

Temporal Topic Analysis with Endogenous and Exogenous Processes

Abstract

We consider the problem of modeling temporal textual data taking endogenous and exogenous processes into account. Such text documents arise in real world applications, including job advertisements and economic news articles, which are influenced by the fluctuations of the general economy. We propose a hierarchical Bayesian topic model which imposes a "group-correlated" hierarchical structure on the evolution of topics over time incorporating both processes, and show that this model can be estimated from Markov chain Monte Carlo sampling methods. We further demonstrate that this model captures the intrinsic relationships between the topic distribution and the time-dependent factors, and compare its performance with latent Dirichlet allocation (LDA) and the structural topic model (STM). The model is applied to two collections of documents to illustrate its empirical performance: online job advertisements from DirectEmployers Association and journalists' postings on BusinessInsider.com.

Content areas: topic modeling; hierarchical models; temporal models ID: 1308

1 Introduction

Many organizations nowadays provide portals for job posting and job search, such as [glassdoor.com](#) from Glassdoor, [indeed.com](#) from Recruit, and [my.jobs](#) from DirectEmployers Association. Our work is inspired by data collected from the portal [my.jobs](#), a website where job seekers can apply to the posted job openings through a provided link. The data collected from the website includes user clickstreams (users create accounts on the site) and attributes of job advertisements, such as their description, location, company name, and posted date.

In this paper, we investigate the relationship between economic fluctuations and the related changes in job advertisements, which can reveal the economic conditions of different time periods. More generally, this question is about the influence of any exogenous process on textual data with temporal dimensions. We adopt the perspective that the documents are organized into a certain number of topics, and study the impact of the exogenous process on the topic distribution, i.e. the relative topic proportions. Given a corpus of text documents with time stamps and a related exogenous process, the problem is to find a relationship between the topics discussed and the exogenous process. This setting is natural

in an economic context; for instance, changes in macroeconomic indicators have an impact on government reports and *Wall Street Journal* news articles. Meanwhile, we also notice that for most temporal documents, the topic proportions change over time, which indicates an endogenous process of topic evolution.

With the goal of establishing topic dependency on the endogenous and exogenous processes, LDA-type topic models are especially suitable. The latent Dirichlet allocation (LDA) (Blei, Ng, and Jordan 2003) is the original model. Since then, a large number of variants have been proposed, many of which can be found in Blei (2011).

Meanwhile, there has been relatively limited discussion on modeling time-dependent documents when there are relevant simultaneous exogenous processes. Many time-dependent topic models without the exogenous component have been proposed, such as the Topics over Time (ToT) model (Wang and McCallum 2006) and the dynamic topic model (DTM) (Blei and Lafferty 2006), to name a few. However, to the best of our knowledge, none of these papers incorporate the effect of exogenous processes. On the other hand, the structural topic model (STM) (Roberts, Stewart, and Airoldi 2015) considers the effect of metadata, i.e. the attributes specified for each document, on the topic distribution. While STM can be applied for mining time-dependent textual data with exogenous covariates, it does not explicitly consider the time factor or the endogenous topic evolution processes of time-stamped documents.

Our approach to this problem is to incorporate both endogenous and exogenous processes into a topic model. For the endogenous part of our paper, we impose a Markovian structure on the topic distribution over time, similar to Blei and Lafferty (2006) and Dubey et al. (2014). For the exogenous process, we incorporate it into the topic distribution in each period, adjusting the endogenous topic evolution process. In this way, our model is essentially a stick-breaking truncation of a "group-correlated" hierarchical Dirichlet process. Our model has the following contributions: (i) it addresses the question of measuring the influence of exogenous processes on the topics in related documents, (ii) it incorporates both endogenous and exogenous aspects, and (iii) it demonstrates that text mining can also have useful implications in the realm of economics, which, from the authors' perspective, is a relatively new finding.

Section 2 offers a brief review on the topic modeling techniques related to our model. Section 3 develops our hierarchical Bayesian model and describes how to make posterior inferences with a variant of the Markov chain Monte Carlo (MCMC) technique. Section 4 studies the online job advertisements from DirectEmployers Association and journalists' postings in finance on BusinessInsider.com with our

proposed method, providing a comparison of performance with the standard LDA and STM. Section 5 suggests possible directions for the future and concludes the paper.

2 Review of Time-Dependent Topic Modeling

We first introduce the standard model of LDA (Blei, Ng, and Jordan 2003). Suppose that there is a collection of documents d_i , $i = 1, \dots, N$ and words $\{x_{i,j}\}_{j=1}^{J_i}$ within each document d_i indexed by a common dictionary containing V words, where N is the number of documents, and J_i is the number of words in d_i . The LDA model is as follows,

$$\begin{cases} \theta_i \stackrel{iid}{\sim} Dir(\alpha), \phi_k \stackrel{iid}{\sim} Dir(\beta), \\ z_{i,j} | \theta_i \stackrel{iid}{\sim} Cat(\theta_i), x_{i,j} | z_{i,j} \sim Cat(\phi_{z_{i,j}}). \end{cases} \quad (1)$$

Here $i = 1, \dots, N$, $j = 1, \dots, J_i$, $k = 1, \dots, K$; θ_i is the length- K per-document topic distribution for d_i , ϕ_k is the length- V per-topic word distribution for the k -th topic, z_i^j is the topic for the j -th word in d_i , and K is the number of topics. $Dir(\cdot)$ denotes the Dirichlet distribution and $Cat(\cdot)$ denotes the categorical distribution, a special case of the multinomial distribution when $n_{obs} = 1$.

