

Learning Topic Models by Belief Propagation

Jia Zeng, *Member, IEEE*, William K. Cheung, *Member, IEEE*, and Jiming Liu, *Fellow, IEEE*

Abstract—Latent Dirichlet allocation (LDA) is an important hierarchical Bayesian model for probabilistic topic modeling, which attracts worldwide interests and touches on many important applications in text mining, computer vision and computational biology. This paper represents LDA as a factor graph within the Markov random field (MRF) framework, which enables the classic loopy belief propagation (BP) algorithm for approximate inference and parameter estimation. Although two commonly-used approximate inference methods, such as variational Bayes (VB) and collapsed Gibbs sampling (GS), have gained great successes in learning LDA, the proposed BP is competitive in both speed and accuracy as validated by encouraging experimental results on four large-scale document data sets. Furthermore, the BP algorithm has the potential to become a generic learning scheme for variants of LDA-based topic models. To this end, we show how to learn two typical variants of LDA-based topic models, such as author-topic models (ATM) and relational topic models (RTM), using BP based on the factor graph representation.

Index Terms—Latent Dirichlet allocation, topic models, belief propagation, message passing, factor graph, Bayesian networks, Markov random fields, hierarchical Bayesian models, Gibbs sampling, variational Bayes.



1 INTRODUCTION

Latent Dirichlet allocation (LDA) [1] is a three-layer hierarchical Bayesian model (HBM) that can infer probabilistic word clusters called topics from the document-word (document-term) matrix. LDA has no exact inference methods because of loops in its graphical representation. Variational Bayes (VB) [1] and collapsed Gibbs sampling (GS) [2] have been two commonly-used approximate inference methods for learning LDA and its extensions, including author-topic models (ATM) [3] and relational topic models (RTM) [4]. Other inference methods for probabilistic topic modeling include expectation-propagation (EP) [5] and collapsed VB inference (CVB) [6]. The connections and empirical comparisons among these approximate inference methods can be found in [7]. Recently, LDA and HBMs have found many important real-world applications in text mining and computer vision (e.g., tracking historical topics from time-stamped documents [8] and activity perception in crowded and complicated scenes [9]).

This paper represents LDA by the factor graph [10] within the Markov random field (MRF) framework [11]. From the MRF perspective, the topic modeling problem can be interpreted as a labeling problem, in which the objective is to assign a set of semantic topic labels to explain the observed nonzero elements in the document-word matrix. MRF solves the labeling problem existing widely in image analysis and computer vision by two important concepts: *neighborhood systems* and *clique potentials* [12] or *factor functions* [11]. It assigns the best

topic labels according to the *maximum a posteriori* (MAP) estimation through maximizing the posterior probability, which is in nature a prohibited combinatorial optimization problem in the discrete topic space. However, we often employ the smoothness prior [13] over neighboring topic labels to reduce the complexity by encouraging or penalizing only a limited number of possible labeling configurations.

The factor graph is a graphical representation method for both directed models (e.g., hidden Markov models (HMMs) [11, Chapter 13.2.3]) and undirected models (e.g., Markov random fields (MRFs) [11, Chapter 8.4.3]) because factor functions can represent both conditional and joint probabilities. In this paper, the proposed factor graph for LDA describes the same joint probability as that in the three-layer HBM, and thus it is not a new topic model but interprets LDA from a novel MRF perspective. The basic idea is inspired by the collapsed GS algorithm for LDA [2], [14], which integrates out multinomial parameters based on Dirichlet-Multinomial conjugacy and views Dirichlet hyperparameters as the pseudo topic labels having the same layer with the latent topic labels. In the collapsed hidden variable space, the joint probability of LDA can be represented as the product of *factor functions* in the factor graph. By contrast, the undirected model “harmonium” [15] encodes a different joint probability from LDA and probabilistic latent semantic analysis (PLSA) [16], so that it is a new and viable alternative to the directed models.

The factor graph representation facilitates the classic loopy belief propagation (BP) algorithm [10], [11], [17] for approximate inference and parameter estimation. By designing proper *neighborhood system* and *factor functions*, we may encourage or penalize different local labeling configurations in the neighborhood system to realize the topic modeling goal. The BP algorithm operates well on the factor graph, and it has the potential to become a

-
- J. Zeng is with the School of Computer Science and Technology, Soochow University, Suzhou 215006, China. He is also with the Shanghai Key Laboratory of Intelligent Information Processing, China. To whom correspondence should be addressed. E-mail: j.zeng@ieee.org.
 - W. K. Cheung and J. Liu are with the Department of Computer Science, Hong Kong Baptist University, Kowloon Tong, Hong Kong.

generic learning scheme for variants of LDA-based topic models. For example, we also extend the BP algorithm to learn ATM [3] and RTM [4] based on the factor graph representations. Although the convergence of BP is not guaranteed on general graphs [11], it often converges and works well in real-world applications.

The factor graph of LDA also reveals some intrinsic relations between HBM and MRF. HBM is a class of directed models within the Bayesian network framework [14], which represents the causal or conditional dependencies of observed and hidden variables in the hierarchical manner so that it is difficult to factorize the joint probability of hidden variables. By contrast, MRF can factorize the joint distribution of hidden variables into the product of factor functions according to the Hammersley-Clifford theorem [18], which facilitates the efficient BP algorithm for approximate inference and parameter estimation. Although learning HBM often has difficulty in estimating parameters and inferring hidden variables due to the causal coupling effects, the alternative factor graph representation as well as the BP-based learning scheme may shed more light on faster and more accurate algorithms for HBM.

The remainder of this paper is organized as follows. In Section 2 we introduce the factor graph interpretation for LDA, and derive the loopy BP algorithm for approximate inference and parameter estimation. Moreover, we discuss the intrinsic relations between BP and other state-of-the-art approximate inference algorithms. Sections 3 and 4 present how to learn ATM and RTM using the BP algorithms. Section 5 validates the BP algorithm on four document data sets. Finally, Section 6 draws conclusions and envisions future work.

2 BELIEF PROPAGATION FOR LDA

The probabilistic topic modeling task is to assign a set of semantic topic labels, $\mathbf{z} = \{z_{w,d}^k\}$, to explain the observed nonzero elements in the document-word matrix $\mathbf{x} = \{x_{w,d}\}$. The notations $1 \leq k \leq K$ is the topic index, $x_{w,d}$ is the number of word counts at the index $\{w,d\}$, $1 \leq w \leq W$ and $1 \leq d \leq D$ are the word index in the vocabulary and the document index in the corpus. Table 1 summarizes some important notations.

Fig. 1 shows the original three-layer graphical representation of LDA [1]. The document-specific topic proportion $\theta_d(k)$ generates a topic label $z_{w,d,i}^k \in \{0,1\}$, $\sum_{k=1}^K z_{w,d,i}^k = 1$, which in turn generates each observed word token i at the index w in the document d based on the topic-specific multinomial distribution $\phi_k(w)$ over the vocabulary words. Both multinomial parameters $\theta_d(k)$ and $\phi_k(w)$ are generated by two Dirichlet distributions with hyperparameters α and β , respectively. For simplicity, we consider only the smoothed LDA [2] with the fixed symmetric Dirichlet hyperparameters α and β . The plates indicate replications. For example, the document d repeats D times in the corpus, the word tokens w_n repeats N_d times in the document d , the vocabulary size is W , and there are K topics.

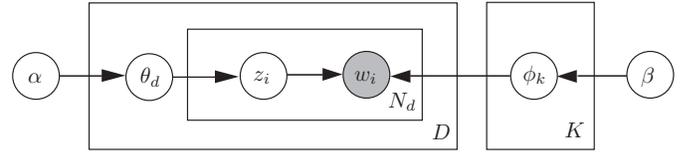


Fig. 1. Three-layer graphical representation of LDA [1].

TABLE 1
Notations

$1 \leq d \leq D$	Document index
$1 \leq w \leq W$	Word index in vocabulary
$1 \leq k \leq K$	Topic index
$1 \leq a \leq A$	Author index
$1 \leq c \leq C$	Link index
$\mathbf{x} = \{x_{w,d}\}$	Document-word matrix
$\mathbf{z} = \{z_{w,d}^k\}$	Topic labels for words
$\mathbf{z}_{-w,d}$	Labels for d excluding w
$\mathbf{z}_{w,-d}$	Labels for w excluding d
\mathbf{a}_d	Coauthors of the document d
$\mu_{\cdot,d}(k)$	$\sum_w x_{w,d} \mu_{w,d}(k)$
$\mu_{w,\cdot}(k)$	$\sum_d x_{w,d} \mu_{w,d}(k)$
θ_d	Factor of the document d
ϕ_w	Factor of the word w
η_c	Factor of the link c
$f(\cdot)$	Factor functions
α, β	Dirichlet hyperparameters

2.1 Factor Graph Representation

We begin by transforming the directed graph of Fig. 1 into a two-layer factor graph, of which a representative fragment is shown in Fig. 2. The notation, $z_{w,d}^k = \sum_{i=1}^{x_{w,d}} z_{w,d,i}^k / x_{w,d}$, denotes the average topic labeling configuration over all word tokens $1 \leq i \leq x_{w,d}$ at the index $\{w,d\}$. We define the *neighborhood system* of the topic label $z_{w,d}$ as $\mathbf{z}_{-w,d}$ and $\mathbf{z}_{w,-d}$, where $\mathbf{z}_{-w,d}$ denotes a set of topic labels associated with all word indices in the document d except w , and $\mathbf{z}_{w,-d}$ denotes a set of topic labels associated with the word indices w in all documents except d . The factors θ_d and ϕ_w are denoted by squares, and their connected variables $z_{w,d}$ are denoted by circles. The factor θ_d connects the neighboring topic labels $\{z_{w,d}, \mathbf{z}_{-w,d}\}$ at different word indices within the same document d , while the factor ϕ_w connects the neighboring topic labels $\{z_{w,d}, \mathbf{z}_{w,-d}\}$ at the same word index w but in different documents. We absorb the observed word w as the index of ϕ_w , which is similar to absorbing the observed document d as the index of θ_d . Because the factors can be parameterized functions [11], both θ_d and ϕ_w can represent the same multinomial parameters with Fig. 1.

