

Neural Attention-Aware Hierarchical Topic Model

Yuan Jin¹ He Zhao¹ Ming Liu² Lan Du^{1*} Wray Buntine¹

¹Faculty of Information Technology, Monash University, Australia
{yuan.jin, ethan.zhao, lan.du, wray.buntine}@monash.edu

²School of Information Technology, Deakin University, Australia
m.liu@deakin.edu.au

Abstract

Neural topic models (NTMs) apply deep neural networks to topic modelling. Despite their success, NTMs generally ignore two important aspects: (1) only document-level word count information is utilized for the training, while more fine-grained sentence-level information is ignored, and (2) external semantic knowledge regarding documents, sentences and words are not exploited for the training. To address these issues, we propose a variational autoencoder (VAE) NTM model that jointly reconstructs the sentence and document word counts using combinations of bag-of-words (BoW) topical embeddings and pre-trained semantic embeddings. The pre-trained embeddings are first transformed into a common latent topical space to align their semantics with the BoW embeddings. Our model also features hierarchical KL divergence to leverage embeddings of each document to regularize those of their sentences, thereby paying more attention to semantically relevant sentences. Both quantitative and qualitative experiments have shown the efficacy of our model in 1) lowering the reconstruction errors at both the sentence and document levels, and 2) discovering more coherent topics from real-world datasets.

1 Introduction

Topic models are a family of powerful techniques that can effectively discover human-interpretable topics from unstructured corpora for text analysis purposes. Among them, Bayesian topic models, based on the latent Dirichlet allocation (LDA) (Blei et al., 2003), have been the mainstream for nearly two decades. They usually adopt count/bag-of-words (BoW) representations for text content and model the generation of the BoW data with a probabilistic structure of latent variables. These variables follow pre-specified distributions under Bayes' theorem, and are learned through Bayesian inference

such as variational inference (VI) and Monte Carlo Markov chain (MCMC) sampling. Despite their success, conventional Bayesian topic models, however, are known to lack flexibility in their model structures and scalability to large volumes of data.

To address the above limitations, increasing effort has been made in leveraging deep neural networks (DNNs) for topic modelling, which leads to the so-called neural topic models (NTMs) (Zhao et al., 2021a). Most of these models follow the framework of variational auto-encoders (VAEs) (Kingma and Welling, 2014; Rezende et al., 2014) and adopt an encoder-decoder architecture, in which the encoder transforms the BoW data of each document into the corresponding document-topical embeddings, and the decoder attempts to map these embeddings back to the same data. With a moderate increase in model complexity, NTMs have largely outperformed conventional topic models on BoW data reconstruction and topic interpretability/coherence (Miao et al., 2016; Srivastava and Sutton, 2017; Ding et al., 2018; Zhou et al., 2020; Zhao et al., 2021b).

With this being said, most NTMs only exploit the BoW information of internal documents while ignoring (1) the sentence-level BoW information of these documents, and (2) the external (semantic) information regarding the documents, sentences and words (e.g., extracted from other larger relevant corpora). These limitations have hindered the further performance improvement for NTMs. Therefore, in this paper, we propose a new NTM that address these limitations. It jointly reconstructs the BoW data of each document and their sentences with combinations of both internal BoW topical embeddings and external pre-trained semantic embeddings. To do this, we design an *internal BoW data encoder* and an *external knowledge encoder* to respectively transform the BoW data and the pre-trained embeddings of the same documents and sentences into a shared latent topical space. The

Corresponding author

resulting internal and external topical embeddings are then combined to decode the BoW data.

To address the BoW data sparsity (Zhao et al., 2019) at the sentence level, which has been a problem for many topic models, our model enforces hierarchical KL divergence on pairs of sentences and corresponding documents with respect to their topical embeddings derived from both the BoW data and the external knowledge. The intent is that both types of topical embeddings for each sentence should be governed by the same types of embeddings of their parent documents. Furthermore, the hierarchical KL for a sentence is weighted by its semantic relations to its parent document, so that a document’s topical information is more influential to more semantically representative sentences. Our contribution can be summarized as follows:

- We propose a VAE-based neural topic model, which encodes internal BoW information and external semantic knowledge specific to word, sentence and document levels into the same latent topical space to refine topic quality.
- Our model imposes attention-weighted hierarchical KL divergence on pairs of sentences and documents to smooth the learning of the topical embeddings from sparse BoW data of sentences.
- We demonstrate that our model is effective in BoW data reconstruction at both the document and sentence levels. It also improves the internal and external coherence of the discovered topics.

2 Related Work

To our best knowledge, most state-of-the-art NTMs, like ours, are based on the VAE framework. For a detailed discussion of these NTMs, we refer readers to Zhao et al. (2021a). Here, we only discuss the two lines of research that are most relevant to ours.

NTMs with pre-trained language models Recently, pre-trained transformer-based language models such as BERT (Devlin et al., 2019) have started to draw the attention of the topic modelling community thanks to their capability of generating contextualized word embeddings with rich semantic information that is absent from the BoW data. Thus, an emerging trend focuses on incorporating these contextualized word embeddings into NTMs.

Based on the popular VAE framework of Srivastava and Sutton (2017), Bianchi et al. (2021) proposed a contextualized topic model (CTM) that incorporates pre-trained document embeddings generated by sentence-transformers (Reimers and

Gurevych, 2019). However, unlike ours, CTM ignores the other levels of pre-trained knowledge and BoW information. Chaudhary et al. (2020) proposed to combine an NTM with a fine-tuned BERT model by concatenating the topic distribution and the learned BERT embedding of a document as the features for document classification. Hoyle et al. (2020) proposed BAT, a NTM framework “taught” by external knowledge distilled from a pre-trained BERT model. BERT predicts probabilities for each word of a document which are then averaged to generate a pseudo BoW vector for the document. The BAT framework can be instantiated with various existing NTMs such as Scholar (Card et al., 2018) (i.e. BAT+Scholar), which imposes a logistic Normal as the variational posteriors for document embeddings in the VAE, and W-LDA (Nan et al., 2019) (i.e. BAT+W-LDA), which replaces the KL divergence with the maximum mean discrepancy (Gretton et al., 2012).

