic coherence result is listed in Table 2, where the size of top words set T in each topic ranges from 5 to 20. As shown in Table 2, the average topic coherence of WNTM is slightly higher than that of LDA when $T = 5$ and 10, while slightly lower than that of LDA when $T = 20$. From the results, we can see that learning topics by grouping words co-occurred in a small range of context can benefit the top 10 words’ coherence in each topic, but when $T = 20$, the coherence of top words might need more plentiful document-level word co-occurrence information to maintain. However, with the increasing of distance between two words, the relation between them becomes less relevant. Because of this, the difference between WNTM and LDA is not likely to be obvious. According to the results of Mann-Whitney U test, two models gain no statistically significant improvement than each other. Therefore, we can conclude that WNTM can produce similar high-quality topics as LDA dose on normal texts. Note that the word-by-document space has no sparsity problem in normal texts, so LDA can utilize the rich contextual information in each document to learn high quality topics. Thus, for normal texts, learning topics from word-by-document space and word-word space makes little difference in topic quality.

5.2 Word similarity tasks

The average topic coherence is an intrinsic measure used to evaluate the quality of all topics. Higher topic coherence often indicates better topic quality, but it does not guarantee a

Table 2 The average topic coherence of LDA and WNTM on the Wikipedia data set. A larger value means more coherent topics.

T	5	10	20
LDA	-13.9 ± 0.3	-69.7 ± 1.0	-327.4 ± 3.3
WNTM	-13.8 ± 0.2	-69.5 ± 0.6	-329.7 ± 2.4

better performance on extrinsic tasks. The study of Henderson and Eliassi-Rad (2009) revealed that LDA is better than LSA (Latent Semantic Analysis) (Deerwester et al 1990) at learning descriptive topics, while LSA is better than LDA at creating a compact semantic representation of words and documents and outperforms LDA in extrinsic tasks (Stevens et al 2012). Hence, we carry out experiments to compare the performance of WNTM and other methods on extrinsic tasks such as word similarity tasks and document classification, which would help further illustrate each models' effectiveness in creating compact semantic representation of words and documents. In this section, we compare models' ability of learning semantic representation of words on short and normal texts, respectively. To begin with, we will introduce how to calculate the semantical similarity between two words.

5.2.1 Semantic representation of a word

For LDA, DSTM and BTM, the conditional topic distribution for a word w can be defined as its semantic representation

$$s_w = [p(z_1|w), p(z_2|w), \dots, p(z_k|w)],$$

where k is the number of topics. $p(z_k|w)$ is easy to obtain after Gibbs sampling stage is completed,

$$p(z_k|w) = \frac{n_{w|z_k}}{n_w}, \quad (5)$$

where $n_{w|z_k}$ stands for how many times w has been assigned with topic k during the sampling and n_w means the total occurrence of w in given corpus. With respect to WNTM, we do not have to calculate $p(z_k|w)$. The θ_w can be directly used as the semantic representation of word w , since WNTM learns the topic distribution over words as matrix Θ . Therefore, we use Θ 's row vector θ_w corresponding to w as its semantic representation, which is denoted as

$$s_w = \theta_w = [\Theta_{w,1}, \Theta_{w,2}, \dots, \Theta_{w,k}].$$

Since word representations are actually distributions over topics, we can measure the distance of two words via the Jensen-Shannon divergence

$$JS(s_i, s_j) = \frac{1}{2}D_{kl}(s_j \| m) + \frac{1}{2}D_{kl}(s_i \| m), \quad (6)$$

where s_i and s_j are the semantic representations of words i and j , $m = \frac{1}{2}(s_i + s_j)$ and $D_{kl}(p||q) = \sum_i p_i \ln \frac{p_i}{q_i}$ is the Kullback-Leibler divergence. If we take word representations as space vectors, cosine similarity can also be used to measure the distance of two words. It can be defined as

$$Cosine(s_i, s_j) = \frac{s_i \cdot s_j}{\|s_i\| \|s_j\|}. \quad (7)$$




Fig. 3 Ranked correlation results of word similarity task on Weibo data. The result demonstrates that WNTM outperforms baseline models on creating better semantic representations of words on short texts.

