
Sawtooth Factorial Topic Embeddings Guided Gamma Belief Network

Zhibin Duan^{*1} Dongsheng Wang^{*1} Bo Chen¹ Chaojie Wang¹
Wenchao Chen¹ Yewen Li¹ Jie Ren¹ Mingyuan Zhou²

Abstract

Hierarchical topic models such as the gamma belief network (GBN) have delivered promising results in mining multi-layer document representations and discovering interpretable topic taxonomies. However, they often assume in the prior that the topics at each layer are independently drawn from the Dirichlet distribution, ignoring the dependencies between the topics both at the same layer and across different layers. To relax this assumption, we propose sawtooth factorial topic embedding guided GBN, a deep generative model of documents that captures the dependencies and semantic similarities between the topics in the embedding space. Specifically, both the words and topics are represented as embedding vectors of the same dimension. The topic matrix at a layer is factorized into the product of a factor loading matrix and a topic embedding matrix, the transpose of which is set as the factor loading matrix of the layer above. Repeating this particular type of factorization, which shares components between adjacent layers, leads to a structure referred to as sawtooth factorization. An auto-encoding variational inference network is constructed to optimize the model parameter via stochastic gradient descent. Experiments on big corpora show that our models outperform other neural topic models on extracting deeper interpretable topics and deriving better document representations.

1. Introduction

Probabilistic topic models, such as latent Dirichlet allocation (LDA) (Blei et al., 2003) and Poisson factor analysis (PFA) (Zhou et al., 2012), have the ability to discover the underlying

semantic themes from a collection of documents, achieving great success in text analysis. In general, a topic model represents each document as a mixture of latent topics, each of which describes an interpretable semantic concept. While being widely used, vanilla topic models assume that topics are independent and there are no structures among them, which limits those models' ability to explore any hierarchical thematic structures. To remove its limitation, a series of hierarchical extensions, including nonparametric Bayesian hierarchical prior based topic models (Blei et al., 2010; Paisley et al., 2014), deep PFA (Gan et al., 2015; Henao et al., 2015), gamma belief network (GBN) (Zhou et al., 2015; Cong et al., 2017), and Dirichlet belief network (DirBN) (Zhao et al., 2018a), have been proposed. Commonly, these models learn directed acyclic graph (DAG)-structured hierarchical topics, which assumes that the topics in the upper layers are more general/abstract than those in the lower layers. Consequently, revealing hierarchical relations between topics provides the user an intuitive way to better understand text data.

With the development of deep neural networks (DNN), there is a growing interest in developing neural topic models (NTMs). Specifically, most neural topic models are based on variational auto-encoders (VAEs) (Kingma & Welling, 2013; Rezende et al., 2014), which employ a variational inference network (encoder) to approximate the posterior distribution and are equipped with a decoder to reconstruct the document's Bag-of-Words (BOW) representation (Miao et al., 2016; Srivastava & Sutton, 2017; Card et al., 2017). However, most NTMs rely on Gaussian latent variables, which often fail to well approximate the posterior distributions of sparse and nonnegative document latent representations. To address this limitation, Zhang et al. (2018) develop Weibull hybrid autoencoding inference (WHAI) for deep LDA, which infers posterior samples via a hybrid of stochastic-gradient MCMC and autoencoding variational Bayes. As a hierarchical neural topic model, WHAI shows attractive qualities in multi-layer document representation learning and hierarchical explainable topic discovery. Compared with traditional Bayesian probabilistic topic models, these NTMs usually enjoy better flexibility and scalability, which are important for modeling large-scale data and performing downstream tasks (Zhang et al., 2019; Wang et al.,

^{*}Equal contribution ¹National Laboratory of Radar Signal Processing, Xidian University, Xi'an, China. ²McCombs School of Business The University of Texas at Austin, Austin, TX 78712, USA. Correspondence to: Bo Chen <bchen@mail.xidian.edu.com>.

2020c; chen et al., 2020; Wang et al., 2020a; Duan et al., 2021; Zhao et al., 2021).

Despite their attractive performance, existing hierarchical topic models such as GBN often assume in the prior that the topics at each layer are independently drawn from the Dirichlet distribution, ignoring the dependencies between the topics both at the same layer and across different layers. To relax this assumption, we propose the Sawtooth Factorial Topic Embeddings Guided GBN (SawETM), a deep generative model of documents that captures the dependencies and semantic similarities between the topics in the embedding space. Specifically, as sketched in Fig. 1, both the words and hierarchical topics are first converted into the shared embedding space. Then we develop the Sawtooth Connection technique to capture the dependencies between the topics at different layers, where the factor loading at layer l is the factor score at layer $l - 1$, which enables the hierarchical topics to be coupled together across all layers. Our work is inspired by both GBN (Zhou et al., 2015), a multi-stochastic-layer hierarchical topic model, and the embedding topic models (Dieng et al., 2019; 2020), which represent the words and single layer topic as embedding vectors. The proposed Sawtooth Connector is a novel method that combines the advantages of both models for hierarchical topic modeling.

We further note that previous work on NTMs has been restricted to shallow models with one or three layers of stochastic latent variables, which could limit their ability. Generally, due to the well-known component collapsing problem of VAEs (Sønderby et al., 2016), constructing a deep latent variable model is challenging work. As discussed in Child (2020), the hierarchical VAEs with a sufficient depth can not only learn arbitrary orderings over observed variables but also learn more effective latent variable distributions, if such distributions exist. Moving beyond text modeling, the recent development on image generation has shown its promising performance and outstanding generation ability (Maaløe et al., 2019; Vahdat & Kautz, 2020; Child, 2020). Inspire by their work, we carefully design the inference network of SawETM in a deep hierarchical VAE framework to improve the model’s ability of modeling textual data. In particular, we propose the integration of a skip-connected deterministic upward path and a stochastic path to approximate the posterior of the latent variables and obtain hierarchical representations of a document. We also provide customized training strategies to build deeper neural topic models. To the best of our knowledge, SawETM is the first neural topic model that well supports a deep network structure (e.g., 15).

Our main contributions are summarized as follows:

- To move beyond the independence assumption between the topics of two adjacent layers in most hierarchical topic models, the Sawtooth Connection technique is

developed to extend GBN by capturing the dependencies and semantic similarities between the topics in the embedding space.

- To avoid posterior collapse, We carefully design a residual upward-downward inference network for SawETM to improve the model’s ability of modeling count data and approximating sparse, non-negative and skewed document latent representations.
- Overall, SawETM, a novel hierarchical NTM equipped with a flexible training algorithm, is proposed to infer multi-layer document representations and discover topic hierarchies in both the embedding space and vocabulary space. Experiments on big corpora show that our models outperform other NTMs on extracting deeper interpretable topics and deriving better document representation.

