

Tracking Brand-Associated Polarity-Bearing Topics in User Reviews

Runcong Zhao^{1,2}, Lin Gui¹, Hanqi Yan², Yulan He^{1,2,3}

¹King's College London, United Kingdom, ²University of Warwick, United Kingdom,

³The Alan Turing Institute, United Kingdom

runcong.zhao@warwick.ac.uk, yulan.he@kcl.ac.uk

Abstract

Monitoring online customer reviews is important for business organizations to measure customer satisfaction and better manage their reputations. In this paper, we propose a novel dynamic Brand-Topic Model (dBTM) which is able to automatically detect and track brand-associated sentiment scores and polarity-bearing topics from product reviews organized in temporally ordered time intervals. dBTM models the evolution of the latent brand polarity scores and the topic-word distributions over time by Gaussian state space models. It also incorporates a meta learning strategy to control the update of the topic-word distribution in each time interval in order to ensure smooth topic transitions and better brand score predictions. It has been evaluated on a dataset constructed from MakeupAlley reviews and a hotel review dataset. Experimental results show that dBTM outperforms a number of competitive baselines in brand ranking, achieving a good balance of topic coherence and uniqueness, and extracting well-separated polarity-bearing topics across time intervals.¹

1 Introduction

With the increasing popularity of social media platforms, customers tend to share their personal experience towards products online. Tracking customer reviews online could help business organizations to measure customer satisfaction and better manage their reputations. Monitoring brand-associated topic changes in reviews can be done through the use of dynamic topic models (Blei and Lafferty, 2006; Wang et al., 2008; Dieng et al., 2019). Approaches such as the dynamic Joint Sentiment-Topic (dJST) model (He et al., 2014) are able to extract polarity-bearing topics evolved over time by assuming the dependency of the sentiment-topic-word distributions across time

slices. They require the incorporation of word prior polarity information, however, and assume topics are associated with discrete polarity categories. Furthermore, they are not able to infer brand polarity scores directly.

A recently proposed Brand-Topic Model (BTM) (Zhao et al., 2021) is able to automatically infer real-valued brand-associated sentiment scores from reviews and generate a set of sentiment-topics by gradually varying its associated sentiment scores from negative to positive. This allows users to detect, for example, strongly positive topics or slightly negative topics. BTM, however, assumes all documents are available prior to model learning and cannot track topic evolution and brand polarity changes over time.

In this paper, we propose a novel framework inspired by Meta-Learning, which is widely used for distribution adaptation tasks (Suo et al., 2020). When training the model on temporally ordered documents divided into time slice, we assume that extracting polarity-bearing topics and inferring brand polarity scores in each time slice can be treated as a new sub-task and the goal of model learning is to learn to adapt the topic-word distributions associated with different brand polarity scores in a new time slice. We use BTM as the base model and store the parameters learned in a memory. At each time slice, we gauge model performance on a validation set based on the model-generated brand ranking results. The evaluation results are used for early stopping and dynamically initializing model parameters in the next time slice with meta learning. The resulting model is called dynamic Brand Topic Modeling (dBTM).

The final outcome from dBTM is illustrated in Figure 1, in which it can simultaneously track topic evolution and infer latent brand polarity score changes over time. Moreover, it also enables the generation of fine-grained polarity-bearing topics in each time slice by gradually varying brand polarity scores. In essence, we can observe topic

¹Data and code are available at <https://github.com/BLPXSPG/dBTM>.

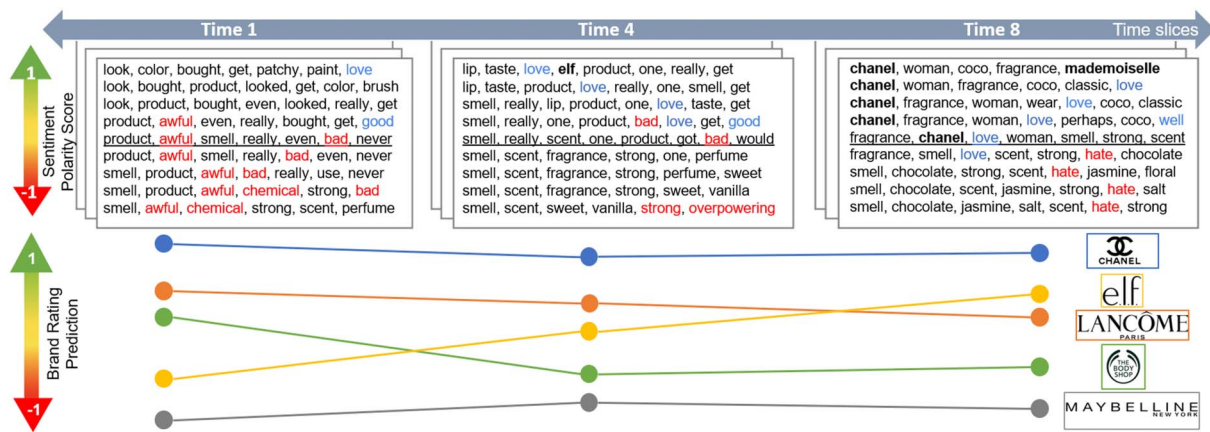


Figure 1: Brand-associated polarity-bearing topics tracking by our proposed model. We show top words from an example topic extracted in time slice 1, 4, and 8 along the horizontal axis. In each time slice, we can see a set of topics generated by gradually varying their associated sentiment scores from -1 (negative) to 1 (positive) along the vertical axis. For easy inspection, positive words are highlighted in blue while negative ones in red. We can observe in Time 1, negative topics are mainly centred on the complaint of the chemical smell of a perfume, while positive topics are about the praise of the look of a product. From Time 1 to Time 8, we can also see the evolving aspects in negative topics moving from complaining about the strong chemical of perfume to overpowering sweet scent. In the lower part of the figure, we show the inferred polarity scores of three brands. For example, Chanel is generally ranked higher than Lancôme, which in turn scores higher than The Body Shop.

transitions in two dimensions, either along a discrete time dimension, or along a continuous brand polarity score dimension.

We have evaluated dBTM on a review dataset constructed from MakeupAlley,² consisting of over 611K reviews spanning over 9 years, and a hotel review dataset sampled from HotelRec (Antognini and Faltings, 2020), containing reviews of the most popular 25 hotels over 7 years. We compare its performance with a number of competitive baselines and observe that it generates better brand ranking results, predicts more accurate brand score time series, and produces well-separated polarity-bearing topics with more balanced topic coherence and diversity. More interestingly, we have evaluated dBTM in a more difficult setup, where the supervised label information, that is, review ratings, is only supplied in the first time slice, and afterwards, dBTM is trained in an unsupervised way without the use of review ratings. dBTM under such a setting can still produce brand ranking results across time slices more accurately compared to baselines trained under the supervised setting. This is a desirable property as dBTM, initially trained on a small set of labeled data, can self-adapt its parameters with streaming data in an unsupervised way.

²<https://www.makeupalley.com/>.

Our contributions are three-fold:

- We propose a new model, called dBTM, built on the Gaussian state space model with meta learning for dynamic brand topic and polarity score tracking;
- We develop a novel meta learning strategy to dynamically initialize the model parameters at each time slice in order to better capture rating score changes, which in turn generates topics with a better overall quality;
- Our experimental results show that dBTM trained with the supervision of review ratings at the initial time slice can self-adapt its parameters with streaming data in an unsupervised way and yet still achieve better brand ranking results compared to supervised baselines.