Fig. 2 describes the same joint probability with Fig. 1 if we properly design the factor functions. The bipartite factor graph is inspired by the collapsed GS [2], [14] algorithm, which integrates out parameter variables $\{\theta, \phi\}$ in Fig. 1 and treats the hyperparameters $\{\alpha, \beta\}$ as pseudo

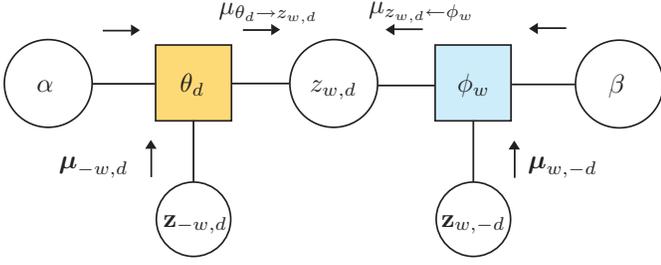


Fig. 2. Factor graph of LDA and message passing.

topic counts having the same layer with hidden variables \mathbf{z} . Thus, the joint probability of the *collapsed* hidden variables can be factorized as the product of factor functions. This collapsed view has been also discussed within the mean-field framework [19], inspiring the zero-order approximation CVB (CVB0) algorithm [7] for LDA. So, we speculate that all three-layer LDA-based topic models can be collapsed into the two-layer factor graph, which facilitates the BP algorithm for efficient inference and parameter estimation. However, how to use the two-layer factor graph to represent more general multi-layer HBM still remains to be studied.

Based on Dirichlet-Multinomial conjugacy, integrating out multinomial parameters $\{\theta, \phi\}$ yields the joint probability [14] of LDA in Fig. 1,

$$P(\mathbf{x}, \mathbf{z} | \alpha, \beta) \propto \prod_{d=1}^D \prod_{k=1}^K \frac{\Gamma(\sum_{w=1}^W x_{w,d} z_{w,d}^k + \alpha)}{\Gamma[\sum_{k=1}^K (\sum_{w=1}^W x_{w,d} z_{w,d}^k + \alpha)]} \times \prod_{w=1}^W \prod_{k=1}^K \frac{\Gamma(\sum_{d=1}^D x_{w,d} z_{w,d}^k + \beta)}{\Gamma[\sum_{d=1}^D (\sum_{w=1}^W x_{w,d} z_{w,d}^k + \beta)]}, \quad (1)$$

where $x_{w,d} z_{w,d}^k = \sum_{i=1}^{x_{w,d}} z_{w,d,i}^k$ recovers the original topic configuration over the word tokens in Fig. 1. Here, we design the *factor functions* as

$$f_{\theta_d}(\mathbf{x}_{\cdot,d}, \mathbf{z}_{\cdot,d}, \alpha) = \prod_{k=1}^K \frac{\Gamma(\sum_{w=1}^W x_{w,d} z_{w,d}^k + \alpha)}{\Gamma[\sum_{k=1}^K (\sum_{w=1}^W x_{w,d} z_{w,d}^k + \alpha)]}, \quad (2)$$

$$f_{\phi_w}(\mathbf{x}_{w,\cdot}, \mathbf{z}_{w,\cdot}, \beta) = \prod_{k=1}^K \frac{\Gamma(\sum_{d=1}^D x_{w,d} z_{w,d}^k + \beta)}{\Gamma[\sum_{d=1}^D (\sum_{w=1}^W x_{w,d} z_{w,d}^k + \beta)]}, \quad (3)$$

where $\mathbf{z}_{\cdot,d} = \{z_{w,d}, \mathbf{z}_{-w,d}\}$ and $\mathbf{z}_{w,\cdot} = \{z_{w,d}, \mathbf{z}_{w,-d}\}$ denote subsets of the variables in Fig. 2. Therefore, the joint probability (1) of LDA can be re-written as the product of factor functions [11, Eq. (8.59)] in Fig. 2,

$$P(\mathbf{x}, \mathbf{z} | \alpha, \beta) \propto \prod_{d=1}^D f_{\theta_d}(\mathbf{x}_{\cdot,d}, \mathbf{z}_{\cdot,d}, \alpha) \prod_{w=1}^W f_{\phi_w}(\mathbf{x}_{w,\cdot}, \mathbf{z}_{w,\cdot}, \beta). \quad (4)$$

Therefore, the two-layer factor graph in Fig. 2 encodes exactly the same information with the three-layer graph for LDA in Fig. 1. In this way, we may interpret LDA within the MRF framework to treat probabilistic topic

modeling as a labeling problem.

2.2 Belief Propagation (BP)

The BP [11] algorithms provide exact solutions for inference problems in tree-structured factor graphs but approximate solutions in factor graphs with loops. Rather than directly computing the conditional joint probability $p(\mathbf{z} | \mathbf{x})$, we compute the conditional marginal probability, $p(z_{w,d}^k = 1, x_{w,d} | \mathbf{z}_{-w,-d}^k, \mathbf{x}_{-w,-d})$, referred to as *message* $\mu_{w,d}(k)$, which can be normalized using a local computation, i.e., $\sum_{k=1}^K \mu_{w,d}(k) = 1, 0 \leq \mu_{w,d}(k) \leq 1$. According to the Markov property in Fig. 2, we obtain

$$p(z_{w,d}^k, x_{w,d} | \mathbf{z}_{-w,-d}^k, \mathbf{x}_{-w,-d}) \propto p(z_{w,d}^k, x_{w,d} | \mathbf{z}_{-w,d}^k, \mathbf{x}_{-w,-d}) p(z_{w,d}^k, x_{w,d} | \mathbf{z}_{w,-d}^k, \mathbf{x}_{w,-d}), \quad (5)$$

where $-w$ and $-d$ denote all word indices except w and all document indices except d , and the notations $\mathbf{z}_{-w,d}$ and $\mathbf{z}_{w,-d}$ represent all possible neighboring topic configurations. From the message passing view, $p(z_{w,d}^k, x_{w,d} | \mathbf{z}_{-w,d}^k, \mathbf{x}_{-w,-d})$ is the neighboring message $\mu_{\theta_d \rightarrow z_{w,d}}(k)$ sent from the factor node θ_d , and $p(z_{w,d}^k, x_{w,d} | \mathbf{z}_{w,-d}^k, \mathbf{x}_{w,-d})$ is the other neighboring message $\mu_{\phi_w \rightarrow z_{w,d}}(k)$ sent from the factor node ϕ_w . Notice that (5) uses the smoothness prior in MRF, which encourages only K smooth topic configurations within the neighborhood system. Using the Bayes' rule and the joint probability (1), we can expand Eq. (5) as

$$\begin{aligned} \mu_{w,d}(k) &\propto \frac{p(\mathbf{z}_{\cdot,d}, \mathbf{x}_{\cdot,d})}{p(\mathbf{z}_{-w,d}^k, \mathbf{x}_{-w,-d})} \times \frac{p(\mathbf{z}_{w,\cdot}, \mathbf{x}_{w,\cdot})}{p(\mathbf{z}_{w,-d}^k, \mathbf{x}_{w,-d})}, \\ &\propto \frac{\sum_{-w} x_{-w,d} z_{-w,d}^k + \alpha}{\sum_{k=1}^K (\sum_{-w} x_{-w,d} z_{-w,d}^k + \alpha)} \times \\ &\quad \frac{\sum_{-d} x_{w,-d} z_{w,-d}^k + \beta}{\sum_{w=1}^W (\sum_{-d} x_{w,-d} z_{w,-d}^k + \beta)}, \end{aligned} \quad (6)$$

where the property, $\Gamma(x+1) = x\Gamma(x)$, is used to cancel the common terms in both nominator and denominator [14]. We find that Eq. (6) updates the message on the variable $z_{w,d}^k$ if its neighboring topic configuration $\{\mathbf{z}_{-w,d}^k, \mathbf{z}_{w,-d}^k\}$ is known. However, due to uncertainty, we know only the neighboring messages rather than the precise topic configuration. So, we replace topic configurations by corresponding messages in Eq. (6) and obtain the following message update equation,

$$\mu_{w,d}(k) \propto \frac{\mu_{-w,d}(k) + \alpha}{\sum_k [\mu_{-w,d}(k) + \alpha]} \times \frac{\mu_{w,-d}(k) + \beta}{\sum_w [\mu_{w,-d}(k) + \beta]}, \quad (7)$$

where

$$\mu_{-w,d}(k) = \sum_{-w} x_{-w,d} \mu_{-w,d}(k), \quad (8)$$

$$\mu_{w,-d}(k) = \sum_{-d} x_{w,-d} \mu_{w,-d}(k). \quad (9)$$

Messages are passed from variables to factors, and in turn from factors to variables until convergence or the

input : $\mathbf{x}, K, T, \alpha, \beta$. output : θ_d, ϕ_w . $\mu_{w,d}^1(k) \leftarrow$ initialization and normalization; for $t \leftarrow 1$ to T do $\mu_{w,d}^{t+1}(k) \propto \frac{\mu_{w,d}^t(k) + \alpha}{\sum_k [\mu_{w,d}^t(k) + \alpha]} \times \frac{\mu_{w,-d}^t(k) + \beta}{\sum_w [\mu_{w,-d}^t(k) + \beta]}$; end $\theta_d(k) \leftarrow [\mu_{\cdot,d}(k) + \alpha] / \sum_k [\mu_{\cdot,d}(k) + \alpha]$; $\phi_w(k) \leftarrow [\mu_{w,\cdot}(k) + \beta] / \sum_w [\mu_{w,\cdot}(k) + \beta]$

Fig. 3. The synchronous BP for LDA.

maximum number of iterations is reached. Notice that we need only pass messages for $x_{w,d} \neq 0$. Because \mathbf{x} is a very sparse matrix, the message update equation (7) is computationally fast by sweeping all nonzero elements in the sparse matrix \mathbf{x} .

