

Jointly Modeling Topics and Intents with Global Order Structure

Bei Chen[†], Jun Zhu^{†*}, Nan Yang[‡], Tian Tian[†], Ming Zhou[‡], Bo Zhang[†]

[†]Dept. of Comp. Sci. & Tech., State Key Lab of Intell. Tech. & Sys.,
Center for Bio-Inspired Computing Research, TNLList, Tsinghua University, Beijing, 100084, China

[‡]Microsoft Research Asia, Beijing, 100080, China

{chenbei12@mails., dcszj@, tiant13@mails., dcszb}@tsinghua.edu.cn; {nanya, mingzhou}@microsoft.com

Abstract

Modeling document structure is of great importance for discourse analysis and related applications. The goal of this research is to capture the document intent structure by modeling documents as a mixture of topic words and rhetorical words. While the topics are relatively unchanged through one document, the rhetorical functions of sentences usually change following certain orders in discourse. We propose GMM-LDA, a topic modeling based Bayesian unsupervised model, to analyze the document intent structure cooperated with order information. Our model is flexible that has the ability to combine the annotations and do supervised learning. Additionally, entropic regularization can be introduced to model the significant divergence between topics and intents. We perform experiments in both unsupervised and supervised settings, results show the superiority of our model over several state-of-the-art baselines.

Introduction

People often organize utterances into meaningful and coherent documents, conforming to certain conventions and underlying structures. For example, scripts (Frermann, Titov, and Pinkal 2014), scientific papers (Ó Séaghdha and Teufel 2014), official mails, news and encyclopedia articles all have relatively fixed discourse structure and exhibit recurrent patterns. Learning the document structure is of great importance for discourse analysis and has various applications, such as text generation (Prasad et al. 2005) and text summarization (Louis, Joshi, and Nenkova 2010).

There are two important aspects of document structure learning: topic modeling and rhetorical structure modeling. Topic modeling assumes multiple topics often exist within a domain. It aims to discover the latent semantics of the documents, with many popular models such as Latent Dirichlet Allocation (LDA) (Blei, Ng, and Jordan 2003), in which each document is posited as an admixture over an underlying set of topics, and each word is draw from a specific topic. Rhetorical structure modeling aims to uncover the underlying organization of documents. Inspired by the discourse theory (Mann and Thompson 1988), each sentence in a document can be assigned a rhetorical function, or called *in-*

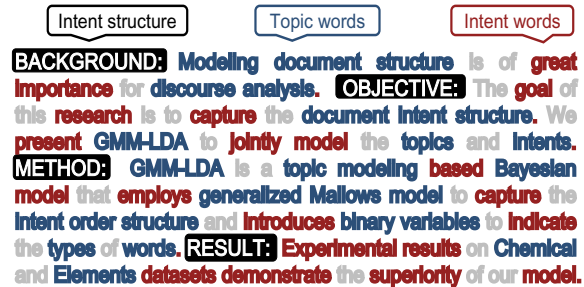


Figure 1: Demonstration of the intent and topic words. Stop words (in gray) can be removed by preprocessing.

tent. For example, the sentences in a scientific paper may have different intents such as “background”, “objective”, “method” and “result”. Fig 1 presents an example of the intent structure of an abstract.

Document-level topics and sentence-level intents usually show contradictory characteristics. For example, it is often sensible to assume that the topics are relatively unchanged through one document (e.g., in LDA), while the sentences’ intents usually change following certain order in discourse. Furthermore, each document often follows a progression of nonrecurring coherent intents (Halliday and Hasan 1976), and the sentences with the same intent tend to appear within the same block of a document. Based on these observations, it’s natural to think that jointly considering these two incompatible structures can help to model document better.

In this paper, we present a hierarchical Bayesian model to discover both the topic structure and rhetorical structure of documents by jointly considering topics and the above order structure in discourse. To this end, we assume that all the words can be divided into two types: *topic word* and *intent word*. Specifically, topic words in a document are relevant to the document’s topic and spread throughout the document, while intent words mainly contribute to the rhetorical functions of the sentences. Following the example in Fig. 1, the words “document”, “discourse” and “intent” are likely to be topic words; they indicate the specific research domain of this document. Meanwhile, the words “result”, “dataset” and “demonstrate” are likely to be intent words and a sentence with these words may have the intent structure label “result”.

*Corresponding author.

We introduce a binary variable for each word to indicate its type, and model the topic structure and intent structure respectively using topic models. Inspired by the generalized Mallows model (GMM) (Fligner and Verducci 1986), we incorporate the intent order structure using GMM-Multi prior, which not only conforms to our intuition of nonrecurring coherent intents but also captures the global effects of the orders. To further improve the expressive ability, we present two important variants of GMM-LDA. One is to incorporate the known intent labels of sentences for supervised learning, and the other is to incorporate entropic regularization to better separate the words into two types under the regularized Bayesian inference framework (Zhu, Chen, and Xing 2014). Finally, we represent experiments on real datasets to demonstrate the effectiveness of our methods over several state-of-the-art baselines. To the best of our knowledge, we present the first model for topics and intents simultaneously, where the intent order structure is described globally by using GMM-Multi prior. In the rest paper, we first present GMM-LDA, followed by the supervised version and entropic regularization. Then we present experimental results with analysis. Finally, we discuss related work and conclude.

