

Dynamic and Static Topic Model for Analyzing Time-Series Document Collections

Rem Hida[†] Naoya Takeishi^{‡†} Takehisa Yairi[†] Koichi Hori[†]

[†]Department of Aeronautics and Astronautics, The University of Tokyo

{hida, yairi, hori}@ailab.t.u-tokyo.ac.jp

[‡]RIKEN Center for Advanced Intelligence Project, Tokyo, Japan

naoya.takeishi@riken.jp

Abstract

For extracting meaningful topics from texts, their structures should be considered properly. In this paper, we aim to analyze structured time-series documents such as a collection of news articles and a series of scientific papers, wherein topics evolve along time depending on multiple topics in the past, and are also related to each other at each time. To this end, we propose a *dynamic and static topic model*, which simultaneously considers the dynamic structures of the temporal topic evolution and the static structures of the topic hierarchy at each time. We show the results of experiments on collections of scientific papers, in which the proposed method outperformed conventional models. Moreover, we show an example of extracted topic structures, which we found helpful for analyzing research activities.

1 Introduction

Probabilistic topic models such as latent Dirichlet allocation (LDA) (Blei et al., 2003) have been utilized for analyzing a wide variety of datasets such as document collections, images, and genes. Although vanilla LDA has been favored partly due to its simplicity, one of its limitations is that the output is not necessarily very understandable because the priors on the topics are independent. Consequently, there has been a lot of research aimed at improving probabilistic topic models by utilizing the inherent *structures* of datasets in their modeling (see, e.g., Blei and Lafferty (2006); Li and McCallum (2006); see Section 2 for other models).

In this work, we aimed to leverage the dynamic and static structures of topics for improving the modeling capability and the understandability of

topic models. These two types of structures, which we instantiate below, are essential in many types of datasets, and in fact, each of them has been considered separately in several previous studies. In this paper, we propose a topic model that is aware of both of these structures, namely *dynamic and static topic model* (DSTM).

The underlying motivation of DSTM is twofold. First, a collection of documents often has *dynamic structures*; i.e., topics evolve along time influencing each other. For example, topics in papers are related to topics in past papers. We may want to extract such dynamic structures of topics from collections of scientific papers for summarizing research activities. Second, there are also *static structures* of topics such as correlation and hierarchy. For instance, in a collection of news articles, the “sports” topic must have the “baseball” topic and the “football” topic as its subtopic. This kind of static structure of topics helps us understand the relationship among them.

The remainder of this paper is organized as follows. In Section 2, we briefly review related work. In Section 3, the generative model and the inference/learning procedures of DSTM are presented. In Section 4, the results of the experiments are shown. This paper is concluded in Section 5.

2 Related Work

Researchers have proposed several variants of topic models that consider the dynamic or static structure. Approaches focusing on the dynamic structure include dynamic topic model (DTM) (Blei and Lafferty, 2006), topic over time (TOT) (Wang and McCallum, 2006), multiscale dynamic topic model (MDTM) (Iwata et al., 2010), dependent Dirichlet processes mixture model (D-DPMM) (Lin et al., 2010), and infinite dynamic topic model (iDTM) (Ahmed and Xing, 2010).

D^t	number of documents at epoch t
n_d^t	number of words in the d -th doc. at epoch t
$w_{d,i}^t$	the i -th word in the d -th doc. at epoch t
K	total number of subtopics
S	number of supertopics
$y_{d,i}^t$	supertopic of $w_{d,i}^t$
$z_{d,i}^t$	subtopic of $w_{d,i}^t$
${}^1\theta_d^t$	multinomial distribution over supertopics for the d -th doc. at epoch t
${}^2\theta_{d,s}^t$	multinomial distribution over subtopics for the d -th doc. in s -th supertopic at epoch t
ϕ_k^t	multinomial distribution over words for the k -th subtopic at epoch t
${}^2\alpha_s^t$	static structure weight (prior of ${}^2\theta_{d,s}^t$)
β^t	dynamic structure weight between topics at time $t-1$ and those at epoch t

Table 1: Notations in the proposed model.

These methods have been successfully applied to a temporal collection of documents, but none of them take temporal dependencies between multiple topics into account; i.e., in these models, only a single topic contributes to a topic in the future.