5.2.2 Word similarity tasks

Word similarity tasks are widely used to evaluate distributional semantic spaces. Topics learned by topic models can be viewed as the knowledge about semantic distributions of words. If a topic model learns topics accurately, then we can expect similar words, such as “man” and “woman”, to be represented with similar semantic representations. In this paper, we use the word similarity task presented by Wang et al (2011) to evaluate the ability of word semantic modeling of two models on Weibo data set. Regarding to normal texts, we use word similarity tasks designed by Rubenstein and Goodenough (1965) and Finkelstein et al (2002) on Wikipedia data set. In each task, the semantic similarity between given pairs of words were evaluated by human. The word similarity task introduced by Wang et al. contains 240 pairs of Chinese words and each pair’s semantic relatedness is rated from 0-10, in which a higher score reflects a more semantically similar word pair. The task introduced by Finkelstein was constituted by 353 pairs of English words and Rubenstein and Goodenough’s task contains 65 English words. Each pair was also given a human rate indicating the pair’s semantic closeness.

We evaluate topic models by calculating the similarity between each pair of words through two similarity measurements (JS is short for Jensen-Shannon divergence and Cosine is short for Cosine similarity) in the evaluate set and then compare the model’s ratings with human ratings by the ranked correlation. Intuitively higher correlation indicates better word semantic modeling. The number of topics is set to 100 for both models and the ranked correlation result on Weibo data is illustrated in Fig. 3.

From the result, we can find that both DSTM and BTM outperform LDA, no matter using JS as the measurement or Cosine. This result indicates that applying sparse constraint to topic models or performing topic modeling on dense word space both can improve the quality of learned word semantic representations. WNTM outperforms all other three models on both JS and Cosine significantly (p -value < 0.001 by Mann-Whitney U test). The result of word similarity task demonstrates that WNTM is not only good at inducing more coherence topics, but also capable of learning better word semantic representations.

The result of Wikipedia data is illustrated in Fig. 4. We surprisingly find that WNTM still outperforms LDA on both similarity measurements. The performance gap between two models on normal texts shrinks compared to the result on short texts, since the document-level data sparsity problem is gone. However, WNTM’s result is more stable than LDA

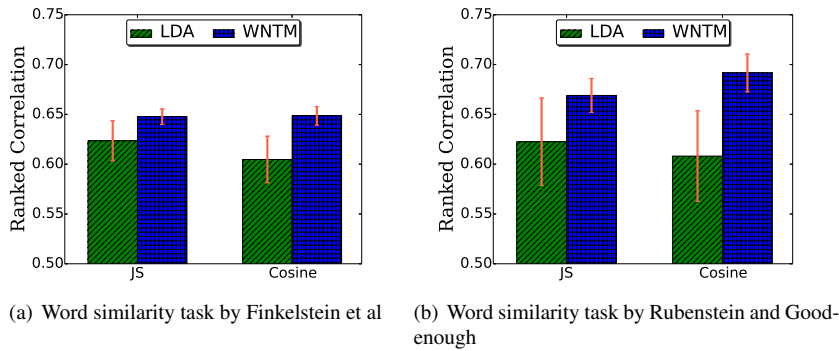


Fig. 4 Ranked correlation results of word similarity task on Wikipedia data set. The result shows that WNTM even outperforms LDA on creating better semantic representations of words on normal texts.

Table 3 Label counts for news titles and contents used in document classification tasks. Here Ent represents Entertainment.

(a) Label counts for news titles				(b) Label counts for news content			
Label	Count	Label	Count	Label	Count	Label	Count
Finance	31,414	Car	6,532	Finance	133,480	Car	18,675
Sports	25,414	IT	2,321	Sports	115,946	IT	10,650
Society	14,889	Military	1,733	Society	70,743	Military	8,706
Ent	11,208	House	1,410	Ent	53,335	House	6,407
Lady	8,128	Culture	983	Olympics	34,767	Health	2,340
Olympics	7,117	Health	962	Lady	31,689	Culture	2,334
Education	6,594			Education	19,482		

by possessing much lower deviations. Although sharing similar topic coherence results in normal texts, the two models' ability in creating compact semantic representation for words are different and WNTM performs better than LDA in word similarity tasks.

5.3 Document classification

In order to compare these models' ability in learning semantic representation of documents on short and normal texts, we evaluate them by performing document classifications on news titles and news content respectively in this section. To illustrate WNTM's advantage on imbalanced texts as compared to LDA, we further conduct classification experiments by tuning the heterogeneity of topic distributions.