2 Related Work

Our work is related to the following research:

2.1 Dynamic Topic Models

Topic models such as the Latent Dirichlet Allocation (LDA) model (Blei et al., 2003) is one of the most successful approaches for the statistical analysis of document collections. Dynamic topic models aim to analyse the temporal evolution of

topics in large document collections over time. Early approaches built on LDA include the dynamic topic model (DTM) (Blei and Lafferty, 2006), which uses the Kalman filter to model the transition of topics across time, and the continuous time dynamic topic model (Wang et al., 2008), which replaced the discrete state space model of the DTM with its continuous generalization. More recently, DTM has been combined with word embeddings in order to generate more diverse and coherent topics in document streams (Dieng et al., 2019).

Apart from the commonly used LDA, Poisson factorization can also be used for topic modeling, in which it factorizes a document-word count matrix into a product of a document-topic matrix and a topic-word matrix. It can be extended to analyse sequential count vectors such as a document corpus which contains a single word count matrix with one column per time interval, by capturing dependence among time steps by a Kalman filter (Charlin et al., 2015), neural networks (Gong and Huang, 2017), or by extending a Poisson distribution on the document-word counts as a non-homogeneous Poisson process over time (Hosseini et al., 2018).

While the aforementioned models are typically used in the unsupervised setting, the Joint Sentiment-Topic (JST) model (Lin and He, 2009; Lin et al., 2012) incorporated the polarity word prior into model learning, which enables the extraction of topics grouped under different sentiment categories. JST is later extended into a dynamic counterpart, called dJST, which tracks both topic and sentiment shifts over time (He et al., 2014) by assuming that the sentiment-topic word distribution at the current time is generated from the Dirichlet distribution parameterised by the sentiment-topic word distributions at previous time intervals.

2.2 Market/Brand Topic Analysis

LDA and its variants have been explored for marketing research. Examples include user interests detection by analyzing consumer purchase behavior (Gao et al., 2017; Sun et al., 2021), the tracking of the competitors in the luxury market among given brands by mining the Twitter data (Zhang et al., 2015), and identify emerging app issues from user reviews (Yang et al., 2021). Matrix factorization, which is able to extract the global information, is also used to be applied in

product recommendation (Zhou et al., 2020) and review summarization (Cui and Hu, 2021). The interaction between topics and polarities can be modeled by the incorporation of approximations by sampling based methods (Lin and He, 2009) with sentiment prior knowledge such as sentiment lexicon (Lin et al., 2012). But such prior knowledge would be highly domain-specific. Seed words with known polarities or seed words generated by morphological information (Brody and Elhadad, 2010) is another common method to obtain topic polarity. But those methods are focused on analyzing the polarity of existing topics. More recently, the Brand-Topic Model built on Poisson factorization was proposed (Zhao et al., 2021), which can infer brand polarity scores and generate fine-grained polarity-bearing topics. The detailed description of BTM can be found in Section 3.

2.3 Meta Learning

Meta learning, or learning to learn, can be broadly categorized into metric-based learning and optimization-based learning. Metric-based learning aims to learn a distance function between training instances so that it can classify a test instance by comparing it with the training instances in the learned embedding space (Sung et al., 2018). Optimization-based learning usually splits the labeled samples into *training* and *validation* sets. The basic idea is to fine-tune the parameters on the *training* set to obtain the updated parameters. These are then evaluated on the *validation* set to get the error, which is converted as a loss value for optimizing the original parameters (Finn et al., 2017; Jamal and Qi, 2019). Meta learning has been explored in many tasks, including text classification (Geng et al., 2020), topic modeling (Song et al., 2020), knowledge representation (Zheng et al., 2021), recommender systems (Neupane et al., 2021; Dong et al., 2020; Lu et al., 2020), and event detection (Deng et al., 2020). Especially, the meta learning based methods have achieved significant successes in distribution adaptation (Suo et al., 2020; Yu et al., 2021). We propose a meta learning strategy here to learn how to automatically initialize model parameters in each time slice.

3 Preliminary: Brand Topic Model

The Brand-Topic Model (BTM) (Zhao et al., 2021), as shown in the middle part of Figure 2,




Figure 2: The overall architecture of the dynamic Brand-Topic Model (dBTM), which extends the Brand-Topic Model (BTM) shown in the upper box to deal with streaming documents. In particular, at time slice t , the document-topic distribution θ^t is initialized by a vanilla Poisson factorization model, the evolution of the latent brand-associated polarity scores x^t and the polarity-associated topic-word offset η^t is modeled by two separate Gaussian state space models. The topic-word distribution β^t has its prior set based on the trend of the model performance on brand ranking results in the previous two time slices. Lines colored in gray indicate parameters are linked by Gaussian state space models, while those colored in green indicate forward calculations.

is trained on review documents paired with their document-level sentiment class labels (e.g., ‘Positive’, ‘Negative’, and ‘Neutral’). It can automatically infer real-valued brand-associated polarity scores and generate fine-grained sentiment-topics in which a continuous change of words under a certain topic can be observed with a gradual change of its associated sentiment. It was partly inspired by the Text-Based Ideal Point (TBIP) model (Vafa et al., 2020), which aims to model the generation of text via Poisson factorization. In particular, for the input bag of words data, the count for term v in document d is formulated as term count c_{dv} , which is assumed to be sampled

from a Poisson distribution $c_{dv} \sim \text{Poisson}(\lambda_{dv})$ where the rate parameter λ_{dv} can be factorized as:

$$\lambda_{dv} = \sum_k \theta_{dk} \beta_{kv} \quad (1)$$

Here, θ_{dk} denotes the per-document topic intensity, and β_{kv} represents the topic-word distribution. We have $\theta \in \mathbb{R}_+^{D \times K}$, $\beta \in \mathbb{R}_+^{K \times V}$, where D is the total number of documents in a corpus, K is the topic number, and V is the vocabulary size. Then, brand-polarity score x_{b_d} and topic-word offset η_{kv} are added to the model:

$$\lambda_{dv} = \sum_k \theta_{dk} \exp(\log \beta_{kv} + x_{b_d} \eta_{kv}) \quad (2)$$

x_{b_d} is the brand polarity score for document d of brand b and we have $\eta \in \mathbb{R}_+^{K \times V}$, $x \in \mathbb{R}$. The model-normalized brand polarity assignment to $[-1, 1]$ in its output for demonstration purposes.

The intuition behind the above formulation is that the latent variable x_{b_d} , which captures the brand polarity score, can be either positive or negative. If a word tends to frequently occur in reviews with positive polarities, but the polarity score of the current brand is negative, then the occurrence count of such a word would be reduced by making x_{b_d} and η_{kv} to have opposite signs.

A Gamma prior is placed on θ and β , with a, b, c, d being hyperparameters, while a normal prior is placed over the brand polarity score x and the topic-word count offset η .

$$\theta_{dk} \sim \text{Gamma}(a, b), \quad \beta_{kv} \sim \text{Gamma}(c, d), \\ x_{b_d} \sim \mathcal{N}(0, 1), \quad \eta_{kv} \sim \mathcal{N}(0, \mathbf{I})$$

BTM makes use of Gumbel-Softmax (Jang et al., 2017) to construct document features for sentiment classification. This is because directly sampling word counts from the Poisson distribution is not differentiable. Gumbel-Softmax, which is a gradient estimator with the reparameterization trick, is used to enable back-propagation of gradients. More details can be found in Zhao et al. (2021).