NTMs for modelling document structures Although document structures, i.e., the structured relationships between documents, paragraphs, and sentences have been modelled in conventional Bayesian topic models (e.g., in Du et al. (2012); Balikas et al. (2016a); Jiang et al. (2019)), they have not been carefully studied in NTMs to our knowledge. The closest work to our idea is Nallapati et al. (2017), which proposes an NTM that samples a topic for each sentence of an input document and then generates the word sequence of the sentence with an RNN conditioned on the sentence’s topic. However, this work focuses on sentence generation instead of topic modelling.

3 Preliminaries

In this section, we introduce the VAE framework of neural variational topic models, based on which our model will be developed. Table 1 details the notations and symbols used throughout the paper.

Problem Formulation Consider a corpus that consists of I documents where the i^{th} ($1, \dots, I$) document is represented as a V -dimensional vector of word counts, w_i^D , also known as the bag-of-words (BoW) data. Here, V is the size of the vocabulary from the corpus, and w_{iv}^D is the number of times the v^{th} ($1, \dots, V$) word occurs in the i -th document. Topic modelling assumes that there exist K topics that can be used to describe each document. Its goal is to recover these topics from the BoW data

Symbols	Description
I, V, K, L, J_i, M	size of corpus, vocabulary, topics, hidden neurons, number of sentences in document $i \in (1, \dots, I)$ and pre-trained embedding dimension
N	a vector of the lengths of each document
w_i^D, w_{ij}^S	BoW vectors of the i^{th} document and its j^{th} sentence. ($\mathbb{Z}_{\geq 0}^V$)
$\mathbf{W}^D, \mathbf{W}_i^S$	BoW matrices over the I documents ($\mathbb{Z}_{\geq 0}^{I \times V}$) and the J_i sentences of document i ($\mathbb{Z}_{\geq 0}^{J_i \times V}$)
$\mathbf{X}^W, \mathbf{X}_i^S, \mathbf{X}^D$	matrices of pre-trained non-contextual word embeddings ($\mathbb{R}^{V \times M}$), pre-trained sentence embeddings ($\mathbb{R}^{J_i \times M}$) and document embeddings ($\mathbb{R}^{I \times M}$)
$q^h(\cdot)$	the VAE encoder for $h \in \{\text{“B”}, \text{“E”}\}$ that generate the topical embeddings of words, sentences and documents
$f^h(\cdot)$	a shared MLP of the encoder for h that maps \mathbb{R}^V to \mathbb{R}^L
$l_1^h(\cdot), l_2^h(\cdot)$	linear layers over the shared MLP of the encoder for h that maps \mathbb{R}^L to \mathbb{R}^K
Φ^h	word-topic embedding matrix over each word specific to the encoder for h , ($\mathbb{R}^{V \times K}$)
s_{ij}^h, \mathbf{S}_i^h	sentence-topic embedding of the j^{th} ($1, \dots, J_i$) sentence in document i generated by the encoder for h , (\mathbb{R}^K); sentence-topic matrix, ($\mathbb{R}^{J_i \times K}$)
z_i^h, \mathbf{Z}^h	topical embedding of document i generated by the encoder for h , ($\mathbb{R}_{\geq 0}^K$); document-topic matrix, ($\mathbb{R}_{\geq 0}^{I \times K}$)
α_{ij}^S	attention weight of sentence j for document i , (\mathbb{R})

Table 1: List of Notation used. Here, h is either “B” or “E” representing respectively the internal BoW data and the external semantic knowledge.

of the corpus \mathbf{W}^D . In this paper, we further formulate the topic modelling problem at the level of sentences; that is the i^{th} document comprising a total of J_i sentences. In this case, the document can be alternatively represented as a word count matrix \mathbf{W}_i^S , an additional level of BoW information we aim to leverage. The j^{th} ($1, \dots, J_i$) row of this matrix accommodates the V -dimensional BoW vector w_{ij}^S of the j^{th} sentence from the i^{th} document.

Neural Variational Topic Model Most traditional topic models are graphical models with exponential family model parameters. Therefore, they yield tractable inference for the posterior distributions of the model parameters. However, the inference is limited in expressiveness, thereby less capable of capturing the true underlying generative process and distributions. To solve this problem, neural variational inference is introduced into topic models, named neural variational topic models (NVTM), where more expressive posterior dis-

tributions are constructed for the model parameters using neural networks. A typical NVTM is learned by maximizing the Evidence Lower Bound (ELBO) of the marginal likelihood of the BoW data \mathbf{W}^D of each document with respect to their topical embeddings \mathbf{Z}^B :

$$\mathbb{E}_{\mathbf{Z}^B \sim q^B(\mathbf{W}^D)} [\log p(\mathbf{W}^D | \mathbf{Z}^B)] - \text{KL}[q(\mathbf{Z}^B | \mathbf{W}^D) || p(\mathbf{Z}^B)] \quad (1)$$

where $\log p(\mathbf{W}^D | \mathbf{Z}^B)$, $q(\mathbf{Z}^B | \mathbf{W}^D)$ and $p(\mathbf{Z}^B)$ are respectively the BoW data likelihood, the variational posterior distribution and the prior distribution of \mathbf{Z}^B . A key characteristic of NVTM is to model the first two terms as decoder and encoder networks. The encoder network $q^B(\mathbf{W}^D)$ models the posterior distribution $q(\mathbf{Z}^B | \mathbf{W}^D)$ by taking in the BoW data \mathbf{W}^D to parameterize diagonal Gaussian distributions¹ over \mathbf{Z}^B . Specifically, for document i , its encoding process is formulated as:

$$\begin{aligned} \mu_i^B &:= l_1^B(f^B(w_i^D)) \\ \text{diag}(\Sigma_i^B) &:= l_2^B(f^B(w_i^D)) \\ z_i^B \sim q^B(w_i^D) &:= \mathcal{N}(\mu_i^B, \text{diag}(\Sigma_i^B)\mathbb{I}) \\ \text{Output with } z_i^B &:= \text{Softmax}(z_i^B) \end{aligned} \quad (2)$$

