

Keyword-based Topic Modeling and Keyword Selection

Xingyu Wang, Lida Zhang, Diego Klabjan
Northwestern University
Evanston, Illinois

xingyuwang2017,lidazhang2018@u.northwestern.edu,d-klabjan@northwestern.edu

ABSTRACT

Certain type of documents such as tweets are collected by specifying a set of keywords. As topics of interest change with time it is beneficial to adjust keywords dynamically. The challenge is that these need to be specified ahead of knowing the forthcoming documents and the underlying topics. The future topics should mimic past topics of interest yet there should be some novelty in them. We develop a keyword-based topic model that dynamically selects a subset of keywords to be used to collect future documents. The generative process first selects keywords and then the underlying documents based on the specified keywords. The model is trained by using a variational lower bound and stochastic gradient optimization. The inference consists of finding a subset of keywords where given a subset the model predicts the underlying topic-word matrix for the unknown forthcoming documents. We compare the keyword topic model against a benchmark model using viral predictions of tweets combined with a topic model. The keyword-based topic model outperforms this sophisticated baseline model by 67%.

CCS CONCEPTS

• **Computing methodologies** → **Learning latent representations**.

KEYWORDS

topic modeling, social media

ACM Reference Format:

Xingyu Wang, Lida Zhang, Diego Klabjan. . Keyword-based Topic Modeling and Keyword Selection. In . ACM, New York, NY, USA, 12 pages.

1 INTRODUCTION

Keywords are an important means to capture most significant information in a text document. They are used in information retrieval, text summarization, social media streaming, etc. Twitter offers an Application Programming Interface (API) for streaming of tweets. The API supports many options to filter tweets based on several streaming endpoints, such as language, location, users, and keywords. Filtering by keywords is a very popular and efficient way to capture tweets based on target interest and domain. With the keyword filter, a tweet is captured if it includes any keyword from the keyword list. Many companies use the names of the company or products to monitor customer feedback, provided customer support, conduct focus groups, and engage in similar activities. However, this approach of collecting feedback has limitations. Since the keywords are selected in advance and then kept static for an extended

period of time, a useful tweet containing valuable user feedback but not including any of the selected keywords cannot be captured. For example, consider a bank using its name as the keyword to monitor online feedback about their new credit cards. If a Twitter user posts an image of a credit card as a tweet and comments, “I dislike the rewards of this credit card,” then this tweet would not be captured, even though it contains very useful information. Such cases frequently occur in replies to posted tweets or comments of retweets because the company or product names might already be mentioned in the original tweet and are not captured by the keyword. If the set of keywords is expanded by including “credit card,” then too many tweets would be captured. There is even a bigger fundamental flaw. In the near future, there might be a relevant event or activity coming up leading to a new topic that would not be captured by prior keywords. This calls for having a prediction model for the forthcoming topics and dynamically adjusting the keywords to anticipate emerging topics. Indeed, this research results from a collaboration with a large global company having challenges updating keywords on a regular basis.

As input the problem consists of a set of documents with a temporal dimension in a corpus and a set of keywords used to collect each document. Each document includes at least one keyword. The documents collected in a time window form a set of topics. The problem is to select a set of new keywords that are going to be used to collect documents in the immediate future. The future topics which have to be predicted from the keywords implied by the yet-to-be-collected documents should not deviate too much from the recent topics, but, on the other hand, should capture possible emerging topics.

Word frequency is a straight forward approach to extract new keywords. TFIDF is an early word extraction algorithm comparing word frequency in a document and in the corpus to determine the importance of a word. There are extensions to this notion based on TFIDF that apply extra features, such as the length of the words, however, tweets are short (maximum 280 character posts) and thus it is hard to find meaningful words that frequently appear in certain documents but have low frequency in others. With TFIDF, it is also easy to pick background words that seem useful but are too broad to be keywords, such as “today,” “money,” etc. With such broad words as keywords, it is easy to filter short tweets that include these words but are off-topic. Thus we need a different algorithm and model to select meaningful keywords applicable to short text media like Twitter. In addition, TFIDF and its derived algorithms select keywords based on historical data, so the selected keywords can only represent past information. However, new keywords need to be used in the future instead of being based on the past. To capture useful information, we need to select keywords that should encapsulate upcoming topics and trends. Therefore, a generative model is necessary for our problem in addition to an inference part

that is going to predict upcoming topics based on sets of keywords. Unfortunately, TFIDF does not provide such a generative model.

To select meaningful keywords that can be used to capture future relevant topics, we propose a generative model with a topic model as an information filter where priors depend on candidate sets of keywords. There are some keyword extraction models based on topics, such as Topical PageRank [15] and context-sensitive Topical PageRank [28]. These models calculate word ranking scores using the graph-based Topical PageRank and extract new keywords from the word ranking scores. They build word graphs and use topics to measure the importance of the words. In such models topics are fixed, and the keywords within the chosen topics are subsequently selected. Since the topics are fixed, these models do not meet our requirement of anticipating and predicting future active keywords together with resulting topics. In our proposed model, both topics and keywords have equal importance and are jointly modeled with the generative model. The generative model captures past interactions between the keywords and topics and thus in inference given a set of keywords it predicts the resulting topics. A challenge in inference is to balance the fact that future topics should be similar to the current topics, yet at the same time they should capture emerging topics.

Different from Topical PageRank and context-sensitive Topical PageRank, which are limited to pre-selected and fixed topics, we model topics by a topic model (Latent Dirichlet Allocation (LDA), [3]). The topics are modeled in time and depending on the underlying keywords, and thus can reflect events and trends, such as a new product release. We assume that the set of all documents is grouped into subsets (e.g. all documents collected in a week) where each subset of documents is generated based on its own set of keywords. The keywords pertaining to the different subsets can overlap. We model the interaction between keywords and topics through priors at the document-topic and topic-word levels. The keyword selection process has its own simple generative process whose output forms the input to neural networks that in turn output the priors to LDA. The model has three prior parameters γ , λ , and ψ which are all trainable. In each updating period (e.g. a week) corresponding to a subset of documents and the underlying keywords, the selected keywords are first sampled from prior parameter γ . The generated keywords work together with prior parameter ψ (weights in a neural network) to generate the document-topic distribution η through a neural network, and with the other prior parameter λ (also weights in a neural network) through another neural network to generate the topic-word distribution ϕ indicating the word distribution for each topic. Regarding word generation in document d , first latent topic variable $z_d^{k_w}$ is selected based on the generated set of keywords and next the keyword belonging to this document is sampled from the topic-word distribution ψ (which also depends on the generated keywords through a neural network). The latter sampling within the topic selected is based on a discrete probability conditioned on topic $z_d^{k_w}$ and the topic-word distribution. For each remaining word in the document, latent topic $z_{w,d}$ is sampled from the topic-word distribution ψ , and a word is chosen from discrete probability conditioned on topic $z_{w,d}$ (this step follows standard LDA). In this way, every word and keyword for each document are

sampled. The process guarantees that each generated document has at least one word from the generated set of keywords.

The model includes the standard LDA parameters which are augmented by the parameters pertaining to the generative process behind keyword generation and it also includes neural network parameters. The parameters are trained by a combination of expectation maximization (EM), variational inference, and stochastic gradient optimization. Since the generative process behind keywords and a part of the incumbent LDA adjusted model are using discrete distributions, we also resort to the Gumbel softmax trick in training. In inference, we need to decide which keywords to use for immediate future streaming based on the generative model.

The proposed model is evaluated on three sets of tweets with the benchmark algorithm being a sophisticated algorithm based on NLP techniques (also designed in this work). In terms of accuracy appropriately defined the improvement is 74%, 61% over the benchmark algorithm if 2 or 3 keywords are recommended, respectively. If a random subset of keywords is selected among all candidate words, then the improvement is 3-fold.

In this paper, we propose a novel generative probabilistic model for keyword selection, Keyword Latent Dirichlet Allocation (KLDA). The main contributions of our work as as follows:

- (1) A topic based generative model is proposed for the future document keyword selection.
- (2) Neural networks are introduced in producing priors.
- (3) An inference model that recommends keywords to use in the future trading off the facts that future topics should be similar to the current topics, yet they should include novelty in order not to miss emerging topics.

In Section 2 we review the literature while in Section 3 the problem is formally stated. In Section 4, we describe the details of KLDA, including the generative process and inference. In Section 5 we introduce the datasets and discuss all of the experimental results.