4 Dynamic Brand Topic Model (dBTM)

To track brand-associated topic dynamics in customer reviews, we split the documents into time slices where the time period of each slice can be set arbitrarily at, for example, a week, a month, or

a year. In each time slice, we have a stream of M documents $\{d_1, \dots, d_M\}$ ordered by their publication timestamps. A document d at time slice t is input as a Bag-of-Words representation. We extend BTM to deal with streaming documents by assuming that documents at the current time slice are influenced by documents at past. The resulting model is called dynamic Brand-Topic Model (dBTM), with its architecture illustrated in Figure 2.

4.1 Initialization

In the original BTM model, the latent variables to be inferred include the document-topic distribution θ , topic-word distribution β , the brand-associated polarity score x , and the polarity-associated topic-word offset η . At time slice 0, we represent all documents in this slice as a document-word count matrix. We then perform Poisson factorization with coordinate-ascent variational inference (Gopalan et al., 2015) to derive θ and β (see Eq. (1)). The topic-word count offset η and the brand polarity score x are sampled from a standard normal distribution.

4.2 State-Space Model

At time slice t , we can model the evolution of the latent brand-associated polarity scores x^t and the polarity-associated topic-word offset η^t over time by a Gaussian state space model:

$$x^t | x^{t-1} \sim \mathcal{N}(x^{t-1}, \sigma_x^2 \mathbf{I}) \quad (3)$$

$$\eta^t | \eta^{t-1} \sim \mathcal{N}(\eta^{t-1}, \sigma_\eta^2 \mathbf{I}) \quad (4)$$

For the topic-word distribution β , a similar Gaussian state-space model is adopted except that log-normal distribution is used:

$$\beta^t | \beta^{t-1} \sim \mathcal{LN}(\beta^{t-1}, \sigma_\beta^2 \mathbf{I}) \quad (5)$$

While topic-word distribution could be inherited from a previous time slice, the document-topic distribution θ^t needs to be re-initialized at the start of each time slice since there is a different set of documents at each time slice. We propose to run a simple Poisson factorization to derive the initial values of $\theta_{(p)}^t$ before we do the model adaption at each time slice: Here, the topic-word distribution in the previous time slice $\beta_{(p)}^{t-1}$ becomes the prior of the topic-word distribution in the current time slice $\beta_{(p)}^t$ as defined in Eq. (5).

We use the subscript (p) to denote that the parameters are derived in the Poisson factorization initialization stage at the start of each time slice.

Essentially, at each time slice t , we initialize the document-topic distribution θ^t of the BTM model as $\theta_{(p)}^t$, which is obtained by performing Poisson factorization on the document-word count matrix in t . For the topic-word distribution, within BTM, we can set β^t to be inherited from β^{t-1} as defined in Eq. (5), but additionally, we also have $\beta_{(p)}^t$, which is obtained by directly performing Poisson factorization of the document-word count matrix in the current time slice. In what follows, we will present how we initialize the value of β^t through meta learning.

4.3 Meta Learning

We notice that although parameters in each time interval are linked with parameters in the previous time interval by Gaussian state-space models, the results generated at each time interval are not stable. Inspired by meta learning, we consider latent brand score prediction and sentiment topic extraction at each time interval as a new sub-task, and propose a learning strategy to dynamically initialize model parameters in each interval based on the brand rating prediction performance on the validation set of the previous interval. In particular, we set aside 10% of the training data in each time interval as the validation set and compare the model-inferred brand ranking result with the gold standard one using the Spearman's rank correlation coefficient. By default, the topic-word distribution in the current interval, β^t , would have its prior set to β^{t-1} learned in the previous time interval. However, if the brand ranking result in the previous interval is poor, then β^t would be initialized as a weighted interpolation of β^{t-1} and the topic-word distribution obtained from the Poisson factorization initialization stage in the current interval $\beta_{(p)}^t$. The rationale is that if the model performs poorly in the previous interval, then its learned topic-word distribution should have less impact on the parameters in the current interval. More concretely, we first evaluate the brand ranking result returned by the model at time slice $t - 1$ on the validation set at $t - 1$:

$$\rho^{t-1} = \text{SpearmanRank}(\hat{r}^{t-1}, r^{t-1}) \quad (6)$$

where \hat{r}^{t-1} denotes the derived brand ranking result based on the model predicted latent brand

polarity scores, $\hat{\mathbf{x}}^{t-1}$, at time slice $t - 1$, \mathbf{r}^{t-1} is the gold-standard brand ranking, and ρ^{t-1} is the Spearman’s rank correlation coefficient. To check if the brand ranking result gets worse or not, we compare it with the brand ranking evaluation result, ρ^{t-2} , in the earlier interval. In particular, we first take Fisher’s z-transformation $z_{\rho^{t-1}}$ of ρ^{t-1} , which is assumed following a Gaussian distribution:

$$z_{\rho^{t-1}} \sim \mathcal{N}\left(\ln\left(\frac{1 + \rho^{t-1}}{1 - \rho^{t-1}}\right)^{0.5}, \frac{1}{B - 3}\right) \quad (7)$$

where B denotes the total number of brands. Then we compute the Cumulative Distribution Function (CDF) of the above normal distribution, denoted as $\Phi_{z_{\rho^{t-1}}}$, and calculate $\Phi_{z_{\rho^{t-1}}}(\rho^{t-2})$, which essentially returns $\Pr(z_{\rho^{t-1}} \leq \rho^{t-2})$. Lower value of $\Phi_{z_{\rho^{t-1}}}(\rho^{t-2})$ indicates that the model at $t - 1$ generates a better brand rank result than that in the previous time slice $t - 2$. This is equivalent to performing a hypothesis test in which we compare the rank evaluation result ρ^{t-1} with ρ^{t-2} to test if the model at $t - 1$ performs better than that at $t - 2$. The hypothesis testing result can be used to set the weight γ^t to determine how to initialize the topic-word distribution at t , β^t :

$$\gamma^t = \max(0.05, \Phi_{z_{\rho^{t-1}}}(\rho^{t-2})) \quad (8)$$

$$\beta^t = (1 - \gamma^t)\beta^{t-1} + \gamma^t\beta_{(p)}^t \quad (9)$$

The above equations state that if the model trained at $t - 1$ generates a better brand ranking result than that in the previous time slice significantly (p -value > 0.05), then we are more confident to initialize β^t largely based on β^{t-1} according to the estimated probability of $\Pr(z_{\rho^{t-1}} > \rho^{t-2}) = 1 - \gamma^t$. Otherwise, we will have to re-initialize β^t mostly based on the topic-word distribution obtained from the Poisson factorization initialization stage in the current interval $\beta_{(p)}^t$.