Here, $f^B : \mathbb{R}^V \rightarrow \mathbb{R}^L$ is a multi-layer perceptron (MLP) for computing a shared hidden layer output from w_i^D , where L is the number of hidden neurons. The functions $l_1^B, l_2^B : \mathbb{R}^L \rightarrow \mathbb{R}^K$ are two linear layers for respectively predicting the posterior mean and variance vectors: $\mu_i^B, \text{diag}(\Sigma_i^B) \in \mathbb{R}^K$ that generate the document-topical embedding z_i^B for document i . The symbol $\mathbb{I} \in \mathbb{R}^{K \times K}$ denotes an identity matrix used to create the diagonal Gaussian. Finally, it is common for NTMs to transform the topical embedding z_i^B into a probability distribution of topics for each document using Softmax.

As for the NVTM’s decoder network, it reconstructs \mathbf{W}^D using the document-topic matrix \mathbf{Z}^B and the word-topic matrix Φ^B as follows:

$$\mathbf{W}^D \sim \text{Multi}(\text{Softmax}(\mathbf{Z}^B(\Phi^B)^T), \mathbf{N}) \quad (3)$$

where \mathbf{N} is a vector of document lengths and Φ^B , capturing the latent topics of each word in the vocabulary, can also be viewed as the weight matrix of the (output layer) of the decoder network.

4 Our Model

Our proposed model, **Neural Attention-aware Hierarchical Topic Model (NAHTM)**, comprises

¹In this paper, we follow the original VAE’s parametrization of the posterior distribution as a diagonal Gaussian distribution.




Figure 1: Basic Architecture of NAHTM. Each colour highlights a level of information, including the word-level (red), sentence-level (orange) and document-level (green) internal BoW and external semantic information. The blue dash line highlights the hierarchical KL (HKL) constraints over the sentence-document pairs.

(1) two types of encoders, *internal BoW data encoder* and *external knowledge encoder*, that capture latent topics of documents, sentences and words from both internal and external sources, respectively; (2) an attention-aware hierarchical KL divergence that regularizes topical embeddings of sentences with those of their documents. Figure 1 shows the basic architecture of NAHTM.

4.1 Internal BoW Data Encoder

This encoder, $q^B(\cdot)$, aims to capture the document- and sentence-level BoW data information. The encoding process of the former has been specified in eq. (2), which yields the posterior distributions $q^B(\mathbf{W}^D)$ for the document-topic matrix \mathbf{Z}^B .

As for encoding the sentence-level BoW data, it is motivated by the fact that sentences convey complete logical statements organized by topics as documents. The difference, as argued by the past research on LDA models, is that sentences are more concise with shorter text and focused topics (Balikas et al., 2016a,b; Amoualian et al., 2017). NAHTM encodes the sentence-level BoW data in the same way as it encodes the document-level data except for the final activation function. More specifically, each document is now viewed as a corpus, while each sentence is viewed as a (short) document. For document i , its BoW data \mathbf{W}_i^S , over its sentences, is encoded as $q^B(\mathbf{W}_i^S)$ with the same encoding process as in eq. (2). Then, the sentence-topical embedding matrix \mathbf{S}_i^B for document i is generated as: $\mathbf{S}_i^B \sim q^B(\mathbf{W}_i^S)$.

Based on the well-grounded argument that a

sentence should be bound in topics, NAHTM forces \mathbf{S}_i^B to be sparse over topics by using Sparsemax (Martins and Astudillo, 2016; Lin et al., 2019) which projects real-valued embeddings into sparse probability vectors: $\mathbf{S}_i^B := \text{Sparsemax}(\mathbf{S}_i^B)$. Specifically, for the j^{th} sentence of document i , its embedding s_{ij}^B is converted by Sparsemax as:

$$\text{Sparsemax}(s_{ij}^B) := \text{argmin}_c \|c - s_{ij}^B\|^2 \quad (4)$$

where c lies on the $(K - 1)$ -dimensional probability simplex. In other words, Sparsemax projects s_{ij}^B from the Euclidean space onto the probability simplex.

4.2 External Knowledge Encoder

External semantic knowledge, extracted by pre-trained language models (e.g. BERT (Devlin et al., 2019)) from large general corpora, provides rich prior information regarding the contexts and semantic relatedness of instances for each entity (i.e. document, sentence and word) modelled by NTMs. The language models account for ordering patterns of the entities (i.e. sentence and word orderings), which are complementary to the orderless BoW information captured by the NTMs. Incorporating such knowledge into the NTMs can potentially help them better infer sentences' true topics in scenarios which cannot be distinguished by the BoW information. For example, a pair of sentences with the same word counts might have very different topics due to different word orderings. Meanwhile, another pair without any word overlap might still have strongly correlated topics due to a next-sentence or entailment relationship. Hence, NTMs can be guided to better infer topics of documents as well as the entire set of topics underlying the corpus.