For the static structure, several models including correlated topic model (CTM) (Lafferty and Blei, 2006), pachinko allocation model (PAM) (Li and McCallum, 2006), and segmented topic model (STM) (Du et al., 2010) have been proposed. CTM models the correlation between topics using the normal distribution as the prior, PAM introduces the hierarchical structure to topics, and STM uses paragraphs or sentences as the hierarchical structure. These models can consider the static structure such as correlation and hierarchy between topics. However, most of them lack the dynamic structure in their model; i.e., they do not premise temporal collections of documents.

One of the existing methods that is most related to the proposed model is the hierarchical topic evolution model (HTEM) (Song et al., 2016). HTEM captures the relation between evolving topics using a nested distance-dependent Chinese restaurant process. It has been successfully applied to a temporal collection of documents for extracting structure but does not take multiple topics dependencies into account either.

In this work, we built a new model to overcome the limitation of the existing models, i.e., to examine both the dynamic and static structures simultaneously. We expect that the proposed model can be applied to various applications such as topic trend analysis and text summarization.




Figure 1: Graphical model of the proposed model for epochs $t-1$ and t .

3 Dynamic and Static Topic Model

In this section, we state the generative model of the proposed method, DSTM. Afterward, the procedure for inference and learning is presented. Our notations are summarized in Table 1.

3.1 Generative Model

In the proposed model, DSTM, the dynamic and static structures are modeled as follows.

Dynamic Structure We model the temporal evolution of topic-word distribution by making it proportional to a weighted sum of topic-word distributions at the previous time (epoch), i.e.,

$$\phi_k^t \sim \text{Dirichlet} \left(\sum_{k'=1}^K \beta_{k,k'}^t \phi_{k'}^{t-1} \right), \quad (1)$$

where ϕ_k^t denotes the word distribution of the k -th topic at the t -th time-epoch, and $\beta_{k,k'}^t$ is a weight that determines the dependency between the k -th topic at epoch t and the k' -th topic at epoch $t-1$.

Static Structure We model the static structure as a hierarchy of topics at each epoch. We utilize the supertopic-subtopic structure as in PAM (Li and McCallum, 2006), where the priors of topics (subtopics) are determined by their supertopic.

Generative Process In summary, the generative process at epoch t is as follows.

1. For each subtopic $k = 1, \dots, K$,
 - (a) Draw a topic-word distribution $\phi_k^t \sim \text{Dirichlet}(\sum_{k'} \beta_{k,k'}^t \phi_{k'}^{t-1})$.
2. For each document $d = 1, \dots, D^t$,
 - (a) Draw a supertopic distribution ${}^1\theta_d^t \sim \text{Dirichlet}({}^1\alpha^t)$.

- (b) For each supertopic $s = 1, \dots, S$,
- i. Draw a subtopic distribution ${}^2\theta_{d,s}^t \sim \text{Dirichlet}({}^2\alpha_s^t)$.
- (c) For each word $i = 1, \dots, n_d^t$,
- i. Draw a supertopic-word assignment $y_{d,i}^t \sim \text{Multinomial}({}^1\theta_d^t)$.
 - ii. Draw a subtopic-word assignment $z_{d,i}^t \sim \text{Multinomial}({}^2\theta_{d,y_{d,i}^t}^t)$.
 - iii. Draw a word-observation $w_{d,i}^t \sim \text{Multinomial}(\phi_{z_{d,i}^t}^t)$.

Note that the above process should be repeated for every epoch t . The corresponding graphical model is presented in Figure 1.

3.2 Inference and Learning

Since analytical inference for DSTM is intractable, we resort to a stochastic EM algorithm (Andrieu et al., 2003) with the collapsed Gibbs sampling (Griffiths and Steyvers, 2004). However, such a strategy is still much costly due to the temporal dependencies of ϕ . Therefore, we introduce a further approximation; we surrogate $\phi_{k'}^{t-1}$ in Eq. (1) by its expectation $\hat{\phi}_{k'}^{t-1} = \mathbb{E}[\phi_{k'}^{t-1}]$. This compromise enables us to run the EM algorithm for each epoch in sequence from $t = 1$ to $t = T$ without any backward inference. In fact, such approximation technique is also utilized in the inference of MDTM (Iwata et al., 2010).

Note that the proposed model has a moderate number of hyperparameters to be set manually, and that they can be tuned according to the existing know-how of topic modeling. This feature makes the proposed model appealing in terms of inference and learning.