The Dirichlet process is a class of randomized probability measures and can be applied for non-parametric modeling of mixture models. Denoting the concentration parameter by α and the mean probability measure by H , a realization G from the Dirichlet process can be written as $G \sim DP(\gamma, H)$. With the stick-breaking notation (Sethuraman 1994), we have

$$G = \sum_{k=1}^{\infty} b_k \delta_{\varphi_k}, \quad (2)$$

where δ_{φ_k} is a ‘‘delta’’ probability measure with all the probability mass placed at φ_k , $\varphi_k \stackrel{iid}{\sim} H$, $b_k = b'_k \prod_{i=1}^{k-1} (1 - b'_i)$, $b'_k \stackrel{iid}{\sim} Beta(1, \gamma)$, $k = 1, 2, \dots$. We write $\mathbf{b} = (b_1, b_2, \dots) \sim Stick(\gamma)$. More properties of the Dirichlet process can be found in Ferguson (1973).

A hierarchical Dirichlet process (HDP) was proposed in the context of text modeling by Teh et al. (2005). The following hierarchical structure is assumed,

$$\begin{cases} G_0 | \gamma \sim DP(\gamma, H), \\ G_1, \dots, G_N | (\alpha, G) \stackrel{iid}{\sim} DP(\alpha, G), \\ \phi_{i,j} | G_i \stackrel{iid}{\sim} G_i, x_{i,j} \sim Cat(\phi_{i,j}). \end{cases} \quad (3)$$

Here $i = 1, \dots, N$, $j = 1, \dots, J_i$. The V -dimensional random vectors G_0, G_1, \dots, G_N are ‘‘random word distributions,’’ each of which is a draw from a Dirichlet process in (3). Moreover, each draw from a random word distribution is a V -dimensional fixed vector $\phi_{i,j}$; it is the word distribution for $x_{i,j}$. The posterior inference can be achieved by different strategies of Gibbs sampling.

There are mainly two approaches in the literature of measuring endogenous topic evolution processes. One approach is to impose a finite mixture structure on the topic distribution: a dynamic hierarchical Dirichlet process (dHDP) (Ren,

Dunson, and Carin 2008) was proposed by adding a temporal dimension, and its variation was further applied on topic modeling with a stick-breaking truncation of Dirichlet processes (Pruteanu-Malinici et al. 2010). The other approach imposes a Markovian structure. For instance, the dynamic topic model (DTM) (Blei and Lafferty 2006) is as follows,

$$\begin{cases} \phi_{t,k} | \phi_{t-1,k} \sim N(\phi_{t-1,k}, \sigma^2 I), \\ \alpha_t | \alpha_{t-1} \sim N(\alpha_{t-1}, \delta^2 I), \\ \theta_{t,i} | \alpha_t \stackrel{iid}{\sim} N(\alpha_t, a^2 I), \\ z_{t,i,j} | \theta_{t,i} \stackrel{iid}{\sim} Cat(\exp(\theta_{t,i})), \\ x_{t,i,j} | z_{t,i,j} \sim Cat(\exp(\phi_{t,z_{t,i,j}})). \end{cases} \quad (4)$$

Here $t = 1, \dots, T$, $i = 1, \dots, N_t$, $j = 1, \dots, J_{t,i}$, $k = 1, \dots, K$ ($t \geq 2$ for the first two equations); T is the number of time periods, N_t is the number of documents in the t -th period, and $J_{t,i}$ is the number of words in the i -th document in the t -th period; the rest are similarly defined as in LDA. One major difference between DTM and LDA is that the topic distributions $\theta_{t,i}$ and word distributions $\phi_{t,k}$ are in log-scale in DTM. A variational Kalman filtering was proposed for the posterior inference. As this Markovian approach is simpler for both interpretation and posterior inference, we apply a more generalized version of it to specify the endogenous process in our model.