Another advantage of external knowledge is that it can potentially alleviate the data sparsity problem under the BoW modelling, especially at the sentence level. Since sentences have much shorter text compared to documents, therefore, their BoW data is also much sparser with significantly fewer word co-occurrences within each sentence and word overlaps in between. To make the learning less affected by the sparse data, external knowledge can be leveraged to calibrate it with prior information regarding the sentences. NAHTM incorporates three levels of external knowledge in the form of the following pre-trained embeddings for words, sentences and documents.

External word embeddings X^W are output by the embedding layer of the pre-trained transformer

model for each word in the vocabulary. Here, \mathbf{X}^W are non-contextualized and untrainable embeddings whose dimension is predefined by the pre-trained transformer and thus, not necessarily equal to the number of topics. The reason for using the non-contextualized word embedding is that the word-level external information is expected to be “injected” correspondingly into Φ^B , the factorized and non-contextualized word-topical embeddings in our topic model.

External sentence embeddings \mathbf{X}_i^S for the sentences of document i are the aggregated results of the outputs from the last transformer encoder (layer) of the pre-trained model. In this case, the inputs to the pre-trained model consist of the word sequences for each sentence. The outputs are the contextualized embeddings of the words in the sequences. In this paper, we use sentence-transformers², a programmable framework that provides a variety of pre-trained transformer models for computing sentence embeddings. We adopt the default aggregation strategy of sentence-transformers which is to average all the (output) contextualized word embeddings over the sentence.

External document embeddings \mathbf{X}^D can be constructed as either the unweighted or the weighted average of all the sentence embeddings for the same document. Specifically, for the latter case, \mathbf{x}_i^D can be obtained as follows:

$$\begin{aligned} y_{ij}^S &:= \sigma_1(\mathbf{x}_{ij}^S \Theta_1 + \mathbf{b}_1) \theta_2 \\ \alpha_i^S &:= \sigma_2(\mathbf{y}_i^S) \\ \mathbf{x}_i^D &:= \sum_{j=1}^{J_i} \alpha_{ij}^S \mathbf{x}_{ij}^S \end{aligned} \quad (5)$$

where $\alpha_i^S = [\alpha_{i1}^S, \dots, \alpha_{iJ_i}^S]^T$ is the attention (weight) vector over each sentence. Its j^{th} element α_{ij}^S is the normalized weight of the j^{th} sentence in representing document i , i.e. $\sum_{j=1}^{J_i} \alpha_{ij}^S = 1$; $\mathbf{y}_i^S = [y_{i1}^S, \dots, y_{iJ_i}^S]^T$ is the corresponding unnormalized attention vector; $\Theta_1 \in \mathbb{R}^{M \times M}$, $\theta_2 \in \mathbb{R}^M$ and $\mathbf{b}_1 \in \mathbb{R}^M$ are the attention parameters with M being the (pre-defined) dimension of the pre-trained embedding; σ_1, σ_2 are the activation functions. For σ_1 , we use the Tanh function for all the experiments. For σ_2 , it can be either the conventional Softmax function or the Sparsemax function which induces sparsity over α_i^S such that unimportant sentences tends to have zero weights in

representing the document.

Mapping $\mathbf{X}^W, \mathbf{X}^D, \mathbf{X}_i^S$ to the topical space is done subsequently by the external knowledge encoder to align the dimension and semantic meanings of external embeddings with the topical embeddings. More specifically, each level of the external knowledge data $\mathbf{X} \in \{\mathbf{X}^W, \mathbf{X}^D, \mathbf{X}_i^S\}$ is encoded into the corresponding posterior distributions $\mathcal{N}(l_1^E(f^E(\mathbf{X})), l_2^E(f^E(\mathbf{X})))$. Here, all the symbols have the same meanings as those in eq. (2) except that they are dedicated to the external knowledge encoder. Correspondingly, we denote the topical embeddings generated by this encoder for the word, sentence and document-level external knowledge respectively as: $\Phi^E \sim q^E(\mathbf{X}^W) := \mathcal{N}_{\mathbf{X}^W}^E$, $\mathbf{S}_i^E \sim q^E(\mathbf{X}_i^S) := \mathcal{N}_{\mathbf{X}_i^S}^E$ and $\mathbf{Z}^E \sim q^E(\mathbf{X}^D) := \mathcal{N}_{\mathbf{X}^D}^E$. NAHTM combines the internal and external topical embedding matrices as follows:

$$\begin{aligned} \Phi^{\text{Comb}} &:= \Phi^B + \gamma_1 \Phi^E, \\ \mathbf{S}_i^{\text{Comb}} &:= \text{Sparsemax}(\mathbf{S}_i^B + \gamma_2 \mathbf{S}_i^E), \\ \mathbf{Z}^{\text{Comb}} &:= \text{Softmax}(\mathbf{Z}^B + \gamma_3 \mathbf{Z}^E) \end{aligned} \quad (6)$$

where γ_1, γ_2 and γ_3 are the hyper-parameters that control the influence of external knowledge in calibrating the internal one at the different levels.

4.3 Attention-Aware KL Divergence

NAHTM makes use of the hierarchical structure of documents by setting the posterior distributions of the document topical embeddings as the priors to their sentences’ topical embeddings. More specifically, for document i , its dedicated KL divergence term is $Q_i^D := \beta_0 \left(\text{KL}[q(\mathbf{z}_i^B | \mathbf{w}_i^D) || p(\mathbf{z}_i^B)] + \text{KL}[q(\mathbf{z}_i^E | \mathbf{x}_i^D) || p(\mathbf{z}_i^E)] \right)$, while for the j^{th} sentence of document i , the dedicated KL term is $Q_{ij}^S := \beta_1 \left(\text{KL}[q(\mathbf{s}_{ij}^B | \mathbf{w}_{ij}^S) || q(\mathbf{z}_i^B | \mathbf{w}_i^D)] + \text{KL}[q(\mathbf{s}_{ij}^E | \mathbf{x}_{ij}^S) || q(\mathbf{z}_i^E | \mathbf{x}_i^D)] \right)$. Here, β_0 and β_1 are the hyper-parameters that control the regularization strengths of the corresponding KL terms in the ELBO. Note that all the embeddings involved in the above KL terms are *unnormalized*; in other words, they have not been transformed by the Softmax/Sparsemax function.