Unsupervised GMM-LDA Model

We consider the following document structure learning problem. We are given a corpus $\mathcal{D} = \{\mathbf{s}_d\}_{d=1}^D$ with D documents, where a document \mathbf{s}_d is a sequence of N_d sentences denoted by $\mathbf{s}_d = (\mathbf{w}_{d1}, \mathbf{w}_{d2}, \dots, \mathbf{w}_{dN_d})$ and a sentence \mathbf{w}_{ds} is a bag of N_{ds} words denoted by $\mathbf{w}_{ds} = \{w_{dsm}\}_{m=1}^{N_{ds}}$. The size of the vocabulary is V . There are T topics and K intents in total. Topics are document-level, while intents are sentence-level. Our goal of document structure learning is to model the intent and the topic simultaneously and assign an intent label to each sentence in the corpus.

We build our models on the following assumptions, 1) **Type**: Each word in the corpus is either an intent word or a topic word; 2) **Order**: The intents of sentences within a document change following certain orders and the orders are similar within a domain; and 3) **Coherence**: The same intent does not appear in disconnected portions of a document.

To characterize the type assumption, we associate each word with a binary variable to indicate whether it is an intent word or a topic word. For the order and coherence assumptions, we introduce a GMM-Multi prior to model the intent order structure. We start with the description of GMM-Multi prior, then GMM-LDA in detail with its inference method.

GMM-Multi Prior for Intent Ordering

According to the coherence assumption, the same intent could not be assigned to the unconnected portions of a document. To satisfy this assumption, we introduce a GMM-Multi prior over possible intent permutations. GMM-Multi, motivated by (Chen et al. 2009), is an extension to the generalized Mallows model (Fligner and Verducci 1986). It concentrates probability mass on a small set of similar permutations around a canonical ordering π_0 , which confirms to the intuition that the intent orders are similar within a domain. The *inversion representation* of permutations is used

instead of the direct order sequence: If we set the canonical ordering π_0 as the identity permutation $(1, 2, \dots, K)$, then any permutation π can be denoted as a $(K-1)$ -dimensional vector $\mathbf{v} = (v_1, v_2, \dots, v_{K-1})$, where v_k is the count of the numbers in π that are before k and greater than k . For instance, permutation $\pi = (2, 1, 4, 5, 3)$ can be represented as $\mathbf{v} = (1, 0, 2, 0)$, where $v_3 = 2$ as there are two numbers, 4 and 5, that are before 3 and greater than 3. v_K is omitted as it is always zero. Obviously, a one-to-one correlation exists between these two kinds of permutation representations. Furthermore, each element of \mathbf{v} is independent of each other. The marginal distribution over v_k is

$$\text{GMM}_k(v_k; \rho_k) = \frac{e^{-\rho_k v_k}}{\psi_k(\rho_k)}, \quad (1)$$

where $\psi_k(\rho_k) = \frac{1 - \exp(-\frac{K-k+1}{\rho_k})}{1 - \exp(-\rho_k)}$ is the normalization factor and $\rho_k > 0$ is the dispersion parameter. Then the probability mass function with parameters $\boldsymbol{\rho} = (\rho_1, \rho_2, \dots, \rho_{K-1})$ can be written as the product over all the components, so that the number of parameters grows linearly with the number of intents. Due to the exponential form, the conjugate prior for ρ_k is also an exponential distribution with two parameters :

$$\text{GMM}_0(\rho_k; v_{k,0}, \nu_0) \propto e^{(-\rho_k v_{k,0} - \log \psi_k(\rho_k)) \nu_0}. \quad (2)$$

Intuitively, ν_0 is the number of previous trials and $v_{k,0}$ is the average number in previous trials for v_k . Let ρ_0 be the prior value for each ρ_k . By setting the maximum likelihood estimate of ρ_k to be ρ_0 , we can obtain $v_{k,0} = \frac{1}{\exp(\rho_0) - 1} - \frac{K-k+1}{\exp((K-k+1)\rho_0) - 1}$.

The GMM-Multi prior is defined over sentence intents \mathbf{z}_d via a generative process on a given inversion representation of permutation \mathbf{v}_d . Firstly, we draw a bag of intents \mathbf{u}_d (N_d elements) and each element $u_{ds} \sim \text{Multinomial}(\boldsymbol{\lambda})$. Then $\sum_s \mathbb{I}(u_{ds} = k)$ represents the number of sentences of intent k in this document. Secondly, we obtain the permutation of intents $\boldsymbol{\pi}_d = \text{Compute-}\pi(\mathbf{v}_d)$, where $\text{Compute-}\pi$ transforms \mathbf{v}_d into the intent permutation $\boldsymbol{\pi}_d$. And finally, we obtain the intent structure of the sentences $\mathbf{z}_d = \text{Compute-}z(\mathbf{u}_d, \boldsymbol{\pi}_d)$, so that z_{ds} is the intent label for sentence \mathbf{w}_{ds} . $\text{Compute-}z$ is the algorithm to obtain the intent structure \mathbf{z}_d using the bag of intents \mathbf{u}_d and the permutation $\boldsymbol{\pi}_d$. It arranges all the intent labels in \mathbf{u}_d as the order of $\boldsymbol{\pi}_d$ with the same intent labels appearing together. For example, when $\mathbf{u}_d = \{1, 1, 2, 3, 3, 3, 5\}$ and $\boldsymbol{\pi}_d = (2, 1, 4, 5, 3)$, we can obtain $\mathbf{z}_d = (2, 1, 1, 5, 3, 3, 3)$. Note that not all the K intent labels should appear in \mathbf{z}_d ; it depends on the intent labels in \mathbf{u}_d . As in the example, \mathbf{u}_d does not contain 4, so 4 does not appear in the intent structure \mathbf{z}_d .