2. Related work

The proposed model in this paper marries a hierarchical neural topic model with word embedding, resulting in a deep generative framework for text modeling. The related work can be roughly divided into two categories, one is the research on constructing neural topic model and the other is on leveraging word embedding for topic models.

Neural topic models Most existing NTMs can be regarded as an extension of Bayesian topic models like LDA within the VAE framework for text modeling, where the latent variables z can be viewed as topic proportions. NTMs usually utilize a single-layer network as their decoder, e.g., $\phi(z) = \text{softmax}(\mathbf{W}z)$ (Srivastava & Sutton, 2017), where \mathbf{W} is a learnable weights between topics and words. Different NTMs may place different distributions on latent variable z , such as Gaussian and Dirichlet distributions (Miao et al., 2016; Burkhardt & Kramer, 2019; Nan et al., 2019; Wang et al., 2020b). Different from these models that only focus on a single layer latent variable model, Zhang et al. (2018) propose WHAI, a hierarchical neural topic model that employs a Weibull upward-downward variational encoder to infer multi-layer document latent representations and use GBN as a decoder. All of these works focus on relatively shallow models with one or three layers of stochastic latent variables. We note that our work is based on WHAI, but different from it, we propose a new decoder to capture the dependencies of topics and a new powerful encoder to approximate the posteriors, which result in a deeper neural topic model.

Word embedding topic models Word embeddings can capture word semantics at low-dimensional continuous space and are well studied in neural language models (Bengio et al., 2003; Mikolov et al., 2013a;b; Levy & Goldberg,




Figure 1. (a) Gamma belief network and (b) overview of the proposed SawETM and its corresponding hierarchical upward and downward encoder networks

2014). Due to their ability to capture semantic information, there is a growing interest in applying word embeddings to topic models. Pre-trained word embeddings, such as GloVe (Pennington et al., 2014) and word2vec (Mikolov et al., 2013b), can serve as complementary information to guide topic discovery, which is effective to alleviate the sparsity issue in topic models (Zhao et al., 2017; 2018b; Li et al., 2016). Dieng et al. (2020) propose an embedding-based topic model (ETM), which directly models the similarity in its generative process, rather than via a Dirichlet distribution. What should be noted is that our model is related to ETM in modeling the correlation between words and topics using their semantic similarities. Different from ETM, the Sawtooth Connection can be seen as injecting the learned knowledge information of a lower layer to a higher layer, which alleviates the sparsity issue in the higher layer.

3. The proposed model

In this section, we develop SawETM for text analysis, which aims at mining document’s multi-layer representations and exploring topic hierarchies. The motivation for designing SawETM focuses on tackling two main challenges: (i) a hierarchical decoder to construct the dependencies between topics; (ii) designing expressive neural networks to approximate the posterior distribution more accurately. Below, we will first describe the decoder and encoder of SawETM, and then discuss the model’s properties. Finally, we will provide details of model inference and stable training techniques.

3.1. Document decoder: Sawtooth Factorial Topic Embeddings Guided GBN

To explore hierarchical thematic structure from a collection of documents, SawETM adapts GBN of Zhou et al. (2015) as its generative module (decoder). Different from GBN, where the interactions between the topics of two adjacent

layers are modeled as a Dirichlet distribution, SawETM utilizes the Sawtooth Connection (SC) technique to couple hierarchical topics across all layers in the shared embedding space. Specifically, assuming the observations are multi-variate count vectors $x_j \in \mathbb{Z}^{K_0}$, the generative model of SawETM with L layers, from top to bottom, can be expressed as

$$\begin{aligned} \theta_j^{(L)} &\sim \text{Gam}(r, c_j^{(L+1)}), \\ \theta_j^{(l)} &\sim \text{Gam}(\Phi^{(l+1)} \theta_j^{(l+1)}, c_j^{(l+1)}), \quad l = 1, \dots, L-1 \\ \Phi_k^{(l)} &= \text{Softmax}(\alpha^{(l-1)T} \alpha_k^{(l)}), \quad l = 1, \dots, L \\ x_j &\sim \text{Pois}(\Phi^{(1)} \theta_j^{(1)}) \end{aligned} \quad (1)$$

where, the count vector x_j (e.g., the bag-of-words of document j) is factorized as the product of the factor loading matrix $\Phi^{(1)}$ (topics), and gamma distributed factor scores $\theta_j^{(1)}$ (topic proportions), under the Poisson likelihood; and the hidden units $\theta_j^{(l)} \in \mathbb{R}_+^{K_l}$ of layer l are further factorized into the product of the factor loading $\Phi^{(l+1)} \in \mathbb{R}_+^{K_l \times K_{l+1}}$ and hidden units of the next layer to infer a multi-layer latent representations. K_l denotes the topic numbers at layer l . the vector $\alpha_k^{(l)} \in \mathbb{R}^D$ is a distributed representation of the k^{th} topic at layer l in the semantic space of words. Especially, $\alpha^{(0)} \in \mathbb{R}^{D \times V}$ is the word embedding matrix. $\Phi_k^{(l)}$ captures the relationship between topics of two adjacent layers and is calculated by the SC technique. In detail, SC first calculates the semantic similarities between the topics of two adjacent layers by the inner product of their embedding vectors and then applies softmax normalization to make sure the sum of each column of $\Phi^{(l)}$ is equal to one. Note that, the factor loading at layer l is the factor score at layer $l-1$, which constructs the dependency between the parameters of the adjacent layers. Repeating this process, the hierarchical topic can be coupled together across all layers.

To verify the role of SC, we also consider two simple variants for ablation study, formulated as

$$\Phi_k^{(l)} = \text{Softmax}(\mathbf{W}_k^{(l)}), \quad l = 1, \dots, L \quad (2)$$

$$\Phi_k^{(l)} = \text{Softmax}(\alpha^{(l)} \beta_k^{(l)}), \quad l = 1, \dots, L \quad (3)$$

where, Eq. (2) directly models the dependence via the learnable parameters \mathbf{W} , which is a common choice in most NTMs. Eq. (3) employs a similar decomposition form as SC but does not share topic embeddings between two adjacent layers. We refer to the first variant as a deep neural topic model (DNTM), and the second as a deep embedding topic model (DETM). Note that when the number of layers is set to one, the DETM has the same structure as SawETM.