Based on messages, we can estimate the multinomial parameters θ and ϕ by the expectation-maximization (EM) algorithm [14]. The E-step infers the normalized message $\mu_{w,d}(k)$. Using the Dirichlet-Multinomial conjugacy and Bayes' rule, we express the marginal Dirichlet distributions on parameters as follows,

$$p(\theta_d) = \text{Dir}(\theta_d | \mu_{\cdot,d}(k) + \alpha), \quad (10)$$

$$p(\phi_w) = \text{Dir}(\phi_w | \mu_{w,\cdot}(k) + \beta). \quad (11)$$

The M-step maximizes (10) and (11) with respect to θ_d and ϕ_w , resulting in the following point estimates of multinomial parameters,

$$\theta_d(k) = \frac{\mu_{\cdot,d}(k) + \alpha}{\sum_k [\mu_{\cdot,d}(k) + \alpha]}, \quad (12)$$

$$\phi_w(k) = \frac{\mu_{w,\cdot}(k) + \beta}{\sum_w [\mu_{w,\cdot}(k) + \beta]}. \quad (13)$$

In this paper, we consider only fixed hyperparameters $\{\alpha, \beta\}$. Interested readers can figure out how to estimate hyperparameters based on inferred messages in [14].

To implement the BP algorithm, we must choose either the synchronous or the asynchronous update schedule to pass messages [20]. Fig. 3 shows the synchronous message update schedule. At each iteration t , each variable uses the incoming messages in the previous iteration $t - 1$ to calculate the current message. Once every variable computes its message, the message is passed to the neighboring variables and used to compute messages in the next iteration $t + 1$. An alternative is the asynchronous message update schedule. It updates the message of each variable immediately. The updated message is immediately used to compute other neighboring messages at each iteration t . The asynchronous update schedule often passes messages faster across variables, which causes the BP algorithm converge faster than the synchronous update schedule. Another termination condition for convergence is that the change of the multinomial parameters [1] is less than a predefined threshold λ , for example, $\lambda = 0.00001$ [21].

2.3 An Alternative View of BP

We may also adopt one of the BP instantiations, the sum-product algorithm [11], to infer $\mu_{w,d}(k)$. For convenience, we will not include the observation $x_{w,d}$ in the formulation. Fig. 2 shows the message passing from two factors θ_d and ϕ_w to the variable $z_{w,d}$, where the arrows denote the message passing directions. Based on the smoothness prior, we encourage only K smooth topic configurations without considering all other possible configurations. The message $\mu_{w,d}(k)$ is proportional to the product of both incoming messages from factors,

$$\mu_{w,d}(k) \propto \mu_{\theta_d \rightarrow z_{w,d}}(k) \times \mu_{\phi_w \rightarrow z_{w,d}}(k). \quad (14)$$

Eq. (14) has the same meaning with (5). The messages from factors to variables are the sum of all incoming messages from the neighboring variables,

$$\mu_{\theta_d \rightarrow z_{w,d}}(k) = f_{\theta_d} \prod_{-w} \mu_{-w,d}(k) \alpha, \quad (15)$$

$$\mu_{\phi_w \rightarrow z_{w,d}}(k) = f_{\phi_w} \prod_{-d} \mu_{w,-d}(k) \beta, \quad (16)$$

where α and β can be viewed as the pseudo-messages, and f_{θ_d} and f_{ϕ_w} are the factor functions that encourage or penalize the incoming messages.

In practice, however, Eqs. (15) and (16) often cause the product of multiple incoming messages close to zero [12]. To avoid arithmetic underflow, we use the sum operation rather than the product operation of incoming messages because when the product value increases the sum value also increases,

$$\prod_{-w} \mu_{-w,d}(k) \alpha \propto \sum_{-w} \mu_{-w,d}(k) + \alpha, \quad (17)$$

$$\prod_{-d} \mu_{w,-d}(k) \beta \propto \sum_{-d} \mu_{w,-d}(k) + \beta. \quad (18)$$

Such approximations as (17) and (18) transform the sum-product to the sum-sum algorithm, which resembles the relaxation labeling algorithm for learning MRF with good performance [12].

The normalized message $\mu_{w,d}(k)$ is multiplied by the number of word counts $x_{w,d}$ during the propagation, i.e., $x_{w,d} \mu_{w,d}(k)$. In this sense, $x_{w,d}$ can be viewed as the weight of $\mu_{w,d}(k)$ during the propagation, where the bigger $x_{w,d}$ corresponds to the larger influence of its message to those of its neighbors. Thus, the topics may be distorted by those documents with greater word counts. To avoid this problem, we may choose another weight like term frequency (TF) or term frequency-inverse document frequency (TF-IDF) for *weighted* belief propagation. In this sense, BP can not only handle *discrete* data, but also process *continuous* data like TF-IDF. The MRF model in Fig. 2 can be extended to describe both discrete and continuous data in general, while LDA in Fig. 1 focuses only on generating discrete data.

In the MRF model, we can design the factor functions arbitrarily to encourage or penalize local topic labeling configurations based on our prior knowledge. From

Fig. 1, LDA solves the topic modeling problem according to three intrinsic assumptions:

- 1) Co-occurrence: Different word indices within the same document tend to be associated with the same topic labels.
- 2) Smoothness: The same word indices in different documents are likely to be associated with the same topic labels.
- 3) Clustering: All word indices do not tend to associate with the same topic labels.

The first assumption is determined by the document-specific topic proportion $\theta_d(k)$, where it is more likely to assign a topic label $z_{w,d}^k = 1$ to the word index w if the topic k is more frequently assigned to other word indices in the document d . Similarly, the second assumption is based on the topic-specific multinomial distribution $\phi_k(w)$, which implies a higher likelihood to associate the word index w with the topic label $z_{w,d}^k = 1$ if k is more frequently assigned to the same word index w in other documents except d . The third assumption avoids grouping all word indices into one topic through normalizing $\phi_k(w)$ in terms of all word indices. If most word indices are associated with the topic k , the multinomial parameter ϕ_k will become too small to allocate the topic k to these word indices.

According to the above assumptions, we design f_{θ_d} and f_{ϕ_w} over messages as

$$f_{\theta_d}(\boldsymbol{\mu}_{\cdot,d}, \alpha) = \frac{1}{\sum_k [\boldsymbol{\mu}_{-w,d}(k) + \alpha]}, \quad (19)$$

$$f_{\phi_w}(\boldsymbol{\mu}_{w,\cdot}, \beta) = \frac{1}{\sum_w [\boldsymbol{\mu}_{w,-d}(k) + \beta]}. \quad (20)$$

Eq. (19) normalizes the incoming messages by the total number of messages for all topics associated with the document d to make outgoing messages comparable across documents. Eq. (20) normalizes the incoming messages by the total number of messages for all words in the vocabulary to make outgoing messages comparable across vocabulary words. Notice that (15) and (16) realize the first two assumptions, and (20) encodes the third assumption of topic modeling. The similar normalization technique to avoid partitioning all data points into one cluster has been used in the classic normalized cuts algorithm for image segmentation [22]. Combining (14) to (20) will yield the same message update equation (7). To estimate parameters θ_d and ϕ_w , we use the joint marginal distributions (15) and (16) of the set of variables belonging to factors θ_d and ϕ_w including the variable $z_{w,d}$, which produce the same point estimation equations (12) and (13).

2.4 Simplified BP (siBP)

We may simplify the message update equation (7). Substituting (12) and (13) into (7) yields the approximate message update equation,

$$\mu_{w,d}(k) \propto \theta_d(k) \times \phi_w(k), \quad (21)$$

```
function [phi, theta] = siBP(X, K, T, ALPHA, BETA)

% X is a W*D sparse matrix.
% W is the vocabulary size.
% D is the number of documents.
% The element of X is the word count 'xi'.
% 'wi' and 'di' are word and document indices.
% K is the number of topics.
% T is the number of iterations.
% mu is a matrix with K rows for topic messages.
% phi is a K*W matrix.
% theta is a K*D matrix.
% ALPHA and BETA are hyperparameters.
%
% normalize(A,dim) returns the normalized values
% (sum to one) of the elements along different
% dimensions of an array.

[wi,di,xi] = find(X);
% random initialization
mu = normalize(rand(K,nnz(X)),1);
% simplified belief propagation
for t = 1:T
    for k = 1:K
        md(k,:) = accumarray(di,xi'.*mu(k,:));
        mw(k,:) = accumarray(wi,xi'.*mu(k,:));
    end
    theta = normalize(md+ALPHA,1); %Eq.(9)
    phi = normalize(mw+BETA,2); %Eq.(10)
    mu = normalize(theta(:,di).*phi(:,wi),1); %Eq.(18)
end

return
```

Fig. 4. The MATLAB code for siBP.

which includes the current message $\mu_{w,d}(k)$ in both numerator and denominator in (7). In many real-world topic modeling tasks, a document often contains many different word indices, and the same word index appears in many different documents. So, at each iteration, Eq. (21) deviates slightly from (7) after adding the current message to both numerator and denominator. Such slight difference may be enlarged after many iterations in Fig. 3 due to accumulation effects, leading to different estimated parameters. Intuitively, Eq. (21) implies that if the topic k has a higher proportion in the document d , and it has the a higher likelihood to generate the word index w , it is more likely to allocate the topic k to the observed word $x_{w,d}$. This allocation scheme in principle follows the three intrinsic topic modeling assumptions in the subsection 2.3. Fig. 4 shows the MATLAB code for the simplified BP (siBP).

2.5 Relationship to Previous Algorithms

Here we discuss some intrinsic relations between BP with three state-of-the-art LDA learning algorithms such as VB [1], GS [2], and zero-order approximation CVB (CVB0) [7], [19] within the unified message passing framework. The message update scheme is an instantiation of the E-step of EM algorithm [23], which has been widely used to infer the marginal probabilities of hidden variables in various graphical models according to the maximum-likelihood estimation [11] (e.g., the E-step inference for GMMs [24], the forward-backward algorithm for HMMs [25], and the probabilistic relaxation labeling

algorithm for MRF [26]). After the E-step, we estimate the optimal parameters using the updated messages and observations at the M-step of EM algorithm.