4.4 Parameter Inference

We use the mean-field variational distribution to approximate the posterior distribution of latent variables, $\theta, \beta, \eta, \mathbf{x}$, given the observed document-word count data \mathbf{c} by maximizing the Evidence Lower-Bound (ELBO):

$$\mathcal{L}_{\text{ELBO}} = \mathbb{E}_{q_\phi}[\log p(\theta, \beta, \eta, \mathbf{x})] + \log p(\mathbf{c}|\theta, \beta, \eta, \mathbf{x}) - \log q_\phi(\theta, \beta, \eta, \mathbf{x}) \quad (10)$$

where

$$q_\phi(\theta, \beta, \eta, \mathbf{x}) = \prod_{d,k,b} q(\theta_d)q(\beta_k)q(\eta_k)q(x_b) \quad (11)$$

In addition, for each document d , we construct its representation z_d by sampling word counts using Gumbel softmax from the aforementioned learned parameters, which is fed to a sentiment classifier to predict a class distribution \hat{y}_d . We also perform adversarial learning by inverting the sign of the inferred polarity score of the brand associated with document d and produce the adversarial representation \tilde{z}_d . This is also fed to the same sentiment classifier which generates another predicted class distribution \tilde{y}_d . We train the model by minimizing the Wasserstein distance between the prediction and the actual class distributions. The final loss function is the combination of the ELBO and the Wasserstein distance losses:

$$\mathcal{L} = -\mathcal{L}_{\text{ELBO}} + \frac{1}{M} \sum_{d=1}^M (\mathcal{L}_{\text{WD}}(\hat{y}_d, y_d) + \mathcal{L}_{\text{WD}}(\tilde{y}_d, \bar{y}_d)) \quad (12)$$

where $\mathcal{L}_{\text{WD}}(\cdot)$ denotes the Wasserstein distance, y_d is the gold-standard class distribution, and \bar{y}_d is the class distribution derived from the inverted document rating. By inverting the document rating, we essentially balance the document rating distributions that for each positive document, we also create a synthetic negative document, and vice versa.

5 Experimental Setup

Datasets Popular datasets such as Yelp and Amazon products (Ni et al., 2019) and Multi-Domain Sentiment dataset (Blitzer et al., 2007) are constructed by randomly selecting reviews from Amazon or Yelp without considering their distributions over various brands and across different time periods. Therefore, we construct our own dataset by crawling reviews from top 25 brands from MakeupAlley, a review website on beauty products. Each review is accompanied with a rating score, product type, brand, and post time. We consider reviews with the ratings of 1 and 2 as the negative class, those with the rating of 3 as the neutral class, and the remaining with the ratings of 4 and 5 as the positive class, following the label setting in BTM. The entire dataset contains

Algorithm 1: Training procedure of dBTM

Input : Number of topics K , number of brands B , time slice $t \in \{0, 1, 2, \dots, T\}$, a stream of document-word count matrix $\mathcal{C} = \{c^0, \dots, c^T\}$

Output: Document-topic intensity $\{\theta^t\}_{t=0}^T$, Topic-word matrix $\{\beta^t\}_{t=0}^T$, brand scores $\{x_1^t, \dots, x_B^t\}_{t=0}^T$

- 1 **Initialization:**
- 2 Initialize θ^0, β^0 by Poisson factorization $c^0 \sim \text{Poisson}(\theta^0 \beta^0)$, $s\mathbf{x}^0 \sim \mathcal{N}(0, \mathbf{I})$, $\boldsymbol{\eta}^0 \sim \mathcal{N}(0, \Sigma_{\boldsymbol{\eta}}^2 \mathbf{I})$, $\gamma^0 = 0$
- 3 Update parameters by minimizing the loss defined in Eq. (12)
- 4 Derive the brand ranking $\hat{\mathbf{r}}^0$ on the validation set based on the inferred brand polarity score $\{x_b^0\}_{b=1}^B$
- 5 Calculate the Spearman's rank correlation coefficient $\rho^0 = \text{SpearmanRank}(\hat{\mathbf{r}}^0, \mathbf{r}^0)$
- 6 Set the weight $\gamma^1 = \max(0.05, \Phi_{z_{\rho^0}}(0))$
- 7 **Training:**
- 8 **for** $t = 1$ **to** T **do**
- 9 **Pre-training:** $c^t \sim \text{Poisson}(\theta_{(p)}^t \beta_{(p)}^t)$
- 10 **Per epoch initialization:**
- 11 $\mathbf{x}^t \sim \mathcal{N}(\mathbf{x}^{t-1}, \sigma_x^2 \mathbf{I})$, $\boldsymbol{\eta}^t \sim \mathcal{N}(\boldsymbol{\eta}^{t-1}, \sigma_{\boldsymbol{\eta}}^2 \mathbf{I})$, $\theta^t = \theta_{(p)}^t$, $\beta^t = (1 - \gamma^t) \beta^{t-1} + \gamma^t \beta_{(p)}^t$
- 12 **for** $i = 0$ **to** *maximum iterations* **do**
- 13 Update parameters by minimizing the loss defined in Eq. (12)
- 14 **if** *checkpoint* **then**
- 15 Derive the brand ranking $\hat{\mathbf{r}}^t$ on the validation set based on the inferred brand polarity score $\{x_b^t\}_{b=1}^B$
- 16 Calculate the Spearman's rank correlation coefficient $\rho^t = \text{SpearmanRank}(\hat{\mathbf{r}}^t, \mathbf{r}^t)$
- 17 **if** $(\Phi_{z_{\rho^t}}(\rho^{t-1}) > 0.95)$ **then**
- 18 | break # *Null hypothesis is rejected by the upper quartile according to Eq. (7)*
- 19 **end**
- 20 Set the weight $\gamma^{t+1} = \max(0.05, \Phi_{z_{\rho^t}}(\rho^{t-1}))$
- 21 **end**
- 22 **end**
- 23 **end**

Dataset	MakeupAlley-Beauty Reviews
No. of documents per class	
Neg / Neu / Pos	114,837 / 88,710 / 407,581
No. of brands	25
Total no. of documents	611,128
No. of time Slices	9
Average review length (#words)	123
Average no. of documents per slice	~ 68k
Vocabulary size	~ 4500
Dataset	HotelRec Reviews
No. of documents per class	
Neg / Neu / Pos	14,600 / 20,629 / 150,265
No. of hotels	25
Total no. of documents	185,496
No. of time Slices	7
Average review length (#words)	204
Average no. of documents per slice	~ 26k
Vocabulary size	~ 7000

Table 1: Dataset statistics of the reviews.

611,128 reviews spanning over 9 years (2005 to 2013). We treat each year as a time slice and split reviews into 9 time slices. The average review length is 123 words. Besides the MakeupAlley-Beauty, we also run our experiments on HotelRec (Antognini and Faltings, 2020), by selecting reviews from the top 25 hotels over 7 years (2012 to 2018). The statistics of our datasets are shown in Table 1. It can be observed that the dataset is imbalanced with positive reviews being over

triple the size of negative ones for MakeupAlley-Beauty and nearly 10 times for HotelRec.

Models for Comparison We conduct experiments using the following models:

- **Dynamic Joint Sentiment-Topic (dJST) model** (He et al., 2014), built on LDA, can detect and track polarity-bearing topics from text with the word prior sentiment knowledge incorporated. In our experiments, the MPQA subjectivity lexicon³ is used to derive the word prior sentiment information.
- **Text-Based Ideal Point (TBIP)** (Vafa et al., 2020), an unsupervised Poisson factorization model which can infer latent brand sentiment scores.
- **Brand Topic Model (BTM)** (Zhao et al., 2021), a supervised Poisson factorization model extended from TBIP with the incorporation of document-level sentiment labels.
- **dBTM**, our proposed dynamic Brand Topic model in which the model is trained with the document-level sentiment labels at each time slice.

³<https://mpqa.cs.pitt.edu/lexicons/>.