2 LITERATURE REVIEW

2.1 Keyword Extraction

Keywords are an efficient and common way of filtering streaming textual data (in particular on Twitter). Several prior works deal with keyword extraction. Turney [23] proposes a supervised learning approach for this problem, classifying phrases as positive and negative in each document. Many later algorithms have been developed based on the Turney’s approach, with additional features and data sources. Hulth [9] introduces part-of-speech (POS) tags to represent the data, Yih et al. [26] analyze information from many resources, including meta-tags, URL, and query frequency, and Liu et al. [13] utilize a rich set of features such as lexical. However, deciding the importance of a word based on its frequency does not capture the resulting topics. In addition, supervised learning methods require accurate learning labels which are usually not available. The goal is to periodically set new keywords for future upcoming documents, so it is hard to get labels for a massive amount of documents. Besides, we do not have a clear target for what the keywords should be. Therefore, unsupervised learning is necessary in our keyword selection problem.

Graph-based keyword ranking is a popular and effective strategy to extract keywords. PageRank [19] is proposed by Page and Brin to estimate the importance of Web pages. Relying on PageRank, Mihalcea and Tarau [18] propose a keyword extraction method based on their TextPage model. TextPage is widely used in later keyword extraction algorithms [14, 24]. However, these graph-based models do not consider the semantics of keywords. Thus, Liu et al. and Zhao et al. [15, 28] propose topic based PageRank focusing on keyword extraction in pre-selected and fixed topics. However, we aim to select keywords without knowing the topics a priori and thus we rely on the concept of predicting future topics.

2.2 Topic Models

Topic models discover semantic relationships in documents by maintaining latent topics of documents. Latent Semantic Indexing [20] is an early topic model that maps words and documents to a representation in a semantic space using Singular Value Decomposition assuming that the words with the same semantics have similar meanings. Hofmann [8] propose a generative model Probabilistic Latent Semantic Indexing that contains a statistical foundation to build document indexing and information retrieval. It obtains the probabilities of different topics by generating every word from a topic and thus each document may contain multiple topics. Based on this, Blei et al. [3] introduce LDA to overcome overfitting of Latent Semantic Indexing models. A survey of topic models can be found in [4].

LDA is a generative model for detecting topics. It represents each topic with the word distribution and models every document as a mixture of multiple latent topics. Many models consider other factors, and extend LDA in different dimensions. Blei et al. [2] consider topic evolution in sequential documents and extends the topic model to the dynamic version, and Hoffman et al. [7] extrapolate LDA from local analysis to online learning in streams. In image problems, Wang et al. [25] introduce image information into LDA priors to encode the spatial structure among visual cues, Li et al. [5] extend LDA to natural scene categories, and Barnard et al. [1] also apply LDA parameters on images for matching words and pictures. Lin et al. and Mass et al. [12, 16] generalize LDA in sentiment analyses, and Zhang et al., Steyvers et al., and Rosen-Zvi et al. [21, 22, 27] add user information into LDA priors. Similar to these works, we introduce keywords into the LDA model by means of a flexible prior. The approaches are fundamentally different because of the dependence of topics on the underlying keywords.

3 MODEL AND ALGORITHM

We first present notation and the setting. We assume that, from a finite vocabulary set $\mathcal{V} = \{W_1, W_2, \dots, W_V\}$ (where V represents the cardinality of the vocabulary), we are given a subset $\mathcal{Q} \subset \mathcal{V}$ called the keyword set, which is the focus of this work. We use Q to denote the cardinality of \mathcal{Q} . A document d is a set of words $\{W_{d,1}, W_{d,2}, \dots, W_{d,n_d}\}$ where $n_d \in \mathbb{Z}$ is the size of the set, $W_{d,i} \in \mathcal{V}$ for all $i = 1, 2, \dots, n_d$, and there exists at least one $kw(d) \in \{1, 2, \dots, n_d\}$ such that $W_{d,kw(d)} \in \mathcal{Q}$. Note that (1) we require each document to contain at least one keyword; (2) there can be more than one keyword in a document; (3) we adopt the bag-of-words assumption throughout, thus the exact order of the words in

the document is of no importance. A corpus \mathcal{W} is a sequence of documents (d_1, d_2, \dots, d_D) where we assume that the index reflects the temporal dimension, i.e., d_1 is the oldest document and d_D the most recent document. We denote by K the number of topics which is a fixed hyperparameter.

As input (training data) we are given \mathcal{W} together with \mathcal{Q} . The main goal of keyword selection is to find a subset $\mathcal{Q}_{next} \subseteq \mathcal{Q}$ such that yet-to-be-collected documents d_{D+1}, d_{D+2}, \dots based on keywords \mathcal{Q}_{next} have similar topics as the topics of the most recent documents in the corpus. ‘‘Similarity’’ is defined later in Section 4; it needs to capture the desire to have identical topics yet to also include some emerging relevant topics.

As an example with tweets, let us assume that at the beginning of each week we specify a set of keywords. During a week we collect all tweets with respect to these keywords. Training data consists of all past tweets collected together with the keywords used. The problem is to select a set of keywords for the next week tweets. In this example, \mathcal{W} corresponds to all historical tweets and \mathcal{Q} can be the set of all keywords used in prior weeks possibly augmented by additional potential keywords. The keywords for the next week are denoted by \mathcal{Q}_{next} . The intent is that the topics of the tweets in the next week would be similar to the topics in most recent weeks but they should also capture some emerging relevant topics.

We approach this problem from the perspective of topic modeling. Each week of tweets yields a set of topics where the document-topic and topic-word probabilities depend on keywords used. We model this dependency through priors of the underlying Dirichlet distributions. This yields a relationship between a potential set of keywords and topics for the next week.

We first introduce the generative model and then the solution methodology based on variational inference.

3.1 Generative Model

We propose the KLDA model (illustrated in Figure 1 left) to explain the generative process of a document (labeled with subscript d and containing N words) with keywords, the steps of which can be summarized as follows.

- (1) Keyword generation: Sample $z_d^{kw} \sim p_{kw}(\cdot | \theta)$
- (2) Document keywords: $d = \{w : (z_d^{kw})_w = 1\}$
- (3) Topic weights: Sample $\eta_d \sim \text{Dir}(f_\psi(z_d^{kw}))$
- (4) Topic-word matrix parameters: $\beta_d = g_\lambda(z_d^{kw})$
- (5) Let $n_d^{kw} = |\{w : (z_d^{kw})_w = 1\}|$
- (6) For each of the remaining $N - n_d^{kw}$ words $W_{d,w}$
 - (a) Topic label: Sample $z_{d,w} \sim \text{Multin}(\eta_d)$
 - (b) Generate word: Sample $W_{d,w} \sim p_w(\cdot | z_{d,w}; \beta_d)$
 - (c) Add word to document: $d = d \cup \{W_{d,w}\}$

Notation ‘‘Dir,’’ ‘‘Multin’’ represent the Dirichlet and multinomial distributions, respectively. The trainable parameters are θ, ψ, λ . In the exposition for simplicity we assume that N is fixed, but everything can be easily extended to the case when N is distributed based on the Poisson distribution.

Step 1 determines which of the words in the candidate keyword set \mathcal{Q} are actually used as keywords for the current document. Binary $z_d^{kw} \in \{0, 1\}^{\mathcal{Q}} \setminus \{0\}$ represents the inclusion of

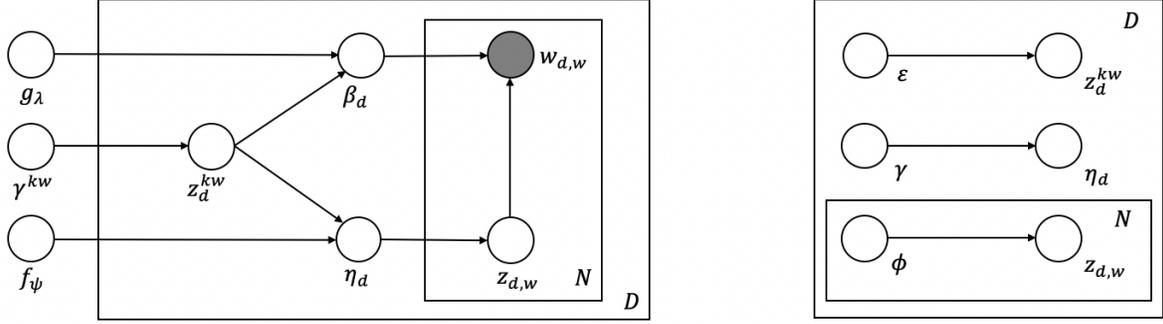


Figure 1: (Left) Plate representation of KLDA. (Right) Graphical model Representation for the variational distribution approximating posterior of KLDA. Subscript d is used for variables that are realized at document level, while w for variables generated for each observed word in the corpus.

the corresponding keyword in d (1 if the keyword is in d and 0 otherwise). These keywords become part of the document (step 2). Note that $z_d^{kw} \neq 0$ which means that each document must contain at least one keyword. We assume an energy-based distribution for keywords, where the unnormalized probability is $\tilde{p}_{k,w}(z|\theta) = \exp(z^T \theta - c(\mathbf{1}^T z - 1))$ with c being a hyperparameter. The partition function $Z(\theta) = \sum_{\tilde{z}} \tilde{p}_{k,w}(\tilde{z}|\theta)$ is a summation over all possible combinations of $\tilde{z} \in \{0, 1\}^Q \setminus \{0\}$. As the result, the normalized probability is $p_{k,w}(z|\theta) = \tilde{p}_{k,w}(z|\theta)/Z(\theta)$. In the energy function $\tilde{p}_{k,w}(z|\theta)$ term $z^T \theta$ encourages the model to generate keywords that are aligned with θ , which represents the underlying probability of each keyword while term $c(\mathbf{1}^T z - 1)$ penalizes z if too many keywords are included, which provides a regularization and prohibits the generation of a large number of keywords for a single document.