VB is a variational message passing method [27] that uses a set of factorized variational distributions $q(\mathbf{z})$ to approximate the joint distribution (1) by minimizing the Kullback-Leibler (KL) divergence between them. Employing the Jensen’s inequality makes the approximate variational distribution an adjustable lower bound on the joint distribution, so that maximizing the joint probability is equivalent to maximizing the lower bound by tuning a set of variational parameters. The lower bound $q(\mathbf{z})$ is also an MRF in nature that approximates the joint distribution (1). Because there is always a gap between the lower bound and the true joint distribution, VB introduces bias when learning LDA. The variational message update equation is

$$\mu_{w,d}(k) \propto \frac{\exp[\Psi(\boldsymbol{\mu}_{\cdot,d}(k) + \alpha)]}{\exp[\Psi(\sum_k [\boldsymbol{\mu}_{\cdot,d}(k) + \alpha])]} \times \frac{\boldsymbol{\mu}_{w,\cdot}(k) + \beta}{\sum_w [\boldsymbol{\mu}_{w,\cdot}(k) + \beta]}, \quad (22)$$

which resembles the synchronous BP (7) but with two major differences. First, VB uses complicated digamma functions $\Psi(\cdot)$, which not only introduces bias [7] but also slows down the message updating. Second, VB uses a different variational EM schedule. At the E-step, it simultaneously updates both variational messages and parameter of θ_d until convergence, holding the variational parameter of ϕ fixed. At the M-step, VB updates only the variational parameter of ϕ .

The message update equation of GS is

$$\mu_{w,d,i}(k) \propto \frac{n_{\cdot,d}^{-i}(k) + \alpha}{\sum_k [n_{\cdot,d}^{-i}(k) + \alpha]} \times \frac{n_{w,\cdot}^{-i}(k) + \beta}{\sum_w [n_{w,\cdot}^{-i}(k) + \beta]}, \quad (23)$$

where $n_{\cdot,d}^{-i}(k)$ is the total number of topic labels k in the document d except the topic label on the current word token i , and $n_{w,\cdot}^{-i}(k)$ is the total number of topic labels k of the word w except the topic label on the current word token i . Eq. (23) resembles the asynchronous BP implementation (7) but with two subtle differences. First, GS randomly samples the current topic label $z_{w,d,i}^k = 1$ from the message $\mu_{w,d,i}(k)$, which truncates all K -tuple message values to zeros except the sampled topic label k . Such information loss introduces bias when learning LDA. Second, GS must sample a topic label for each word token, which repeats $x_{w,d}$ times for the word index $\{w, d\}$. The sweep of the entire word tokens rather than word index restricts GS’s scalability to large-scale document repositories containing billions of word tokens.

CVB0 is exactly equivalent to our asynchronous BP implementation but based on word tokens. Previous empirical comparisons [7] advocated the CVB0 algorithm for LDA within the approximate mean-field framework [19] closely connected with the proposed BP. Here we clearly explain that the superior performance of CVB0 has been

largely attributed to its asynchronous BP implementation from the MRF perspective. Our experiments also support that the message passing over word indices instead of tokens will produce comparable or even better topic modeling performance but with significantly smaller computational costs.

Eq. (21) also reveals that siBP is a probabilistic matrix factorization algorithm that factorizes the document-word matrix, $\mathbf{x} = [x_{w,d}]_{W \times D}$, into a matrix of document-specific topic proportions, $\boldsymbol{\theta} = [\theta_d(k)]_{K \times D}$, and a matrix of vocabulary word-specific topic proportions, $\boldsymbol{\phi} = [\phi_w(k)]_{K \times W}$, i.e., $\mathbf{x} \sim \boldsymbol{\phi}^T \boldsymbol{\theta}$. We see that the larger number of word counts $x_{w,d}$ corresponds to the higher likelihood $\sum_k \theta_d(k) \phi_w(k)$. From this point of view, the multinomial principle component analysis (PCA) [28] describes some intrinsic relations among LDA, PLSA [16], and non-negative matrix factorization (NMF) [29]. Eq. (21) is the same as the E-step update for PLSA except that the parameters θ and ϕ are smoothed by the hyperparameters α and β to prevent overfitting.

VB, BP and siBP have the computational complexity $\mathcal{O}(KDW_dT)$, but GS and CVB0 require $\mathcal{O}(KDN_dT)$, where W_d is the average vocabulary size, N_d is the average number of word tokens per document, and T is the number of learning iterations.

3 BELIEF PROPAGATION FOR ATM

Author-topic models (ATM) [3] depict each author of the document as a mixture of probabilistic topics, and have found important applications in matching papers with reviewers [30]. Fig. 5A shows the generative graphical representation for ATM, which first uses a document-specific uniform distribution u_d to generate an author index $a, 1 \leq a \leq A$, and then uses the author-specific topic proportions θ_a to generate a topic label $z_{w,d}^k = 1$ for the word index w in the document d . The plate on θ indicates that there are A unique authors in the corpus. The document often has multiple coauthors. ATM randomly assigns one of the observed author indices to each word in the document based on the document-specific uniform distribution u_d . However, it is more reasonable that each word $x_{w,d}$ is associated with an author index $a \in \mathbf{a}_d$ from the multinomial rather than uniform distribution, where \mathbf{a}_d is a set of author indices of the document d . As a result, each topic label takes two variables $z_{w,d}^{a,k} = \{0, 1\}$, $\sum_{a,k} z_{w,d}^{a,k} = 1, a \in \mathbf{a}_d, 1 \leq k \leq K$, where a is the author index and k is the topic index attached to the word.

We transform Fig. 5A to the factor graph representation of ATM in Fig. 5B. As with Fig. 2, we absorb the observed author index $a \in \mathbf{a}_d$ of the document d as the index of the factor $\theta_{a \in \mathbf{a}_d}$. The notation \mathbf{z}_{-w}^a denotes all labels connected with the authors $a \in \mathbf{a}_d$ except those for the word index w . The only difference between ATM and LDA is that the author $a \in \mathbf{a}_d$ instead of the document d connects the labels $z_{w,d}^a$ and \mathbf{z}_{-w}^a . As a result, ATM encourages topic smoothness among labels $z_{w,d}^a$ attached to the same author a instead of the same document d .

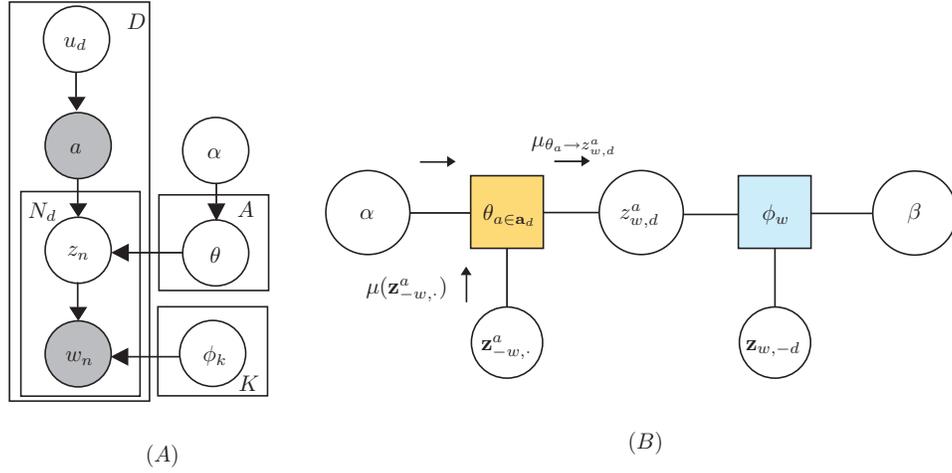


Fig. 5. (A) The three-layer graphical representation [3] and (B) two-layer factor graph of ATM.

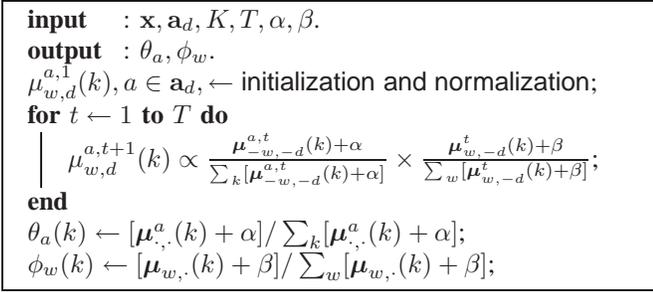


Fig. 6. The synchronous BP for ATM.

3.1 Inference and Parameter Estimation

Unlike passing the K -tuple message $\mu_{w,d}(k)$ in Fig. 3, the BP algorithm for learning ATM passes the $|\mathbf{a}_d| \times K$ -tuple message vectors $\mu_{w,d}^a(k), a \in \mathbf{a}_d$ through the factor $\theta_{a \in \mathbf{a}_d}$ in Fig. 5B, where $|\mathbf{a}_d|$ is the number of authors in the document d . Nevertheless, we can still obtain the K -tuple word topic message $\mu_{w,d}(k)$ by marginalizing the message $\mu_{w,d}^a(k)$ in terms of the author variable $a \in \mathbf{a}_d$ as follows,

$$\mu_{w,d}(k) = \sum_{a \in \mathbf{a}_d} \mu_{w,d}^a(k). \quad (24)$$

Since Figs. 2 and 5B have the same right half part, the message passing equation from the factor ϕ_w to the variable $z_{w,d}$ and the parameter estimation equation for ϕ_w in Fig. 5B remain the same as (7) and (13) based on the marginalized word topic message in (24). Thus, we only need to derive the message passing equation from the factor $\theta_{a \in \mathbf{a}_d}$ to the variable $z_{w,d}^a$ in Fig. 5B. Because of the topic smoothness prior, we design the factor function as follows,

$$f_{\theta_a} = \frac{1}{\sum_k [\mu_{-w,-d}^a(k) + \alpha]}, \quad (25)$$

where $\mu_{-w,-d}^a(k) = \sum_{-w,-d} x_{w,d}^a \mu_{w,d}^a(k)$ denotes the sum of all incoming messages attached to the author

index a and the topic index k excluding $x_{w,d}^a \mu_{w,d}^a(k)$. Likewise, Eq. (25) normalizes the incoming messages attached the author index a in terms of the topic index k to make outgoing messages comparable for different authors $a \in \mathbf{a}_d$. Similar to (15), we derive the message passing $\mu_{f_{\theta_a} \rightarrow z_{w,d}^a}$ through adding all incoming messages evaluated by the factor function (25).