3.2. Document encoder: upward and downward encoder networks

The goal of encoder is to reparameterize the variational posterior distribution $q(\theta_j)$, which denotes topic proportions. While the gamma distribution, satisfying the nonnegative constraint and encouraging sparsity, appears to be a natural choice, it is not reparameterizable and therefore not amenable to gradient descent based optimization. Here we consider Weibull distribution because $i)$, latent variable $x \sim \text{Weibull}(k, \lambda)$ can be easily reparameterized as:

$$x = \lambda(-\ln(1 - \varepsilon))^{1/k}, \quad \varepsilon \sim \text{Uniform}(0, 1). \quad (4)$$

$ii)$, The KL divergence from the gamma to Weibull distributions has an analytic expression (Zhang et al., 2018):

$$KL(\text{Weibull}(k, \lambda) \parallel \text{Gamma}(\alpha, \beta)) = \frac{\gamma\alpha}{k} - \alpha \log \lambda +$$

$$\log k + \beta \lambda \Gamma(1 + \frac{1}{k}) - \gamma - 1 - \alpha \log(\beta) + \log \Gamma(\alpha).$$

where γ is the Euler-Mascheroni constant. What's more, designing a capacious inference network is necessary for deep VAEs training. Inspired by the series of studies about deep generative models (Sønderby et al., 2016; Zhang et al., 2018; Vahdat & Kautz, 2020; Child, 2020), we develop an upward-downward inference network, which contains a bottom-up residual deterministic path and a top-down stochastic path.

Upward-downward inference model: Like most neural topic models, the exact posterior distribution for $\theta^{(l)}$ is intractable and needs to be approximated. Following VAEs (Kingma & Welling, 2013; Rezende et al., 2014), we define the variational distribution as $q(\theta_j | x_j)$, which need to be flexible enough to well approximate the true posterior distribution. The variational distribution is factorized with a bottom-up structure as

$$q(\theta_j | x_j) = q(\theta_j^{(1)} | x_j) \prod_{l=1}^{L-1} q(\theta_j^{(l+1)} | \theta_j^{(l)}) \quad (5)$$

Here we emphasize that this hierarchical structure makes the top stochastic latent variables tend to collapse into the prior, especially when the layer is large (Sønderby et al., 2016). To address this problem, SawETM first parameters a skip-connected deterministic upward path to obtain the latent representations of input x :

$$\begin{aligned} \mathbf{h}_j^{(l)} &= \mathbf{h}_j^{(l-1)} + \text{MLP}(\mathbf{h}_j^{(l-1)}) \\ \hat{\mathbf{k}}_j^{(l)} &= \text{Relu}(\text{Linear}(\mathbf{h}_j^{(l)})) \\ \hat{\lambda}_j^{(l)} &= \text{Relu}(\text{Linear}(\mathbf{h}_j^{(l)})) \end{aligned}$$

where $\mathbf{h}^{(0)} = \mathbf{x}$, MLP is a two layer fully connected network, Linear is a single linear layer, and Relu applies non-linear activation function. SawETM combines the obtained

latent features with the prior from the stochastic up-down path to construct the variational posterior:

$$\begin{aligned} q(\theta_j^{(l)} | \Phi^{(l+1)}, \mathbf{h}_j^{(l)}, \theta_j^{(l+1)}) &= \text{Weibull}(\mathbf{k}_j^{(l)}, \lambda_j^{(l)}) \quad (6) \\ \mathbf{k}_j^{(l)} &= \text{Softplus}(\text{Linear}(\Phi^{(l+1)} \theta_j^{(l+1)} \oplus \hat{\mathbf{k}}_j^{(l)})) \\ \lambda_j^{(l)} &= \text{Softplus}(\text{Linear}(\Phi^{(l+1)} \theta_j^{(l+1)} \oplus \hat{\lambda}_j^{(l)})) \end{aligned}$$

where \oplus denotes the concatenation at topic dimension, and Softplus applies $\log(1 + \exp(\cdot))$ nonlinearity to each element to ensure positive Weibull shape and scale parameters. The Weibull distribution is used to approximate the gamma-distributed conditional posterior, and its parameters $\mathbf{k}_j^{(l)} \in \mathbb{R}^{K_l}$ and $\lambda_j^{(l)} \in \mathbb{R}^{K_l}$ are both inferred by combining the bottom-up likelihood information and the prior information from the generative distribution using the neural networks. The inference network is structured as

$$q(\theta | x) = q(\theta^{(L)} | x) \prod_{l=1}^{L-1} q(\theta^{(l)} | \theta^{(l+1)}, x). \quad (7)$$

compared with Eq. (5), the residual upward pass in SawETM allows all the latent variables to have a deterministic dependency on input x , thus the top stochastic latent variables could receive efficient information and will be empirically less likely to collapse. Note that, compared with the inference network in WHAI (Zhang et al., 2018), we construct a more powerful residual network structure to better approximate the true posteriors.

3.3. Model properties

SawETM inherits the good properties of both deep topic model and word embeddings, as described below.

Semantic topic taxonomies: The loading matrices $\Phi^{(l)}$ in Eq. (1) capture the semantic correlations of the topics of adjacent layers. Using the law of total expectation, we have

$$\mathbb{E} \left[\mathbf{x}_j | \theta_j^{(l)}, \left\{ \Phi^{(t)}, c_j^{(t)} \right\}_{t=1}^l \right] = \left[\prod_{t=1}^l \Phi^{(t)} \right] \frac{\theta_j^{(l)}}{\prod_{t=2}^l c_j^{(t)}}. \quad (8)$$

Therefore, $\left[\prod_{t=1}^{l-1} \Phi^{(t)} \right] \Phi^{(l)}$ is naturally interpreted as the projection of topic $\Phi^{(l)}$ to the vocabulary space, providing us with a way to visualize the topics at different semantic layers. The topics in the bottom layers are more specific and become increasingly more general when moving upward, as shown in Fig.4.

Hierarchical topics in the same embedding space: In SawETM, both words and hierarchical topics are represented with embedding vector (e.g., $\alpha^{(l)} \in \mathbb{R}^{D \times K_l}$), and the topic can be defined by the Sawtooth Connection (e.g. the l^{th} layer topic is defined as $\Phi^{(l)} = \text{softmax}(\alpha^{(l-1)T} \alpha^{(l)})$).

The first advantage is that the learned words and hierarchical topics can be projected into the same embedding space, which is shown in Fig. 3. And the second advantage is that SawETM can establish dependencies between different layers, which can be seen as the learned knowledge information of a lower layer can be injected into a higher layer. The intuition is there is semantic relation between the same layer topics, such as the topic about ‘basketball’ have a strong relation with the topic about ‘game’, which should be considered by the higher layer topics. Note that, other hierarchical topic models such as GBN usually assume the hierarchical topic is independent and ignore this semantic structure (Blei et al., 2007), while SawETM try to capture this structure in the embedding space.