Time Slice	dJST		TBIP		BTM		O-dBTM		dBTM	
	Corr	p-value	Corr	p-value	Corr	p-value	Corr	p-value	Corr	p-value
<i>MakeupAlley-Beauty</i>										
1	-0.249	0.230	-0.567	0.003	0.552	0.004	0.454	0.023	0.402	0.046
2	-0.437	0.029	0.527	0.007	0.488	0.013	0.459	0.021	0.438	0.029
3	-0.327	0.111	-0.543	0.005	-0.384	0.058	0.504	0.010	0.523	0.007
4	-0.127	0.545	-0.431	0.032	-0.428	0.033	0.448	0.025	0.453	0.023
5	0.112	0.596	-0.347	0.089	0.402	0.047	0.438	0.028	0.394	0.051
6	-0.118	0.573	-0.392	0.053	0.432	0.031	0.402	0.047	0.433	0.031
7	-0.203	0.330	0.400	0.048	0.417	0.038	0.400	0.048	0.402	0.047
8	-0.552	0.004	0.348	0.089	0.363	0.074	0.359	0.078	0.364	0.074
<i>HotelRec</i>										
1	0.097	0.645	0.121	0.565	-0.508	0.009	0.356	0.081	0.285	0.168
2	-0.242	0.244	0.443	0.027	-0.337	0.100	0.196	0.347	0.382	0.059
3	-0.112	0.596	-0.392	0.053	0.318	0.121	0.419	0.037	0.355	0.082
4	-0.362	0.076	0.276	0.181	0.301	0.144	0.349	0.087	0.315	0.126
5	-0.045	0.829	0.292	0.156	0.225	0.279	0.323	0.115	0.364	0.074
6	0.222	0.285	0.298	0.148	0.306	0.137	0.294	0.154	0.312	0.130

Table 2: Brand ranking results generated by various models trained on time slice t and tested on time slice $t + 1$. We report the correlation coefficients $corr$ and its associated two-sided p -values.

- **O-dBTM**, a variant of our model that is only trained with the supervised review-level sentiment labels in the first time slice (denoted as the 0-th time slice). In the subsequent time slices, it is trained under the unsupervised setting. In such a case, we no longer have a gold-standard brand ranking in time slices other than the 0-th one. Instead of directly calculating the Spearman’s rank correlation coefficient, we measure the difference of the brand ranking results in neighboring time slices and use it to set the weight γ^t in Eq. (8).

Parameter Setting Frequent bigrams and trigrams⁴ are added as features in addition to unigrams for document representations. In our experiments, we train the models using the data from the current time slice and test the model performance on the full data from the next time slice. During training, we set aside 10% of data in each time slice as the validation set. For hyperparameters, we set the batch size to 256, the maximum training steps to 50,000, the topic number to 50.⁵ It is worth noting that since topic dynamics are

⁴Frequent but less informative n -grams such as ‘*actually bought*’ were filtered out using NLTK.

⁵The topic number is set empirically based on the validation set in the 0-th time slice.

not explicitly modeled in the static models such as TBIP and BTM, their topics extracted in different time slices are not directly linked.

6 Experimental Results

In this section, we present the experimental results in comparison with the baseline models in brand rating, topic coherence/uniqueness measures, and qualitative evaluation of generated topics. For fair comparison, baselines are trained based on all previous time slices and predict on the current time slice.

6.1 Brand Rating

TBIP, BTM, and dBTM can infer each brand’s associated polarity score automatically. For dJST, we derive the brand rating by aggregating the label distribution of its associated review documents and then normalizing over the total number of brand-related reviews. The average of the document-level ratings of a brand b at a time slice t is used as the ground truth of the brand rating x_b^t . We evaluate two aspects of the brand ratings:

Brand Ranking Results We report in Table 2 the brand ranking results measured by the Spearman’s correlation coefficient, showing the correlation of predicted brand rating and the ground

Maybelline New York

— Real Score




Figure 3: The rating time series for ‘Maybelline New York’. The rating scores are normalized in the range of $[-1, 1]$ with positive values denoting positive sentiment and negative ones for negative sentiment. In each subfigure, the dashed curve shows the actual rating scores.

truth, along with the associated two-sided p -values of the Spearman’s correlations

Topic model variants, such as dJST, TBIP, and BTM, produced brand ranking results either positively or negatively correlated with the true ranking results. We can see the correlation of BTM has switched between positive correlated and negative rated between time slices. With Gaussian state space models, our proposed model dBTM and its variant O-dBTM generate more stable results. On MakeupAlley-Beauty, dBTM gives the best results in 4 out of 8 time slices. Interestingly, O-dBTM with the supervised information supplied in only the first time slice outperforms the static models such as BTM in 3 out of 8 time slices, showing the effectiveness of our proposed architecture in tracking brand score dynamics. Similar conclusions can be drawn on HotelRec that O-dBTM gives superior performance compared to BTM on 5 out of 6 time slices. Both O-dBTM and dBTM outperform the other baselines except TBIP in time slice 2.

In summary, in dBTM, the brand rating score is treated as a latent variable (i.e., x_{b_d} in Eq. (2)) and is directly inferred from the data. On the contrary, models such as dJST, which require post-processing to derive brand rating scores by aggregating the document-level sentiment labels, are inferior to dBTM. This shows the advantage of our proposed dBTM over traditional dynamic topic models in brand ranking.

Brand Rating Time Series The brand rating time series aims to compare the ability of models to track the trend of brand rating. For easy comparison, we normalize the ratings produced by each model, so that the plot only reflects the fluctuation of ratings over time. Figure 3 shows the brand rating on the brand ‘Maybelline New York’ generated on the test set of MakeupAlley-Beauty by various models across time slices. It can be observed that the brand ratings generated by TBIP and BTM do not correlate well with the actual rating scores. dJST shows a better aligned rating

Time Slice	dJST			TBIP			BTM			O-dBTM			dBTM		
	coh	uni	quality	coh	uni	quality	coh	uni	quality	coh	uni	quality	coh	uni	quality
<i>MakeupAlley-Beauty</i>															
1	-3.087	0.564	0.183	-3.653	0.861	0.236	-3.836	0.862	0.225	-3.486	0.820	0.235	-3.685	0.833	0.226
2	-3.008	0.513	0.170	-4.043	0.850	0.210	-3.867	0.864	0.223	-3.360	0.807	0.240	-3.642	0.829	0.228
3	-3.286	0.552	0.168	-3.949	0.843	0.214	-3.716	0.851	0.229	-3.369	0.787	0.234	-3.611	0.823	0.228
4	-3.004	0.515	0.172	-3.629	0.808	0.223	-3.837	0.846	0.220	-3.457	0.771	0.223	-3.549	0.799	0.225
5	-3.112	0.560	0.180	-4.168	0.838	0.201	-4.023	0.839	0.208	-3.412	0.793	0.232	-3.523	0.818	0.232
6	-3.139	0.542	0.173	-4.100	0.841	0.205	-3.976	0.846	0.213	-3.433	0.761	0.222	-3.577	0.814	0.228
7	-3.269	0.521	0.159	-4.049	0.854	0.211	-3.675	0.845	0.230	-3.330	0.772	0.232	-3.667	0.825	0.225
8	-3.060	0.560	0.183	-3.942	0.843	0.214	-3.715	0.837	0.225	-3.589	0.789	0.220	-3.546	0.818	0.231
Average	-3.120	0.541	0.173	-3.942	0.842	0.214	-3.831	0.849	0.222	-3.430	0.788	0.230	-3.600	0.820	0.228
<i>HotelRec</i>															
1	-3.749	0.615	0.164	-4.024	0.767	0.191	-3.935	0.851	0.216	-4.051	0.812	0.201	-3.716	0.818	0.220
2	-4.020	0.633	0.158	-3.577	0.753	0.211	-3.960	0.813	0.205	-3.851	0.803	0.209	-3.696	0.809	0.219
3	-3.667	0.593	0.162	-3.905	0.817	0.209	-4.078	0.844	0.207	-3.861	0.819	0.212	-3.854	0.820	0.213
4	-4.008	0.644	0.161	-3.747	0.808	0.216	-3.946	0.859	0.218	-3.637	0.814	0.224	-3.681	0.794	0.216
5	-3.751	0.691	0.184	-4.057	0.800	0.197	-3.953	0.823	0.208	-3.705	0.804	0.217	-3.547	0.817	0.230
6	-3.916	0.697	0.178	-3.770	0.810	0.215	-4.061	0.855	0.210	-3.510	0.800	0.228	-3.705	0.821	0.222
Average	-3.852	0.645	0.168	-3.847	0.793	0.206	-3.989	0.841	0.211	-3.769	0.809	0.215	-3.700	0.813	0.220