Next in step 3 we generate a probability vector η_d of topics based on the Dirichlet distribution. Its prior f_ψ is a model (parameterized by ψ) that takes as input a binary vector of dimension Q , and outputs a vector with Q positive numbers that can be used as a prior for a Dirichlet distribution. In our model the prior depends on the topics generated z_d^{kw} .

In the next step 4 we generate a set of parameters β_d that govern the topic-word distribution of the selected topic $z_{d,w}$. These parameters are specified by function g_λ which is a model (parameterized by λ) that takes as input a binary vector of dimension Q and outputs a $K \times V$ matrix with each row corresponding to a distribution that is used as a topic-word matrix in the LDA model.

In step 5, n_d^{kw} is the number of keywords selected. Step 6 generates all of the remaining keywords in d . In step 6a a topic is then selected based on the multinomial distribution with probabilities η_d . A word is finally sampled in step 6b based on a probability depending on the generated topic and topic-word priors, and added to the document in the last step. Distribution $p_w(\cdot | z; \beta)$ is a multinomial distribution conditioned on the one-hot vector z , meaning that for index k with $z_k = 1$, the distribution is a multinomial distribution parametrized by the k -th row of β .

The generative process implies the joint probability distribution

$$p(\mathcal{W}, \mathbf{z}^{kw}, \eta, \mathbf{z} | \theta, \psi, \lambda) = \prod_{d \in \mathcal{W}} p_{k,w}(z_d^{kw} | \theta) \cdot p(\eta_d | z_d^{kw}, \psi) \cdot \prod_{w=1}^{n_d} p(z_{d,w} | \eta_d) \cdot p_w(W_{d,w} | z_{d,w}, \beta_d(\lambda)).$$

One peculiarity of this generative process is that on the surface the topic-word matrix seem to depend on document d . However each document generated on the same set of keywords z^{kw} is subject to the same topic-word matrix. This is consistent with our goal that a set of keywords implies a topic-word matrix. On the other hand, it is consistent with the standard LDA where there is a single set of keywords and thus just one such matrix not depending on documents.

We describe the specific choices of f, g and the procedures for the updates of their parameters in the next section.

3.2 Variational Inference with EM Algorithm

Due to the intractability of the proposed probabilistic model, we resort to variational inference combined with the EM algorithm for model training. The variational distribution for latent variables in KLDA is shown in Figure 1 right and it reads

$$q(\mathbf{z}^{kw}, \eta, \mathbf{z} | \epsilon, \gamma, \phi) = \prod_{d \in \mathcal{W}} q(z_d^{kw} | \epsilon_d) \cdot q(\eta_d | \gamma_d) \cdot \prod_w q(z_{d,w} | \phi_{d,w}).$$

In the study we assume that $q(z_d^{kw} | \epsilon_d)$ is the Bernoulli distribution with success rate ϵ_d which agrees with the assumption that multiple keywords can be presented in a document. Distribution $q(z_{d,w} | \phi_{d,w})$ of a given document d is multinomial with parameters $\phi_{d,w}$. Finally, $q(\eta_d | \gamma_d)$ is Dirichlet with prior parameter γ_d . For topic weights η_d and topic labels $z_{d,w}$, we adopt the approach used in the standard LDA model (see [3]). In the variational distribution, η_d ($d = 1, 2, \dots, D$) is sampled from a Dirichlet distribution parametrized by the d -th row of γ , a $D \times K$ matrix with positive entries. Latent variable $z_{d,w}$ ($w = 1, 2, \dots, n_d$) is sampled from a

multinomial distribution parametrized by a probability vector $\phi_{d,w}$ of length K .

The introduction of keywords entails a new approach on the variational distribution that is amenable to both the categorical nature of z_d^{kw} and the framework of the EM algorithm. Here we use the Gumbel-softmax trick [10] to generate keywords z_d^{kw} in the variational distribution. Specifically, the trainable parameter ϵ is a $D \times Q$ matrix, where each element $\epsilon_{d,j} \in [0, 1]$ represents the probability that the j -th candidate keyword is used as a keyword for document d . The Gumbel-softmax trick is used to approximate the categorical distribution of z_d^{kw} in the following way where $s_r(x)$ is 0 if $x < r$ and 1 otherwise.

For a fixed temperature parameter $\tau > 0$ and for any given document index $d = 1, 2, \dots, D$, we sample Q^2 samples $u_{d,j,k}$ from uniform distribution on $[0, 1]$, and we set $g_{d,j,k} = -\log(-\log(u_{d,j,k}))$ for $j = 1, \dots, Q, k = 1, \dots, Q$ to represent a standard Gumbel sample. We then compute

$$y_{d,j,k}^{kw} = \frac{\exp((\log(\epsilon_{d,j,k}) + g_{d,j,k})/\tau)}{\sum_{i=1}^Q \exp((\log(\epsilon_{d,i,k}) + g_{d,i,k})/\tau)} \quad (1)$$

and collect the vector $y_{d,k}^{kw} = (y_{d,j,k}^{kw})_{j=1}^Q$. We then use $y_{d,k}^{kw} = s_{0.5}(\sum_{k=1}^Q y_{d,k}^{kw})$ (here s is applied coordinate-wise) to approximate z_d^{kw} . Note that if $y_{d,k}^{kw}$ are binary, then $y_{d,k}^{kw} = s_{1.0}(\sum_{k=1}^Q y_{d,k}^{kw}) = s_{0.5}(\sum_{k=1}^Q y_{d,k}^{kw})$. The Gumbel-softmax trick provides a differentiable, smooth, empirically efficient and stable gradient estimation approach for categorical distributions, which can recover the original discrete distribution when the temperature is annealed down to 0. For training of KLDA, we use a sequence of pre-determined, decreasing temperatures. In each iteration we first sample y_d^{kw} from the Gumbel distribution described above, and then feed y_d^{kw} into f and g for evaluating the loss function or computing a gradient during training.

With the aforementioned variational distribution, we apply the EM algorithm to maximize the Evidence Lower Bound (ELBO) (see [7] for a detailed description based on the original LDA model). In our context, the probability of the corpus \mathcal{W} in KLDA is $p(\mathcal{W}|\theta, \psi, \lambda)$, the joint distribution of the corpus with all latent variables is denoted by $p(\mathcal{W}, z^{kw}, \eta, z|\theta, \psi, \lambda)$, and the probability of latent variables under the variational distributions is represented as $q(z^{kw}, \eta, z)$. Then we have

$$\begin{aligned} \log p(\mathcal{W}|\theta, \psi, \lambda) &\geq \mathcal{L}(\theta, \psi, \lambda, \epsilon, \gamma, \phi) \\ &:= \mathbb{E}_q[\log p(\mathcal{W}, z^{kw}, \eta, z|\theta, \psi, \lambda)] \\ &\quad - \mathbb{E}_q[\log q(z^{kw}, \eta, z|\epsilon, \gamma, \phi)], \end{aligned}$$

where \mathbb{E}_q denotes the expectation operator with respect to variational distribution q . Note that the introduction of the Gumbel-softmax trick affects the variational distribution and requires sampling from the Gumbel distribution at each iteration, and for large-scale cases the training of the model needs to be executed in an online fashion, using only a batch of the entire data in each iteration. Therefore, the actual training is performed based on an estimator $\hat{\mathcal{L}}$ of $\mathcal{L}(\theta, \psi, \lambda, \epsilon, \gamma, \phi)$ depending on the batch sampled

in each iteration, and the sample y_d^{kw} through the Gumbel-softmax trick.