3.4. Inference and estimation

Similar to VAEs, the training objective of SawETM is the maximization of an Evidence Lower Bound (ELBO):

$$\begin{aligned} \mathcal{L} = & \sum_{j=1}^J \mathbb{E}_{q(\theta_j)} [\ln p(\mathbf{x}_j | \Phi^{(1)}, \theta_j^{(1)})] \\ & - \sum_{j=1}^J \sum_{l=1}^L \text{KL}(q(\theta_j^{(l)}) || p(\theta_j^{(l)} | \Phi^{(l+1)}, \theta_j^{(l+1)})) \end{aligned} \quad (9)$$

where, $\Phi^{(l)} = \text{softmax}(\alpha^{(l-1)T} \alpha^{(l)})$, $q(\theta_j^{(l)})$ is the Weibull variational distribution in Eq. (6), and $p(\theta_j^{(l)})$ is the gamma prior distribution in Eq. (1). The first term is the expected log-likelihood or reconstruction error, while the second term is the Kullback–Leibler (KL) divergence that constrains $q(\theta_j^{(l)})$ to be close to its prior $p(\theta_j^{(l)})$ in the generative model. Thanks to the analytic KL expression and easy reparameterization of the Weibull distribution, the gradient of the ELBO with respect to $\{\alpha^{(l)}\}_{l=0}^L$ and other parameters in the inference network can be accurately evaluated. As describe in Algorithm. 1, the encoder parameters Ω and decoder parameters Ψ in SawETM are updated by SGD, which makes faster inference at both train and test stages compared to Gibbs Sampling. This also helps differ SawETM from WHAI, which updates the global parameters by SG-MCMC and is limited to update local and global parameters alternately.

3.5. Stable training

It is a great optimization challenge to train a deep hierarchical VAE in practice, due to the well-known posterior collapse and unbounded KL divergence in the objective (Razavi et al., 2019; Child, 2020). Here, we propose three approaches for stabilizing the training. We emphasize that all these approaches are applied in WHAI and other hierarchical neural topic model variants for fair comparison when we perform experiments.

Algorithm 1 Upward-Downward Autoencoding Variational Inference for SawETM

```

Set mini-batch size  $m$  and the number of layer  $L$ 
Initialize the encoder parameters  $\Omega$  and decoder parameters  $\Psi$ ;
for iter = 1, 2, ... do
  Randomly select a mini-batch of  $m$  documents to form a subset  $X = \{\mathbf{x}_i\}_{1,m}$ ;
  Draw random noise  $\{\epsilon_i^l\}_{i=1,l=1}^{m,L}$  from uniform distribution;
  Calculate  $\nabla_{\Omega, \Psi} L(\Omega, \Psi; X, \{\epsilon_i^l\}_{i=1,l=1}^{m,L})$  according to Eq. (10), and update encoder parameters  $\Omega$  and decoder parameter  $\Psi$  jointly;
end for

```

Shape parameter skipping of Weibull distribution: As shown in Eq. (4), when the sampled noise ϵ is close to 1, e.g., 0.98, and the Weibull shape parameter k is less than 1e-3, the x will be extremely huge, which could destabilize the training process. In practice, we constrain the shape parameter k such that $k \geq 0.1$ to avoid extreme value. A similar setting can be found in Fan et al. (2020), who view k as a hyperparameter and set it as 0.1.

Warm-up: The variational training criterion in Eq. (10) contains the likelihood term $p(\mathbf{x}_j | \Phi^{(1)}, \theta_j^{(1)})$ and the variational regularization term. During the early training stage, the variational regularization term causes some of the latent units to become inactive before their learning useful representation (Sønderby et al., 2016). We solve this problem by first training the parameters only using the reconstruction error and then adding the KL loss gradually with a temperature coefficient:

$$\begin{aligned} \mathcal{L} = & \sum_{j=1}^J \mathbb{E}_{q(\theta_j)} [\ln p(\mathbf{x}_j | \Phi^{(1)}, \theta_j^{(1)})] \\ & - \beta \sum_{j=1}^J \sum_{l=1}^L \text{KL}(q(\theta_j^{(l)}) || p(\theta_j^{(l)} | \Phi^{(l+1)}, \theta_j^{(l+1)})) \end{aligned} \quad (10)$$

where β is increased from 0 to 1 during the first N training epochs. This idea has been considered in previous VAE-based algorithms (Raiko et al., 2007; Higgins et al., 2016).

Gradient clipping: Optimizing the unbounded KL loss often causes the sharp gradient during training (Child, 2020), we address this by clipping gradient with a large L2-norm above a certain threshold, which we set 20 in all experiments. This technique can be easily implemented and allows networks to train smoothly.

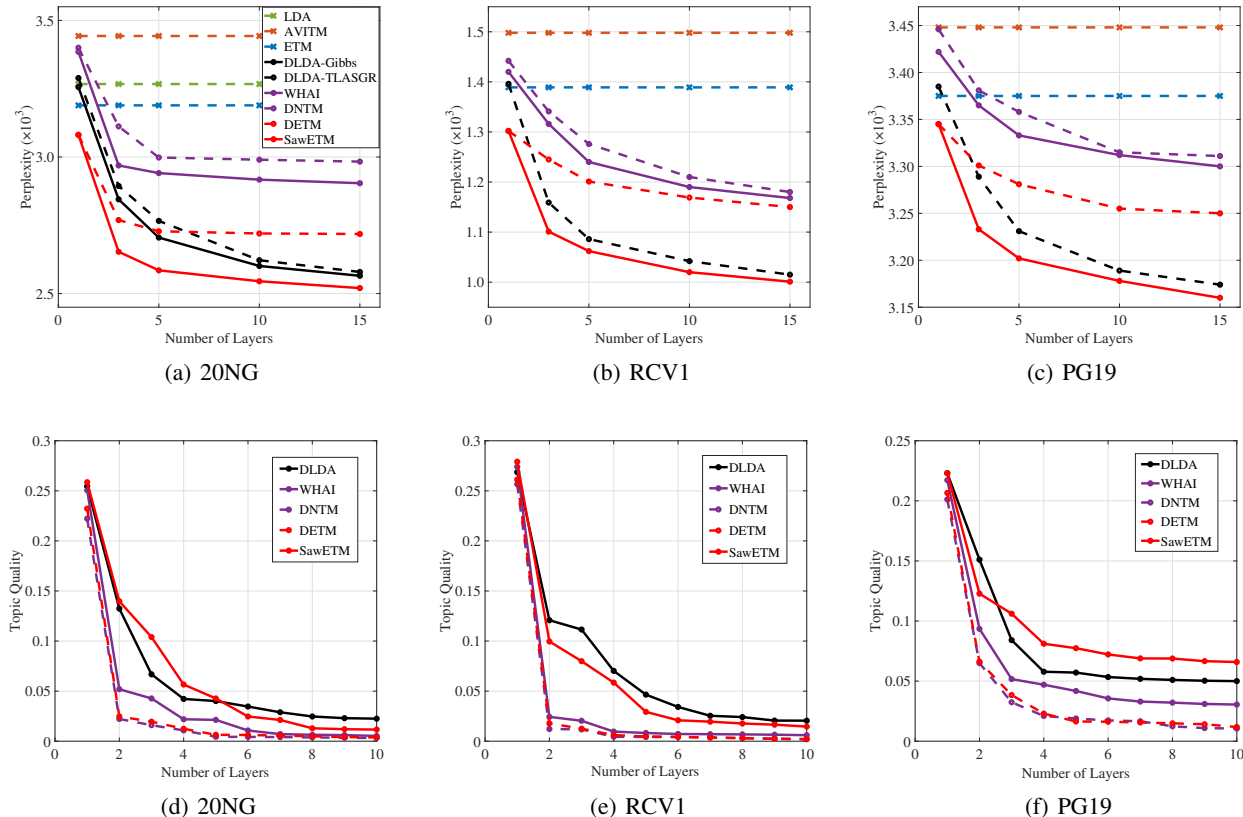


Figure 2. (a)-(c): Comparison of per-heldout-word perplexity (the lower the better). (d)-(f): Comparison of topic quality (the higher the better).