The structural topic model (STM) (Roberts, Stewart, and Airoldi 2015) measures the effect of metadata of each document with the logistic normal distribution. Their model for each document d_i is as follows,

$$\begin{cases} \theta_i | (X_i \gamma, \Sigma) \sim LogisticNormal(X_i \gamma, \Sigma), \\ p(\phi_{i,k}) \propto \exp(m + \kappa_k + \kappa_{g_i} + \kappa_{kg_i}), \\ z_{i,j} | \theta_i \stackrel{iid}{\sim} Cat(\theta_i), x_{i,j} | z_{i,j} \sim Cat(\phi_{i,z_{i,j}}), \end{cases} \quad (5)$$

where $i = 1, \dots, N$, $j = 1, \dots, J_i$; X_i is the metadata matrix, γ is a coefficient vector, Σ is the covariance matrix, $\phi_{i,k}$ is the word distribution for d_i and the k -th topic, m is a baseline log-word distribution, κ_k , and κ_{g_i} and κ_{kg_i} are the topic, group, and interaction effects; the rest are defined similarly to LDA. This model explicitly considers exogenous factors, and can be applied to find the relationship between topic distributions and exogenous processes. Below we adopt a slightly more general approach, incorporating both endogenous and exogenous factors.

3 Model and Algorithm

3.1 Motivation: A Group-Correlated Hierarchical Dirichlet Process

We formulate our problem as in DTM: we are given time periods $t = 1, \dots, T$, documents from each period $d_{t,i}$, $i = 1, \dots, N_t$, $t = 1, \dots, T$, and the indices of words $\{x_{t,i,j}\}_{j=1}^{J_{t,i}}$ within each document $d_{t,i}$ from the first word to the last. The words are indexed by a dictionary containing V words in total. We begin with a hierarchical Dirichlet process in time 1: let $G_1 | \gamma \sim DP(\gamma, H)$, $G_1 | (\alpha_1, G_1) \sim$

$DP(\alpha_1, G_1)$, where G_1 is a baseline random word distribution for time 1, and G_{1i} is the random word distribution for document d_{1i} . For G_2, \dots, G_T , we have the following Markovian structure,

$$p(G_t|G_{t-1}) \propto \exp[-d(G_t, G_{t-1})], \quad t = 2, \dots, T. \quad (6)$$

Here $d(\cdot, \cdot)$ is some distance between two probability measures. This completes our endogenous process. To take an exogenous process $\{\mathbf{y}_t\}_{t=1}^T$ into account, we assume the following

$$\tilde{G}_t = \mathcal{M}(G_t, \mathbf{y}_t), \quad t = 1, \dots, T, \quad (7)$$

where \mathcal{M} maps the endogenous baseline random word distribution G_t to the realized baseline random word distribution \tilde{G}_t for time t , considering the influence of $\{\mathbf{y}_t\}_{t=1}^T$. Therefore, we further assume that each per-document random word distribution $G_{t,i}$ is sampled with mean \tilde{G}_t rather than G_t . The final model is as follows,

$$\begin{cases} G_1 | \gamma \sim DP(\gamma, H), \\ p(G_t | G_{t-1}) \propto \exp[-d(G_t, G_{t-1})], \\ \tilde{G}_t = \mathcal{M}(G_t, \mathbf{y}_t), \\ G_{t,i} | (\alpha_t, \tilde{G}_t) \stackrel{iid}{\sim} DP(\alpha_t, G_t), \\ \phi_{t,i,j} | G_{t,i} \stackrel{iid}{\sim} G_{t,i}, \quad x_{t,i,j} \sim Cat(\phi_{t,i,j}). \end{cases} \quad (8)$$

Here $t = 1, \dots, T, i = 1, \dots, N_t, j = 1, \dots, J_{t,i}$ ($t \geq 2$ for the first line). Throughout this paper, our model is fully conditional on $\{\mathbf{y}_t\}_{t=1}^T$, i.e. we assume $\{\mathbf{y}_t\}_{t=1}^T$ to be fixed; this has an intuitive explanation, as our temporal documents represent a very small portion of the underlying environment, i.e. the exogenous process, so their influence on $\{\mathbf{y}_t\}_{t=1}^T$ is almost negligible.

3.2 A Group-Correlated Temporal Topic Model: Some Simplifications

Below we consider a stick-breaking truncation of the model above, since posterior inference of the exact model can be intricate. With the stick-breaking expression of G_1 in (8), we have

$$\begin{cases} \phi_1, \phi_2, \dots, \stackrel{iid}{\sim} H, \\ \boldsymbol{\pi}_1 = (\pi_{11}, \pi_{12}, \dots) \sim Stick(\gamma), \\ G_1 = \sum_{k=1}^{\infty} \pi_{1k} \delta_{\phi_k}. \end{cases} \quad (9)$$

Here we set $d(\cdot, \cdot) = +\infty$ if the two probability measures have different supports; this necessitates that all periods share the same topics. Our intent is that the topics should remain the same to investigate their relationships with endogenous and exogenous processes; otherwise, changes in topics can blur the relationships and possibly result in overfitting. We apply the total variation distance $d(p, q) = \lambda \cdot \int |p - q| d\mu$ with $\lambda > 0$, although many others can also be applied and lead to, for instance, a log-normal model in DTM, or a normal model (Dubey et al. 2014; Zhang, Kim, and Xing 2015). We have the following,

$$\begin{cases} G_t = \sum_{k=1}^{\infty} \pi_{tk} \delta_{\phi_k}, \\ \pi_{tk} = \pi_{t-1k} + Lap(\lambda), \\ \boldsymbol{\pi}_t = (\pi_{t1}, \pi_{t2}, \dots). \end{cases} \quad (10)$$