Generative Process of GMM-LDA

Now we present GMM-LDA, an unsupervised Bayesian generative model, as illustrated in Fig. 2. GMM-LDA simultaneously models the topics (the blue part in Fig. 2) and the intents (the red part in Fig. 2). The binary variable b_{dsm} denotes the type of word w_{dsm} : if $b_{dsm} = 1$, then w_{dsm} is a topic word; if $b_{dsm} = 0$, then w_{dsm} is an intent word. Each




Figure 2: The graphical structure of GMM-LDA.

sentence has a specific intent label z_{ds} , while each document s_d has a topic mixing distribution θ_d . For topic words, there is a document-specific topic model. It is a hierarchical Bayesian model that posits each document as an admixture of T topics, where each topic β_t is a multinomial distribution over a V -word vocabulary, drawn from a language model $\beta_t \sim \text{Dirichlet}(\beta_0)$ ($t \in [T]$). For intent words, there is a rhetorical language model, in which the intent structure z_d of document s_d is generated from a bag of intent labels u_d and an intent permutation π_d follows the GMM-Multi prior. The total number of intents is K , and each intent α_k is also a multinomial distribution over vocabulary, drawn from another language model $\alpha_k \sim \text{Dirichlet}(\alpha_0)$ ($k \in [K]$).

For each document s_d , the generating process is

1. Draw a topic proportion $\theta_d \sim \text{Dirichlet}(\theta_0)$.
2. Obtain the intent structure $z_d \sim \text{GMM-Multi}(\rho, \lambda)$, so that z_{ds} is the intent label for sentence w_{ds} .
3. For each word w_{dsm} in document d ,
 - (a) Draw an indicator $b_{dsm} \sim \text{Bernoulli}(\gamma)$.
 - (b) If $b_{dsm} = 0$, then w_{dsm} is from intent part: draw $w_{dsm} \sim \text{Multinomial}(\alpha_{z_{ds}})$.
 - (c) if $b_{dsm} = 1$, then w_{dsm} is from topic part: draw a topic $t_{dsm} \sim \text{Multinomial}(\theta_d)$, and draw $w_{dsm} \sim \text{Multinomial}(\beta_{t_{dsm}})$.

For fully-Bayesian GMM-LDA, we assume the following priors: $\gamma \sim \text{Beta}(\gamma_0)$, $\lambda \sim \text{Dirichlet}(\lambda_0)$, which is a distribution to express how likely each intent label is to appear regardless of positions, and each component of ρ follows $\rho_k \sim \text{GMM}_0(\rho_0, \nu_0)$, $k \in [K - 1]$. The variables subscripted with 0 are fixed prior hyperparameters.

Collapsed Gibbs Sampling

Let $\mathcal{M} = \{\mathbf{u}, \rho, \mathbf{v}, \mathbf{t}, \mathbf{b}\}$ denote all the variables to be learned during training. Then with Bayes' theorem, the posterior distribution of GMM-LDA is

$$q(\mathcal{M}|\mathcal{D}) \propto p_0(\mathcal{M})p(\mathcal{D}|\mathcal{M}), \quad (3)$$

where $p_0(\mathcal{M})$ is the prior and $p(\mathcal{D}|\mathcal{M})$ is the likelihood. Then it can be learned with a collapsed Gibbs sampler due to the conjugate priors. We naturally split the variables into four parts, namely \mathbf{u} , ρ , \mathbf{v} and $\{\mathbf{t}, \mathbf{b}\}$, then sample them using their posterior distributions respectively.

For \mathbf{u} : For each document s_d , u_d is a bag of intent labels with N_d elements. We resample each element u_{ds} via:

$$q(u_{ds} = x|\mathcal{D}, \mathcal{M}_{-u_{ds}}) \propto p_0(u_{ds} = x)p(s_d|\mathcal{M}, \mathcal{D}_{-s_d}) \propto (f_{u=x}^{-ds} + \lambda_0) \prod_{k=1}^K \frac{\Gamma(f_{0,k,\cdot}^{-d} + V\alpha_0)}{\Gamma(f_{0,k,\cdot} + V\alpha_0)} \prod_{v=1}^V \frac{\Gamma(f_{0,k,v} + \alpha_0)}{\Gamma(f_{0,k,v}^{-d} + \alpha_0)},$$

where the subscript “-” denotes that some elements are omitted from a set, e.g., $\mathcal{M}_{-u_{ds}}$ is the set \mathcal{M} except u_{ds} and \mathcal{D}_{-s_d} is the set of all documents in \mathcal{D} except s_d . Let $\mathbf{u} = \{u_d\}_{d=1}^D$ denote all the intent labels in the corpus. $f_{u=x}^{-ds}$ is the count of times that intent label x appears in \mathbf{u} except u_{ds} . Let w_v denote the v -th word in the vocabulary¹, where $v \in [V]$. $f_{0,k,v}$ counts the times that w_v appears as an intent word in the sentences with intent label k . Then, $f_{0,k,\cdot}^{-d} = \sum_{v=1}^V f_{0,k,v}^{-d}$. The superscript “- d ” indicates that the frequency is calculated over all documents except s_d .