E-step In E-step, the supertopic/subtopic assignments are sampled. Given the current state of all variables except $y_{d,i}^t$ and $z_{d,i}^t$, new values for them should be sampled according to

$$p(y_{d,i}^t = s, z_{d,i}^t = k \mid w^t, y^t, z^t, \Phi^{t-1}, {}^1\alpha^t, {}^2\alpha^t, \beta^t) \propto \frac{n_{d,s}^t + {}^1\alpha_s^t}{n_{d,i}^t + \sum_{s=1}^S {}^1\alpha_s^t} \cdot \frac{n_{d,s,k}^t + {}^2\alpha_{s,k}^t}{n_{d,s}^t + \sum_{k=1}^K {}^2\alpha_{s,k}^t} \cdot \frac{n_{k,v}^t + \sum_{k'=1}^K \beta_{k,k'}^t \hat{\phi}_{k',v}^{t-1}}{n_{k,i}^t + \sum_{k'=1}^K \beta_{k,k'}^t}, \quad (2)$$

where $n_{k,v}^t$ denotes the number of tokens assigned to topic k for word v at epoch t , $n_k^t = \sum_v n_{k,v}^t$, and $n_{d,s}^t$ and $n_{d,s,k}^t$ denote the number of tokens in document d assigned to supertopic s and subtopic

	NIPS	Drone
Date	1987–1999	2009–2016
# Documents	1,740	1,035
# Vocabulary	11,443	3,442
# Tokens	2,271,087	68,305

Table 2: Summary of the datasets.

k (via s), at epoch t respectively. Moreover, $n_{k,i}^t$ denotes the count yielded excluding the i -th token.

M-step In M-step, ${}^2\alpha^t$ and β^t are updated using the fixed-point iteration (Minka, 2000).

$$({}^2\alpha_{s,k}^t)^* = {}^2\alpha_{s,k}^t \frac{\sum_{d=1}^{D^t} \Psi(n_{d,s,k}^t + {}^2\alpha_{s,k}^t) - \Psi({}^2\alpha_{s,k}^t)}{\sum_{d=1}^{D^t} \Psi(n_{d,s}^t + {}^2\alpha_s^t) - \Psi({}^2\alpha_s^t)}, \quad (3)$$

$$(\beta_{k,k'}^t)^* = \beta_{k,k'}^t \frac{\sum_v \hat{\phi}_{k',v}^{t-1} B_{k',v}^t}{\Psi(n_k^t + \sum_{k'} \beta_{k,k'}^t) - \Psi(\sum_{k'} \beta_{k,k'}^t)}. \quad (4)$$

Here, Ψ is the digamma function, ${}^2\alpha_s^t = \sum_k {}^2\alpha_{s,k}^t$, and

$$B_{k',v}^t = \Psi\left(n_{k',v}^t + \sum_{k'} \beta_{k,k'}^t \hat{\phi}_{k',v}^{t-1}\right) - \Psi\left(\sum_{k'} \beta_{k,k'}^t \hat{\phi}_{k',v}^{t-1}\right).$$

Overall Procedure The EM algorithm is run for each epoch in sequence; at epoch t , after running the EM until convergence, $\hat{\phi}_{k,v}^t$ is computed by

$$\hat{\phi}_{k,v}^t = \frac{n_{k,v}^t + \sum_{k'} \beta_{k,k'}^t \hat{\phi}_{k',v}^{t-1}}{n_k^t + \sum_{k'} \beta_{k,k'}^t},$$

and then this value is used for the EM at the next epoch $t + 1$. Moreover, see Supplementary A for the computation of the statistics of the other variables.

4 Experiments

4.1 Datasets

We used two datasets comprising technical papers: **NIPS** (Perrone et al., 2016) and **Drone** (Liew et al., 2017). **NIPS** is a collection of the papers that appeared in NIPS conferences. **Drone** is a collection of abstracts of papers on unmanned aerial vehicles (UAVs) and was collected from related conferences and journals for surveying recent developments in UAVs. The characteristics of those datasets are summarized in Table 2. See Supplementary B for the details of data preprocessing.