Here $t = 2, \dots, T, k = 1, 2, \dots$, and $Lap(\lambda)$ denotes a Laplacian distribution with scale parameter λ . For the exogenous part, we consider specifying the relationship between $\boldsymbol{\pi}_t$ and $\tilde{\boldsymbol{\pi}}_t = (\tilde{\pi}_{t1}, \tilde{\pi}_{t2}, \dots)$ such that $\tilde{G}_t = \sum_{k=1}^{\infty} \tilde{\pi}_{tk} \delta_{\phi_k}$. We let

$$\tilde{\boldsymbol{\pi}}_t = \boldsymbol{\pi}_t + \boldsymbol{\eta} \cdot \mathbf{y}_t, \quad t = 1, \dots, T, \quad \mathbf{1}' \cdot \boldsymbol{\eta} = \mathbf{0}. \quad (11)$$

Here $\boldsymbol{\eta}$ is a $K \times p$ matrix which indicates the relationship between the topic distribution $\tilde{\boldsymbol{\pi}}_t$ and the length- p vector \mathbf{y}_t . However, we notice that $\tilde{\boldsymbol{\pi}}_t$ and $\boldsymbol{\pi}_t$ are of infinite length, which creates difficulty in our inference. Therefore we adopt a stick-breaking truncation approach, i.e. we only consider $\{\phi_k\}_{k=1}^K$ in our model; the probability weights for $\{\phi_k\}_{k=K+1}^{\infty}$ in $\boldsymbol{\pi}_t$ will be added into π_{tK} . We note that a number of papers in topic modeling have put this approach into practice (Pruteanu-Malinici et al. 2010; Wang, Paisley, and Blei 2011).

It has been shown (Pruteanu-Malinici et al. 2010) that when the truncation level K is large, we may as well replace the distribution of $\boldsymbol{\pi}_1$ with $\boldsymbol{\pi}_1 \sim Dir(\gamma \boldsymbol{\pi}_0)$, where $\gamma = 1$, $\boldsymbol{\pi}_0 = (1/K, \dots, 1/K)$. We also let $H = Dir(\beta, \dots, \beta)$ as in the paper by Teh et al. (2005). We summarize our model as

$$\begin{cases} \phi_1, \dots, \phi_K \stackrel{iid}{\sim} Dir(\beta, \dots, \beta), \\ \boldsymbol{\pi}_1 \sim Dir(\gamma \boldsymbol{\pi}_0), \quad \pi_{tk} = \pi_{t-1k} + Lap(\lambda), \\ \tilde{\boldsymbol{\pi}}_t = \boldsymbol{\pi}_t + \boldsymbol{\eta} \cdot \mathbf{y}_t, \\ \theta_{t,i} | (\alpha_t, \tilde{\boldsymbol{\pi}}_t) \stackrel{iid}{\sim} Dir(\alpha_t \tilde{\boldsymbol{\pi}}_t), \\ z_{t,i,j} | \theta_{t,i} \stackrel{iid}{\sim} Cat(\theta_{t,i}), \quad x_{t,i,j} \sim Cat(\phi_{z_{t,i,j}}). \end{cases} \quad (12)$$

Here $t = 1, \dots, T, i = 1, \dots, N_t, j = 1, \dots, J_{t,i}$ ($t \geq 2$ for the second line). The last two lines above are derived as in Teh et al. (2005). We note that here ϕ_k is the per-topic word distribution, $\theta_{t,i}$ is the per-document topic distribution, and $z_{t,i,j}$ is the actual topic for each word; they have the same meaning as in LDA.

We also note that π_{tk} and $\tilde{\pi}_{tk}$ may be negative from our model construction. This is not really an issue in terms of computation, since we always set the overall likelihood to 0 if there exists $\pi_{t,k} \notin \mathbb{R}^+$ or $\tilde{\pi}_{t,k} \notin \mathbb{R}^+$.

We name our model a ‘‘group-correlated temporal topic model’’ (GCLDA). Here a ‘‘group’’ stands for all the documents within the same time period. We use the term ‘‘correlated’’ because the baseline topic distributions $\{\boldsymbol{\pi}_t\}_{t=1}^T$ for each period, controlling for $\{\mathbf{y}_t\}_{t=1}^T$, are endogenously correlated; meanwhile, the realized baseline topic distributions $\{\tilde{\boldsymbol{\pi}}_t\}_{t=1}^T$ for each period are also correlated with the given exogenous process $\{\mathbf{y}_t\}_{t=1}^T$.

3.3 Sampling the posterior: An MCMC Approach

Direct estimation of the Bayesian posterior is often intractable since the closed-form expression, if it exists, can be difficult to integrate and thus, many approaches to approximate the posterior have been proposed. Monte Carlo methods, which draw a large number of samples from the posterior as its approximation, are particularly helpful. In this paper, we adopt the Markov chain Monte Carlo (MCMC) approach which constructs samples from a Markov chain and

is asymptotically exact. Below we provide the Metropolis-within-Gibbs sampling approach tailored to our situation, which is a variant of the general MCMC approach. It only requires specifying the full conditionals of the unknown variables, which is covered below.