Table 3: Topic coherence (coh) and uniqueness (uni) measures of the results generated by various models. We also combine the two scores to derive the overall quality of the extracted topics.

trend, but its prediction missed some short-term changes such as the peak of brand rating at time slice 7. By contrast, dBTM correctly predicts the general trend of the brand rating. The weakly-supervised O-dBTM is able to follow the general trend but misses some short-term changes such as the upward trend from the time slice 1 to 2, and from the slice 6 to 7.

6.2 Topic Evaluation Results

We use the top 10 words of each topic to calculate the context-vector-based topic coherence scores (Röder et al., 2015) as well as topic uniqueness (Nan et al., 2019) which measures the ratio of word overlap across topics. We want to achieve balanced topic coherence and diversity. As such, topic coherence and topic diversity are combined to give an overall quality measure of topics (Dieng et al., 2020). Since the results for topic coherence is negative in our experiment, that is, smaller absolute values are better, we define the overall quality of a topic as $q = \frac{\text{topic uniqueness}}{|\text{topic coherence}|}$. Table 3 shows the topic evaluation results. In general, there is a trade-off between topic coherence and topic diversity. On average, dJST has the highest coherence but the lowest uniqueness scores, while TBIP has quite high uniqueness but the lowest coherence values. Both O-dBTM and dBTM achieve a good balance between coherence and

uniqueness and outperform other models in overall quality.

6.3 Example Topics across Time Periods

We illustrate some representative topics generated by dBTM in various time slices. For easy inspection, we retrieve a representative sentence from the corpus for each topic. For a sentence, we derive its representation by averaging the GloVe embeddings of its constituent words. For a topic, we also average the GloVe embeddings of its associated top words, but weighted by the topic-word probabilities. The sentence with the highest cosine similarity is selected.

Example of generated topics relating to ‘Eye Products’ and ‘Skin Care’ from MakeupAlley-Beauty is shown in Figure 4. We can observe that for the topic ‘Eye Products’, the top words of negative comments for ‘eye cleanser’ evolve from the reaction of skin (e.g., ‘sting’, ‘burned’) to the cleaning ability (e.g., ‘remove’, ‘residue’). We could also see that the positive topics gradually change from praising the ability of the product for ‘dark circle’ in time slice 1 to the quality of eye shadow in time slice 4 and eye primer in time slice 8. Moreover, we observe the brand name *M.A.C.* in the positive topic in time slice 4, which aligns with its ground truth rating. For the topic ‘Skin Care’, it can be observed that negative topics gradually move from the complaint of a skin

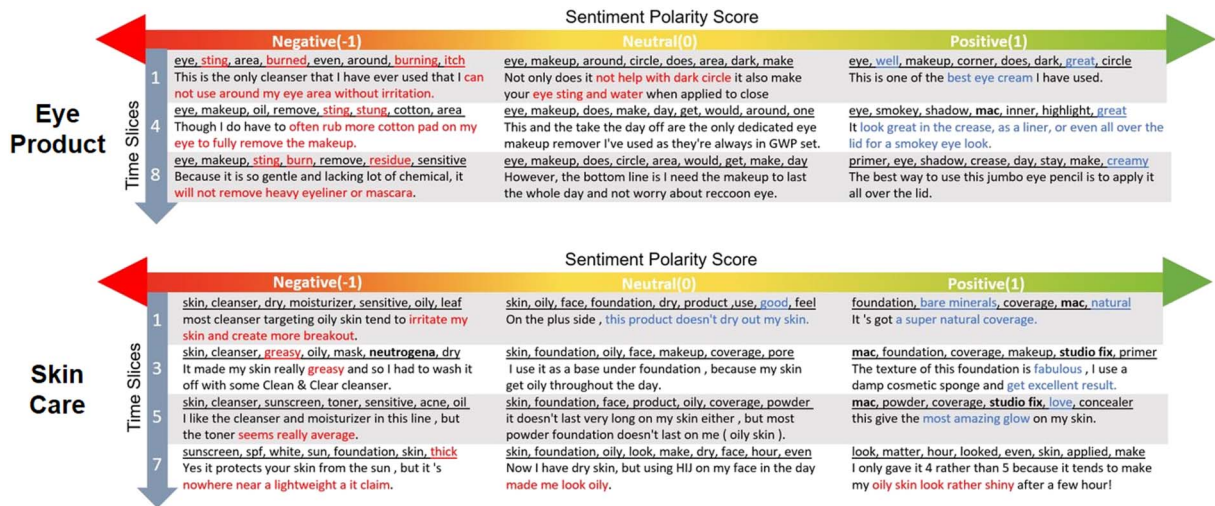


Figure 4: Example of generated topics shown as a list of top associated words (underlined) in different time slices from the MakeupAlley dataset. For easy inspection, we also show the most representative sentence under each topic. The negative, neutral and positive topics in each time slice are generated by varying the brand polarity score from -1 to 0 , and to 1 . Positive words/phrases are highlighted in blue, negative words/phrases are in red, while brand names are in bold.

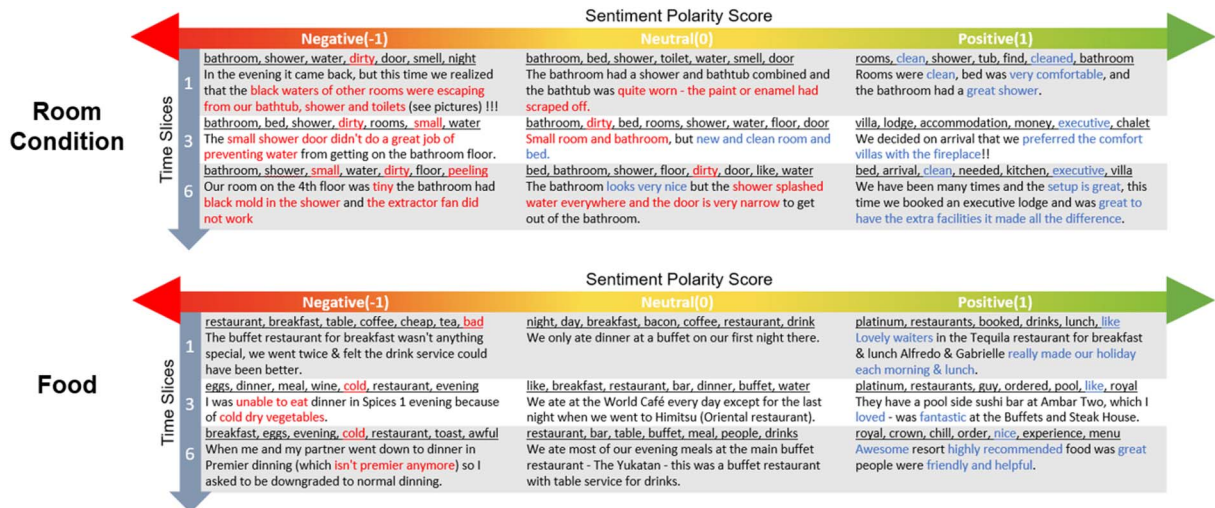


Figure 5: Example of generated topics shown as a list of top associated words (underlined) in different time slices from the HotelRec dataset. The representative sentence for each topic is also shown for easy inspection.