For ρ : We update each ρ_k from its posterior distribution:

$$q(\rho_k|\mathcal{D}, \mathcal{M}_{-\rho_k}) = \text{GMM}_0\left(\rho_k; \frac{\sum_{d=1}^D v_{d,k} + v_{k,0}\nu_0}{D + \nu_0}, D + \nu_0\right),$$

Since the normalization constant is unknown, it is intractable to sample directly from GMM_0 . Fortunately, slice sampling (Neal 2003) can be used to solve this problem.

For \mathbf{v} : For the inversion representation v_d of document s_d , each $v_{d,k}$ can be resampled independently from its posterior distribution:

$$q(v_{d,k} = v|\mathcal{D}, \mathcal{M}_{-v_{d,k}}) \propto p_0(v_{d,k} = v)p(s_d|\mathcal{M}, \mathcal{D}_{-s_d}) \propto p_0(v_{d,k} = v) \prod_{k=1}^K \frac{\Gamma(f_{0,k,\cdot}^{-d} + V\alpha_0)}{\Gamma(f_{0,k,\cdot} + V\alpha_0)} \prod_{v=1}^V \frac{\Gamma(f_{0,k,v} + \alpha_0)}{\Gamma(f_{0,k,v}^{-d} + \alpha_0)},$$

where $p_0(v_{d,k} = v) = \text{GMM}_k(v_{d,k} = v; \rho_k)$ is the prior.

For \mathbf{t}, \mathbf{b} : Since t_{dsm} has meaningful value only when $b_{dsm} = 1$, we jointly sample b_{dsm} and t_{dsm} for the topics. The joint distributions are

$$q(b_{dsm} = 1, t_{dsm} = t|\mathcal{D}, \mathcal{M}_{-\{b_{dsm}, t_{dsm}\}}) \propto (f_{b=1}^{-dsm} + \gamma_0) \frac{f_{1,t,v}^{-dsm} + \beta_0}{f_{1,t,\cdot}^{-dsm} + V\beta_0} \frac{f_{1,d,t}^{-dsm} + \theta_0}{f_{1,d,\cdot}^{-dsm} + T\theta_0},$$

$$q(b_{dsm} = 0, t_{dsm} = \emptyset|\mathcal{D}, \mathcal{M}_{-\{b_{dsm}, t_{dsm}\}}) \propto (f_{b=0}^{-dsm} + \gamma_0) \frac{f_{0,z_{ds},v}^{-dsm} + \alpha_0}{f_{0,z_{ds},\cdot}^{-dsm} + V\alpha_0},$$

where $f_{b=1}$ is the number of topic words in the corpus and $f_{b=0}$ is the number of intent words. $f_{1,t,v}$ counts all the times that w_v appears in the corpus with indicator variable value 1 and topic label t . Then, $f_{1,t,\cdot} = \sum_{v=1}^V f_{1,t,v}$. $f_{1,d,t}$ is the number of words in document s_d with indicator variable value 1 and topic label t . Then $f_{1,d,\cdot} = \sum_{t=1}^T f_{1,d,t}$. The superscript “- dsm ” indicates that the frequency is calculated except w_{dsm} . According to the joint distribution of b_{dsm} and t_{dsm} , we compute the $T + 1$ non-zero probabilities, then do normalization and sample from the resulting multinomial.

¹ w_v is different from w_{dsm} . w_{dsm} is a word in sentence w_{ds} .




Figure 3: The graphical expression of computing π_0 .

Supervised GMM-LDA

GMM-LDA is unsupervised; it learns the document structure without human annotation. As suggested by existing supervised topic models (McAuliffe and Blei 2008; Wang, Blei, and Li 2009; Zhu, Ahmed, and Xing 2012), the predictive power can be greatly improved with a small amount of labeled documents. Here, we present a supervised GMM-LDA which combines the known intent labels of sentences during learning.

We consider the setting where a part of the documents in the corpus are labeled, that is, each sentence is assigned an intent label. Our goal is to develop a supervised model to learn the intent structures for the remaining unlabeled documents. The simplest way to leverage the label information is to use the labels directly during learning instead of sampling them. However, the GMM describes the intent order structure in a global way, which makes the process more complicated. In the unsupervised case, the canonical intent permutation is the identity ordering $(1, 2, \dots, K)$ as the intent numbers are completely symmetric and not linked to any meaningful intent label. However, the true intent labels are already known in the supervised case. Then a challenging problem is to determine the canonical permutation π_0 . Moreover, for the part of labeled documents, we have the true intent structure z_d , which is unnecessary to draw from the GMM-Multi prior. So the next challenging problem is how to leverage the known intent structure z_d to help for learning. That is, how to obtain u_d and π_d using z_d . Below, we discuss in detail on how to solve these problems.