We consider sampling the following variables $\mathbf{Z} = \{z_{t,i,j}\}_{j=1}^{J_{t,i}} \{N_t\}_{t=1}^T$, $\{\alpha_t\}_{t=1}^T$, $\{\tilde{\pi}_t\}_{t=1}^T$, $\boldsymbol{\eta}$, λ . We integrate out $\{\boldsymbol{\theta}_{t,i}\}_{i=1}^{N_t}$ and $\{\phi_k\}_{k=1}^K$ to speed up calculation, using a collapsed Gibbs approach. Following Griffiths and Steyvers (2004), conditioning on all other variables listed for sampling,

$$p(z_{t,i,j} = k | rest) \propto (C_{t,i,k}^{(-1)} + \alpha_t \tilde{\pi}_{tk}) \cdot \frac{C_{x_{t,i,k},k}^{(-1)} + \beta}{C_k^{(-1)} + V\beta}. \quad (13)$$

Here $C_{t,i,k}^{(-1)}$ is the count of elements in $\mathbf{Z} \setminus \{z_{t,i,j}\}$ which belong to $d_{t,i}$ and has values equal to k ; $C_{x_{t,i,k},k}^{(-1)}$ is the count of elements in $\mathbf{Z} \setminus \{z_{t,i,j}\}$ whose values are k and corresponding words are $x_{t,i,j}$; $C_k^{(-1)}$ is the count of elements in $\mathbf{Z} \setminus \{z_{t,i,j}\}$ whose values are k . Also following Griffiths and Steyvers (2004), we have for α_t and $\tilde{\pi}_t$

$$\begin{aligned} & p(\alpha_t, \tilde{\pi}_t | rest) \propto p(\mathbf{Z} | \alpha_t, \tilde{\pi}_t) \cdot p(\boldsymbol{\pi}_{1,\dots,T} | \gamma, \boldsymbol{\pi}_0) \cdot p(\alpha_t) \\ & \propto \left[\frac{\Gamma(\alpha_t)}{\prod_{k=1}^K \Gamma(\alpha_t \tilde{\pi}_{tk})} \right]^{N_t} \cdot \prod_{i=1}^{N_t} \frac{\prod_{k=1}^K \Gamma(C_{t,i,k} + \alpha_t \tilde{\pi}_{tk})}{\Gamma(J_{t,i} + \alpha_t)} \\ & \cdot \prod_{k=1}^K \pi_{1k}^{\gamma \pi_{0k-1}} \cdot \exp \left[-\lambda \left(1_{t>1} \cdot \sum_{k=1}^k |\pi_{t+1k} - \pi_{tk}| \right. \right. \\ & \left. \left. + 1_{t<T} \cdot \sum_{k=1}^k |\pi_{tk} - \pi_{t-1k}| \right) \right] \cdot p(\alpha_t). \quad (14) \end{aligned}$$

Here the difference between $C_{t,i,k}$ and $C_{t,i,k}^{(-1)}$ is to replace $\mathbf{Z} \setminus \{z_{t,i,j}\}$ with \mathbf{Z} . We also view $\boldsymbol{\pi}_t$, π_{tk} , etc. as functions of the other parameters; specifically, $\pi_{tk} = \tilde{\pi}_{tk} - \boldsymbol{\eta}_k \mathbf{y}_t$, where $\boldsymbol{\eta}_k$ is the k -th row of $\boldsymbol{\eta}$. This involves a transformation of variables; however, the related Jacobian determinant $\det(J) = 1$, so (14) is still valid. For parameter $\boldsymbol{\eta}$, we have

$$\begin{aligned} p(\boldsymbol{\eta} | rest) & \propto \prod_{t=2}^T p(\boldsymbol{\pi}_t | \boldsymbol{\pi}_{t-1}, \lambda) \cdot p(\boldsymbol{\eta}) \\ & \propto \exp \left(-\lambda \cdot \sum_{t=2}^T \sum_{k=1}^K |\pi_{tk} - \pi_{t-1k}| \right) \cdot p(\boldsymbol{\eta}). \quad (15) \end{aligned}$$

Finally, for parameter λ , we have

$$\begin{aligned} p(\lambda | rest) & \propto \prod_{t=2}^T p(\boldsymbol{\pi}_t | \boldsymbol{\pi}_{t-1}, \lambda) \cdot p(\lambda) \\ & \propto \lambda^{(T-1)K} \exp \left(-\lambda \cdot \sum_{t=2}^T \sum_{k=1}^K |\pi_{tk} - \pi_{t-1k}| \right) \cdot p(\lambda). \quad (16) \end{aligned}$$

We note that (13) and (16) are full conditionals, and we can easily derive the full conditionals of α_t , π_{tk} , and η_k from