cleanser to the thickness of a sunscreen, while positive topics are about the praise of the coverage of the M.A.C. foundation more consistently over time. The results show that dBTM can generate well-separated polarity-bearing topics and it also allows the tracking of topic changes over time.

Example of generated topics relating to 'Room Condition' and 'Food' from HotelRec is shown in Figure 5. We can see that for the topic 'Room Condition', top words gradually shift from the expression of cleanliness (e.g., 'clean' in positive and 'dirty' in negative comments) to the description of the type and size of the rooms (e.g., 'ex-

ecutive' and 'villa' in positive reviews, and the concern of 'small' room size in negative comments). For the topic 'Food', the concerned food changes across time from drinks (e.g., 'coffee', 'tea') to meals (e.g., 'eggs', 'toast'). Negative reviews mainly focus on the concern of food quality, (e.g., 'cold'), while positive reviews contain a general praise of food and services (e.g., 'like', 'nice').

6.4 Ablation Study

We investigate the contribution of the meta learning component (i.e., Eq. (8) and (9)) by conducting an ablation study and the results are shown in

Time Slice	dBTM				dBTM (no meta learning)			
	cor	coh	uni	quality	cor	coh	uni	quality
<i>MakeupAlley-Beauty</i>								
1	0.402	-3.685	0.833	0.226	0.435	-3.972	0.861	0.217
2	0.438	-3.642	0.829	0.228	0.189	-3.704	0.840	0.227
3	0.523	-3.611	0.823	0.228	-0.162	-3.873	0.828	0.214
4	0.453	-3.549	0.799	0.225	-0.042	-3.745	0.849	0.227
5	0.394	-3.523	0.818	0.232	0.086	-3.990	0.832	0.209
6	0.433	-3.577	0.814	0.228	-0.029	-3.958	0.856	0.216
7	0.402	-3.667	0.825	0.225	-0.042	-3.587	0.842	0.235
8	0.364	-3.546	0.818	0.231	-0.125	-3.920	0.847	0.216
<i>HotelRec</i>								
1	0.285	-3.716	0.818	0.220	0.222	-3.559	0.791	0.222
2	0.382	-3.696	0.809	0.219	0.210	-3.796	0.801	0.211
3	0.355	-3.854	0.820	0.213	0.285	-3.763	0.790	0.210
4	0.315	-3.681	0.794	0.216	0.408	-3.597	0.817	0.227
5	0.364	-3.547	0.817	0.230	0.362	-3.657	0.793	0.217
6	0.312	-3.705	0.821	0.222	0.262	-3.694	0.809	0.219

Table 4: Results of dBTM with and without the meta learning component.

Table 4. We can observe that in general, removing meta learning leads to a significant reduction in brand ranking correlations across all time slices for the MakeupAlley-Beauty dataset. In terms of topic quality, we observe reduced coherence scores, but slightly increased uniqueness scores without meta learning, leading to an overall reduction of topic quality scores in most time slices.

For HotelRec, we can see that removing meta learning also leads to a reduction in brand ranking results, but the impact is smaller compared to MakeupAlley-Beauty. For topic quality, we observe increased coherence but worse uniqueness, resulting in slightly worse topic quality results without meta learning in most time slices. One main reason is that unlike makeup brands where new products are introduced over time, leading to the change of discussed topics in reviews, the topic-word distribution does not change much across different time slices for hotel reviews. Therefore, the results are less impacted with or without meta learning.

6.5 Training Time Complexity

All experiments were run on a single GeForce 1080 GPU with 11GB memory. The training time for each model across time slices is shown in Figure 6. It can be observed that with the increasing number of time slices, the training time of dJST and BTM grows quickly. Both TBIP and dBTM take significantly less time to train. TBIP simply performs Poisson factorization independently in each time slice and fails to track topic/sentiment changes over time. On the contrary, our proposed dBTM and O-dBTM are able to




Figure 6: Training time of models across time slices.

monitor topic/sentiment evolution and yet take even less time to train compared to TBIP. One main reason is that dBTM and O-dBTM can automatically adjust the number of iterations with our proposed meta learning and hence can be trained more efficiently.

7 Conclusion

We have presented dBTM, which is able to automatically detect and track brand-associated topics and sentiment scores. Experimental evaluation based on the reviews from MakeupAlley and HotelRec demonstrates the superiority of dBTM over previous models in brand ranking and dynamic topic extraction. The variant of dBTM, O-dBTM, trained with document-level sentiment labels in the first time slice only, outperforms baselines in brand ranking and achieves the best overall result in topic quality evaluation. This shows the effectiveness of the proposed architecture in modeling the evolution of brand scores and topics across time intervals.

Our model currently only considers review ratings, but real-world applications potentially involve additional factors (e.g., user preference). A possible solution is to explore simultaneous modeling of user preferences to extract personalised brand polarity topics.

Acknowledgments

This work was supported in part by the UK Engineering and Physical Sciences Research

Council (grant no. EP/T017112/1, EP/V048597/1, EP/X019063/1). YH is supported by a Turing AI Fellowship funded by the UK Research and Innovation (grant no. EP/V020579/1).