5.3.1 Evaluation on news corpus

To explore the effectiveness of WNTM in document semantic modeling on short texts and normal texts, we first evaluate its performance on document classification tasks of news titles and news content, which are extracted from the news corpus provided by Sogou.com⁶. After preprocessing, we obtain 508,554 news titles with label distributions listed in Table 3, and

⁶ <http://www.sogou.com/labs/dl/ca.html>

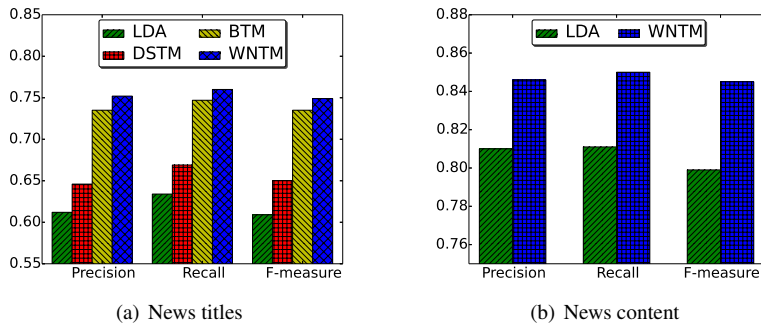


Fig. 5 Classification performance comparison between WNTM and baseline models on short and normal texts, respectively.


59,348 distinct words in total. The average number of token in news titles is 5.5. We also obtain 118,705 news reports with label distributions listed in Table 3, and 76,114 distinct words in total. The average number of token in each report is 175.9.

Taking topic model as a method of dimensionality reduction, we can reduce a document into a fixed set of topics, which can be features for document categorization. For each topic model trained on 100 topics, we perform 10-fold cross-validation on news titles (or content). In each fold, we randomly split news titles (or content) into training and test subsets with the ratio 9:1, and classified them by using LIBLINEAR⁷. The macro precision, recall and F-measure for news titles are shown in Fig. 5(a). The result for news content is shown in Fig. 5(b).


From the result, we can find DSTM outperforms LDA in classifying short texts, which indicates the sparse constraint applied to LDA can alleviate its data sparsity problem. However, the improvement made by DSTM is limited compared to BTM and WNTM, which indicates resorting to perform topic modeling on word space is more reliable than adding sparse constraint to topic models. WNTM outperforms BTM and DSTM significantly (p -value < 0.001 by Mann-Whitney U test), similar to word similarity task. Regarding to normal texts, LDA’s classification result becomes acceptable since the data is no longer sparse. However, WNTM still outperforms LDA, which indicates that WNTM is better than LDA in recognizing the resemblance of documents even for normal texts.

Fig. 6 shows the confusion matrix of LDA and WNTM in the news contents classification. For example, LDA confuses “House” with “Sports”, “Finance” and “Lady”, where “House” and “Finance” may be hard to distinguish because they are strongly relevant, while obviously “House” has few connections with “Sports” and “Lady”. Consequently, LDA almost misclassifies all documents under label “House”. The similar situation also happens to “Culture” and “Health”. As listed in Table 3, the significant confusion of those labels mainly comes from their low numbers of documents and in fact they represent exactly the rare topics in the employed corpus. However, as shown in Fig. 6(b), WNTM still achieves a relative better performance as compared to LDA. It indicates that WNTM gains outstanding improvement on the performance of identifying rare labels, which evidently justifies the conjecture that WNTM can identify more rare topics than LDA as the text is imbalanced.

⁷ <http://www.csie.ntu.edu.tw/~cjlin/liblinear/>



(a) LDA



(b) WNTM

Fig. 6 The confusion matrix for news content classification. The result from news title is similar and hence it is not reported here.

5.3.2 Document classification on imbalanced texts

To compare performance between WNTM and LDA in classifying documents containing rare topics further, we investigate the variation of their classification results by continuously tuning the imbalance of texts. We first build a balanced data set from news content introduced previously. This balanced data set includes news from five classes, namely “Education”, “Car”, “Finance”, “Lady” and “Ent”. Within each class, there are 1,000 documents allocated equally. Second, we take various numbers of documents away from “Car”, which is randomly selected, to build data sets with different levels of imbalance. Specifically, we build 8 groups of imbalanced data sets, each group contains 1,000 documents belong to classes except “Car”, which contains d_c documents, where d_c is tunable parameter and we let $d_c = 800, 600, 400, 200, 100, 80, 60$ and 40 . As d_c ranges from 800 to 40, the imbal-