In the remainder, we abuse the notation of q to denote the revised variational distribution where sampling of z^{kw} is performed through the Gumbel trick. First, note that

$$\begin{aligned} \hat{\mathcal{L}} &= \mathbb{E}_\epsilon \left[\mathbb{E}_q[\log p(\mathcal{W}, y^{kw}, \eta, z|\theta, \psi, \lambda)] \right. \\ &\quad \left. - \log q(y^{kw}, \eta, z|\epsilon, \gamma, \phi) | y^{kw} \right] \end{aligned}$$

where \mathbb{E}_ϵ is the expectation with respect to the random vectors y^{kw} sampled from the Gumbel distribution parametrized by ϵ . Now we consider each terms one by one, conditioned on y^{kw} . We have

$$\begin{aligned} &\mathbb{E}_q \left[\log p(\mathcal{W}, y^{kw}, \eta, z|\theta, \psi, \lambda) | y^{kw} \right] \\ &\approx \sum_{d=1}^D \log p_{k_w}(y_d^{kw}|\theta) \end{aligned} \quad (2)$$

$$+ \sum_{d=1}^D \mathbb{E}_q[\log p(\eta_d | y_d^{kw}, \psi)] \quad (3)$$

$$+ \sum_{d=1}^D \sum_{w=1}^{n_d} \mathbf{1}\{W_{d,w} \notin I(y_d^{kw})\} \cdot \mathbb{E}_q[\log p(z_{d,w} | \eta_d)] \quad (4)$$

$$+ \sum_{d=1}^D \sum_{w=1}^{n_d} \mathbf{1}\{W_{d,w} \notin I(y_d^{kw})\} \cdot \mathbb{E}_q[\log p_w(W_{d,w} | z_{d,w}, \beta_d)] \quad (5)$$

where $I(y) = \{i : y_i > h\}$ for hyperparameter h set to 0.5 in experiments. The indicator functions in (4) and (5) are capturing the fact that if a word has already been generated as a keyword, we do not need to generate this word again as a regular word; we use hyperparameter h for y_d^{kw} to determine if a word has been generated as a keyword. Note that this is an approximation only because we are using the Gumbel trick and thus values of y^{kw} are not strictly 0 or 1.

A direct calculation gives an explicit form for each term (2)-(5). For term (2), we have

$$\log p_{k_w}(y|\theta) = \log \tilde{p}_{k_w}(y|\theta) - \log Z(\theta)$$

where the unnormalized probability \tilde{p}_{k_w} and partition function Z are defined in the previous section. For term (3), we have

$$\mathbb{E}_q \log p(\eta_d | y_d^{kw}, \psi) = \tilde{h}(\gamma_d, f_d)$$

where $f_d = f_\psi(y_d^{kw})$ with γ being the parameter for the variational distribution, and

$$\begin{aligned} \tilde{h}(u, v) &= \sum_i (v_i - 1) [\Psi(u_i) - \Psi(\sum_j u_j)] \\ &\quad + \log \Gamma(\sum_i v_i) - \sum_i \log \Gamma(v_i) \end{aligned}$$

and Ψ denoting the Digamma function. For term (4) given word with index w in document d we have

$$\mathbb{E}_q[\log p(z_{d,w} | \eta_d)] = \sum_{k=1}^K \phi_{d,w,k} \left(\Psi(\gamma_{d,k}) - \Psi\left(\sum_{l=1}^K \gamma_{d,l}\right) \right)$$

with ϕ coming from the variational distribution. For term (5) we derive

$$\begin{aligned} & \mathbb{E}_q[\log p_w(W_{d,w}|z_{d,w}, \beta_d)] \\ &= \sum_{v=1}^V \sum_{k=1}^K \mathbf{1}\{W_{d,w} = v\} \phi_{d,w,k} \left(\Gamma(\beta_{d,k,v}) - \Gamma\left(\sum_{j=1}^V \beta_{d,k,j}\right) \right) \end{aligned}$$

where ϕ is the parameter for z in the variational distribution, and β is the output of model g .

Also, for the variational distribution we have

$$\begin{aligned} & \mathbb{E}_q[\log q(y^{\text{kw}}, \eta, z | \epsilon, \gamma, \phi)] \\ & \approx \sum_{d=1}^D \log q(y_d^{\text{kw}} | \epsilon_d) \end{aligned} \quad (6)$$

$$+ \sum_{d=1}^D \mathbb{E}_q[\log q(\eta_d | \gamma_d)] \quad (7)$$

$$+ \sum_{d=1}^D \sum_{w=1}^{n_d} \mathbf{1}\{W_{d,w} \notin I(y_d^{\text{kw}})\} * \mathbb{E}_q[\log q(z_{d,w} | \phi_{d,w})]. \quad (8)$$

Specifically, in term (6), given the original Bernoulli distribution in Figure 1 right, we have

$$\log q(y|\epsilon) = \sum_{j=1}^Q y_j \log(\epsilon_j) + (1 - y_j) \log(1 - \epsilon_j)$$

for vectors y, ϵ of length Q ; for term (7) we derive

$$\mathbb{E}_q[\log q(\eta|\gamma_d)] = \tilde{h}(\gamma_d, \gamma_d);$$

and for term (8) we have

$$\mathbb{E}_q[\log q(z_{d,w}|\phi_{d,w})] = \sum_{k=1}^K \phi_{d,w,k} \log \phi_{d,w,k}.$$

We next consider the case where we can only access a batch $\tilde{\mathcal{W}} \subset \mathcal{W}$ of the entire corpus. Denote the number of documents in $\tilde{\mathcal{W}}$ as $B = |\tilde{\mathcal{W}}|$, and the total number of words of the documents in this batch as N_B . Then we can use the following estimator with

adjusted weights for each term:

$$\begin{aligned} & \hat{\mathcal{L}} \quad (9) \\ &= \frac{D}{B} \sum_{d \in \tilde{\mathcal{W}}} \log p_{k_w}(y_d^{\text{kw}} | \theta) \\ &+ \frac{D}{B} \sum_{d \in \tilde{\mathcal{W}}} \mathbb{E}_q[\log p(\eta_d | y_d^{\text{kw}}, \psi)] \\ &+ \frac{V}{N_B} \sum_{d \in \tilde{\mathcal{W}}} \sum_{w=1}^{n_d} \mathbf{1}\{W_{d,w} \notin I(y_d^{\text{kw}})\} \cdot \mathbb{E}_q[\log p(z_{d,w} | \eta_d)] \\ &+ \frac{V}{N_B} \sum_{d \in \tilde{\mathcal{W}}} \sum_{w=1}^{n_d} \mathbf{1}\{W_{d,w} \notin I(y_d^{\text{kw}})\} \\ &\quad \cdot \mathbb{E}_q[\log p_w(W_{d,w} | z_{d,w}, \beta_d)] \\ &- \frac{D}{B} \sum_{d \in \tilde{\mathcal{W}}} \log q(y_d^{\text{kw}} | \epsilon_d) \\ &- \frac{D}{B} \sum_{d \in \tilde{\mathcal{W}}} \mathbb{E}_q[\log q(\eta_d | \gamma_d)] \\ &- \frac{V}{N_B} \sum_{d \in \tilde{\mathcal{W}}} \sum_{w=1}^{n_d} \mathbf{1}\{W_{d,w} \notin I(y_d^{\text{kw}})\} \cdot \mathbb{E}_q[\log q(z_{d,w} | \phi_{d,w})]. \end{aligned}$$

At an E-step, the optimal solution with respect to γ and ϕ can be solved through the iterative updates for any $d \in \tilde{\mathcal{W}}$ and $k = 1, 2, \dots, K$ as

$$\begin{aligned} \phi_{d,w,k} &\propto \beta_{d,k,W_{d,w}} \exp(\Psi(\gamma_{d,k}) - \Psi(\sum_{j=1}^K \gamma_{d,j})) \\ &\quad \text{for all } W_{d,w} \notin I(y_d^{\text{kw}}); \\ \gamma_{d,k} &= (f_\psi(y_d^{\text{kw}}))_k + \sum_{w=1}^{n_d} \phi_{d,w,k}. \end{aligned}$$

The derivation can be found in [3]. For all the other parameters, the updates can be conducted through gradient descent with respect to the estimated log probability $\hat{\mathcal{L}}$. Note that ϵ can also be updated in this fashion since y_d^{kw} is essentially a function of ϵ (and random uniform samples u).