The assumption behind the above hierarchical structure of KL divergence terms is straightforward: topics of sentences should be somewhat similar to those of their documents. In this case, the degree of the topical similarity constraint enforced into the

²<https://github.com/UKPLab/sentence-transformers>

learning of NAHTM is controlled by β_1 . For example, if β_1 becomes smaller, the similarity constraint is going to be weakened accordingly.

Customising β_1 with attention weights is an alternative method we propose that allows for adaptive control on the regularization strengths of the KL divergence terms specific to individual sentence-document pairs. The intent of this method is that the semantic relevance of sentences towards the document, as revealed by the external knowledge, should also be indicative of their topical relevance in the corpus. There are two strategies for implementing this method.

In **strategy 1**, NAHTM integrates each unnormalized attention weight y_{ij}^S , computed from eq. (5), into the sentence-document KL terms with respect to the pre-trained sentence embedding as:

$$Q_{ij}^S := \beta_1 \sigma(y_{ij}^{S'}) \left(\text{KL}[q(\mathbf{s}_{ij}^B | \mathbf{w}_{ij}^S) || q(\mathbf{z}_i^B | \mathbf{w}_i^D)] + \text{KL}[q(\mathbf{s}_{ij}^E | \mathbf{x}_{ij}^S) || q(\mathbf{z}_i^E | \mathbf{x}_i^D)] \right) + \lambda_0 \|y_{ij}^{S'} - y_{ij}^S\|^2 \quad (7)$$

where $y_{ij}^{S'}$ is a latent variable to be learned to control the KL terms specific to the sentence j of document i , and is constrained to be close to y_{ij}^S ; λ_0 is a hyper-parameter that controls the strength of the constraint; $\sigma(y_{ij}^{S'})$ is the corresponding normalized result by either the Softmax or the Sparsemax function. Eq. (7) is essentially a ‘‘soft’’ version of the strategy that directly uses the normalized attention weight as the controlling parameter, i.e. $\beta_1 \alpha_{ij}^S$, for the KL terms across the J_i sentences.

In **strategy 2**, instead of learning the attention weights y_{ij}^S , they are first pre-computed based on the pre-trained sentence and document embeddings by $y_{ij}^S := -\|\mathbf{x}_{ij}^S - \mathbf{x}_i^D\|$ where $\mathbf{x}_i^D := \frac{1}{J_i} \sum_{j=1}^{J_i} \mathbf{x}_{ij}^S$. In this case, the unnormalized attention weight y_{ij}^S for sentence j is the negative Euclidean distance between its embedding and the document embedding which is the centroid across all the sentences. The further away the embedding \mathbf{x}_{ij}^S is from the centroid \mathbf{x}_i^D , the smaller the weight. Therefore, the control over strengths of the KL constraints is now adaptive only towards the external knowledge. The rest of the KL computation follows exactly eq. (7). For a ‘‘hard’’ version of this strategy, we can subsequently normalize the pre-computed y_{ij}^S to obtain α_{ij}^S and then, directly use them as the controlling parameters (same as in strategy 1).

4.4 Training Objective

In summary, the training objective of NAHTM is:

$$\mathbb{E}_{\mathcal{Z}^{\text{Comb}}} [\log p_{\Phi^{\text{Comb}}}(\mathbf{W}^D | \mathcal{Z}^{\text{Comb}})] + \gamma_4 \sum_i^I \mathbb{E}_{\mathcal{S}_i^{\text{Comb}}} [\log p_{\Phi^{\text{Comb}}}(\mathbf{W}_i^S | \mathcal{S}_i^{\text{Comb}})] - \sum_i^I \sum_j^J Q_{ij}^S - \sum_i^I Q_i^D - \sum_v^V Q_v^W \quad (8)$$

where the likelihood $p_{\Phi^{\text{Comb}}}$ is modelled by the decoder network as in eq. (3) with the weight matrix now being Φ^{Comb} ; γ_4 controls the influence of the sentence-level training loss; Q_i^D and Q_{ij}^S are respectively the document- and sentence-level KL terms specified in Section 4.3; Q_v^W consists of the regularization and KL terms respectively for the internal and external embeddings for each word v from the vocabulary; $Q_v^W := \lambda_1 \|\Phi^B\| + \beta_2 \text{KL}[q(\phi_v^E | \mathbf{x}_v^W) || p(\phi_v^E)]$. Again, λ_1 and β_2 are the controlling hyper-parameters for the respective terms.

5 Experimental Setup

Data We evaluate the efficacy of NAHTM using four real-world corpora from a variety of domains, Wikitext-103 (Nan et al., 2019), 20NewsGroup (Srivastava and Sutton, 2017), COVID-19 open research dataset (CORD)³ and NIPS⁴ datasets. Table 2 summarises the key statistics of these datasets. For the Wikitext-103 dataset, we only use the introduction part of each document, named the WikiIntro dataset, to examine the efficacy of our model on short text. For the CORD dataset, we randomly sampled 20,000 documents from its original corpus for our experiments, which is named CORD20K.

We adopt the same training-validation-testing split ratios and preprocessing steps from the original papers for the Wikitext-103 (i.e. 70%-15%-15%) and 20NewsGroup (i.e. 48%-12%-40%) datasets. To construct their vocabularies for topic modelling, we follow the same stemming/lemmatization and stopword removing procedures of Hoyle et al. (2020).