4. Experiments

In this paper, SawETM is proposed for extracting deep latent features and analyzing documents unsupervised. So we evaluate the effectiveness of the proposed model by unsupervised learning tasks in this section. Specifically, four widely-used performance measure of topic models are used in the experiments, which include perplexity, topic quality, and document clustering accuracy/normalized mutual information metric. The experiments are conducted on four real-world datasets including both regular and big corpora. In order to further understand the proposed model and verify our motivation, we visually inspect the learned word and topic embeddings as well as topic hierarchies. Our code is available at <https://github.com/BoChenGroup/SawETM>.

4.1. Experimental settings

Datasets. We run our experiments on four widely used benchmark corpora including **R8**, 20Newsgroups (**20NG**), Reuters Corpus Volume I (**RCV1**), and **PG-19**. The R8 dataset is a subset of the Reuters 21578 dataset, which con-

sists of 7,674 documents from 8 different review groups. R8 is partitioned into a training set of 5,485 ones and a testing set of 2,189 ones. The 20NG dataset, with a vocabulary size 36,534, consists of 18,774 documents from 20 different news groups and its average document length is 221. 20NG is split into a training set of 11,314 ones and a testing set of 7,532 one. The RCV1 dataset, with a vocabulary size 50,000, consists of 804,414 documents (Lewis et al., 2004) and its average document length is 140. The PG-19 dataset is extracted from Project Gutenberg (Rae et al., 2019) and contains 28,752 book. We first build a vocabulary with 20,000 words from this dataset, and then split each book with 1024 tokens, which result in 613,386 documents. For data processing, we first preprocess all the datasets by cleaning and tokenizing text, then removed stop words and low frequency words appearing less than 5 times, and finally select the N most frequent words to build a vocabulary. Note that, the R8 and 20NG datasets are used for document clustering experiments as there are ground-truth labels.

Baselines. We compare SawETM with basic Bayesian topic models and neural topic models: 1. LDA Group, single layer

topic models, including Latent Dirichlet Allocation (**LDA**) (Blei et al., 2003), which is a basic probability topic model; LDA with Products of Experts (**AVITM**) (Srivastava & Sutton, 2017), which replaces the mixture model in LDA with a product of experts and uses the variational inference update parameters; LDA with word embeddings (**ETM**) (Dieng et al., 2020), a generative model of documents that marries traditional topic models with word embeddings. 2. **DLDA Group**, hierarchical topic models, including Deep latent allocation inferred by Gibbs sampling (**DLDA-Gibbs**) (Zhou et al., 2015) and by TLASGR-MCMC (**DLDA-TLASGR**) (Cong et al., 2017); and Weibull hybrid autoencoding inference model for Deep latent allocation (**WHAI**) (Zhang et al., 2018), which employs a deep variational encoder to infer hierarchical document representations and updates the parameters by a hybrid of stochastic-gradient MCMC and variational inference. 3. **Variants Group**, variants of SawETM, including deep neural topic model (**DNTM**) as introduced in Eq. 2, which directly models the dependence via the learnable parameters W ; and deep embedding topic models (**DETM**) as introduced in Eq. 3, which employs a similar decomposition form as SC but does not share topic embeddings between two adjacent layers. These two Variant models use the same encoder with SawETM. For all the baseline models, we use their official default parameters with best-reported settings.

Setting. The hyperparameter settings used for the GBN group are similar to the ones used in Zhang et al. (2018). For the hierarchical topic models, the network structures of 15-layer models are set as [256, 224, 192, 160, 128, 112, 96, 80, 64, 56, 48, 40, 32, 16, 8]. For the embedding-based topic models such as ETM, DETM, and SawETM, we set the embedding size as 100. For the NTMs, we set the hidden size as 256. For optimization, the Adam optimizer (Kingma & Ba, 2014) is utilized with a learning rate of $1e^{-2}$. The mini-batch size is set as 200 in all experiments. All experiments are performed on Nvidia GTX 8000 GPU and coded with PyTorch.

4.2. Per-heldout-word perplexity

Per-heldout-word perplexity is a widely-used performance measure of topic models. Similar to Zhou et al. (2015), for each corpus, we randomly select 80% of the word token from each document to form a training matrix T , holding out the remaining 20% to form a testing matrix Y . We use T to train the model and calculate the per-held-word perplexity as

$$\exp \left\{ -\frac{1}{y_{..}} \sum_{v=1}^V \sum_{n=1}^N y_{vn} \ln \frac{\sum_{s=1}^S \sum_{k=1}^{K^1} \phi_{vk}^{(1)s} \theta_{kn}^{(1)s}}{\sum_{s=1}^S \sum_{v=1}^V \sum_{k=1}^{K^1} \phi_{vk}^{(1)s} \theta_{kn}^{(1)s}} \right\},$$

where S is the total number of collected samples and $y_{..} = \sum_{v=1}^V \sum_{n=1}^N y_{vn}$.

Fig. 2 (a)-(c) show how the perplexity changes as a function of the number of layers for various models over three different datasets. For both RCV1 and PG19, which are too large to run Gibbs sampling, we omit DLDA-Gibbs and only include DLDA-TLASGR for comparison. In the LDA group, ETM gets the best performance compared with LDA and AVITM, which can be attributed to the powerful word embeddings decoder (Dieng et al., 2020). LDA has the better performance compared with AVITM, which is not surprising as this batch algorithm can sample from the true posteriors given enough Gibbs sampling iterations. But these models are limited to single-layer shallow models, which can't benefit from the deep structure and results in the gap with the second group of models.