Multiplying two messages from factors $\theta_{a \in \mathbf{a}_d}$ and ϕ_w yields the message update equation as follows,

$$\mu_{w,d}^a(k) \propto \frac{\mu_{-w,-d}^a(k) + \alpha}{\sum_k [\mu_{-w,-d}^a(k) + \alpha]} \times \frac{\mu_{w,-d}(k) + \beta}{\sum_w [\mu_{w,-d}(k) + \beta]}. \quad (26)$$

Notice that the $|\mathbf{a}_d| \times K$ -tuple message $\mu_{w,d}^a(k), a \in \mathbf{a}_d$ is normalized in terms of all combinations of $\{a, k\}, a \in \mathbf{a}_d, 1 \leq k \leq K$. Based on the normalized messages, the author-specific topic proportion $\theta_a(k)$ can be estimated from the sum of all incoming messages including $\mu_{w,d}^a$ evaluated by the factor function f_{θ_a} as follows,

$$\theta_a(k) = \frac{\mu_{\cdot,\cdot}^a(k) + \alpha}{\sum_k [\mu_{\cdot,\cdot}^a(k) + \alpha]}. \quad (27)$$

As a summary, Fig. 6 shows the synchronous BP algorithm for learning ATM. The difference between Fig. 3 and Fig. 6 is that Fig. 3 considers the author index a as the label for each word. At each iteration, the computational complexity is $\mathcal{O}(KDW_d A_d T)$, where A_d is the average number of authors per document.

4 BELIEF PROPAGATION FOR RTM

Network data, such as citation and coauthor networks of documents [30], [31], tag networks of documents and images [32], hyperlinked networks of web pages, and social networks of friends, exist pervasively in data mining and machine learning. The probabilistic relational topic modeling of network data can provide both useful predictive models and descriptive statistics [4].

In Fig. 7A, relational topic models (RTM) [4] represent entire document topics by the mean value of the docu-

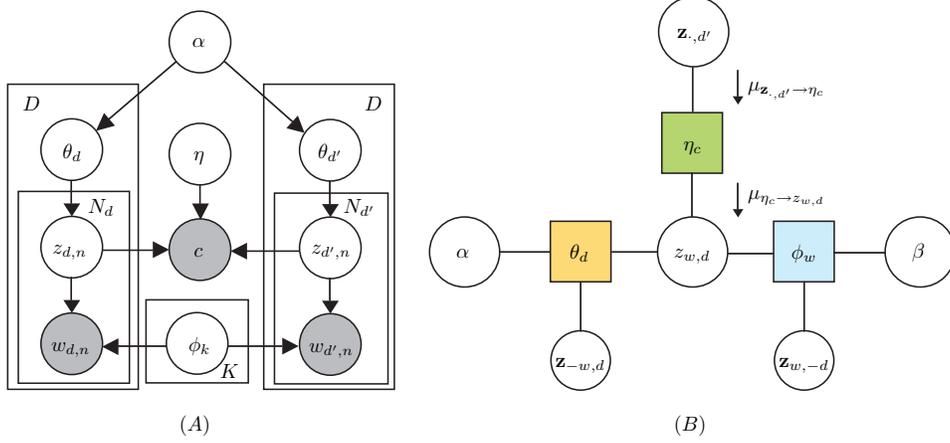


Fig. 7. (A) The three-layer graphical representation [4] and (B) two-layer factor graph of RTM.

```

input :  $w, c, \xi, K, T, \alpha, \beta$ .
output :  $\theta_a, \phi_w$ .
 $\mu_{w,d}^1(k) \leftarrow$  initialization and normalization;
for  $t \leftarrow 1$  to  $T$  do
   $\mu_{w,d}^{t+1}(k) \propto$ 
   $[(1 - \xi)\mu_{\theta_d \rightarrow z_{w,d}}^t(k) + \xi\mu_{\eta_c \rightarrow z_{w,d}}^t(k)] \times \mu_{\phi_w \rightarrow z_{w,d}}^t(k)$ ;
end
 $\theta_d(k) \leftarrow [\mu_{.,d}(k) + \alpha] / \sum_k [\mu_{.,d}(k) + \alpha]$ ;
 $\phi_w(k) \leftarrow [\mu_{w,.}(k) + \beta] / \sum_w [\mu_{w,.}(k) + \beta]$ ;

```

Fig. 8. The synchronous BP for RTM.

ment topic proportions, and use Hadamard product of mean values $\bar{z}_d \circ \bar{z}_{d'}$ from two linked documents $\{d, d'\}$ as link features, which are learned by the generalized linear model (GLM) η to generate the observed binary citation link variable $c = 1$. Besides, all other parts in RTM remain the same as LDA.

We transform Fig. 7A to the factor graph Fig. 7B by absorbing the observed link index $c \in \mathbf{c}, 1 \leq c \leq C$ as the index of the factor η_c . Each link index connects a document pair $\{d, d'\}$, and the factor η_c connects word topic labels $z_{w,d}$ and $\mathbf{z}_{.,d'}$ of the document pair. Besides encoding the topic smoothness, RTM explicitly describes the topic structural dependencies between the pair of linked documents $\{d, d'\}$ using the factor function $f_{\eta_c}(\cdot)$.

4.1 Inference and Parameter Estimation

In Fig. 7, the messages from the factors θ_d and ϕ_w to the variable $z_{w,d}$ are the same as LDA in (15) and (16). Thus, we only need to derive the message passing equation from the factor η_c to the variable $z_{w,d}$.

We design the factor function $f_{\eta_c}(\cdot)$ for linked documents as follows,

$$f_{\eta_c}(k|k') = \frac{\sum_{\{d,d'\}} \mu_{.,d}(k) \mu_{.,d'}(k')}{\sum_{\{d,d'\},k'} \mu_{.,d}(k) \mu_{.,d'}(k')}, \quad (28)$$

which depicts the likelihood of topic label k assigned to the document d when its linked document d' is

associated with the topic label k' . Notice that the designed factor function does not follow the GLM for link modeling in the original RTM [4] because the GLM makes inference slightly more complicated. However, similar to the GLM, Eq. (28) is also able to capture the topic interactions between two linked documents $\{d, d'\}$ in document networks. Instead of smoothness prior encoded by factor functions (19) and (20), it describes arbitrary topic dependencies $\{k, k'\}$ of linked documents $\{d, d'\}$.

Based on the factor function (28), we resort to the sum-product algorithm to calculate the message,

$$\mu_{\eta_c \rightarrow z_{w,d}}(k) = \sum_{d'} \sum_{k'} f_{\eta_c}(k|k') \mu_{.,d'}(k'), \quad (29)$$

where we use the sum rather than the product of messages from all linked documents d' to avoid arithmetic underflow. The standard sum-product algorithm requires the product of all messages from factors to variables. However, in practice, the direct product operation cannot balance the messages from different sources. For example, the message $\mu_{\theta_d \rightarrow z_{w,d}}$ is from the neighboring words within the same document d , while the message $\mu_{\eta_c \rightarrow z_{w,d}}$ is from all linked documents d' . If we pass the product of these two types of messages, we cannot distinguish which one influences more on the topic label $z_{w,d}$. Hence, we use the weighted sum of two types of messages,

$$\mu(z_{w,d} = k) \propto [(1 - \xi)\mu_{\theta_d \rightarrow z_{w,d}}(k) + \xi\mu_{\eta_c \rightarrow z_{w,d}}(k)] \times \mu_{\phi_w \rightarrow z_{w,d}}(k), \quad (30)$$

where $\xi \in [0, 1]$ is the weight to balance two messages $\mu_{\theta_d \rightarrow z_{w,d}}$ and $\mu_{\eta_c \rightarrow z_{w,d}}$. When there are no link information $\xi = 0$, Eq. (30) reduces to (7) so that RTM reduces to LDA. Fig. 8 shows the synchronous BP algorithm for learning RTM. Given the inferred messages, the parameter estimation equations remain the same as (12) and (13). The computational complexity at each iteration is $\mathcal{O}(K^2CDW_dT)$, where C is the total number of links in the document network.

TABLE 2
Summarization of four document data sets

Data sets	D	A	W	C	N_d	W_d
CORA	2410	2480	2961	8651	57	43
MEDL	2317	8906	8918	1168	104	66
NIPS	1740	2037	13649	—	1323	536
BLOG	5177	—	33574	1549	217	149

5 EXPERIMENTS

We use four large-scale document data sets:

- 1) CORA [33] contains abstracts from the CORA research paper search engine in machine learning area, where the documents can be classified into 7 major categories.
- 2) MEDL [34] contains abstracts from the MEDLINE biomedical paper search engine, where the documents fall broadly into 4 categories.
- 3) NIPS [35] includes papers from the conference “Neural Information Processing Systems”, where all papers are grouped into 13 categories. NIPS has no citation link information.
- 4) BLOG [36] contains a collection of political blogs on the subject of American politics in the year 2008. where all blogs can be broadly classified into 6 categories. BLOG has no author information.

Table 2 summarizes the statistics of four data sets, where D is the total number of documents, A is the total number of authors, W is the vocabulary size, C is the total number of links between documents, N_d is the average number of words per document, and W_d is the average vocabulary size per document.