As for the CORD20K and NIPS datasets, we set the split ratio to be 60%-20%-20%, and use the most frequent 10,000 words (with stemming and stop-words removed) as the vocabulary. Specifically, we apply WordNet lemmatizer and English stopword list (both from the NLTK⁵ toolkit) to

³<https://www.kaggle.com/allen-institute-for-ai/CORD-19-research-challenge>

⁴<https://www.kaggle.com/benhamner/nips-papers>

⁵<https://www.nltk.org/>

preprocess their text. Finally, we use the first 50 sentences of each document from the two datasets for the experiments.

Furthermore, to extract the sentence embedding from the external pre-trained transformer, we have not done any pre-processing (neither stemming/lemmatization nor stopword removing) on the sentence text for all the datasets, as required by the sentence-transformers library. Doing so guarantees that the pre-trained language model can capture the full contextual information within each sentence.

	I	V	Avg N
20NewsGroup	17,992	1,995	87.1
WikiIntro	28,532	20,000	124.11
CORD20K	20,000	10,000	1127.01
NIPS	7,241	10,000	1139.26

Table 2: Summary of the four datasets in terms of the document number I , vocabulary size V and average document length Avg N (first 50 sentences).

Evaluation Metrics We seek to perform two types of evaluation on our model. The first one is the ability to discover a set of latent topics that are meaningful and useful to human. To achieve this, we look at topic coherence with the normalized mutual pointwise information (NPMI) (Aletras and Stevenson, 2013; Lau et al., 2014), which is positively correlated with human judgments on topic quality. Specifically, we calculate the NPMI scores on both the internal corpora from Table 2, and a large external corpus. For each calculation, we first select the top 10 words under each topic based on the (sorted) values in each column of the word embedding matrix Φ^{Comb} . Then, we calculate and average the internal NPMI scores of each topic as in (Bouma, 2009) with a sliding window of size 10 and the training data as the reference corpus. As for the external NPMI, it is calculated using Palmetto⁶ over a large external English Wikipedia dump.

The second metric we use is perplexity, a popular criterion that evaluates how well topic models fit the BoW data, which is calculated over the testing data as⁷: $\exp(-\frac{1}{I} \sum_{i=1}^I \frac{1}{N_i} \sum_{n=1}^{N_i} p(\mathbf{w}_i^{\text{D}}))$ where N_i is the length of document i , and $p(\mathbf{w}_i^{\text{D}})$ is the log-likelihood of the model on the document’s BoW data. It is approximated as $p(\mathbf{w}_i^{\text{D}}) \approx$

⁶<https://aksw.org/Projects/Palmetto.html>

⁷Our perplexity calculation follows that of (Miao et al., 2016, 2017; Card et al., 2018) but for fair comparisons, the KL divergence is not included for calculating the perplexity as different models weight KL differently.

$p_{\Phi^{\text{Comb}}}(\mathbf{w}_i^{\text{D}}|z_i^{\text{Comb}})$. Here, Φ^{Comb} combines the decoder weights Φ^{B} and the posterior mean for Φ^{E} , while z_i^{Comb} combines the posterior means for z_i^{B} and z_i^{E} , as shown in eq. (6), at the first training step after the model has converged with respect to the validation perplexity.

Baselines We compare NAHTM with various state-of-the-art baselines which can be broadly categorised into 1) BoW-based autoencoder models without external knowledge, which include Scholar (without meta-data) (Card et al., 2018), W-LDA (Nan et al., 2019), GSM (Miao et al., 2017), ETM (Dieng et al., 2020) and RRT (Tian et al., 2020); 2) models with external knowledge learned by pre-trained language models including CTM (Bianchi et al., 2021) and BAT (Hoyle et al., 2020). GSM is similar to Scholar but with a simpler encoder-decoder structure. ETM further factorizes the topic-word distribution matrix into multiplication of topic and word embedding vectors. RRT proposes a new reparameterization trick for Dirichlet distributions over the document embeddings.

As for the implementations and settings of the baselines, we use their official codes and default settings obtained from their official Github repositories, except for GSM whose original code is unavailable. In this case, we re-implement GSM by referring to other credible sources⁸. For BAT, we use its enhanced versions with Scholar (i.e. BAT+Scholar) and W-LDA (i.e. BAT+W-LDA) from its implementation. To allow for fair comparisons, we tune the major parameters for all the models, including the numbers of hidden layers {1, 2, 3} and hidden neurons {100, 300, 600}, the learning rate {0.001, 0.002, 0.005, 0.01} and the batch size {8, 20, 200}. The ranges of the above hyper-parameters are the most common ones set by the baselines in their own implementations. We generally found that 1 hidden layer with 300 neurons, a learning rate of 0.002 and a batch size of 20 yields the best overall perplexity and NPMI performance across the models. For CTM and BAT that incorporate external knowledge as NAHTM does, they use the same pre-trained transformers as NAHTM, including “bert-base-uncased”, “distilbert-base-uncased” and “roberta-base” from the Huggingface Transformers models⁹.

⁸https://github.com/zll17/Neural_Topic_Models

⁹https://huggingface.co/transformers/pretrained_models.html

5.1 NAHTM Settings

The hyper-parameters of NAHTM control the influence of its different components and we use the validation dataset to optimize their values in terms of the validation perplexity. We find the same values generally hold, across the datasets, for $\gamma_1 = 0.01$ and $\gamma_2 = 0.001$ that control the effects of external embeddings for words and documents respectively. On the other hand, γ_3 and γ_4 were respectively tuned over the sets of candidate parameter values $\{0.001, 0.01, 0.1, 1\}$ and $\{0.01, 0.1, 1, 5\}$ to different optimal values for different datasets; β_0 , β_1 and β_2 were all optimized over the value set $\{0.001, 0.01, 0.1, 1\}$, which, compared with the KL annealing approach (Bowman et al., 2016), is much more efficient albeit sub-optimal; λ_0 and λ_1 were both tuned over the value set $\{0.01, 0.1, 1, 5, 10\}$. To allow for fair comparisons with the baselines, we set the number of hidden layers for both the internal BoW and the external knowledge encoders to be 1, the number of neurons to be 300, the batch size to be 20, and the learning rate to be 0.002 for all the experiments.