3.3 Partition Function

Updating of θ requires special care with respect to $\log Z(\theta)$ in (2). The exponential nature of this term also requires approximations in inference. We resort to sampling described next.

For estimating the gradient $\nabla_\theta \log Z(\theta)$ we use the following well-known fact (see [6] for details)

$$\nabla_\theta \log Z(\theta) = \mathbb{E}_{X \sim p_{k_w}} \nabla_\theta \log \tilde{p}_{k_w}(X).$$

During training, the gradient estimator is obtained through the Markov Chain Monte Carlo algorithm. Specifically, we employ Gibbs sampling, and maintain N_{MCMC} different independent chains whose state spaces are $\{0, 1\}^Q \setminus \{0\}$, the possible range of z^{kw} . Before training starts, we initialize all the chains by uniformly randomly sampling a possible z^{kw} . At iteration t of gradient optimization in the M-step, we perform m_t steps of Gibbs sampling for each chain, and denote the current state for each chain by $\tilde{z}_t^{(j)}$

for each $j = 1, 2, \dots, N_{MCMC}$. Here $\{m_t\}_{t \geq 1}$ is a preset sequence of integers. The gradient estimator we use to perform stochastic gradient descent on θ , conditioned on y^{k_w} from the Gumbel trick, is

$$\frac{D}{B} \sum_{d \in \tilde{\mathcal{W}}} \left(\nabla_{\theta} \log \tilde{p}_{k_w}(y_d^{k_w} | \theta) - \frac{\sum_{j=1}^{N_{MCMC}} \log \tilde{p}_{k_w}(z_t^{(j)} | \theta)}{N_{MCMC}} \right). \quad (10)$$

During inference with respect to the $\log Z(\theta)$ term, we employ a simple importance sampling algorithm. Alternative options of MCMC or acceptance-rejection are not appropriate for the following reasons. Since the estimate \hat{Z} is used in the log function as an approximation to $\log(Z)$, even though $\mathbb{E}\hat{Z} = Z$, it holds that $\mathbb{E} \log(\hat{Z}) \neq \log(\mathbb{E}\hat{Z})$ due to the Jensen's inequality, and the gap is affected by the variance of \hat{Z} and thus we need a lot of samples (we use 100,000). For MCMC running that many chains till burn-in is too costly.

To be specific, importance sampling is used to estimate the partition function Z , where we construct the following estimator

$$\hat{Z} = \frac{1}{N_e} \sum_{j=1}^{N_e} \frac{\tilde{p}_{k_w}(X^{(j)})}{h(X^{(j)})}.$$

Values $\{X^{(j)}\}$ are i.i.d samples from distribution $h(\cdot)$ constructed in the following way.

- (1) Sample \tilde{N} from Geometric distribution with mean 1/0.25, and set $N = \min\{\tilde{N}, Q\}$ to correspond to the number of keywords;
- (2) Randomly sample a subset $\tilde{Q} \subset Q$ of size N (the sampling is without replacement);
- (3) Output a vector X of length Q , where $X_l = 1$ for $l \in \tilde{Q}$ and $X_l = 0$ otherwise.

If the distribution of h is relatively close to the distribution of Z , the variance of the estimator should be low. The designed h is trying to match the generative process. Even though we are ignoring the actual dependency on θ for each candidate keyword and use uniform distribution instead, the geometric distribution is capturing the fact that, due to the penalty term $c > 0$ in the definition of \tilde{p}_{k_w} , it is approximately geometrically or exponentially unlikely to have many keywords.

Note that $h(X)$ is simple to evaluate since it only relies on the number of positive elements in X , and only requires the summation of a geometric sequence. In our experiments, we set $N_e = 100,000$ so that the variance of $\log(\hat{Z})$ is small enough when plugging in the estimator \hat{Z} , which further ensures that the bias to $\log Z(\theta)$ is small enough.

3.4 Overall Training Algorithm

The overall algorithm is presented in Algorithm 1. The main loop in step 2 iterates over all EM iterations. In the inner loop 3 we first create independent samples in support of keyword candidates based on Gibbs Sampling described in Section 3.3. The actual samples are derived in step 6 for the select batch samples of documents. The forward step is executed in step 7. Possible several E steps are done in step 8 while a single M-step gradient adjustment is performed in steps 11-15.

Algorithm 1 EM algorithm for KLDA model

Require:

Corpus $\mathcal{W} = \{d_1, d_2, \dots, d_D\}$; Vocabulary \mathcal{V} ; Candidate keyword set $Q \subset \mathcal{V}$;

Learning rates $\{\alpha_t\}_{t \geq 1}$; Temperature $\{\tau_t\}_{t \geq 1}$; MCMC repetitions $\{m_t\}_{t \geq 1}$; Integer for number of MCMC chains N_{MCMC} ; Batch size B ; E-step iterations K_E

- 1: **Initialize:** Parameters ψ, λ for model f and g ; Parameters $\theta, \epsilon, \gamma, \phi$ for the probabilistic model; States for MCMC chains $\{z_0^{(j)}\}_{j=1}^{N_{MCMC}}$
 - 2: **for** $t = 1, 2, 3, \dots$ **do**
 - 3: **for** $j = 1, 2, \dots, N_{MCMC}$ **do** ▷ MCMC
 - 4: Perform m_t Gibbs Sampling steps on $z_{t-1}^{(j)}$ to
 \hookrightarrow obtain $z_t^{(j)}$
 - 5: Uniformly sample a batch $\tilde{\mathcal{W}} \subset \mathcal{W}$ of size B
 - 6: Obtain $y_d^{k_w}$ for any $d \in \tilde{\mathcal{W}}$ via (1) with
 \hookrightarrow temperature τ_t ▷ Gumbel Trick
 - 7: Evaluate $\hat{\mathcal{L}}$ using (9)
 - 8: **for** $\tilde{t} = 1, 2, \dots, K_E$ **do** ▷ E-step
 - 9: $\phi_{d,w,k} \propto \beta_{d,k} w_{d,w} \exp(\Psi(\gamma_{d,k}) - \Psi(\sum_{j=1}^K \gamma_{d,j}))$
 \hookrightarrow for each $W_{d,w} \notin I(y_d^{k_w})$
 \hookrightarrow and for each $d \in \tilde{\mathcal{W}}$;
 - 10: $\gamma_{d,k} = (f_{\psi}(y_d^{k_w}))_k + \sum_{w=1}^{n_d} \phi_{d,w,k} \forall d \in \tilde{\mathcal{W}}$;
 - 11: $\epsilon \leftarrow \epsilon + \alpha_t \nabla_{\epsilon} \hat{\mathcal{L}}$ ▷ M-step
 - 12: $\psi \leftarrow \psi + \alpha_t \nabla_{\psi} \hat{\mathcal{L}}$
 - 13: $\lambda \leftarrow \lambda + \alpha_t \nabla_{\lambda} \hat{\mathcal{L}}$
 - 14: Evaluate $\nabla_{\theta} \hat{\mathcal{L}}$ using (10)
 - 15: $\theta \leftarrow \theta + \alpha_t \nabla_{\theta} \hat{\mathcal{L}}$
-

4 INFERENCE

After training the KLDA model, we obtain model g_{λ} that can output the topic-word matrices given keyword set z , and vector θ that can be used to evaluate the possibility that keyword set z is actually generated by the KLDA procedure. The next goal is to decide the best possible keyword sets for further collection of documents. To this end, we propose two inference metrics that aim to capture two antagonistic purposes of deriving a keyword set. Let us recall that the documents are assumed to be collected in time (which is actually disregarded in training). Let Q_{last} be the most recent set of keywords used to collect documents $\mathcal{D}_{last} = \{d_{D-s}, \dots, d_D\}$ for a hyperparameter s . In practice s should reflect the frequency of keyword updates, e.g., if we plan to update the keywords once every week, then these documents should correspond to the documents in the last week and thus Q_{last} are the keywords used in the previous week.

Given a subset $U \subset Q$ of keywords, then $g_{\lambda}(1_U)$ is a topic-word matrix where $1_U \in \{0, 1\}^Q$ is the binary representation of U . Let $a(U) = \sum_k g_{\lambda}(1_U)_k / K$ is a distribution over vocabulary (with $g_{\lambda}(1_U)_k$ denoting the k -th row of matrix $g_{\lambda}(1_U)$). Let also $\text{freq}(S)$ be the ratio of the number of documents in Q_{last} containing every word in S and s (the number of all documents in Q_{last}).