5.1 BP for LDA

We compare BP with two commonly-used LDA learning algorithms such as VB [1] (Here we use Blei’s implementation of digamma functions)¹ and GS [2]² under the same fixed hyperparameters $\alpha = \beta = 0.01$. We use MATLAB C/C++ MEX-implementations for all these algorithms, and carry out experiments on a common PC with CPU 2.4GHz and RAM 4G. With the goal of repeatability, we have made our source codes and data sets publicly available [37].

To examine the convergence property of BP, we use the entire data set as the training set, and calculate the training perplexity [1] at every 10 iterations in the total of 1000 training iterations from the same initialization. Fig. 9 shows that the training perplexity of BP generally decreases rapidly as the number of training iterations increases. In our experiments, BP on average converges with the number of training iterations $T \approx 170$ when the difference of training perplexity between two successive iterations is less than one. Although this paper does not

theoretically prove that BP will definitely converge to the fixed point, the resemblance among VB, GS and BP in the subsection 2.5 implies that there should be the similar underlying principle that ensures BP to converge on general sparse word vector space in real-world applications. Further analysis reveals that BP on average uses more number of training iterations until convergence than VB ($T \approx 100$) but much less number of training iterations than GS ($T \approx 300$) on the four data sets. The fast convergence rate is a desirable property as far as the online [21] and distributed [38] topic modeling for large-scale corpus are concerned.

The predictive perplexity for the unseen test set is computed as follows [1], [7]. To ensure all algorithms to achieve the local optimum, we use the 1000 training iterations to estimate ϕ on the training set from the same initialization. In practice, this number of training iterations is large enough for convergence of all algorithms in Fig. 9. We randomly partition each document in the test set into 90% and 10% subsets. We use 1000 iterations of learning algorithms to estimate θ from the same initialization while holding ϕ fixed on the 90% subset, and then calculate the predictive perplexity on the left 10% subset,

$$\mathcal{P} = \exp \left\{ - \frac{\sum_{w,d} x_{w,d}^{10\%} \log [\sum_k \theta_d(k) \phi_w(k)]}{\sum_{w,d} x_{w,d}^{10\%}} \right\}, \quad (31)$$

where $x_{w,d}^{10\%}$ denotes word counts in the 10% subset. Notice that the perplexity (31) is based on the marginal probability of the word topic label $\mu_{w,d}(k)$ in (21).

Fig. 10 shows the predictive perplexity (average \pm standard deviation) from five-fold cross-validation for different topics, where the lower perplexity indicates the better generalization ability for the unseen test set. Consistently, BP has the lowest perplexity for different topics on four data sets, which confirms its effectiveness for learning LDA. On average, BP lowers around 11% than VB and 6% than GS in perplexity. Fig. 11 shows that BP uses less training time than both VB and GS. We show only 0.3 times of the real training time of VB because of time-consuming digamma functions. In fact, VB runs as fast as BP if we remove digamma functions. So, we believe that it is the digamma functions that slow down VB in learning LDA. BP is faster than GS because it computes messages for word indices. The speed difference is largest on the NIPS set due to its largest ratio $N_d/W_d = 2.47$ in Table 2. Although VB converges rapidly attributed to digamma functions, it often consumes triple more training time. Therefore, BP on average enjoys the highest efficiency for learning LDA with regard to the balance of convergence rate and training time.

We also compare six BP implementations such as siBP, BP and CVB0 [7] using both synchronous and asynchronous update schedules. We name three synchronous implementations as s-BP, s-siBP and s-CVB0, and three asynchronous implementations as a-BP, a-siBP

1. <http://www.cs.princeton.edu/~blei/lda-c/index.html>

2. http://psixp.ss.uci.edu/research/programs_data/toolbox.htm

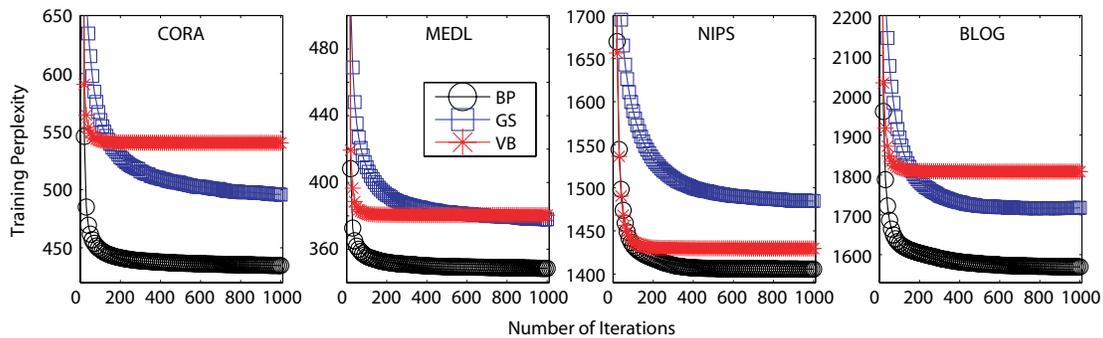


Fig. 9. Training perplexity as a function of number of iterations when $K = 50$ on CORA, MEDL, NIPS and BLOG.

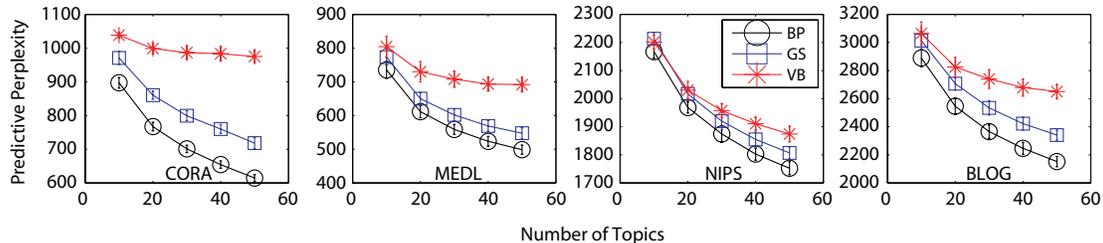


Fig. 10. Predictive perplexity as a function of number of topics on CORA, MEDL, NIPS and BLOG.

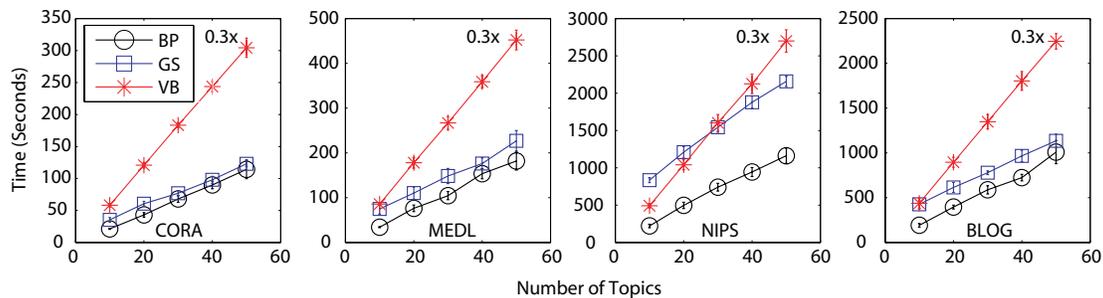


Fig. 11. Training time as a function of number of topics on CORA, MEDL, NIPS and BLOG. For VB, it shows 0.3 times of the real learning time denoted by 0.3x.

and a-CVB0. Because these six belief propagation implementations produce comparable perplexity, we show the relative perplexity that subtracts the mean value of six implementations in Fig. 12. Overall, the asynchronous schedule gives slightly lower perplexity than synchronous schedule because it passes messages faster and more efficiently. Except on CORA set, siBP generally provides the highest perplexity because it introduces subtle biases in computing perplexity at each iteration. The biased message will be propagated and accumulated leading to inaccurate parameter estimation. Although the proposed BP achieves lower perplexity than CVB0 on NIPS set, both of them work comparably well on other sets. But BP is much faster because it computes messages over word indices. The comparable results also confirm our assumption that topic modeling can be efficiently performed on word indices instead of tokens.

To measure the interpretability of a topic model, the *word intrusion* and *topic intrusion* are proposed to involve subjective judgements [39]. The basic idea is to ask volunteer subjects to identify the number of word

intruders in the topic as well as the topic intruders in the document, where intruders are defined as inconsistent words or topics based on prior knowledge of subjects. Fig. 13 shows the top ten words of $K = 10$ topics inferred by VB, GS and BP algorithms on NIPS set. We find no obvious difference with respect to word intrusions in each topic. Most topics share the similar top ten words but with different ranking orders. Despite significant perplexity difference, the topics extracted by three algorithms remains almost the same interpretability at least for the top ten words. This result coincides with [39] that the lower perplexity may not enhance interpretability of inferred topics.

Similar phenomenon has also been observed in MRF-based image labeling problems [20]. Different MRF inference algorithms such as graph cuts and BP often yield comparable results. Although one inference method may find more optimal MRF solutions, it does not necessarily translate into better performance compared to the ground-truth. The underlying hypothesis is that the ground-truth labeling configuration is often less optimal

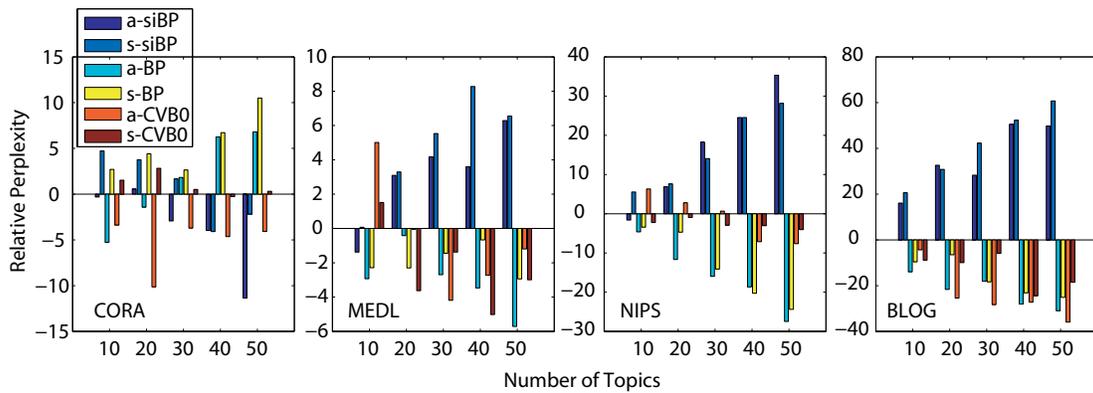


Fig. 12. Relative predictive perplexity as a function of number of topics on CORA, MEDL, NIPS and BLOG.