Table 4 Average precision, recall and F-measure on classifying “Car” news contents in the imbalanced document classification. Results in boldface are those with p -value < 0.001 by Mann-Whitney U test.

d_c	Precision		Recall		F-measure	
	LDA	WNTM	LDA	WNTM	LDA	WNTM
800	0.948±0.009	0.958±0.008	0.881±0.012	0.941±0.010	0.913±0.007	0.950±0.005
600	0.953±0.007	0.957±0.007	0.864±0.011	0.923±0.012	0.906±0.007	0.936±0.006
400	0.956±0.011	0.954±0.006	0.839±0.026	0.911±0.016	0.893±0.015	0.932±0.009
200	0.961±0.012	0.932±0.011	0.792±0.021	0.866±0.019	0.868±0.014	0.898±0.013
100	0.935±0.019	0.909±0.022	0.727±0.043	0.792±0.038	0.818±0.034	0.846±0.025
80	0.903±0.069	0.913±0.023	0.718±0.087	0.773±0.042	0.798±0.077	0.836±0.029
60	0.874±0.039	0.884±0.034	0.619±0.199	0.700±0.100	0.702±0.186	0.789±0.051
40	0.760±0.100	0.820±0.065	0.491±0.192	0.553±0.119	0.600±0.200	0.658±0.104

ance of data set is enhanced accordingly. To avoid the influence of the randomness in taking away documents from “Car”, we randomly sample documents under “Car” 10 times for each group. The average precision and recall on classifying “Car” documents is illustrated in Table 4. From the result, we can find that the average precision of LDA first slightly increases when the number of news under “Car” reducing from 800 to 200, and then decreases dramatically due to the more and more serious imbalance of the data set. On the contrary, the average precision of WNTM decreases slowly with the enhancement of the imbalance. Particularly, when the number of “Car” documents equals to 60 and 40, the average precision of LDA has decreased to 86% and 72% while WNTM’s average precision is 90% and 85%. This result indicates that WNTM is more accurate than LDA in distinguishing documents under rare labels. Note that as d_c varies in the range between 200 and 500, the averaged precision of LDA is better than that of WNTM. However, the recall and F-measure of WNTM outperforms LDA. From those results, we can find that WNTM and LDA have similar trending with the imbalance enhancement, while WNTM always outperforms LDA in terms of the average recall and f-measure. To sum up, we can conclude that WNTM performs better than LDA in dealing with text classification on imbalanced data sets.

From the above experiments, we find WNTM dominates LDA in learning rare topics. Thus it is reasonable to conjecture that WNTM could be a better choice than LDA in detecting newly emerging topics or unexpected events at early stages in social media, like Twitter or Weibo. To illustrate this point, we collect 10,000 micro-blogs from Weibo, and then inject d_m micro-blogs related to the event of “MH370” into them, where d_m ranges from 10 to 100. Note that the 10,000 micro-blogs have no relation with “MH370” since they are posted exactly before the event happens. For different settings of d_m , we train both models 10 times since the result of Gibbs sampling varies each time. After training, we look over the top 20 words of each topic to search the topic word “MH370”, and count how many times it is found for each setting of d_m . From the results we find that WNTM can always identify a topic with word “MH370” contained by its top 20 word-list when $d_m \geq 30$. However, LDA achieves the same result only as $d_m \geq 50$. Thus, WNTM is more sensitive to rare topics and can learn newly emerging topics at earlier stages than LDA. Besides, we also evaluate the quality of the topic related to “MH370” by human review. If a topic’s top words contains more “MH370” related words than others, we evaluate this topic as the best. We list top 15 words of the best “MH370” topic learned by two models for comparison when $d_m = 50$ and 100. It is interesting that when $d_m = 50$, the best “MH370” topic learned by LDA contains 8 words that never appear in the 50 injected “MH370” micro-blogs, while WNTM only has 4. When $d_m = 100$, LDA’s “MH370” topic still contains 4 unrelated words, while WNTM only has 1 error word left. Based on these results, we can conclude that WNTM can detect

event topic earlier than LDA and the quality of the topic is much better than that identified by LDA.