Among these DLDA based models in group two, the DLDA-Gibbs outperforms other models, attributed to the more accurate posterior estimations, while DLDA-TLASGR is a mini-batch algorithm and slightly degraded performance in out-of-sample prediction. As a hierarchical neural topic model, WHAI gets the worse performance compared with DLDA-Gibbs/DLDA-TLASGR. Meanwhile, we can see that WHAI with a single hidden layer clearly outperforms AVITM, indicating that using the Weibull distribution is more appropriate than using the logistic normal distribution to model the document latent representation. Besides, the performance of DLDA-Gibbs/DLDA-TLASGR can be effectively improved by increasing the number of layers, while WHAI fails to improve its performance when the layer size becomes greater than three. This maybe due to all layers of the natural topic models are trained by SGD together, which makes it difficult to learn meaningful prior (Wu et al., 2021). And similar to deep VAE, this phenomenon is called posterior collapse (Sønderby et al., 2016; Maaløe et al., 2019).

With the powerful word embedding decoder and effective Weibull upward-downward variational encoder, DETM of the variants group gets significant performance improvement with single layer. However, it also experience the similar problem with WHAI that no clear performance improvement when the number of layers becomes greater than three. Benefiting from the SC module between different layers, the learned knowledge at lower layers can flow to the upper layer, which can help the higher layers learn meaningful topics, resulting in better prior learned by SawETM. We can see that SawETM further improve the performance with the layer size becomes bigger, and get comparable performance with DLDA-Gibbs/DLDA-TLASGR. Note that, although the improvement of SawETM is not that significant compared with DLDA-Gibbs/DLDA-TLASGR, DLDA-Gibbs/DLDA-TLASGR require iterative sampling to infer latent document representations in the testing stage, while SawETM can infer latent representations via direct projection, which makes it both scalable to large corpora

and fast in out-of-sample prediction. Besides, thanks to SC and the improved encoder, SawETM can significantly outperform other NTMs.

4.3. Topic quality

A good topic model can provide interpretable topics. In this section, we measure the model’s performance in terms of topic interpretable (Dieng et al., 2020). Specifically, topic coherence and topic diversity are combined here to evaluate topic interpretable/quality. Topic coherence is obtained by taking the average Normalized Pointwise Mutual Information (NPMI) of the top 20 words of each topic (Aletras & Stevenson, 2013). It provides a quantitative measure of the interpretability of a topic (Mimno et al., 2011). The second metric is topic diversity (Dieng et al., 2020), which denotes the percentage of unique words in the top 20 words of all topics. Diversity close to 1 means more diverse topics. Topic quality is defined as the product between topic coherence and diversity.

Fig. 2 (d)-(f) show the topic quality results of different layer for various models over three different datasets. Clearly, DLDA performs the best in terms of topic quality especially the higher layer, which is not surprising as all its parameters are updated by Gibbs sampling/TLASGR-MCMC (Cong et al., 2017). Thanks to the use TLASGR-MCMC rather than a simple SGD procedure, WHAI consistently outperforms DNTM and DETM, which update all the parameters by SGD. Although equipped with a powerful word embeddings decoder, the topic quality of DETM clearly decreases as the number of layers increases. Through the above experimental phenomenon, we can find that it is difficult for NTMs to learn meaningful hierarchical topics. This is probably because NTMs often suffer from the posterior collapse problem in VAEs, making it hard to learn deeper semantic structure. However, SawETM achieving comparable performance to DLDA, which clearly outperforms the other deep neural topic models. As discussed in the Sec. 4.2, this improvement come from the SC module. The results of topic quality also agree with the results of perplexity, which are shown in Fig. 2 (a)-(c).

4.4. Document clustering

We consider the multi-class classification task for predicting the clusters for test documents to evaluate the quality of latent document representations extracted by these models. In detail, we use the trained topic models to extract the latent representations of the testing documents and then use k-means to predict the clusters. As shown in Table 1, the accuracy (AC) and normalized mutual information metric (NMI) are used to measure the clustering performance, both of which are the higher the better. Tab. 1 shows the clustering results of all the models on 20NG and R8 dataset. It can be observed that with powerful word embeddings decoder

Table 1. Results of AC and NMI for document clustering task.

Model	Layer	20News		R8	
		AC	NMI	AC	NMI
LDA	1	46.52	45.15	51.41	40.47
AVITM	1	48.31	46.33	52.43	41.20
ETM	1	49.79	48.40	55.34	41.28
PGBN	1	46.62	45.43	51.67	40.76
PGBN	5	48.33	46.51	54.21	41.21
WHAI	1	49.43	46.56	57.86	42.31
WHAI	5	49.51	46.98	60.45	43.98
DNTM	1	49.17	46.32	57.58	42.12
DNTM	5	49.25	46.79	59.93	43.90
DETM	1	50.24	48.69	61.21	43.45
DETM	5	50.33	48.87	61.86	44.12
SawETM	5	51.25	50.77	63.82	45.90

and the Sawtooth Connection, SawETM can extract more expressive document latent representations and outperforms the other models included for comparison.

4.5. Qualitative analysis

One of the most appealing properties of SawETM is interpretability, we can visually inspect the inferred topics at different layers and the inferred connection weights between the topics of adjacent layers. Specifically, we conduct an extensive qualitative evaluation on the quality of the topics discovered by SawETM, including word embedding, topic embedding, and topic hierarchies. In this section, we use a 15-layer SawETM trained on PG-19 for the visualization of embedding space and hierarchical structure experiments.

Visualisation of embedding space The top 10 words from six topics are visualized in Fig. 3(a) by t-SNE visualization (Van der Maaten & Hinton, 2008). As we can see, the words under the same topic stay closer together, and the words under different topics are far apart. Besides, the words under the same topic are semantically more similar, which can demonstrate the meaning of the learned word embeddings. Apart from the visualization of word embeddings, we visualize the topic embeddings. As shown in Fig. 3(b), we select the top 2 sub-topics of the topic at the second layer. We can see that the sub-topic from the same topic have semantic similarity, and are closer in the embedding space. The above experiments show that the proposed SawETM learns not only meaningful word embeddings but also meaningful topic embeddings. More importantly, the learned words and hierarchical topics embedding can be projected into the same space, which can further support our motivations.

Hierarchical structure of topic model: Limited by the paper space, we only visualize topics at the top two layers and the bottom two layers. As shown in Fig. 4, the semantic

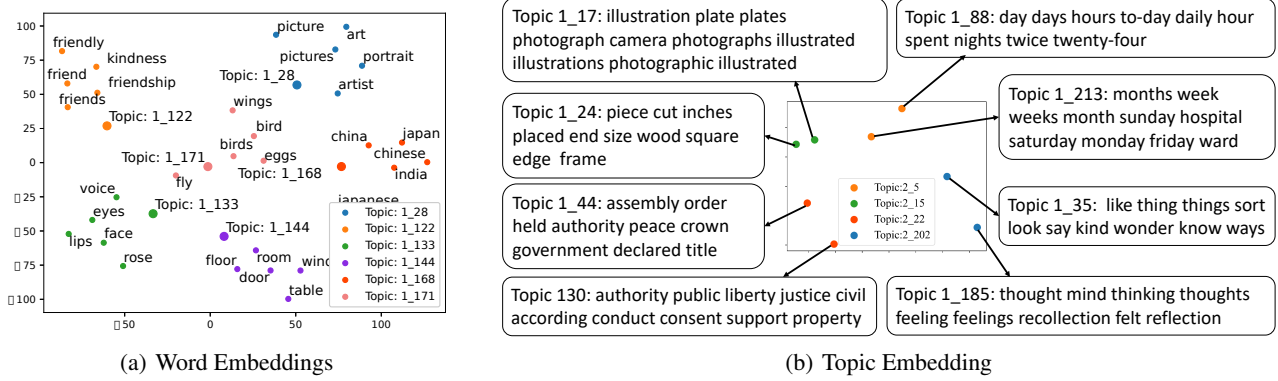


Figure 3. t-SNE visualisation of (a) word embeddings, which we choose the top ten words for each topic at layer one and (b) topic embeddings, which we choose the top two sub topics for each topic at layer two. (Note that the Topic: t_j denotes the j^{th} topic at t^{th} layer.)