(14) and (15). Since each $z_{t,i,j} | rest$ has a categorical distribution, and $\lambda | rest$ has a Gamma distribution with a conjugate prior, they can be updated with Gibbs updates. For α_t , π_{tk} , and η_k , we replace a Gibbs update with a Metropolis update. Specifically, suppose we know $p(par | rest)$ up to a multiplicative constant, where par is any length-1 parameter. We also assume $par = par^{(r)}$ at the r -th iteration. Then at the $(r+1)$ -th iteration,

$$par_{new} \sim q(\cdot | par^{(r)}), \delta_i^{(r)} = \begin{cases} par^{(r+1)} & \text{with prob. } P, \\ par_{new} & \text{with prob. } 1 - P, \end{cases} \quad (17)$$

where $P = \min\{p(par_{new} | rest)/p(par^{(r)} | rest), 1\}$, and $q(\cdot | \cdot)$ is a known conditional probability distribution such that $q(x|y) = q(y|x)$.

This completes our sampling and posterior inference. For the theoretical convergence properties of Metropolis-within-Gibbs samplers, the reader can refer to Robert and Casella (2004) and Roberts and Rosenthal (2006).

4 Case Studies

The proposed model, ‘‘GCLDA,’’ is demonstrated on two data sets: (1) online job advertisements from my.jobs from February to September in 2014, and (2) journalists’ postings in 2014 in the ‘‘Finance’’ section in BusinessInsider.com, an American business and technology news website. Our algorithm has been implemented in Java, and we compare GCLDA with LDA and STM.

4.1 Experiment Settings

We initialize the hyperparameters of LDA as follows: $\boldsymbol{\alpha} = (50/K, \dots, 50/K)$, $\boldsymbol{\beta} = (0.01, \dots, 0.01)$, according to a rule of thumb; this has been carried out in Berry and Kogan (2010) and Sridhar (2015). For GCLDA, we let $\gamma = 1$, $\boldsymbol{\pi}_0 = (1/K, \dots, 1/K)$, $\alpha_t \stackrel{iid}{\sim} \Gamma(1, 1)$, $p(\boldsymbol{\eta}) \propto e^{-0.01 \sum |\eta_k|}$, $\lambda \sim \Gamma(1, 1)$, $\beta = 0.01$. We carry out the Metropolis-within-Gibbs algorithm as described in Section 3.3, and run 5,000 iterations of the Markov chain with 1,000 burn-in samples for GCLDA and LDA; for LDA, we apply the collapsed Gibbs sampling as in Griffiths and Steyvers (2004). The number of topics is set to $K = 50$ for both data sets. For STM, we apply the ‘‘Spectral’’ initialization (Roberts, Stewart, and Tingley 2015) together with other default settings in the R package `stm`. We perform data cleaning, remove the stopwords, stem the documents, and keep most frequent V words in each study. For the job advertisements, $V = 2,000$ and covers 96.2% of all words with repetition, which means that the choice of words in job advertisements is quite narrow; for the journalists’ postings, $V = 3,000$ and covers 93.9% of all words with repetition.

We use perplexity to compare the difference of the prediction power between LDA and GCLDA. The perplexity for N_{test} held-out documents given the training data D is defined as

$$perp = \exp \left\{ -\frac{\sum_{i=1}^{N_{test}} \log p(d_{test,i} | D)}{\sum_{i=1}^{N_{test}} n_{test,i}} \right\} \quad (18)$$

where $d_{test,i}$ represents the i -th held-out document, and $n_{test,i}$ is the number of words in $d_{test,i}$. We expect the perplexity to be small when a model performs well, since this means that under the estimated model, the probability of a word in the testing documents being written *a priori* is large. We apply the “Left-to-right” algorithm (Wallach et al. 2009) and apply point estimates for “ Φ ” and “ αm ” using the training data, as suggested in Section 3 in the same paper.

4.2 My.jobs: Online Job Advertisements

The number of online job advertisements on my.jobs from February to September in 2014 amounts to 17,147,357 in total, and the number of advertisements each day varies greatly. Therefore, we gather a stratified sample of 44,660 advertisements with a roughly equal number of samples for each day, so that we have sampled 0.26% of all the documents in total. The training data set consists of 40,449 advertisements, and the testing data set consists of 4,211 advertisements (9.4% of the sample). For the exogenous variable $\{y_t\}_{t=1}^T$, we use the standardized Consumer Price Index from February to September in 2014, so that $p = 1$, and $T = 8$.

Figure 1 implies that GCLDA better predicts the words in the new documents in terms of perplexity. This is due to the fact that the introduction of the endogenous and exogenous processes allows us to make more accurate inference on the topic distributions of the documents in a given period of time. The standard errors for the perplexity in each period are also shown; we can observe that the difference is quite significant.