			NIPS			Drone		
	static	dynamic	K30 (S15)	K40 (S20)	K50 (S25)	K15 (S3)	K20 (S3)	K25 (S3)
LDA	-	-	1455.6 (16.7)	1407.3 (15.9)	1374.6 (16.8)	1624.3 (191.1)	1634.8 (189.1)	1644.7 (193.0)
PAM	✓	-	1455.1 (18.2)	1407.0 (17.5)	1376.9 (16.7)	1587.4 (185.1)	1589.9 (191.4)	1590.8 (186.8)
DRTM	-	✓	1380.7 (18.5)	1308.6 (17.5)	1253.9 (17.9)	1212.5 (153.2)	1206.1 (148.0)	1201.2 (143.5)
DSTM	✓	✓	1378.7 (16.5)	1301.0 (17.9)	1247.3 (17.2)	1194.2 (148.2)	1180.0 (147.0)	1171.6 (141.4)

Table 3: Means (and standard deviations) of PPLs averaged over all epochs for each dataset with different values of K and S . The proposed method, DSTM, achieved the smallest PPL.

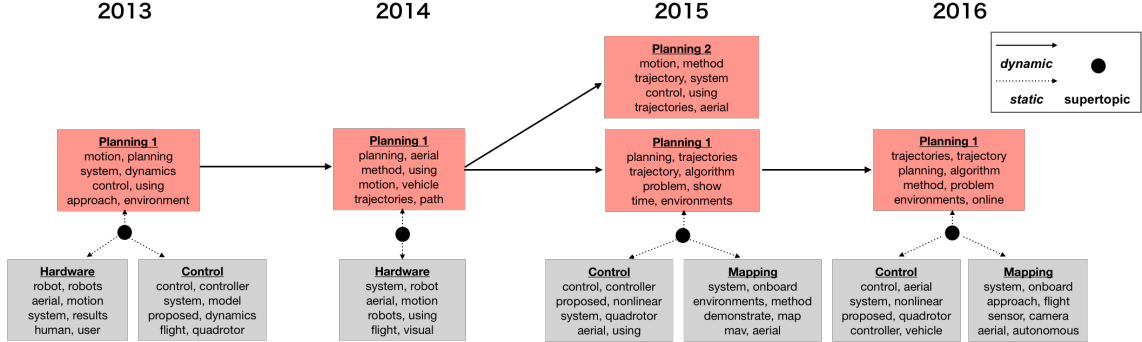


Figure 2: Part of the topic structure extracted from **Drone** dataset using the proposed method. The solid arrows denote the temporal evolution of “planning” topics. The dotted arrows mean that “planning” topics are related to “hardware”, “control”, and “mapping” topics via some supertopics (filled circles).

4.2 Evaluation by Perplexity

First, we evaluate the performance of the proposed method quantitatively using perplexity (PPL):

$$\text{PPL} = \exp \left(- \frac{\sum_{d=1}^D \sum_{w_d^{\text{test}}} \log p(w_d, i | \mathcal{M})}{\sum_{d=1}^D n_d^{\text{test}}} \right).$$

For each epoch, we used 90% of tokens in each document for training and calculated the PPL using the remaining 10% of tokens. We randomly created 10 train-test pairs and evaluated the means of the PPLs over those random trials. We compared the performance of DSTM to three baselines: LDA (Blei et al., 2003), PAM (Li and McCallum, 2006), and the proposed model without the static structure, which we term DRTM. See Supplementary C on their hyperparameter setting.

The means of the PPLs averaged over all epochs for each dataset with different values K are shown in Table 3. In both datasets with every setting of K , the proposed model, DSTM, achieved the smallest PPL, which implies its effectiveness for modeling a collection of technical papers. For clarity, we conducted paired t-tests between the perplexities of the proposed method and those of the baselines. On the differences between DSTM and DRTM, the p-values were 4.2×10^{-2} ($K = 30$), 7.9×10^{-5} ($K = 40$), and 6.4×10^{-7} ($K = 50$) for the **NIPS** dataset, and 1.3×10^{-4}

($K = 15$), 8.8×10^{-5} ($K = 20$), and 4.9×10^{-6} ($K = 25$) for the **Drone** dataset, respectively. It is also noteworthy that DRTM shows more significant improvement relative to LDA than PAM does. This suggests that the dynamic structure with multiple-topic dependencies is essential for datasets of this kind.