For π_0 : As shown in Fig.3, we present an approximate three-step algorithm to obtain the canonical permutation π_0 . **Step 1:** We compute π'_d for each labeled document s_d , where π'_d is the permutation of the $K_d (\leq K)$ intent labels appearing in document s_d . We arrange the K_d intent labels in the order they appear in z_d . However, there exist cases where the same intent label appears in disconnected portions of z_d , which are rare in practice. When encountering these cases, we take the position that the label appears most consecutively in the sequence. If there are two or more of these occurrences, we take the first position (see Fig. 3 for some examples). **Step 2:** We introduce variables $g_{ij}(i, j \in [K])$, where g_{ij} counts the times of the intent label i appearing before j in all π'_d . Then we define a directed graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$, where $\mathcal{V} \in [K]$ is the set of nodes and $\mathcal{E} = \{(i, j) | g_{ij} >= g_{ji}; i, j \in \mathcal{V}\}$ is the set of edges. **Step 3:**

We obtain the π_0 by calculating the topological sequence of \mathcal{G} . If there are circles in \mathcal{G} , we randomly break them. If there are multiple topological sequences, we randomly take one. In our experiments, the real situations are always in accordance with our intuition that only one topological sequence can be obtained from \mathcal{G} .

For u_d and π_d : To combine the label information, we compute u_d and π_d using z_d and π_0 . u_d can be obtained by directly putting all the ordered elements of z_d into a bag. π_d is a permutation of K numbers, which can be obtained by inserting the remaining $K - K_d$ numbers into π'_d . We want π_d to be as close to π_0 as possible, which is consistent with the idea of GMM. Specially, $d(\pi_d, \pi_0)$ is the distance from π_d to π_0 , which defined as the minimum number of swaps of adjacent elements needed to transform π_d into the same order of π_0 . Inspired by the idea of greedy method, we insert the $K - K_d$ numbers into π'_d one by one and each step need to minimize the distance. During supervised learning, when it comes to a labeled document s_d , we directly take the values of π_d and u_d instead of updating them.

GMM-LDA with Entropic Regularization

GMM-LDA jointly models the two incompatible structures of documents by using a binary variable to indicate the type (intent or topic) of each word. It can happen that a same word located in two different positions are assigned with different types, which is somewhat unreasonable. Most of the time, the type of a word can be decided regardless of which position it appears in. For instance, the words “propose” and “experiment” are more likely to be intent words regardless of their positions. In order to model the significant divergence between topics and intents, we introduce the entropy of the words to make our model more descriptive. As we know, in information theory, the entropy of a discrete random variable $X = \{x_1, x_2, \dots, x_n\}$ can explicitly be written as $H(X) = -\sum_{i=1}^n p(x_i) \log p(x_i)$. Similarly, the entropy of a word w_v in the vocabulary can be formulated as:

$$H(b_v) = -\sum_{i=0,1} p(b_v = i) \log p(b_v = i), \quad (4)$$

where $p(b_v = i)$ denotes the probability that word w_v appears as an intent word ($i = 0$) or an topic word ($i = 1$) in the corpus. Lower entropic value means better separation.

As GMM-LDA is under Bayesian inference framework, it is challengeable to incorporate the entropic knowledge. Nevertheless, regularized Bayesian inference framework (Zhu, Chen, and Xing 2014) provides us an alternative interpretation of Bayesian inference (Williams 1980) and can combine domain knowledge flexibly. Specifically, GMM-LDA with entropic regularization can be formulated as:

$$\min_{q(\mathcal{M}) \in \mathcal{P}} \text{KL}(q(\mathcal{M}) || q(\mathcal{M} | \mathcal{D})) + c \sum_{v=1}^V H(b_v), \quad (5)$$

where \mathcal{P} is the space of probability distributions and $c \geq 0$ is a regularization parameter. When $c = 0$, the optimal solution of the Kullback-Leibler divergence $\text{KL}(q || p)$ is $q(\mathcal{M}) = q(\mathcal{M} | \mathcal{D})$, the standard Bayesian posterior distribution as in

Table 1: Results of unsupervised clustering.

Models		Chemical				Elements			
		ARI	Recall	Prec	Fscore	ARI	Recall	Prec	Fscore
$K = 5$	K-means	0.0023	0.7139	0.4257	0.5333	0.0236	0.8107	0.3201	0.4583
	Boilerplate-LDA	0.3953	0.6797	0.6647	0.6720	0.3060	0.6975	0.5091	0.5884
	GMM	0.1942	0.5093	0.5982	0.5501	0.3395	0.7004	0.5348	0.6064
	GMM-LDA (Uniform)	0.2614	0.6886	0.5836	0.6314	0.1911	0.6002	0.4449	0.5106
	GMM-LDA	0.4055	0.7473	0.6699	0.7065	0.3223	0.7176	0.5271	0.6075
	EGMM-LDA	0.4125	0.7553	0.6725	0.7115	0.3567	0.7450	0.5517	0.6339
$K = 10$	K-means	0.0011	0.5779	0.4332	0.4949	0.0504	0.6895	0.3460	0.4606
	Boilerplate-LDA	0.2861	0.5367	0.6955	0.6056	0.3511	0.5885	0.6440	0.6149
	GMM	0.1381	0.3302	0.6268	0.4324	0.3652	0.5692	0.6491	0.6063
	GMM-LDA (Uniform)	0.3206	0.6069	0.6227	0.6138	0.3432	0.5720	0.6146	0.5924
	GMM-LDA	0.4477	0.6747	0.7311	0.7017	0.4013	0.6273	0.6857	0.6549
	EGMM-LDA	0.4546	0.6786	0.7364	0.7063	0.4196	0.6418	0.6865	0.6631

Table 2: Statistics of datasets.