To be more specific about the goal, assuming that a set $Q_{last} \subset Q$ is provided, the goal is to identify another subset $Q_{next} \subseteq Q$ that, in terms of the possible documents generated by keywords contained in Q_{next} , it maintains information in documents \mathcal{D}_{last} corresponding to Q_{last} but also captures new information and provides a natural extension of previous topics.

We construct Q_{next} by considering each word $W_{last} \in Q_{last}$ and try to “extend” it with a different word $W_{next} \in Q$. The set of all derived W_{next} form our final set Q_{next} of keywords. Let W_{last}, W_{next} be fixed in the subsequent discussion. We consider the strength of W_{next} by using the following two metrics.

The first metric aims to capture that $\{W_{next}, W_{last}\}$ can be seen as a good extension of keyword W_{last} , since either of them generates documents that cover similar topics. We need to capture that as long as one of the keywords appears in the document it would be meaningful to collect it. Choices for such extensions W_{next} should satisfy the following criteria: (1) W_{next} should be relatively frequent in \mathcal{D}_{last} (we set the threshold as $f_2 = 0.5\%$), otherwise the mixing of $\{W_{next}, W_{last}\}$ would be extremely similar to using W_{last} only, thus rendering the next criterion meaningless; (2) the KL divergence between the word distributions induced by W_{last} and W_{next} should be small ($KL(a(\{W_{last}\})||a(\{W_{next}\}))$ should be small where KL denotes KL divergence). In our experiments we found that, instead of using KL divergence, most other f -divergence measures can be used and the selected keywords remain largely similar, so other divergence metrics can be used if need be. As long as the first threshold is satisfied, we prefer the keywords that induce as small KL divergence as possible.

Besides this metric, we propose the second metric that is used as a reference and helps filtering out some obviously bad choices before applying the previous metric. We use the term *high-frequency distance* for the quantity defined as follows. Let us assume we have hyperparameters $o_h \in [0, 1]$ (high-frequency threshold), $r \in [0, 1]$ (retention rate), and $\hat{c} > 0$ (penalty weight). We define the following sets

- extension set

$$E = \{w \in \mathcal{V} : a(\{W_{next}\})_w > o_h, \frac{a(\{W_{last}\})_w}{a(\{W_{next}\})_w} < r\}$$

- retention failure set

$$F = \{w \in \mathcal{V} : a(\{W_{last}\})_w > o_h, \frac{a(\{W_{next}\})_w}{a(\{W_{last}\})_w} < r\}$$

and set the high-frequency distance as $R(W_{last}, W_{next}) = |E| - \hat{c}|F|$. Intuitively, this quantity translates the idea of “maintaining old information but also extending to new information” into a value that is based on high-frequency words in both cases. In our experiments we set o_h to be a value that ensures that the high-frequency sets

$$\{w : a(\{W_{next}\})_w > o_h\}, \{w : a(\{W_{last}\})_w > o_h\}$$

are of size around 100 (the value o_h is set as $1/1,000$), and we set $\hat{c} = 1, r = 1/2$. Note that if $W_{last} = W_{next}$, then the two sets are empty and thus $|E| - \hat{c}|F| = 0$.

The actual inference procedure is presented in Algorithm 2. Note that $S \subset C$. It can conceptually happen that $Q_{next} = \emptyset$ is returned by the algorithm. In such a case which was never observed in the experimental study, we can select $Q_{next} = Q_{last}$.

Algorithm 2 Inference

- 1: $Q_{next} = \emptyset$
 - 2: **for** $W_{last} \in Q_{last}$ **do**
 - 3: $C = \{W_{next} \in Q | R(W_{last}, W_{next}) \geq 0\}$
 - 4: $S = \operatorname{argmin}_{W_{next} \in C: \text{freq}(\{W_{next}\}) > f_2} \{KL(a(\{W_{last}\})||a(\{W_{next}\}))\}$
 - 5: $Q_{next} = Q_{next} \cup S$
 - 6: Remove duplicate words in Q_{next}
-

5 COMPUTATIONAL STUDY

5.1 Training Details

Related to (1), function s is not smooth and it requires Q^2 samples. In order to circumvent this in the computational experiments we instead sample Q binary Gumble approximations as follows. For any given document index $d = 1, 2, \dots, D$, and keyword index $j = 1, 2, \dots, Q$, we sample two independent copies $u_{d,j,1}, u_{d,j,0}$ from the uniform distribution on $[0, 1]$, and we set $g_{d,j,i} = -\log(-\log(u_{d,j,i}))$ for $i = 0, 1$. We then compute

$$y_{d,j}^{kw} = \frac{\exp((\log(\epsilon_{d,j,1}) + g_{d,j,1})/\tau)}{\sum_{i \in \{0,1\}} \exp((\log(\epsilon_{d,j,i}) + g_{d,j,i})/\tau)} \quad (11)$$

and define the vector $y_d^{kw} = (y_{d,j}^{kw})_{j=1}^Q$ to approximate z_d^{kw} . While this is smooth it has the drawback that it allows $y_d^{kw} = 0$, i.e., no keyword is in document d . Thus this necessitates allowing state 0 in MCMC (see Section 3.3).

For our experiments, we set $K = 5$ topics (by finding a good number of topics using plain LDA), and the batch size $B = 64$. For the distribution of $p_{k,w}$, we set the penalty $c = 2$. For model g (topic-word matrices) we use a 2-layer neural network with 32 neurons in each layer, using LeakyRelu activation with weight 0.01 for the negative quadrant (see [17] for details on LeakyRelu). For model f , we fix it and always use a uniform 1 vector as the output. In other words, in our experiments we use an uninformative, uniform prior for topic weights. We tried to use a parametric f but we observed that the trained model tends to use only one of the K topics and the others have an extremely small probability to be generated. We have tried to add a Lagrangian penalty term to prevent the topic weights from being too extreme but at no avail. For training of model g and gradient ascent on most parameters, we set the learning rates as 0.001 and use the Adam optimizer [11]. For updates of ϵ , however, we use RMSprop (see [29]) with learning rate 0.005 to cope with the fact that the gradient on ϵ is extremely sparse as only the ϵ corresponding to a mini batch of documents is updated at each iteration. We also add an L2 regularization term to loss function (9) with the regularization penalty of 0.1.

For the MCMC steps, we set $m_t = 1000$ if $t \bmod 100 = 0$ and set $m_t = 1$ otherwise. In other words, periodically we perform more Gibbs sampling steps to ensure the quality of the samples after updating θ . For the Gumbel trick, we anneal down the temperature based on the following scheme

$$\tau_t = \max\{0.25, \exp(-0.0001 * t)\},$$

which is partially following the choices in [10] but we set the threshold as 0.25 as we try to further anneal down the temperature and better recover the original categorical distribution.

Since the corpus is given, during training of KLDA, we cannot allow the model to generate a keyword that is not observed in the document. To this end, at each iteration after we obtain y_d^{kw} through the Gumbel trick, we apply a binary mask to all the elements in y_d^{kw} , which multiplies 0 to slots of candidate keywords that do not appear in the document. Therefore, the masked candidate keywords do not affect the evaluation of $\hat{\mathcal{L}}$, and the corresponding ϵ does not need to be updated.

Training of θ is critical to our inference, since it determines the process that generates keywords. In our experiment we apply several tricks to improve the quality of θ . First, we change the weights of the negative MCMC samples for estimating the gradient of the partition function. In (10) we apply a weight c_{neg} and use the revised form

$$\nabla_{\theta} \log \tilde{p}_{kw}(y_d^{kw} | \theta) - c_{\text{neg}} \frac{\sum_{j=1}^{N_{\text{MCMC}}} \log \tilde{p}_{kw}(z_t^{(j)} | \theta)}{N_{\text{MCMC}}}.$$

In our experiment we set $c_{\text{neg}} = 0.1$. Second, we add an L_2 regularization term in $\hat{\mathcal{L}}$ as the squared error between normalized $\exp(\theta)$ and the empirical probability of each candidate keyword in the corpus. We set the coefficient of this L_2 regularization term as 1.

Lastly, to improve the quality of model g , we adopt a pre-training scheme before training the KLDA model. The only difference between pre-training and complete KLDA model training is in the following. In pre-training, there is no generation of keywords z^{kw} ; instead we use the real observation as z^{kw} , meaning that each element in z^{kw} is 1 if the corresponding keyword appears in the document, and 0 otherwise (namely, we skip the stochastic generation of keywords, and deterministically generate a candidate keyword as a keyword if it is present in the document). We use the same hyperparameters and model structures as described above. Note that this pre-training step does not involve updates on θ (since in fact we skip the keyword generation step), and only serves to improve the quality of topic-word matrices output by g . We run pretrain for 2, 500 iterations, and when we start training the full KLDA model, we load the parameters of model g from pretrain to initialize the new g model.