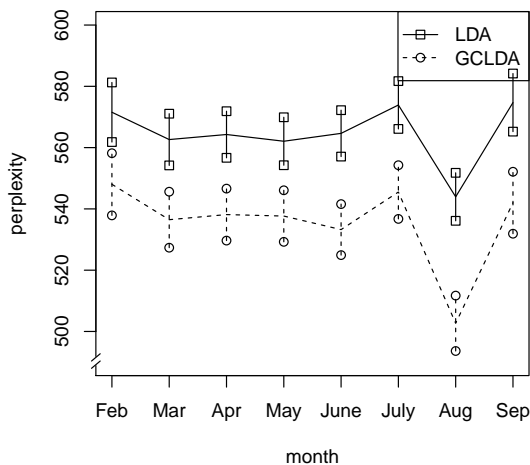


Figure 1: Perplexity results for the job advertisements from February to September in 2014.

The 20 most common topics are presented in Figure 2. The x-axis represent the degree of correlation $\rho = \eta/\pi$ for all topics, i.e. the percent change in the topic proportion given one unit change in the exogenous covariate. Here π and η denote the related component of $\sum \pi_t/T$ and η for each topic. Table 1 lists the highest probability words sorted

by their probabilities from high to low inside the five topics with highest ρ in Figure 2.




Figure 2: The 20 most common topics and their ρ from GCLDA for the job advertisements.

“equal opportunity”:	employment status disabled veteran equal
“health care”:	health care medical service provide center hospital
“software”:	development experience software design application
“secretarial”:	management operations ensure training perform
“nursing”:	care nursing patient required clinical practice medical

Table 1: The highest probability words inside the five topics with highest ρ in Figure 2.

A number of facts can be inferred from Figure 2. The topics with positive ρ are those that have a positive correlation with the growth of the CPI in 2014. We can observe that two of them are supported by the U.S. government spending, namely “equal opportunity” and “health care,” the latter of which is probably related to the Affordable Care Act programs. This suggests that there is a causal relationship between the increase in government spending and the increase in the number of jobs in these categories, and the former was also an underlying factor in the growth of the CPI in 2014. We also observe that “software” and “secretarial” were moving in the same direction of CPI, while some traditional higher-paid job categories, such as engineering and marketing, were not. This partly agrees with some news articles in 2014 in that while the labor market was recovering, there was relatively lower growth in traditional higher-paid job categories (Lowrey 2014; Coy 2014).

We also compare our method with STM. Below is the STM counterpart of Figure 2. We can observe that the topics from GCLDA seem to be more time-related and tend to be more informative of the economy during the period.




Figure 3: The 20 most common topics and their γ from STM for the job advertisements.

4.3 BusinessInsider.com: Financial News Articles

We consider all contributions in the “Finance” section of BusinessInsider.com on all trading days in 2014. There are 15,659 articles in total, which are divided into a training data set containing 12,527 articles and a testing data set containing 3,132 articles (20% of all articles). We increase the proportion of testing documents and let $T = 252$ (all trading days) in order to create a more challenging scenario for GCLDA. We apply the daily price of the Chicago Board Options Exchange Market Volatility Index (VIX) as the exogenous process, measuring the volatility of the U.S. financial market. The other settings are the same as those in Section 4.2. We provide an analysis of the perplexity of LDA and GCLDA in Figure 4. The lines are smoothed by LOESS with a span of 0.15, as there are large fluctuations in perplexity from day to day.

Again we observe that GCLDA generates a lower perplexity for the testing documents over time, therefore improving the fitting of the topic model. From Figure 5, the topics that are strongly positively correlated with the VIX are generally short-term news, such as stock market news and announcements from central banks, as in the topics “FED” (the Federal Reserve), “revenue,” and “stock market,” which are indeed closely related to the stock market. From our analysis, the drop in oil price and the instability in Russia and Ukraine were also major causes of fluctuations in the stock market in 2014. On the other hand, we observe that news about longer-term economic trends is not positively correlated with the VIX, such as “companies” and “labor market.” We posit that GCLDA can be used for finding topics that are major contributors to changes in an exogenous process during a period of time.

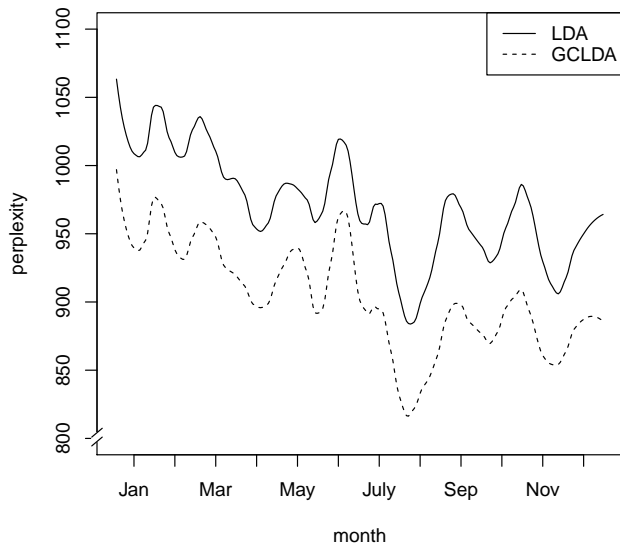


Figure 4: Perplexity results for contributions in the “Finance” section in BusinessInsider.com in 2014.