Corpus	Docs	Sentences	Vocab	Tokens
Chemical	965	9,488	4,981	123,119
Elements	118	1,848	5,301	67,484

Eq. (3). When $c > 0$, the entropic knowledge is imposed as a regularization constraint. With mean-field assumption, the inference of Eq. (5) is similar to that of GMM-LDA. Moreover, the entropic regularization is only relevant to b , which can be combined with both the unsupervised GMM-LDA and the supervised GMM-LDA.

Experimental Results

To demonstrate the efficacy of our models, we evaluate the performance on two tasks: unsupervised clustering and supervised classification. We use two real datasets: 1) **Chemical** (Guo et al. 2010): It contains 965 abstracts of scientific papers about 5 kinds of chemicals, and each abstract focuses on one of the 5 topics. Each sentence is annotated with one of the 7 intent labels: *Background*, *Objective*, *Related Work*, *Method*, *Result*, *Conclusion* and *Future Work*; and 2) **Elements** (Chen et al. 2009): It consists of 118 articles from the English Wikipedia, and each article talks about one of the 118 chemical elements in the periodic table. Each paragraph is annotated with an intent label. We take the 8 most frequently occurring intent labels: *Top-level Segment*, *History*, *Isotopes*, *Applications*, *Occurrence*, *Notable Characteristics*, *Precautions* and *Compounds*, and filter out paragraphs with other labels. Although the intent structure is paragraph-level in Element, while it is sentence-level in Chemical, the word ‘‘sentence’’ is used throughout the paper for simplicity. Tab. 2 summarizes the dataset statistics. Although both datasets are in chemistry domain, they have different characteristics that can be observed from the experimental results. In Chemical, the intent orders are relatively fixed due to the writing conventions of scientific papers, while they are more variable in Elements.

Data preprocessing involves removing a small set of stop words, tokens containing non-alphabetic characters, tokens appearing less than 3 times, tokens of length one and sentences with less than 5 valid tokens. We report the average results over 5 runs, while each run takes a sufficiently large

number of iterations (e.g. 2000) to converge. Statistical significance is measured with *t-test*.

Unsupervised Clustering

Our goal of unsupervised clustering is to learn an intent label for each sentence in the corpus without any true label information. Adjusted Rand Index (ARI) (Vinh, Epps, and Bailey 2010), recall, precision and F-score are used as our evaluation measures. F-score is the harmonic mean of recall and precision. For all the four measures, higher scores are better.

We consider two variants of our model: 1) **GMM-LDA**: Our unsupervised model; and 2) **EGMM-LDA**: GMM-LDA with entropic regularization. For hyperparameters, we set $\theta_0 = 0.1$, $\lambda_0 = 0.1$, $\alpha_0 = 0.1$, $\beta_0 = 0.1$ and $\gamma_0 = 1$, since we find that the results are insensitive to them. ν_0 is set to be 0.1 times the number of documents in the corpus. For EGMM-LDA, we set the regularization parameter c to be 0.1. The baseline methods we use are: 1) **K-means**: The feature used for each sentence is the bag of words. 2) **Boilerplate-LDA**: The model presented in (Ó Séaghdha and Teufel 2014). 3) **GMM**: The intent part of GMM-LDA, which is the content model by (Chen et al. 2009) and can be implemented by fixing all indicator variables b_{dsm} to 0 during learning; and 4) **GMM-LDA (Uniform)**: The model assumes a uniform distribution over all intent permutations and can be implemented by fixing ρ to zero.

Clustering performance: We set the number of topics $T = 10$ for Chemical and $T = 5$ for Elements. Tab. 1 shows the results. For Chemical, our two models outperform all other methods on four measures, with p -values smaller than 0.001 except for Boilerplate-LDA with $K = 5$. The exception of Boilerplate-LDA may be caused by the large variance in the results of multiple runs when K is small. Since the first-order Markov chain is used in Boilerplate-LDA for order structure learning, which only has a local view and is more susceptible to noise. Moreover, first-order Markov chain would select the same intent label for disconnected sentences within a document, which is against our intuition. Our models overcome these problems by using GMM and also achieve better performance. The simpler variants of our models achieve reasonable performance. GMM underperforms GMM-LDA, indicating that modeling topics and intents simultaneously provides a richer and more effective way to document structure learning. The bad performance




Figure 4: Sensitivity of ρ_0 .

of the uniform variant proves the indispensability to model the intent order structure. Moreover, our model yields better results with entropic regularization. As to Elements, the results are similar to that of Chemical. However, we can observe that GMM performs competitively in Elements, which shows the characteristic of this dataset that there is no obvious topic structure. We can also observe that the highest recall scores are obtained by K-means with very low precision scores, since K-means prefers to assign the same label to most of the sentences. It can thus be seen that the task is difficult and the richer models are required.