6 Conclusions

A simple but general approach named WNTM is presented in this paper to facilitate the topic modeling in short and imbalanced texts at an acceptable cost. Different from conventional LDA-like solutions, it explores topics from word co-occurrence networks and successfully alleviates the data sparsity and the topic-document heterogeneity in word-by-document space. Thorough experiments on both short and normal texts suggest that WNTM outperforms baseline methods in tasks of topic coherence, word similarity and feature selection within document classification. Furthermore, the ability of capturing rare topics with high quality indicates that WNTM could be an effective model for detecting newly emerging topics or unexpected events in social media at quite early stages. However, promising new research directions still exist for further study. For example, we would like to investigate the influence imposed on topic results as different means of establishing word co-occurrence networks are taken. Besides, we may also consider of using a word's adjacent words within a given semantic distance to model its topics instead of only the direct neighbors.


Acknowledgements This work was supported by NSFC (Grant Nos. 71501005 and 61421003) and the fund of the State Key Lab of Software Development Environment (Grant No. SKLSDE-2015ZX-05).

References

- Andrzejewski D, Zhu X, Craven M (2009) Incorporating domain knowledge into topic modeling via dirichlet forest priors. In: ICML, pp 25–32
- Arora S, Ge R, Halpern Y, Mimno D, Moitra A, Sontag D, Wu Y, Zhu M (2013) A practical algorithm for topic modeling with provable guarantees. In: ICML, vol 28, pp 280–288
- Blei DM, Lafferty JD (2006) Dynamic topic models. In: ICML, pp 113–120
- Blei DM, McAuliffe JD (2007) Supervised topic models. In: NIPS, pp 121–128
- Blei DM, Ng AY, Jordan MI (2003) Latent dirichlet allocation. *J Mach Learn Res* 3:993–1022
- Cha Y, Cho J (2012) Social-network analysis using topic models. In: SIGIR, pp 565–574
- Chang J, Gerrish S, Wang C, Boyd-graber JL, Blei DM (2009) Reading tea leaves: How humans interpret topic models. In: NIPS, pp 288–296
- Chen M, Jin X, Shen D (2011) Short text classification improved by learning multi-granularity topics. In: IJCAI, pp 1776–1781
- Chen Y, Amiri H, Li Z, Chua TS (2013a) Emerging topic detection for organizations from microblogs. In: SIGIR, pp 43–52
- Chen Z, Mukherjee A, Liu B, Hsu M, Castellanos M, Ghosh R (2013b) Discovering coherent topics using general knowledge. In: CIKM, pp 209–218
- Chua FCT, Asur S (2013) Automatic summarization of events from social media. In: ICWSM
- Deerwester SC, Dumais ST, Landauer TK, Furnas GW, Harshman RA (1990) Indexing by latent semantic analysis. *JASIS* 41(6):391–407
- Fan R, Zhao J, Feng X, Xu K (2014) Topic dynamics in weibo: Happy entertainment dominates but angry finance is more periodic. In: ASONAM, pp 230–233
- Finkelstein L, Gabrilovich E, Matias Y, Rivlin E, Solan Z, Wolfman G, Ruppin E (2002) Placing search in context: The concept revisited. *ACM Trans Inf Syst* 20(1):116–131
- Heinrich G (2005) Parameter estimation for text analysis. Web:<http://www.babylonnet/publications/text-estpdf>
- Henderson K, Eliassi-Rad T (2009) Applying latent dirichlet allocation to group discovery in large graphs. In: SAC, pp 1456–1461
- Hofmann T (1999) Probabilistic latent semantic indexing. In: SIGIR, pp 50–57
- Hong L, Davison BD (2010) Empirical study of topic modeling in twitter. In: SOMA, pp 80–88