Topic 1	<p>recognition image system images network training speech figure model set image images figure object feature features recognition space visual distance recognition image images training system speech set feature features word</p>
Topic 2	<p>network networks units input neural learning hidden output training unit network networks neural input units output learning hidden layer weights network units networks input output neural hidden learning unit layer</p>
Topic 3	<p>analog chip circuit figure neural input output time system network analog circuit chip output figure signal input neural time system analog neural circuit chip figure output input time system signal</p>
Topic 4	<p>time model neurons neuron spike synaptic activity input firing information model neurons cells neuron cell visual activity response input stimulus neurons time model neuron synaptic spike cell activity input firing</p>
Topic 5	<p>model visual figure cells motion direction input spatial field orientation time noise dynamics order results point model system values figure model visual figure motion field direction spatial cells image orientation</p>
Topic 6	<p>function functions algorithm linear neural matrix learning space networks data function functions number set algorithm theorem tree bound learning class function functions algorithm set theorem linear number vector case space</p>
Topic 7	<p>learning network error neural training networks time function weight model function learning error algorithm training linear vector data set space learning error neural network function weight training networks time gradient</p>
Topic 8	<p>data training set error algorithm learning function class examples classification training data set performance classification recognition test class error speech training data set error learning performance test neural number classification</p>
Topic 9	<p>model data models distribution probability parameters gaussian algorithm likelihood mixture model data models distribution gaussian probability parameters likelihood mixture algorithm model data distribution models gaussian algorithm probability parameters likelihood mixture</p>
Topic 10	<p>learning state time control function policy reinforcement action algorithm optimal learning state control time model policy action reinforcement system states learning state time control policy function action reinforcement algorithm model</p>

Fig. 13. Top ten words of $K = 10$ topics of VB (first line), GS (second line), and BP (third line) on NIPS.

than solutions produced by inference algorithms. For example, if we manually label the topics for a corpus, the final perplexity is often higher than that of solutions returned by VB, GS and BP. For each document, LDA provides the equal number of topics K but the ground-truth often uses the unequal number of topics to explain the observed words, which may be another reason why the overall perplexity of learned LDA is often lower than that of the ground-truth. To test this hypothesis, we compare the perplexity of labeled LDA (L-LDA) [40] with LDA in Fig. 14. L-LDA is a supervised LDA that restricts the hidden topics as the observed class labels

of each document. When a document has multiple class labels, L-LDA automatically assigns one of the class labels to each word index. In this way, L-LDA resembles the process of manual topic labeling by human, and its solution can be viewed as close to the ground-truth. For a fair comparison, we set the number of topics $K = 7, 4, 13, 6$ of LDA for CORA, MEDL, NIPS and BLOG according to the number of document categories in each set. Both L-LDA and LDA are trained by BP using 500 iterations from the same initialization. Fig. 14 confirms that L-LDA produces higher perplexity than LDA, which partly supports that the ground-truth often yields the higher perplexity than the optimal solutions of LDA inferred by BP. The underlying reason may be

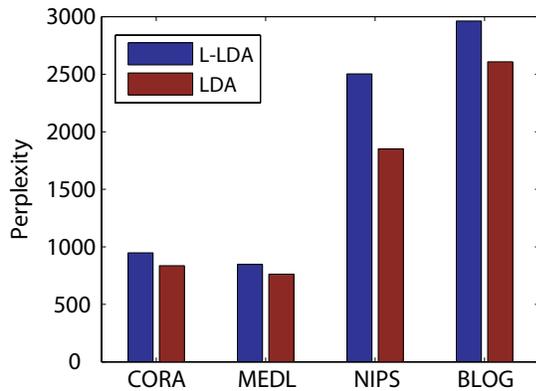


Fig. 14. Perplexity of L-LDA and LDA on four data sets.

that the three topic modeling rules encoded by LDA are still too simple to capture human behaviors in finding topics.

Under this situation, improving the formulation of topic models such as LDA is better than improving inference algorithms to enhance the topic modeling performance significantly. Although the proper settings of hyperparameters can make the predictive perplexity comparable for all state-of-the-art approximate inference algorithms [7], we still advocate BP because it is faster and more accurate than both VB and GS, even if they all can provide comparable perplexity and interpretability under the proper settings of hyperparameters.

5.2 BP for ATM

The GS algorithm for learning ATM is implemented in the MATLAB topic modeling toolbox.³ We compare BP and GS for learning ATM based on 500 iterations on training data. Fig. 15 shows the predictive perplexity (average \pm standard deviation) from five-fold cross-validation. On average, BP lowers 12% perplexity than GS, which is consistent with Fig. 10. Another possible reason for such improvements may be our assumption that all coauthors of the document account for the word topic label using multinomial instead of uniform probabilities.

5.3 BP for RTM

The GS algorithm for learning RTM is implemented in the R package.⁴ We compare BP with GS for learning RTM using the same 500 iterations on training data set. Based on the training perplexity, we manually set the weight $\xi = 0.15$ in Fig. 8 to achieve the overall superior performance on four data sets.

Fig. 16 shows predictive perplexity (average \pm standard deviation) on five-fold cross-validation. On average, BP lowers 6% perplexity than GS. Because the original RTM learned by GS is inflexible to balance information from different sources, it has slightly higher

perplexity than LDA (Fig. 10). To circumvent this problem, we introduce the weight ξ in (30) to balance two types of messages, so that the learned RTM gains lower perplexity than LDA. Future work will estimate the balancing weight ξ based on the feature selection or MRF structure learning techniques.

We also examine the link prediction performance of RTM. We define the link prediction as a binary classification problem. As with [4], we use the Hadamard product of a pair of document topic proportions as the link feature, and train an SVM [41] to decide if there is a link between them. Notice that the original RTM [4] learned by the GS algorithm uses the GLM to predict links. Fig. 17 compares the F-measure (average \pm standard deviation) of link prediction on five-fold cross-validation. Encouragingly, BP provides significantly 15% higher F-measure over GS on average. These results confirm the effectiveness of BP for capturing accurate topic structural dependencies in document networks.

6 CONCLUSIONS

First, this paper has presented the novel factor graph representation of LDA within the MRF framework. Not only does MRF solve topic modeling as a labeling problem, but also facilitate BP algorithms for approximate inference and parameter estimation in three steps:

- 1) First, we absorb $\{w, d\}$ as indices of factors, which connect hidden variables such as topic labels in the neighborhood system.
- 2) Second, we design the proper factor functions to encourage or penalize different local topic labeling configurations in the neighborhood system.
- 3) Third, we develop the approximate inference and parameter estimation algorithms within the message passing framework.

The BP algorithm is easy to implement, computationally efficient, faster and more accurate than other two approximate inference methods like VB [1] and GS [2] in several topic modeling tasks of broad interests. Furthermore, the superior performance of BP algorithm for learning ATM [3] and RTM [4] confirms its potential effectiveness in learning other LDA extensions.

Second, as the main contribution of this paper, the proper definition of neighborhood systems as well as the design of factor functions can interpret the three-layer LDA by the two-layer MRF in the sense that they encode the same joint probability. Since the probabilistic topic modeling is essentially a word annotation paradigm, the opened MRF perspective may inspire us to use other MRF-based image segmentation [22] or data clustering algorithms [17] for LDA-based topic models.

Finally, the scalability of BP is an important issue in our future work. As with VB and GS, the BP algorithm has a linear complexity with the number documents D and the number of topics K . We may extend the proposed BP algorithm for online [21] and distributed [38] learning of LDA, where the former incrementally learns

3. http://psiexp.ss.uci.edu/research/programs_data/toolbox.htm

4. <http://cran.r-project.org/web/packages/lda/>

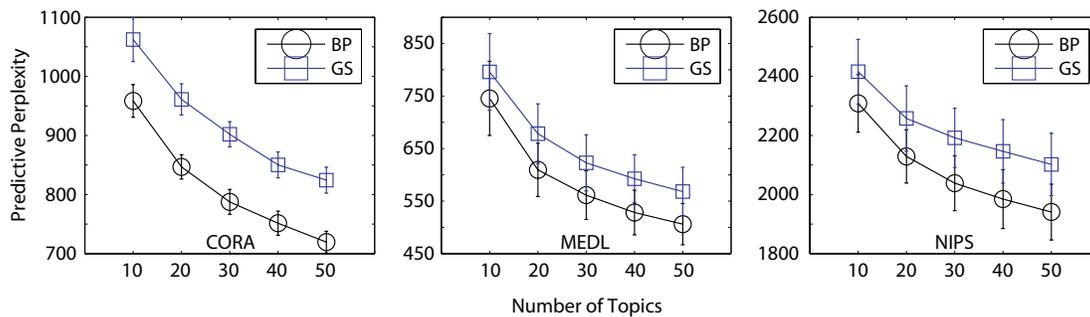


Fig. 15. Predictive perplexity as a function of number of topics for ATM on CORA, MEDL and NIPS.