Figure 5: The 20 most common topics and their ρ from GCLDA for the financial articles.

“FED”:	rate FED inflation policies market federal expected
“revenue”:	quarter billion year revenue million earnings share
“Ukraine”:	Russia Ukraine Moscow gas country president
“stock market”:	market trade stock week day close morning dollar
“energy”:	oil price energies gas production crude supplies

Table 2: The highest probability words inside the five topics with highest ρ in Figure 5.




Figure 6: The 20 most common topics and their γ from STM for the financial articles.

We also compare our method with STM, with Figure 6 being the STM counterpart of Figure 5. Again we observe that the topics from GCLDA seem to be more time-related and tend to be more informative of the stock market during the period. These findings assert our view that GCLDA improves the structure of the topic model and makes it more time-dependent.

5 Conclusion

We have developed a temporal topic model which analyzes time-stamped text documents with known exogenous processes. Our new model, GCLDA, takes both endogenous and exogenous processes into account, and applies Markov chain Monte Carlo sampling for calibration. We have demonstrated that this model better fits temporal documents in terms of perplexity, and extracts well information from job advertisements and financial news articles. We suggest that a possible direction for the future could be analyzing the contents of temporal documents so that they could predict the trends of related exogenous processes.

References

- Berry, M., and Kogan, J. 2010. *Text Mining: Applications and Theory*. Chichester, UK: Wiley.
- Blei, D., and Lafferty, J. 2006. Dynamic topic models. In *23rd International Conference on Machine Learning*.
- Blei, D.; Ng, A.; and Jordan, M. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research* 3:993–1022.
- Coy, P. 2014. America’s low-paying recovery: More jobs than ever, worse wages. *Bloomberg Business*. www.bloomberg.com/bw/articles/2014-08-11/report-new-jobs-in-u-dot-s-dot-offer-lower-wages-than-before-recession.
- Dubey, A.; Ho, Q.; Williamson, S.; and Xing, E. 2014. Dependent nonparametric trees for dynamic hierarchical clustering. In *Advances in Neural Information Processing Systems* 28. MIT Press.
- Ferguson, T. 1973. Bayesian analysis of some nonparametric problems. *Annals of Statistics* 1:209–230.
- Griffiths, T., and Steyvers, M. 2004. Finding scientific topics. *Proceedings of the National Academy of Sciences* 101:5228–5235.
- Lowrey, A. 2014. Recovery has created far more low-wage jobs than better-paid ones. *The New York Times*. www.nytimes.com/2014/04/28/business/economy/recovery-has-created-far-more-low-wage-jobs-than-better-paid-ones.html.
- Pruteanu-Malinici, I.; Ren, L.; Paisley, J.; Wang, E.; and Carin, L. 2010. Hierarchical bayesian modeling of topics in time-stamped documents. *IEEE Trans. Pattern Analysis Machine Intelligence* 32:996–1011.
- Ren, L.; Dunson, D.; and Carin, L. 2008. The dynamic hierarchical dirichlet process. In *25th International Conference on Machine Learning*.
- Robert, C., and Casella, G. 2004. *Monte Carlo Statistical Methods*. New York: Springer, 2nd ed. edition.
- Roberts, G., and Rosenthal, J. 2006. Harris recurrence of metropolis-within-gibbs and trans-dimensional markov chains. *The Annals of Applied Probability* 16:2123–2139.
- Roberts, M.; Stewart, B.; and Airolidi, E. 2015. A model of text for experimentation in the social sciences. scholar.harvard.edu/files/bstewart/files/stm.pdf.
- Roberts, M.; Stewart, B.; and Tingley, D. 2015. Navigating the local modes of big data: The case of topic models. In *Data Analytics in Social Science, Government, and Industry*. New York: Cambridge University Press. forthcoming.
- Sethuraman, J. 1994. A constructive definition of dirichlet priors. *Statistica Sinica* 4:639–650.
- Sridhar, V. 2015. Unsupervised topic modeling for short texts using distributed representations of words. In *Proceedings of NAACL-HLT 2015*.
- Teh, Y.; Jordan, M.; Beal, M.; and Blei, D. 2005. Hierarchical dirichlet processes. *Journal of the American Statistical Association* 101:1566–1582.
- Wallach, H.; Murray, I.; Salakhutdinov, R.; and Mimno, D. 2009. Evaluation methods for topic models. In *26th International Conference on Machine Learning*.
- Wang, X., and McCallum, A. 2006. Topics over time: A non-markov. continuous-time model of topical trends. In *12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*.
- Wang, C.; Paisley, J.; and Blei, D. 2011. Online variational inference for the hierarchical dirichlet process. In *Artificial Intelligence and Statistics*.
- Zhang, H.; Kim, G.; and Xing, E. 2015. Dynamic topic modeling for monitoring market competition from online text and image data. In *21st ACM SIGKDD Conference on Knowledge Discovery and Data Mining*.