References

- Diego Antognini and Boi Faltings. 2020. Hotelrec: A novel very large-scale hotel recommendation dataset. In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 4917–4923.
- 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*, 3:993–1022.
- John Blitzer, Mark Dredze, and Fernando Pereira. 2007. Biographies, bollywood, boom-boxes and blenders: Domain adaptation for sentiment classification. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 440–447.
- Samuel Brody and Noémie Elhadad. 2010. An unsupervised aspect-sentiment model for online reviews. In *Proceedings of the 2010 Annual Conference of the North American Chapter of the Association for Computational Linguistics*, pages 804–812.
- Laurent Charlin, Rajesh Ranganath, James McInerney, and David M. Blei. 2015. Dynamic poisson factorization. In *Proceedings of the 9th ACM Conference on Recommender Systems*, pages 155–162.
- Peng Cui and Le Hu. 2021. Topic-guided abstractive multi-document summarization. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1463–1472. <https://doi.org/10.18653/v1/2021.findings-emnlp.126>
- Shumin Deng, Ningyu Zhang, Jiaojian Kang, Yichi Zhang, Wei Zhang, and Huajun Chen. 2020. Meta-learning with dynamic-memory-based prototypical network for few-shot event detection. In *Proceedings of the Thirteenth ACM International Conference on Web Search and Data Mining*, pages 151–159.
- Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. 2020. Topic modeling in embedding spaces. In *Proceedings of Transactions of the Association for Computational Linguistics*, pages 439–453. <https://doi.org/10.1162/tacla.00325>
- Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. 2019. The dynamic embedded topic model. *arXiv preprint arXiv:1907.05545*.
- Manqing Dong, Feng Yuan, Lina Yao, Xiwei Xu, and Liming Zhu. 2020. MAMO: Memory-augmented meta-optimization for cold-start recommendation. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 688–697. <https://doi.org/10.1145/3394486.3403113>
- Chelsea Finn, Pieter Abbeel, and Sergey Levine. 2017. Model-agnostic meta-learning for fast adaptation of deep networks. In *Proceedings of the 34th International Conference on Machine Learning*, pages 1126–1135.
- Li Gao, Jia Wu, Chuan Zhou, and Yue Hu. 2017. Collaborative dynamic sparse topic regression with user profile evolution for item recommendation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pages 1316–1322. <https://doi.org/10.1609/aaai.v31i1.10726>
- Ruiying Geng, Binhua Li, Yongbin Li, Jian Sun, and Xiaodan Zhu. 2020. Dynamic memory induction networks for few-shot text classification. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 1087–1094. <https://doi.org/10.18653/v1/2020.acl-main.102>
- Chengyue Gong and Win-bin Huang. 2017. Deep dynamic poisson factorization model. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 1665–1673.
- Prem Gopalan, Jake Hofman, and David Blei. 2015. Scalable recommendation with hierarchical poisson factorization. In *Proceedings of the 31th Uncertainty in Artificial Intelligence*, pages 326–335.

- Yulan He, Chenghua Lin, Wei Gao, and Kam-Fai Wong. 2014. Dynamic joint sentiment-topic model. *ACM Transactions on Intelligent Systems and Technology*, 5:1–21. <https://doi.org/10.1145/2594452>
- Seyed Abbas Hosseini, Ali Khodadadi, Keivan Alizadeh, Ali Arabzadeh, Mehrdad Farajtabar, Hongyuan Zha, and Hamid R. Rabiee. 2018. Recurrent poisson factorization for temporal recommendation. *IEEE Transactions on Knowledge and Data Engineering*, 32(1):121–134. <https://doi.org/10.1109/TKDE.2018.2879796>
- Muhammad Abdullah Jamal and Guo-Jun Qi. 2019. Task agnostic meta-learning for few-shot learning. In *Processing of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 11719–11727.
- Eric Jang, Shixiang Gu, and Ben Poole. 2017. Categorical reparameterization with gumbel-softmax. In *International Conference on Learning Representations*.
- Chenghua Lin and Yulan He. 2009. Joint sentiment/topic model for sentiment analysis. In *Proceedings of the ACM Conference on Information and Knowledge Management*, pages 375–384.
- Chenghua Lin, Yulan He, Richard Everson, and Stefan Ruder. 2012. Weakly supervised joint sentiment-topic detection from text. *IEEE Transactions on Knowledge and Data Engineering*, 24(6):1134–1145. <https://doi.org/10.1109/TKDE.2011.48>
- Yuanfu Lu, Yuan Fang, and Chuan Shi. 2020. Meta-learning on heterogeneous information networks for cold-start recommendation. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1563–1573.
- Feng Nan, Ran Ding, Ramesh Nallapati, and Bing Xiang. 2019. Topic modeling with wasserstein autoencoders. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6345–6381.
- Krishna Prasad Neupane, Ervine Zheng, and Qi Yu. 2021. Metaedl: Meta evidential learning for uncertainty-aware cold-start recommendations. In *Proceedings of the 21th IEEE International Conference on Data Mining*, pages 1258–1263. <https://doi.org/10.1109/ICDM51629.2021.00154>
- Jianmo Ni, Jiacheng Li, and Julian McAuley. 2019. Justifying recommendations using distantly-labeled reviews and fine-grained aspects. In *Proceedings of Empirical Methods in Natural Language Processing*, pages 188–197.
- Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. Exploring the space of topic coherence measures. In *Proceedings of the 8th ACM International Conference on Web Search and Data Mining*, pages 399–408.
- Yuanfeng Song, Yongxin Tong, Siqi Bao, Di Jiang, Hua Wu, and Raymond Chi-Wing Wong. 2020. Topicocoean: An ever-increasing topic model with meta-learning. In *Proceedings of the 20th IEEE International Conference on Data Mining*, pages 1262–1267. <https://doi.org/10.1109/ICDM50108.2020.00161>
- Wu-Jiu Sun, Xiao Fan Liu, and Fei Shen. 2021. Learning dynamic user interactions for online forum commenting prediction. In *Proceedings of the 21th IEEE International Conference on Data Mining*, pages 1342–1347.
- Flood Sung, Yongxin Yang, Li Zhang, Tao Xiang, Philip H. S. Torr, and Timothy M. Hospedales. 2018. Learning to compare: Relation network for few-shot learning. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1199–1208. <https://doi.org/10.1109/CVPR.2018.00131>
- Qiuling Suo, Jingyuan Chou, Weida Zhong, and Aidong Zhang. 2020. Tadanet: Task-adaptive network for graph-enriched meta-learning. In *Proceedings of the 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 1789–1799.
- Keyon Vafa, Suresh Naidu, and David M. Blei. 2020. Text-based ideal points. In *Proceedings of the Conference of the Association for Computational Linguistics*, pages 5345–5357. <https://doi.org/10.18653/v1/2020.acl-main.475>
- Chong Wang, David Blei, and David Heckerman. 2008. Continuous time dynamic topic models. In *Proceedings of the 24th Conference*

- on *Uncertainty in Artificial Intelligence*, pages 579–586.
- Tianyi Yang, Cuiyun Gao, Jingya Zang, David Lo, and Michael R. Lyu. 2021. Tour: Dynamic topic and sentiment analysis of user reviews for assisting app release. In *Proceedings of the Web Conference 2021*, pages 708–712. <https://doi.org/10.1145/3442442.3458612>
- Runsheng Yu, Yu Gong, Xu He, Yu Zhu, Qingwen Liu, Wenwu Ou, and Bo An. 2021. Personalized adaptive meta learning for cold-start user preference prediction. In *Proceedings of the 35th AAAI Conference on Artificial Intelligence*, pages 10772–10780. <https://doi.org/10.1609/aaai.v35i12.17287>
- Hao Zhang, Gunhee Kim, and Eric P. Xing. 2015. Dynamic topic modeling for monitoring market competition from online text and image data. In *Proceedings of the 21th International Conference on Knowledge Discovery and Data Mining*, pages 1425–1434. <https://doi.org/10.1145/2783258.2783293>
- Runcong Zhao, Lin Gui, Gabriele Pergola, and Yulan He. 2021. Adversarial learning of poison factorisation model for gauging brand sentiment in user reviews. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics*, pages 2341–2351.
- Wenbo Zheng, Lan Yan, Chao Gou, and Fei-Yue Wang. 2021. Knowledge is power: Hierarchical-knowledge embedded meta-learning for visual reasoning in artistic domains. In *Proceedings of the 27th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 2360–2368. <https://doi.org/10.1145/3447548.3467285>
- Kun Zhou, Yuanhang Zhou, Wayne Xin Zhao, Xiaoke Wang, and Ji-Rong Wen. 2020. Towards topic-guided conversational recommender system. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 4128–4139. <https://doi.org/10.18653/v1/2020.coling-main.365>