6 Results and Discussion

Following the settings from the previous section, we proceed to conduct both quantitative and qualitative experiments to evaluate NAHTM. For the quantitative experiments, we report the results of the average external and internal NPMI, and the average test perplexity over 5 runs with different random seeds for initialization. Moreover, within each run, all the models were learned twice with 50 and 200 topics each time, and their corresponding metric scores have been summarized in Table 3.

It can be observed that NAHTM, equipped with either the β_1 -customising strategy 1 or 2 from Section 4.3 (i.e. NAHTM_{S1} and NAHTM_{S2}), and with uncustomised hierarchical KL (i.e. NAHTM_{HKL}), is overall more coherent, in terms of both types of NPMI, than the baseline models and NAHTM with independent KL terms for sentences and documents (i.e. NAHTM_{KL}). Especially when comparing with BAT+{Scholar,W-LDA} and CTM, which also have leveraged external pre-trained knowledge, NAHTM manages to achieve not only higher topic coherence but also lower perplexity on the test data in general. The only exception is on the NIPS dataset where NAHTM follows BAT+Scholar closely on the perplexity.

In addition to its efficacy on the document-

Models	WikiIntro	20News	NIPS	CORD20K
External NPMI				
Scholar	0.116 / 0.109	0.028 / 0.009	0.013 / -0.017	0.036 / 0.007
W-LDA	0.108 / 0.096	0.034 / 0.012	0.013 / -0.022	0.031 / -0.001
GSM	0.111 / 0.094	0.018 / -0.006	-0.022 / -0.062	0.023 / -0.011
ETM	0.082 / 0.064	0.003 / -0.014	-0.062 / -0.097	0.008 / -0.032
RRT	0.118 / 0.094	0.009 / -0.010	0.011 / -0.031	0.027 / -0.013
BAT+Scholar	0.122 / 0.113	0.039 / 0.017	0.014 / -0.011	0.041 / 0.008
BAT+W-LDA	0.119 / 0.108	0.041 / 0.017	0.014 / -0.010	0.039 / 0.004
CTM	0.125 / 0.110	0.040 / 0.018	0.013 / -0.013	0.036 / 0.010
NAHTM _{KL}	0.132 / 0.114	0.035 / 0.013	0.014 / -0.006	0.039 / 0.008
NAHTM _{HKL}	0.138 / 0.119	0.037 / 0.015	0.014 / -0.005	0.042 / 0.010
NAHTM _{S1}	0.145 / 0.132	0.040 / 0.018	0.015 / -0.003	0.044 / 0.011
NAHTM _{S2}	0.149 / 0.135	0.042 / 0.020	0.016 / 0.007	0.046 / 0.015
Internal NPMI				
Scholar	0.427 / 0.411	0.256 / 0.224	0.285 / 0.256	0.349 / 0.294
W-LDA	0.422 / 0.409	0.249 / 0.220	0.294 / 0.260	0.337 / 0.286
GSM	0.411 / 0.398	0.225 / 0.198	0.289 / 0.242	0.318 / 0.239
ETM	0.393 / 0.384	0.190 / 0.165	0.222 / 0.176	0.247 / 0.188
RRT	0.427 / 0.414	0.254 / 0.218	0.249 / 0.212	0.343 / 0.265
BAT+Scholar	0.439 / 0.421	0.268 / 0.243	0.298 / 0.260	0.361 / 0.303
BAT+W-LDA	0.432 / 0.416	0.264 / 0.240	0.309 / 0.263	0.353 / 0.294
CTM	0.438 / 0.422	0.269 / 0.246	0.295 / 0.254	0.356 / 0.290
NAHTM _{KL}	0.437 / 0.416	0.263 / 0.245	0.302 / 0.262	0.356 / 0.296
NAHTM _{HKL}	0.441 / 0.418	0.266 / 0.249	0.306 / 0.264	0.358 / 0.301
NAHTM _{S1}	0.442 / 0.421	0.273 / 0.251	0.311 / 0.267	0.363 / 0.307
NAHTM _{S2}	0.446 / 0.425	0.279 / 0.256	0.319 / 0.272	0.370 / 0.315
Perplexity				
Scholar	2,071 / 1,938	978 / 936	1,686 / 1,584	1,562 / 1,489
W-LDA	2,243 / 2,091	951 / 927	1,754 / 1,632	1,580 / 1,517
GSM	2,465 / 2,176	1,114 / 1,032	1,796 / 1,726	1,611 / 1,565
ETM	2,184 / 2,106	1,045 / 985	1,722 / 1,639	1,559 / 1,523
RRT	2,274 / 2,123	984 / 955	1,748 / 1,677	1,625 / 1,570
BAT+Scholar	1,963 / 1,872	947 / 924	1,642 / 1,534	1,496 / 1,423
BAT+W-LDA	2,014 / 1,934	936 / <u>903</u>	1,715 / 1,603	1,528 / 1,476
CTM	1,924 / 1,860	972 / 933	1,726 / 1,641	1,568 / 1,519
NAHTM _{KL}	1,987 / 1,898	959 / 941	1,752 / 1,670	1,552 / 1,502
NAHTM _{HKL}	1,949 / 1,882	939 / 916	1,737 / 1,653	1,534 / 1,483
NAHTM _{S1}	1,911 / 1,854	921 / 907	1,704 / 1,634	1,512 / 1,469
NAHTM _{S2}	1,878 / 1,825	905 / 881	1,652 / <u>1,567</u>	1,491 / 1,450