4.3 Analysis of Extracted Structure

We examined the topic structures extracted from the **Drone** dataset using DSTM. In Figure 2, we show a part of the extracted structure regarding planning of the UAV’s path and/or movement. We identified “planning” topics by looking for keywords such as “trajectory” and “motion.” In Figure 2, each node is labeled with eight most probable keywords. Moreover, solid arrows (dynamic relations) are drawn if the corresponding $\beta_{k,k'}^t$ is larger than 200, and dotted arrows (static relations) are drawn between a supertopic and subtopics with the two or three largest values of ${}^2\alpha_{s,k}^t$.

Looking at the dynamic structure, we may see how research interest regarding planning has changed. For example, word “online” first emerges in the “planning” topic in 2016. This is possibly due to the increasing interest in real-time planning problems, which is becoming feasible due to the recent development of on-board computers. In regard to the static structures, for

example, the “planning” topic is related to the “hardware” and “control” topics in 2013 and 2014, whereas it is also related to the “mapping” topic in 2015 and 2016. Looking at these static structures, we may anticipate how research areas are related to each other in each year. In this case, we can anticipate that planning problems are combined with mapping problems well in recent years. Note that we cannot obtain these results unless the dynamic and static structures are considered simultaneously.

5 Conclusion

In this work, we developed a topic model with dynamic and static structures. We confirmed the superiority of the proposed model to the conventional topic models in terms of perplexity and analyzed the topic structures of a collection of papers. Possible future directions of research include automatic inference of the number of topics and application to topic trend analysis in various domains.

Acknowledgments

We thank Chun Fui Liew for sharing his dataset with us and advice.

References

- Amr Ahmed and Eric P. Xing. 2010. [Timeline: A dynamic hierarchical Dirichlet process model for recovering birth/death and evolution of topics in text stream](#). In *Proceedings of the 26th Conference on Uncertainty in Artificial Intelligence*, pages 20–29.
- Christophe Andrieu, Nando de Freitas, Arnaud Doucet, and Michael I. Jordan. 2003. [An introduction to MCMC for machine learning](#). *Machine Learning*, 50(1):5–43.
- David M. Blei and John D. Lafferty. 2006. [Dynamic topic models](#). In *Proceedings of the 23rd International Conference on Machine Learning*, pages 113–120.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent Dirichlet allocation](#). *Journal of Machine Learning Research*, 3(Jan):993–1022.
- Lan Du, Wray Buntine, and Huidong Jin. 2010. [A segmented topic model based on the two-parameter Poisson-Dirichlet process](#). *Machine Learning*, 81(1):5–19.
- Thomas L. Griffiths and Mark Steyvers. 2004. [Finding scientific topics](#). In *Proceedings of the National Academy of Sciences of the United States of America*, volume 101, pages 5228–5235.
- Tomoharu Iwata, Takeshi Yamada, Yasushi Sakurai, and Naonori Ueda. 2010. [Online multiscale dynamic topic models](#). In *Proceedings of the 16th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 663–672.
- John D. Lafferty and David M. Blei. 2006. [Correlated topic models](#). In *Advances in Neural Information Processing Systems*, volume 18, pages 147–154.
- Wei Li and Andrew McCallum. 2006. [Pachinko allocation: DAG-structured mixture models of topic correlations](#). In *Proceedings of the 23rd International Conference on Machine Learning*, pages 577–584.
- Chun Fui Liew, Danielle DeLatte, Naoya Takeishi, and Takehisa Yairi. 2017. [Recent developments in aerial robotics: A survey and prototypes overview](#). arXiv:1711.10085.
- Dahua Lin, Eric Grimson, and John W. Fisher. 2010. [Construction of dependent Dirichlet processes based on Poisson processes](#). In *Advances in Neural Information Processing Systems*, volume 23, pages 1396–1404.
- Thomas Minka. 2000. Estimating a Dirichlet distribution. Technical report, MIT.
- Valerio Perrone, Paul A. Jenkins, Dario Spano, and Yee Whye Teh. 2016. [Poisson random fields for dynamic feature models](#). arXiv:1611.07460.
- Jun Song, Yu Huang, Xiang Qi, Yuheng Li, Feng Li, Kun Fu, and Tinglei Huang. 2016. [Discovering hierarchical topic evolution in time-stamped documents](#). *Journal of the Association for Information Science and Technology*, 67(4):915–927.
- Xuerui Wang and Andrew McCallum. 2006. [Topics over time: A non-Markov continuous-time model of topical trends](#). In *Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 424–433.