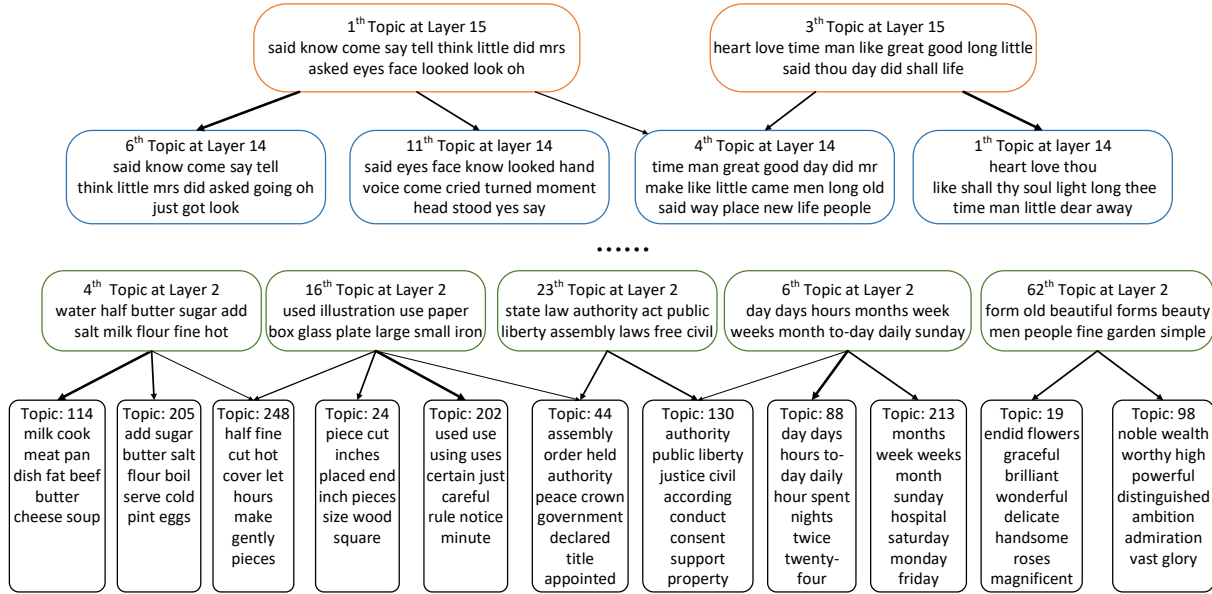


Figure 4. An example of hierarchical topics learned from PG-19 by a 15-layer SawETM, We only show example topics at the top two layers and bottom two layers.

meaning of each topic and the connections between the topics of adjacent layers are highly interpretable. In particular, SawETM can learn meaningful hierarchical topics at higher layers, indicating that it is able to support a deep structure.

5. Conclusion

In this paper, we propose SawETM, a deep neural topic model that captures the dependencies and semantic similarities between the topics at different layers. We design a skip-connected upward-downward inference network to approximate the posterior distribution of a document. Note that with the Sawtooth Connection technique, SawETM provides different views to a deep topic model, and further improves the performance of the neural deep topic model.

As a fully variational deep topic model, SawETM can be optimized by SGD. Extensive experiments have shown that SawETM achieves comparable performance on perplexity, document clustering, and topic quality with the start-of-the-art model. In addition, with learned word and topic embeddings, and topic hierarchies, SawETM can discover interpretable structured topics, which helps to gain a better understanding of text data.

Acknowledgments

Bo Chen acknowledges the support of NSFC (61771361), Shaanxi Youth Innovation Team Project, the 111 Project (No. B18039) and the Program for Oversea Talent by Chinese Central Government.