5.2 Data

We next give a briefing of the data used in the experiments. In order to test the model’s capability under different settings, we use collections of Twitter posts that concern various topics and evolve over time. Specifically, data collection is done through the Twitter streaming API; the entire collection lasted for 11 weeks in year 2019 for each dataset with keywords adapted at the beginning of each week. During this period we maintain 3 sets of tweets revolving around 3 different topics: horse-racing, streaming, and Korean pop music (the first dataset was intended to be a “computer virus” related dataset, but evolved into a horse-racing dataset through the keyword selection benchmark procedure described next). We selected these topics to have variety.

The benchmark algorithm relies on a prediction model that classifies if a tweet is likely to become viral or not. The features of

a tweet are the user identification, text length, number of URLs, number of mentions, number of hashtags and emojis, the number of users the user is following, the number of followers of the user, the number of favorites of the user, and the total number of posts of the user. The label is defined as viral for each tweet that has been retweeted at least ten times in a month and the model is trained based on 3 months of tweets.

In terms of the data collection procedure for any dataset, in the first week we manually pick a set of keywords that are related to these topics and, by querying the API, collect tweets posted in the following 7-day period that contain any of the given keywords. For all subsequent weeks we apply the following algorithm called *viral* to adjust the set of keywords. The algorithm also serves as the main benchmark algorithm.

- We first calibrate an LDA model on all tweets from the last week.
- We next perform inference of the tweets collected in the last week by using the viral prediction model.
- From all tweets predicted to become viral, we select 15 highest count words (excluding words that are not nouns or proper nouns which implies that we first perform part-of-speech tagging) as keyword candidates.
- Let $tweets(k)$ be the set of all tweets in the previous week containing keyword k where k is one of the candidate keywords based on the viral tweets.
- For each keyword k of the 15 keywords and for each tweet in $tweets(k)$ we check if the tweet belongs to one of the topics. This is done by using the LDA probabilities. If over 60% of the tweets in $tweets(k)$ belong to a topic, keyword k is selected and otherwise it is discarded.

The ‘viral’ component produces candidate keywords present in tweets that are likely to become important. The LDA part is necessary to assure that the selected keywords are not deviating too much from the desired topics. For example, in the case with the industry partner, the company had a promotion related to the FIFA World Cup. If the LDA part is removed, then the viral tweets capture everything related to the World Cup which is too broad and most of the material of no interest to the company.

We execute the algorithm for each of the remaining 9 weeks for each one of the 3 datasets. The keywords produced and thus used in the collection process are exhibited in Table 1. We remark that based on this algorithm the initial topic of “computer virus” evolved to horse-racing in the first dataset.

Next, based on the tweets collected in this 10-week span, we generate a set of words that are the union of the following: (1) recommendation from *viral* using the tweets in week 10; (2) high-frequency nouns in tweets not recommended by *viral*; (3) any candidate keyword that has been used in the 10-week data collection period. This candidate keyword set is used to collect tweets for one more week, labeled as week 11, and is utilized as the test set in our experiment. Table 2 lists the keywords for week 11 that are added to the keywords in Table 1 and Table 3 provides basic statistics of the datasets. KLDA for week 11 only recommends a subset of the keywords shown in Tables 1 and 2. At the end of week 10, we assemble the keywords for week 11 and then collect all tweets during week 11 based on all these keywords.

Table 1: Evolution of Keywords, Weeks 1-10

Week	Horse-racing	Streaming	KPOP
1	computer virus, Trojan, invasion, malware	video, stream	idol, album
2	horse, sex, virus, aid, war	league, hero, game, live, playing	ticket, army, follow
3	horse, sex, virus, war	stream, hero, game, live, playing	ticket, army
4	horse	stream, game, live, playing	really, army
5	horse	stream, game, live, playing	army, known, stan
6	horse, happy	stream, game, live, playing	persona, april
7	people, know	stream, live, playing	love, april
8	people, racing	stream, live, playing, time	album, april, army, luv
9	people, racing	stream, live, playing, workshop	army
10	derby, race, racing	stream, game, playing	army

Table 2: Additional Keywords for Week 11

	Recommendations by <i>viral</i>	High-frequency Words
Horse-racing	birthday Kentucky derby shit racing riding people today time today	maximum security town one girl show
Streaming	stream skin map mercy game havana league player night tonight people fun	video guy everyone
KPOP	army idol concert persona stan world day tour people time way	thank everyone album guys

Table 3: Information of Datasets Used in Experiments

	Horse-racing	Streaming	KPOP
#Docs, Weeks 1-10	153,378	135,506	163,013
#Words, Weeks 1-10	1,610,218	1,387,365	1,817,409
#Docs, Week 11	13,433	6,271	12,327
#Words, Week 11	136,973	64,002	12,7454
Vocabulary	172,371	196,299	162,639
Vocabulary(reduced)	10,164	12,033	12,422

We add a few remarks regarding collection and preprocessing of the datasets. First, we apply stemming and removing of the stopwords. Besides, we reduce the vocabulary size by replacing all the low-frequency tokens (appearing less than 10 times in all tweets collected during the first 10 week) by a placeholder token. The purpose is to facilitate model training by preventing rare typos and acronyms, since tweets are more prone to these issues. As we observe in the collected datasets, less than 10% of tokens (words) have appeared 10 or more times in the entire dataset. By substituting the low-frequency tokens with the placeholder token, we manage to reduce the vocabulary to a reasonable size (see Table 3 for details). Lastly, if a token outside of the vocabulary is observed in the tweets collected in week 11, it is also replaced by the placeholder token, thus maintaining a consistent vocabulary between weeks 1-10 and week 11.

5.3 Computational Results

We outline the objective of the computational experiments as follows: for each dataset, different methods yield different sets of recommendations of keywords for week 11; by collecting tweets using candidate keywords recommended by different methods at

week 11, we are able to evaluate the performance of different methods based on tweets in week 11 (given a set of keywords for week 11, we simply select the tweets containing at least one keyword from the set; since such keywords are a subset of all keywords used in the collection process for week 11, no tweet is left out); the comparison of different methods under various topics also demonstrate the versatility and stability of different approaches.

More specifically, we compare KLDA against *viral* and a randomly chosen keyword set. Recall that the inference algorithm based on KLDA generates recommendations for a given keyword by evaluating all the remaining candidate keywords under the aforementioned metric. Based on the score of the metric, we need to specify a threshold: candidate keywords with metric scores meeting the threshold shall be accepted into the recommendation list. In this experiment we compare the results from two different thresholds: we either accept the top 2 words or the top 3 words. Throughout, we use *top 2* or *top 3* to refer to the two different set of criteria, respectively. As for *viral*, recall that the recommendation is generated for the entire set used in week 10 in one shot, instead of extending each keyword in the set as is the case in Algorithm 2. Lastly, to generate a *random* keyword set as another benchmark,

Table 4: Perplexity of Trained Models

	Horse-racing	Streaming	KPOP
LDA on Week 10	-4.86	-4.94	-4.42
Trained on Weeks 1-9, Tested on Week 10	-6.46	-5.92	-6.889
no training at all	-8.30	-7.47	-8.51

Table 5: Recommendation of Keywords from KLDA Model: Individual Keyword Level

Org. Keyword	KLDA (Top 2)	KLDA (Top 3)	Week 11 (Top 2)	Week 11 (Top 3)	Acc. Top 2	Acc. Top 3
derby	racing Kentucky	racing Kentucky time	racing sex	racing sex know	0.50	0.33
race	war riding	war riding horse	horse riding	horse riding racing	0.50	0.67
racing	today kentucky	today kentucky time	sex know	sex know show	0.00	0.00
stream	map	map	none	none		
game	playing tonight	playing tonight stream	playing tonight	playing tonight fun	1.00	0.67
playing	live stream	live stream map	fun tonight	fun tonight stream	0.00	0.33
army	april album	april album	fans concert	fans concert guys	0.00	0.00

we randomly choose 3 keywords used in the week 11 candidate set, with the restriction that only those that are not recommended by Algorithm 2 or *viral* are taken into consideration. The random set generated for the experiment is: Horse-racing: town, trojan, know; Streaming: workshop, everyone, hero; KPOP: know, everyone, luv.