Hyperparameter ρ_0 : ρ_0 controls the variability of the order structure and would be set according to different datasets. The model with large ρ_0 assigns massive probabilities around the canonical permutation and the order structure is relatively fixed, while the model with small ρ_0 relaxes the constraints and has a variety of orders far from the canonical one. For Elements with $K = 10$, we change ρ_0 from 0.125 to 4 and report the F-scores of GMM-LDA in Fig. 4. It can be observed that the performance is stable in a wide range (e.g. $0.5 < \rho_0 < 2$). We set $\rho_0 = 2$ for all the experiments except for Elements with $K = 10$, in which we set $\rho_0 = 1$.

Table 3: Most commonly words of two types on Chemical.

Intent Words			Topic Words		
days	results	increased	rats	protein	expression
other	control	observed	mice	levels	exposure
used	present	significant	DNA	activity	adducts
more	treated	between	gene	effects	treatment
data	suggest	compared	cells	tumor	phenobarbital
study	human	showed	liver	induced	formation

Types of words: To embody the intents and topics in the results, we assume that each word in the vocabulary is either an intent word or a topic word. If a word appears in the corpus more as an intent word than a topic word, we classify it as an intent word; otherwise, it is a topic word. Note that this additional condition is introduced only for the ease of demonstration. In order to see how our models separate these two types, we list 18 most commonly used words of each type in Chemical according to the result of EGMM-LDA. As shown in Tab. 3, we can observe that almost all the intent words has rhetorical functions that can express the intents of sentences, while almost all the topic words are about chemical topics. The good separations justify our assumption and show the effectiveness of our models.

Table 4: Results of supervised classification.

Models	Chemical		Elements	
	ARI	ACC	ARI	ACC
SVM	0.399	0.674	0.326	0.622
sBoilerplate-LDA	0.506	0.709	0.522	0.749
sGMM	0.276	0.530	0.531	0.744
sGMM-LDA	0.510	0.731	0.521	0.752
sEGMM-LDA	0.510	0.730	0.562	0.775

Supervised Classification

Now, we evaluate our supervised models for classifying sentences. For each dataset, we randomly choose 20% documents; annotate their sentences with intent labels; and use them for training. Our goal is to learn the intent labels for the sentences in the remaining 80% documents. We report accuracy (ACC) and the ARI scores to show the improvements compared to the unsupervised learning. We again consider two variants of our model: 1) **sGMM-LDA**: Our supervised model; and 2) **sEGMM-LDA**: sGMM-LDA with entropic regularization. The baseline methods are: 1) **SVM**: We use the bag-of-words features, linear kernel and SVM-Light tools (Joachims 1998). 2) **sBoilerplate-LDA**: The supervised version of Boilerplate-LDA, in which we fix the known labels during learning instead of updating them; and 3) **sGMM**: The intent part of our sGMM-LDA. All the settings are the same as that in the unsupervised learning.

Classification performance: Tab. 4 presents the results. For Chemical, the best accuracies are achieved by our two models ($p < 0.01$), which again proves that our assumption of the two types of words is reasonable and the intent order structure can be better modeled by employing GMM. For Elements, sBoilerplate-LDA and sGMM perform competitively, while EGMM-LDA beats all the other methods ($p < 0.05$). It shows that our model is more robust with entropic regularization. The ARI scores improve a lot compared to that in Tab. 1, indicating that the predictive power can be largely improved with just 20% labeled documents.

Table 5: Results of the canonical intent permutation and the intent words on Chemical with sEGMM-LDA.

No.	Intent Label	High-Frequency Words
0	<i>Background</i>	cancer, studies, carcinogen, environmental, aromatic, known, shown, however, used study, using, investigated, present, examined, used, whether, determine, effect, evaluated
1	<i>Objective</i>	days, treated, diet, single, control, groups, body, followed, water, using, injection
2	<i>Method</i>	increased, significantly, observed, compared, showed, higher, however, respectively
3	<i>Result</i>	polymerase, significantly, results, codes, ICD, comparisons, percent, suggest, enhances
4	<i>Related Work</i>	results, suggest, indicate, findings, study, role, studies, thus, conclusion, important
5	<i>Conclusion</i>	needed, however, needs, required, future, research, investigations, studies, observed
6	<i>Future Work</i>	

Intent words: We can obtain the canonical intent permutation by the known intent labels, at the same time the distribution over vocabulary can be learned. Tab. 5 shows the results on Chemical with sEGMM-LDA. We can observe that the canonical intent order (numbered from 0 to

6) conforms to the convention in scientific writing. Moreover, from the high-frequency words of each intent, we can see that most of the words express the intent labels well. For instance, “study” and “investigated” express the intent *Objective*, while “increased” and “significantly” are for *Result*.

Related Work

From the algorithmic perspective, our work is grounded in topic models, such as Latent Dirichlet Allocation (LDA) (Blei, Ng, and Jordan 2003), which have been widely developed for many NLP tasks. Instead of representing documents as bags of words, many expanded models take specific structural constraints into consideration (Purver et al. 2006; Gruber, Weiss, and Rosen-Zvi 2007). Among different models, our work has a closer relation to the models with order structure. For order modeling, Markov chain can only capture the dependence locally (Ó Séaghdha and Teufel 2014; Barzilay and Lee 2004; Elsner, Austerweil, and Charniak 2007), while the generalized Mallows model (GMM) (Fligner and Verducci 1986) has a global view (Chen et al. 2009; Du, Pate, and Johnson 2015; Cheng, Hühn, and Hüllermeier 2009). A more complete model can be obtained by dividing the words into different types. In early trial of zoneLDA (Varga, Preotiuc-Pietro, and Ciravegna 2012), a type of words are for describing background, which are independent of the category of the sentence. Boilerplate-LDA (Ó Séaghdha and Teufel 2014) also considers two types: document-specific topic words and rhetorical words. Three types of words are learnt by a rule-based method in (Nguyen and Shirai 2015). However, global order structure is not considered in these models. Therefore, jointly modeling topics and intents with global order structure is of great value.