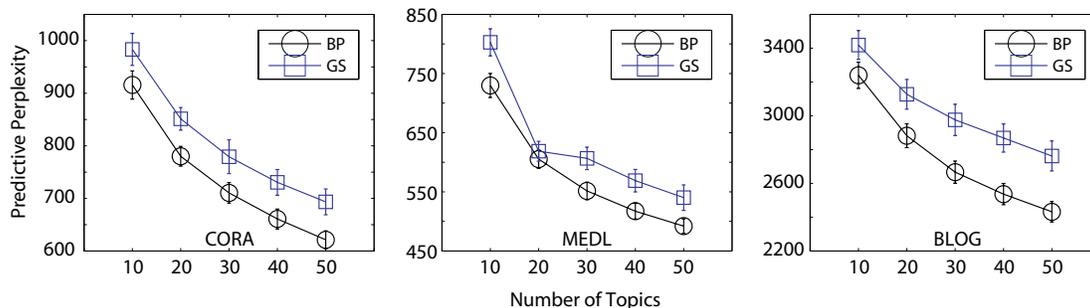


Fig. 16. Predictive perplexity as a function of number of topics for RTM on CORA, MEDL and BLOG.

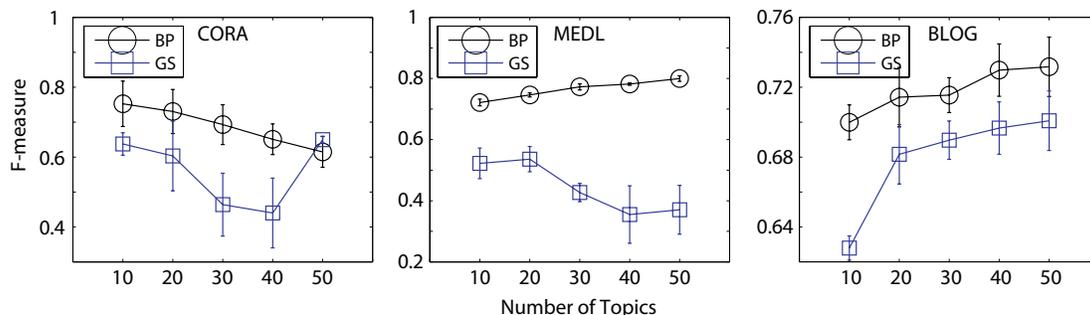


Fig. 17. F-measure of link prediction as a function of number of topics on CORA, MEDL and BLOG.

parts of D documents in data streams and the latter learns parts of D documents on distributed computing units. Since the K -tuple message is often sparse [42], we may also pass only salient parts of the K -tuple messages or only update those informative parts of messages at each learning iteration to speed up the whole message passing process.

ACKNOWLEDGEMENTS

This work is supported by NSFC (Grant No. 61003154), the Shanghai Key Laboratory of Intelligent Information Processing, China (Grant No. I IPL-2010-009), and a grant from Baidu.

REFERENCES

- [1] D. M. Blei, A. Y. Ng, and M. I. Jordan, "Latent Dirichlet allocation," *J. Mach. Learn. Res.*, vol. 3, pp. 993–1022, 2003.
- [2] T. L. Griffiths and M. Steyvers, "Finding scientific topics," *Proc. Natl. Acad. Sci.*, vol. 101, pp. 5228–5235, 2004.
- [3] M. Rosen-Zvi, T. Griffiths, M. Steyvers, and P. Smyth, "The author-topic model for authors and documents," in *UAI*, 2004, pp. 487–494.
- [4] J. Chang and D. M. Blei, "Hierarchical relational models for document networks," *Annals of Applied Statistics*, vol. 4, no. 1, pp. 124–150, 2010.
- [5] T. Minka and J. Lafferty, "Expectation-propagation for the generative aspect model," in *UAI*, 2002, pp. 352–359.
- [6] Y. W. Teh, D. Newman, and M. Welling, "A collapsed variational Bayesian inference algorithm for latent Dirichlet allocation," in *NIPS*, 2007, pp. 1353–1360.
- [7] A. Asuncion, M. Welling, P. Smyth, and Y. W. Teh, "On smoothing and inference for topic models," in *UAI*, 2009, pp. 27–34.
- [8] I. Pruteanu-Malinici, L. Ren, J. Paisley, E. Wang, and L. Carin, "Hierarchical Bayesian modeling of topics in time-stamped documents," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 32, no. 6, pp. 996–1011, 2010.
- [9] X. G. Wang, X. X. Ma, and W. E. L. Grimson, "Unsupervised activity perception in crowded and complicated scenes using hierarchical Bayesian models," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 3, pp. 539–555, 2009.
- [10] F. R. Kschischang, B. J. Frey, and H.-A. Loeliger, "Factor graphs and the sum-product algorithm," *IEEE Transactions on Inform. Theory*, vol. 47, no. 2, pp. 498–519, 2001.
- [11] C. M. Bishop, *Pattern recognition and machine learning*. Springer, 2006.

- [12] J. Zeng and Z.-Q. Liu, "Markov random field-based statistical character structure modeling for handwritten Chinese character recognition," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 30, no. 5, pp. 767–780, 2008.
- [13] S. Z. Li, *Markov Random Field Modeling in Image Analysis*. New York: Springer-Verlag, 2001.
- [14] G. Heinrich, "Parameter estimation for text analysis," University of Leipzig, Tech. Rep., 2008.
- [15] M. Welling, M. Rosen-Zvi, and G. Hinton, "Exponential family harmoniums with an application to information retrieval," in *NIPS*, 2004, pp. 1481–1488.
- [16] T. Hofmann, "Unsupervised learning by probabilistic latent semantic analysis," *Machine Learning*, vol. 42, pp. 177–196, 2001.
- [17] B. J. Frey and D. Dueck, "Clustering by passing messages between data points," *Science*, vol. 315, no. 5814, pp. 972–976, 2007.
- [18] J. M. Hammersley and P. Clifford, "Markov field on finite graphs and lattices," *unpublished*, 1971.
- [19] A. U. Asuncion, "Approximate Mean Field for Dirichlet-Based Models," in *ICML Workshop on Topic Models*, 2010.
- [20] M. F. Tappen and W. T. Freeman, "Comparison of graph cuts with belief propagation for stereo, using identical MRF parameters," in *ICCV*, 2003, pp. 900–907.
- [21] M. Hoffman, D. Blei, and F. Bach, "Online learning for latent Dirichlet allocation," in *NIPS*, 2010, pp. 856–864.
- [22] J. Shi and J. Malik, "Normalized cuts and image segmentation," *IEEE Trans. Pattern Anal. Machine Intell.*, vol. 22, no. 8, pp. 888–905, 2000.
- [23] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the EM algorithm," *Journal of the Royal Statistical Society, Series B*, vol. 39, pp. 1–38, 1977.
- [24] J. Zeng, L. Xie, and Z.-Q. Liu, "Type-2 fuzzy Gaussian mixture models," *Pattern Recognition*, vol. 41, no. 12, pp. 3636–3643, 2008.
- [25] J. Zeng and Z.-Q. Liu, "Type-2 fuzzy hidden Markov models and their application to speech recognition," *IEEE Trans. Fuzzy Syst.*, vol. 14, no. 3, pp. 454–467, 2006.
- [26] J. Zeng and Z. Q. Liu, "Type-2 fuzzy Markov random fields and their application to handwritten Chinese character recognition," *IEEE Trans. Fuzzy Syst.*, vol. 16, no. 3, pp. 747–760, 2008.
- [27] J. Winn and C. M. Bishop, "Variational message passing," *J. Mach. Learn. Res.*, vol. 6, pp. 661–694, 2005.
- [28] W. L. Buntine, "Variational extensions to EM and multinomial PCA," in *ECML*, 2002, pp. 23–34.
- [29] D. D. Lee and H. S. Seung, "Learning the parts of objects by non-negative matrix factorization," *Nature*, vol. 401, pp. 788–791, 1999.
- [30] J. Zeng, W. K. Cheung, C.-H. Li, and J. Liu, "Coauthor network topic models with application to expert finding," in *IEEE/WIC/ACM WI-IAT*, 2010, pp. 366–373.
- [31] J. Zeng, W. K.-W. Cheung, C.-H. Li, and J. Liu, "Multirelational topic models," in *ICDM*, 2009, pp. 1070–1075.
- [32] J. Zeng, W. Feng, W. K. Cheung, and C.-H. Li, "Higher-order Markov tag-topic models for tagged documents and images," *arXiv:1109.5370v1 [cs.CV]*, 2011.
- [33] A. K. McCallum, K. Nigam, J. Rennie, and K. Seymore, "Automating the construction of internet portals with machine learning," *Information Retrieval*, vol. 3, no. 2, pp. 127–163, 2000.
- [34] S. Zhu, J. Zeng, and H. Mamitsuka, "Enhancing MEDLINE document clustering by incorporating MeSH semantic similarity," *Bioinformatics*, vol. 25, no. 15, pp. 1944–1951, 2009.
- [35] A. Globerson, G. Chechik, F. Pereira, and N. Tishby, "Euclidean embedding of co-occurrence data," *J. Mach. Learn. Res.*, vol. 8, pp. 2265–2295, 2007.
- [36] J. Eisenstein and E. Xing, "The CMU 2008 political blog corpus," Carnegie Mellon University, Tech. Rep., 2010.
- [37] J. Zeng, "TMBP: A topic modeling toolbox using belief propagation," *J. Mach. Learn. Res.*, p. arXiv:1201.0838v1 [cs.LG], 2012.
- [38] K. Zhai, J. Boyd-Graber, and N. Asadi, "Using variational inference and MapReduce to scale topic modeling," *arXiv:1107.3765v1 [cs.AI]*, 2011.
- [39] J. Chang, J. Boyd-Graber, S. Gerris, C. Wang, and D. Blei, "Reading tea leaves: How humans interpret topic models," in *NIPS*, 2009, pp. 288–296.
- [40] D. Ramage, D. Hall, R. Nallapati, and C. D. Manning, "Labeled LDA: A supervised topic model for credit attribution in multi-labeled corpora," in *Empirical Methods in Natural Language Processing*, 2009, pp. 248–256.
- [41] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011.
- [42] I. Porteous, D. Newman, A. Ihler, A. Asuncion, P. Smyth, and M. Welling, "Fast collapsed Gibbs sampling for latent Dirichlet allocation," in *KDD*, 2009, pp. 569–577.