References

- Aletras, N. and Stevenson, M. Evaluating topic coherence using distributional semantics. In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013)—Long Papers*, pp. 13–22, 2013.
- Bengio, Y., Ducharme, R., Vincent, P., and Jauvin, C. A neural probabilistic language model. *Journal of machine learning research*, 3(Feb):1137–1155, 2003.
- Blei, D. M., Ng, A. Y., and Jordan, M. I. Latent dirichlet allocation. *Journal of machine Learning research*, 3(Jan): 993–1022, 2003.
- Blei, D. M., Lafferty, J. D., et al. A correlated topic model of science. *The Annals of Applied Statistics*, 1(1):17–35, 2007.
- Blei, D. M., Griffiths, T. L., and Jordan, M. I. The nested chinese restaurant process and bayesian nonparametric inference of topic hierarchies. *Journal of the ACM (JACM)*, 57(2):1–30, 2010.
- Burkhardt, S. and Kramer, S. Decoupling sparsity and smoothness in the dirichlet variational autoencoder topic model. *Journal of Machine Learning Research*, 20(131): 1–27, 2019.
- Card, D., Tan, C., and Smith, N. A. A neural framework for generalized topic models. *arXiv preprint arXiv:1705.09296*, 2017.
- chen, w., Wang, C., Chen, B., Liu, Y., Zhang, H., and Zhou, M. Bidirectional convolutional poisson gamma dynamical systems. In *Advances in Neural Information Processing Systems*, volume 33, 2020.
- Child, R. Very deep vaes generalize autoregressive models and can outperform them on images. *arXiv preprint arXiv:2011.10650*, 2020.
- Cong, Y., Chen, B., Liu, H., and Zhou, M. Deep latent dirichlet allocation with topic-layer-adaptive stochastic gradient riemannian mcmc. *arXiv preprint arXiv:1706.01724*, 2017.
- Dieng, A. B., Ruiz, F. J., and Blei, D. M. The dynamic embedded topic model. *arXiv preprint arXiv:1907.05545*, 2019.
- Dieng, A. B., Ruiz, F. J., and Blei, D. M. Topic modeling in embedding spaces. *Transactions of the Association for Computational Linguistics*, 8:439–453, 2020.
- Duan, Z., Zhang, H., Wang, C., Wang, Z., Chen, B., and Zhou, M. EnsLM: Ensemble language model for data diversity by semantic clustering. In *ACL-IJCNLP 2021: The Joint Conference of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing*, Aug. 2021.
- Fan, X., Zhang, S., Chen, B., and Zhou, M. Bayesian attention modules. *arXiv preprint arXiv:2010.10604*, 2020.
- Gan, Z., Chen, C., Henao, R., Carlson, D., and Carin, L. Scalable deep poisson factor analysis for topic modeling. In *International Conference on Machine Learning*, pp. 1823–1832, 2015.
- Henao, R., Gan, Z., Lu, J., and Carin, L. Deep poisson factor modeling. *Advances in Neural Information Processing Systems*, 28:2800–2808, 2015.
- Higgins, I., Matthey, L., Pal, A., Burgess, C., Glorot, X., Botvinick, M., Mohamed, S., and Lerchner, A. beta-vae: Learning basic visual concepts with a constrained variational framework. 2016.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Kingma, D. P. and Welling, M. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- Levy, O. and Goldberg, Y. Neural word embedding as implicit matrix factorization. *Advances in neural information processing systems*, 27:2177–2185, 2014.
- Lewis, D. D., Yang, Y., Rose, T. G., and Li, F. Rcv1: A new benchmark collection for text categorization research. *Journal of machine learning research*, 5(Apr):361–397, 2004.
- Li, C., Wang, H., Zhang, Z., Sun, A., and Ma, Z. Topic modeling for short texts with auxiliary word embeddings. In *Proceedings of the 39th International ACM SIGIR conference on Research and Development in Information Retrieval*, pp. 165–174, 2016.
- Maaløe, L., Fraccaro, M., Liévin, V., and Winther, O. Biva: A very deep hierarchy of latent variables for generative modeling. In *Advances in neural information processing systems*, pp. 6551–6562, 2019.
- Miao, Y., Yu, L., and Blunsom, P. Neural variational inference for text processing. In *International conference on machine learning*, pp. 1727–1736, 2016.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013a.
- Mikolov, T., Sutskever, I., Chen, K., Corrado, G. S., and Dean, J. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26:3111–3119, 2013b.

- Mimno, D., Wallach, H., Talley, E., Leenders, M., and McCallum, A. Optimizing semantic coherence in topic models. In *Proceedings of the 2011 conference on empirical methods in natural language processing*, pp. 262–272, 2011.
- Nan, F., Ding, R., Nallapati, R., and Xiang, B. Topic modeling with wasserstein autoencoders. *arXiv preprint arXiv:1907.12374*, 2019.
- Paisley, J., Wang, C., Blei, D. M., and Jordan, M. I. Nested hierarchical dirichlet processes. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 37(2):256–270, 2014.
- Pennington, J., Socher, R., and Manning, C. D. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Rae, J. W., Potapenko, A., Jayakumar, S. M., Hillier, C., and Lillicrap, T. P. Compressive transformers for long-range sequence modelling. *arXiv preprint*, 2019. URL <https://arxiv.org/abs/1911.05507>.
- Raiko, T., Valpola, H., Harva, M., and Karhunen, J. Building blocks for variational bayesian learning of latent variable models. *Journal of Machine Learning Research*, 8(1), 2007.
- Razavi, A., Oord, A. v. d., Poole, B., and Vinyals, O. Preventing posterior collapse with delta-vaes. *arXiv preprint arXiv:1901.03416*, 2019.
- Rezende, D. J., Mohamed, S., and Wierstra, D. Stochastic backpropagation and variational inference in deep latent gaussian models. In *International Conference on Machine Learning*, volume 2, 2014.
- Sønderby, C. K., Raiko, T., Maaløe, L., Sønderby, S. K., and Winther, O. Ladder variational autoencoders. In *Advances in neural information processing systems*, pp. 3738–3746, 2016.
- Srivastava, A. and Sutton, C. Autoencoding variational inference for topic models. *arXiv preprint arXiv:1703.01488*, 2017.
- Vahdat, A. and Kautz, J. Nvae: A deep hierarchical variational autoencoder. *arXiv preprint arXiv:2007.03898*, 2020.
- Van der Maaten, L. and Hinton, G. Visualizing data using t-sne. *Journal of machine learning research*, 9(11), 2008.
- Wang, C., Zhang, H., Chen, B., Wang, D., Wang, Z., and Zhou, M. Deep relational topic modeling via graph poisson gamma belief network. *Advances in Neural Information Processing Systems*, 33, 2020a.
- Wang, R., Hu, X., Zhou, D., He, Y., Xiong, Y., Ye, C., and Xu, H. Neural topic modeling with bidirectional adversarial training. *arXiv preprint arXiv:2004.12331*, 2020b.
- Wang, Z., Wang, C., Zhang, H., Duan, Z., Zhou, M., and Chen, B. Learning dynamic hierarchical topic graph with graph convolutional network for document classification. In *International Conference on Artificial Intelligence and Statistics*, pp. 3959–3969. PMLR, 2020c.
- Wu, B., Nair, S., Martin-Martin, R., Fei-Fei, L., and Finn, C. Greedy hierarchical variational autoencoders for large-scale video prediction. *arXiv preprint arXiv:2103.04174*, 2021.
- Zhang, H., Chen, B., Guo, D., and Zhou, M. Whai: Weibull hybrid autoencoding inference for deep topic modeling. *arXiv preprint arXiv:1803.01328*, 2018.
- Zhang, H., Chen, B., Tian, L., Wang, Z., and Zhou, M. Variational hetero-encoder randomized gans for joint image-text modeling. *arXiv preprint arXiv:1905.08622*, 2019.
- Zhao, H., Du, L., and Buntine, W. A word embeddings informed focused topic model. In *Asian Conference on Machine Learning*, pp. 423–438, 2017.
- Zhao, H., Du, L., Buntine, W., and Zhou, M. Dirichlet belief networks for topic structure learning. *arXiv preprint arXiv:1811.00717*, 2018a.
- Zhao, H., Du, L., Buntine, W., and Zhou, M. Inter and intra topic structure learning with word embeddings. In *International Conference on Machine Learning*, pp. 5892–5901, 2018b.
- Zhao, H., Phung, D., Huynh, V., Jin, Y., Du, L., and Buntine, W. Topic modelling meets deep neural networks: A survey. *arXiv preprint arXiv:2103.00498*, 2021.
- Zhou, M., Hannah, L., Dunson, D., and Carin, L. Beta-negative binomial process and poisson factor analysis. In *Artificial Intelligence and Statistics*, pp. 1462–1471, 2012.
- Zhou, M., Cong, Y., and Chen, B. The poisson gamma belief network. *Advances in Neural Information Processing Systems*, 28:3043–3051, 2015.