Conclusion and Future Work

We present GMM-LDA (both unsupervised and supervised) for document structure learning, which simultaneously model topics and intents. The generalized Mallows model is employed to model the intent order globally. Moreover, we consider the entropic regularization to make the model more descriptive. Our results demonstrate the reasonability of our intuitions and the effectiveness of our models. For future work, we are interested in making our models richer by combining local coherence constraints.

Acknowledgments

This work was mainly done when the first author was an intern at Microsoft Research Asia. The work was supported by the National Basic Research Program (973 Program) of China (Nos. 2013CB329403, 2012CB316301), National NSF of China (Nos. 61322308, 61332007), TNList Big Data Initiative, and Tsinghua Initiative Scientific Research Program (Nos. 20121088071, 20141080934).

References

Barzilay, R., and Lee, L. 2004. Catching the drift: Probabilistic content models, with applications to generation and summarization. In *NAACL-HLT*.

- Blei, D.; Ng, A.; and Jordan, M. 2003. Latent dirichlet allocation. *JMLR* 3:993–1022.
- Chen, H.; Branavan, S.; Barzilay, R.; and Karger, D. 2009. Content modeling using latent permutations. *JAIR* 36:129–163.
- Cheng, W.; Hühn, J.; and Hüllermeier, E. 2009. Decision tree and instance-based learning for label ranking. In *ICML*.
- Du, L.; Pate, J.; and Johnson, M. 2015. Topic segmentation with an ordering-based topic model. In *AAAI*.
- Elsner, M.; Austerweil, J.; and Charniak, E. 2007. A unified local and global model for discourse coherence. In *NAACL-HLT*.
- Fligner, M., and Verducci, J. 1986. Distance based ranking models. *JRSS. Series B (Methodological)* 359–369.
- Frermann, L.; Titov, I.; and Pinkal, M. 2014. A hierarchical bayesian model for unsupervised induction of script knowledge. In *EACL*.
- Gruber, A.; Weiss, Y.; and Rosen-Zvi, M. 2007. Hidden topic markov models. In *AISTATS*.
- Guo, Y.; Korhonen, A.; Liakata, M.; Karolinska, I.; Sun, L.; and Stenius, U. 2010. Identifying the information structure of scientific abstracts: an investigation of three different schemes. In *Workshop on Biomedical Natural Language Processing*.
- Halliday, M., and Hasan, R. 1976. Cohesion in english. *London; Martin JR*.
- Joachims, T. 1998. Making large-scale svm learning practical. Technical report, Universität Dortmund.
- Louis, A.; Joshi, A.; and Nenkova, A. 2010. Discourse indicators for content selection in summarization. In *Special Interest Group on Discourse and Dialogue*, 147–156.
- Mann, W., and Thompson, S. 1988. Rhetorical structure theory: Toward a functional theory of text organization. *Text-Interdisciplinary Journal for the Study of Discourse* 8(3):243–281.
- McAuliffe, J., and Blei, D. 2008. Supervised topic models. In *NIPS*.
- Neal, R. 2003. Slice sampling. *Annals of Statistics* 705–741.
- Nguyen, T., and Shirai, K. 2015. Topic modeling based sentiment analysis on social media for stock market prediction. In *ACL*.
- Ó Séaghdha, D., and Teufel, S. 2014. Unsupervised learning of rhetorical structure with un-topic models. In *COLING*.
- Prasad, R.; Joshi, A.; Dinesh, N.; Lee, A.; Miltsakaki, E.; and Webber, B. 2005. The penn discourse treebank as a resource for natural language generation. In *CL Workshop on Using Corpora for NLG*.
- Purver, M.; Griffiths, T.; Körding, K.; and Tenenbaum, J. 2006. Unsupervised topic modelling for multi-party spoken discourse. In *COLING-ACL*, 17–24.
- Varga, A.; Preotiuc-Pietro, D.; and Ciravegna, F. 2012. Unsupervised document zone identification using probabilistic graphical models. In *LREC*.
- Vinh, N.; Epps, J.; and Bailey, J. 2010. Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance. *JMLR* 11:2837–2854.
- Wang, C.; Blei, D.; and Li, F. 2009. Simultaneous image classification and annotation. In *CVPR*.
- Williams, P. 1980. Bayesian conditionalisation and the principle of minimum information. *BJPS* 131–144.
- Zhu, J.; Ahmed, A.; and Xing, E. 2012. Medlda: Maximum margin supervised topic models. *JMLR* 13(1):2237–2278.
- Zhu, J.; Chen, N.; and Xing, E. 2014. Bayesian inference with posterior regularization and applications to infinite latent svms. *JMLR* 15(1):1799–1847.