Table 6: Performance of Different Recommendations

Top 2	KLDA	viral	Random
Horse-racing, Accuracy	0.40	0.20	0.33
Horse-racing, Coverage	0.40	0.40	0.20
Streaming, Accuracy	0.50	0.25	0.00
Streaming, Coverage	0.67	1.00	0.00
KPOP, Accuracy	0.00	0.09	0.00
KPOP, Coverage	0.00	0.50	0.00
Accuracy (Mean)	0.30	0.17	0.11
Coverage (Mean)	0.36	0.63	0.07
Top 3	KLDA	viral	Random
Horse-racing, Accuracy	0.43	0.20	0.33
Horse-racing, Coverage	0.50	0.33	0.17
Streaming, Accuracy	0.60	0.33	0.00
Streaming, Coverage	0.75	1.00	0.00
KPOP, Accuracy	0.00	0.09	0.00
KPOP, Coverage	0.00	0.33	0.00
Accuracy (Mean)	0.34	0.21	0.11
Coverage (Mean)	0.41	0.55	0.06

We first demonstrate the predictability of the trained model in terms of topics of future tweets, thus showing KLDA’s capability for this modelling problem. Specifically, we first train a model using tweets collected from weeks 1 to 9; for tweets in week 10 we train a separate LDA model; we then compare the perplexity of the two models when fitting the tweets in week 10. As shown in Table 4, we see a consistent pattern across all the three datasets: by training a model using tweets from week 1-9, we are able to fit tweets in the next week with very good perplexity, which is both a

significant improvement compared to “no training at all” (a model that is randomly initialized and not trained at all), and close to the performance of the LDA model trained on week 10 directly.

As mentioned before, the KLDA inference Algorithm 2 generates recommendations by inspecting each keyword used in week 10. We next discuss the performance of KLDA at the individual keyword level, the result of which is illustrated in Table 5. The metrics accuracy and coverage are interpreted as follows. Calling the top 2 (or 3) recommendations for week 11 based on the weeks 1-10 model as predictions, and the top 2 (or 3) choices in week 11 based on the two criteria discussed in Section 4 with respect to LDA based on tweets in week 11 as truth, we claim that a prediction generated by KLDA is accurate if it is in truth, and we say one truth is covered by KLDA if it is in the prediction set. We use “none” to indicate that due to bad performance of the high-frequency distance defined in Section 4, all candidate keywords are excluded and there is no good recommendation based on the tweets in week 11. We observe that KLDA demonstrates great accuracy when accepting top 2 choices, and the result further improves when we are allowed to accept top 3 choices.

The comparison of the recommendations for week 11 under the different methods is summarized in Table 6. We observe that under both top 2 and top 3, KLDA yields the highest average accuracy; besides, KLDA also performs the best for 2 of the 3 topics. In the KPOP dataset, which is the only case KLDA is outperformed by *viral*, note that it is expected that *viral* outperforms since it generates a recommendation set of a much larger size, and even so it produces only one accurate prediction out of 11. In general, it is expected *viral* to have higher coverage since it selects many more keywords.

REFERENCES

- [1] Kobus Barnard, Pinar Duygulu, David Forsyth, Nando de Freitas, David M Blei, and Michael I Jordan. Matching words and pictures. *Journal of machine learning research*, 3(Feb):1107–1135, 2003.
- [2] David M Blei and John D Lafferty. Dynamic topic models. In *Proceedings of the 23rd international conference on machine learning*, pages 113–120. ACM, 2006.
- [3] David M Blei, Andrew Y Ng, and Michael I Jordan. Latent Dirichlet allocation. *Journal of machine learning research*, 3(Jan):993–1022, 2003.

- [4] Ali Daud, Juanzi Li, Lizhu Zhou, and Faqir Muhammad. Knowledge discovery through directed probabilistic topic models: a survey. *Frontiers of computer science in China*, 4(2):280–301, 2010.
- [5] Li Fei-Fei and Pietro Perona. A Bayesian hierarchical model for learning natural scene categories. In *Proceedings of the computer vision and pattern recognition conference*, volume 2, pages 524–531. IEEE, 2005.
- [6] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [7] Matthew Hoffman, Francis R Bach, and David M Blei. Online learning for latent Dirichlet allocation. In *Advances in neural information processing systems*, pages 856–864, 2010.
- [8] Thomas Hofmann. Probabilistic latent semantic indexing. In *ACM special interest group on information retrieval forum*, volume 51, pages 211–218. ACM, 2017.
- [9] Anette Hulth. Improved automatic keyword extraction given more linguistic knowledge. In *Proceedings of the 2003 conference on empirical methods in natural language processing*, pages 216–223. ACL, 2003.
- [10] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-softmax. *arXiv preprint arXiv:1611.01144*, 2016.
- [11] Diederik P Kingma and Jimmy Ba. ADAM: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [12] Chenghua Lin and Yulan He. Joint sentiment/topic model for sentiment analysis. In *Proceedings of the 18th ACM conference on information and knowledge management*, pages 375–384. ACM, 2009.
- [13] Fei Liu, Feifan Liu, and Yang Liu. Automatic keyword extraction for the meeting corpus using supervised approach and bigram expansion. In *Spoken language technology workshop*, pages 181–184. IEEE, 2008.
- [14] Feifan Liu, Deana Pennell, Fei Liu, and Yang Liu. Unsupervised approaches for automatic keyword extraction using meeting transcripts. In *Proceedings of human language technologies: The 2009 annual conference of the North American chapter of the ACL*, pages 620–628. ACL, 2009.
- [15] Zhiyuan Liu, Wenyi Huang, Yabin Zheng, and Maosong Sun. Automatic keyphrase extraction via topic decomposition. In *Proceedings of the 2010 conference on empirical methods in natural language processing*, pages 366–376. ACL, 2010.
- [16] Andrew L Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the ACL: Human language technologies-volume 1*, pages 142–150. ACL, 2011.
- [17] Andrew L Maas, Awni Y Hannun, and Andrew Y Ng. Rectifier nonlinearities improve neural network acoustic models. In *Proceedings of the international conference on machine learning*, volume 30, page 3, 2013.
- [18] Rada Mihalcea and Paul Tarau. TextRank: Bringing order into text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*, 2004.
- [19] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
- [20] Christos H Papadimitriou, Prabhakar Raghavan, Hisao Tamaki, and Santosh Vempala. Latent semantic indexing: A probabilistic analysis. *Journal of computer and system sciences*, 61(2):217–235, 2000.
- [21] Michal Rosen-Zvi, Thomas Griffiths, Mark Steyvers, and Padhraic Smyth. The author-topic model for authors and documents. In *Proceedings of the 20th conference on uncertainty in artificial intelligence*, pages 487–494. AUAI Press, 2004.
- [22] Mark Steyvers, Padhraic Smyth, Michal Rosen-Zvi, and Thomas Griffiths. Probabilistic author-topic models for information discovery. In *Proceedings of the tenth ACM SIGKDD international conference on knowledge discovery and data mining*, pages 306–315. ACM, 2004.
- [23] Peter D Turney. Learning algorithms for keyphrase extraction. *Information retrieval*, 2(4):303–336, 2000.
- [24] Xiaojun Wan, Jianwu Yang, and Jianguo Xiao. Towards an iterative reinforcement approach for simultaneous document summarization and keyword extraction. In *Proceedings of the 45th annual meeting of the ACL*, pages 552–559, 2007.
- [25] Xiaogang Wang and Eric Grimson. Spatial latent Dirichlet allocation. In *Advances in neural information processing systems*, pages 1577–1584, 2008.
- [26] Wen-tau Yih, Joshua Goodman, and Vitor R Carvalho. Finding advertising keywords on web pages. In *Proceedings of the 15th international conference on World Wide Web*, pages 213–222. ACM, 2006.
- [27] Wei Zhang and Jianyong Wang. Prior-based dual additive latent Dirichlet allocation for user-item connected documents. In *Proceedings of the international joint conference on artificial intelligence*, volume 15, pages 1405–1411, 2015.
- [28] Wayne Xin Zhao, Jing Jiang, Jing He, Yang Song, Palakorn Achananuparp, Ee-Peng Lim, and Xiaoming Li. Topical keyphrase extraction from Twitter. In *Proceedings of the 49th annual meeting of the ACL: Human language technologies-volume 1*, pages 379–388. ACL, 2011.
- [29] Fangyu Zou, Li Shen, Zequn Jie, Weizhong Zhang, and Wei Liu. A sufficient condition for convergences of ADAM and RMSprop. In *Proceedings of the computer vision and pattern recognition conference*, pages 11127–11135, 2019.