- Jagarlamudi J, Daumé H III, Udupa R (2012) Incorporating lexical priors into topic models. In: EACL, pp 204–213
- Jiang D, Leung KT, Vosecky J, Ng W (2014a) Personalized query suggestion with diversity awareness. In: ICDE, pp 400–411
- Jiang D, Leung KWT, Ng W (2014b) Fast topic discovery from web search streams. In: WWW, pp 949–960
- Jin O, Liu NN, Zhao K, Yu Y, Yang Q (2011) Transferring topical knowledge from auxiliary long texts for short text clustering. In: CIKM, pp 775–784
- Li C, Cheung W, Ye Y, Zhang X, Chu D, Li X (2015) The author-topic-community model for author interest profiling and community discovery. *Knowl Inf Syst* 44(2):359–383
- Lin T, Tian W, Mei Q, Cheng H (2014) The dual-sparse topic model: Mining focused topics and focused terms in short text. In: WWW, pp 539–550
- McCallum A, Mimno D, Wallach HM (2009) Rethinking lda: Why priors matter. In: NIPS, pp 1973–1981
- Mimno D, Wallach HM, Talley E, Leenders M, McCallum A (2011) Optimizing semantic coherence in topic models. In: EMNLP, pp 262–272
- Nigam K, McCallum A, Thrun S, Mitchell T (2000) Text classification from labeled and unlabeled documents using em. *Mach Learn* 39(2-3):103–134
- Peirsman Y, Heylen K, Geeraerts D (2008) Size matters: tight and loose context definitions in english word space models. In: Proceedings of the ESSLLI Workshop on Distributional Lexical Semantics, pp 34–41
- Phan XH, Nguyen LM, Horiguchi S (2008) Learning to classify short and sparse text & web with hidden topics from large-scale data collections. In: WWW, pp 91–100
- Quan X, Liu G, Lu Z, Ni X, Liu W (2010) Short text similarity based on probabilistic topics. *Knowl Inf Syst* 25(3):473–491
- Ramage D, Hall D, Nallapati R, Manning CD (2009) Labeled lda: A supervised topic model for credit attribution in multi-labeled corpora. In: EMNLP, pp 248–256
- Ramage D, Dumais S, Liebling D (2010) Characterizing microblogs with topic models. In: ICWSM
- Rosen-Zvi M, Griffiths T, Steyvers M, Smyth P (2004) The author-topic model for authors and documents. In: UAI, pp 487–494
- Rubenstein H, Goodenough JB (1965) Contextual correlates of synonymy. *Commun ACM* 8(10):627–633
- Rubin TN, Chambers A, Smyth P, Steyvers M (2012) Statistical topic models for multi-label document classification. *Mach Learn* 88(1-2):157–208
- Sahami M, Heilman TD (2006) A web-based kernel function for measuring the similarity of short text snippets. In: WWW, pp 377–386
- Stevens K, Kegelmeyer P, Andrzejewski D, Buttler D (2012) Exploring topic coherence over many models and many topics. In: EMNLP-CoNLL, pp 952–961
- Tang J, Meng Z, Nguyen X, Mei Q, Zhang M (2014) Understanding the limiting factors of topic modeling via posterior contraction analysis. In: ICML, pp 190–198
- Tong Y, Cao CC, Chen L (2014) Tcs: Efficient topic discovery over crowd-oriented service data. In: KDD, pp 861–870
- Wang X, McCallum A (2006) Topics over time: A non-markov continuous-time model of topical trends. In: KDD, pp 424–433
- Wang X, Jia Y, Zhou B, Ding Z, Zheng L (2011) Computing semantic relatedness using chinese wikipedia links and taxonomy. *Journal of Chinese Computer Systems* 32(11):2237–2242
- Weng J, Lim EP, Jiang J, He Q (2010) Twitterank: Finding topic-sensitive influential twitterers. In: WSDM, pp 261–270
- Yan X, Guo J, Lan Y, Cheng X (2013) A biterm topic model for short texts. In: WWW, pp 1445–1456
- Yu L, Asur S, Huberman BA (2011) What trends in chinese social media. arXiv preprint arXiv:11073522
- Yu LL, Asur S, Huberman BA (2013) Dynamics of trends and attention in chinese social media. arXiv preprint arXiv:13120649
- Zhao WX, Jiang J, Weng J, He J, Lim EP, Yan H, Li X (2011) Comparing twitter and traditional media using topic models. In: ECIR, pp 338–349
- Zhou T, Lyu MT, King I, Lou J (2015) Learning to suggest questions in social media. *Knowl Inf Syst* 43(2):389–416

Author Biographies



Yuan Zuo received a B.E. degree from Tianjin University of Technology, Tianjin, China, in 2011. He is currently a Ph.D. student at the School of Computer Science and Engineering, Beihang University, Beijing, China. His research interests include text mining and parallel computing.



Jichang Zhao received his B.E. and Ph.D. degrees from Beihang University in 2008 and 2014. He is currently an assistant professor at the School of Economics and Management, Beihang University. His research interests include online social media, complex networks and data mining.



Ke Xu is a professor at Beihang University, China. He received his B.E., M.E. and Ph.D. degrees from Beihang University in 1993, 1996 and 2000, respectively. His research interests include algorithm and complexity, data mining and complex networks.