Table 3: Results of average external and internal NPMI, and average test perplexity for each model with 50 and 200 topics, whose results are respectively on the left and right side of “/” in the table. Lower perplexity and higher NPMI scores mean better performance. Bold and underlined values indicate the best and second performing models in each dataset/metric setting.

level BoW reconstruction, NAHTM is also able to achieve state-of-the-art reconstruction performance at the sentence level. We illustrate this by treating each sentence (from the beginning 50 sentences) as a short document, and applying all the NTMs models to reconstruct their BoW data. In this case, we focus on the test perplexity of the different models on the sentences, reporting the sentence-level perplexity from NAHTM with its inference jointly performed over the document- and sentence-level log-likelihood under the hierarchical KL constraint). Table 5 shows that at 50 topics, the best variant NAHTM_{S2}, with respect to the document-level BoW reconstruction, remains competitive on the sentence-level reconstruction task. This suggests that NAHTM, with its attention-aware hierarchical KL regularization, can effectively infer both the document- and sentence-level neural topic mod-

	Topics	Models	Topic Words
NIPS	Optimization	BAT+Scholar	<i>problem</i> , siam, np, solve, optimization, loss, min, <i>consider</i> , convex, hard
		NAHTM _{S2}	loss, gradient, minimization, optimization, solution, min, convex, np, descent, objective
NIPS	Neural Networks	BAT+Scholar	layer, neural, deep, convolutional, network, gradient, architecture, <i>visual</i> , <i>input</i> , recurrent
		NAHTM _{S2}	network, neural, layer, deep, backpropagation, recurrent, convolutional, descent, feedforward, gradient
CORD20K	COVID	BAT+Scholar	infection, covid, <i>disease</i> , virus, coronavirus, epidemic, h7n9, vaccine, <i>patient</i> , <i>cases</i>
		NAHTM _{S2}	<i>disease</i> , covid, virus, infected, coronavirus, vaccine, influenza, <i>patient</i> , epidemic, respiratory
CORD20K	Quarantine	BAT+Scholar	quarantine, <i>school</i> , lockdown, <i>unemployed</i> , <i>study</i> , control, province, <i>holiday</i> , social, city
		NAHTM _{S2}	quarantine, lockdown, curfew, <i>unemployment</i> , fatality, distancing, <i>holiday</i> , social, province, <i>inequity</i>

Table 4: Top 10 topic words extracted by NAHTM and BAT+Scholar from the NIPS and CORD20K datasets where the italic words are those either too common or less relevant (compared to the other words) to the topics.

Models	WikiIntro	20News	NIPS	CORD20K
Scholar	965 (36)	648 (54)	861 (45)	3,157 (76)
W-LDA	940 (42)	619 (70)	896 (40)	3,718 (88)
GSM	1,029 (58)	802 (42)	1,055 (66)	3,348 (102)
ETM	944 (19)	710 (22)	965 (32)	2,952 (57)
RRT	981 (74)	747 (95)	1,271 (128)	3,623 (114)
BAT+Scholar	934 (45)	566 (63)	828 (54)	2,256 (68)
BAT+W-LDA	915 (33)	548 (56)	874 (61)	2,744 (93)
CTM	952 (82)	1,025 (51)	1,484 (110)	3,298 (121)
NAHTM _{S2}	892 (78)	605 (84)	788 (94)	<u>2,469</u> (146)

Table 5: With 50 topics, the sentence-level test perplexity results (with standard deviations shown inside brackets) of each model on the BoW data of each of the top 25 sentences within every document.

els it contains, and enables the latter model to be robust against the sparse sentence-level BoW data.

Furthermore, we conduct an ablation study on the importance of different external knowledge components in contributing to the robustness of NAHTM when dealing with the sparse sentence data. We find from Table 6 that the external word embeddings are the most important components for enhancing the performance of NAHTM, while the external sentence embeddings are the least important. Despite this, we can still see that the performance of NAHTM is significantly influenced by all the three types of external knowledge.

Finally, to gain a more intuitive view on how well NAHTM has discovered the underlying topics, we show in Table 4 the top 10 words under each of the four example topics extracted by NAHTM from the NIPS and CORD20K datasets. These four topics are Optimization and Neural Networks from the NIPS dataset, and COVID and Quarantine from the CORD20K dataset. It can be observed that the topics discovered by NAHTM tend to be more coherent and less likely to contain common and irrelevant words which can be found from the topic

Models	WikiIntro	20News	NIPS	CORD20K
NAHTM _{S2}	892 (78)	605 (84)	788 (94)	2,469 (146)
NAHTM _{S2} - \mathbf{X}^W	1,173 (142)	981 (119)	1,035 (182)	3,238 (165)
NAHTM _{S2} - \mathbf{X}^D	976 (67)	711 (79)	890 (86)	2,894 (111)
NAHTM _{S2} - \mathbf{X}^S	950 (42)	683 (60)	848 (62)	2,726 (98)

Table 6: With 50 topics, an ablation study on the effects of removing different external knowledge components on the test perplexity performance of NAHTM.

word lists extracted by BAT+Scholar.

7 Conclusion

In this paper, we have proposed NAHTM, a VAE-based neural topic model with attention-aware hierarchical KL divergence imposed on the pairs of documents and sentences. NAHTM incorporates both the internal BoW data information and the external pre-trained knowledge for refining the topical embeddings of words, sentences and documents. Both quantitative and qualitative experiments have shown the effectiveness of NAHTM on 1) recovering the BoW data at different levels of granularity and 2) discovering coherent topics, by making use of the hierarchical KL constraints on the sentence-document pairs and the external knowledge. As for the future work, we would like to investigate the possibility of combining NAHTM with neural language models for topic-aware language understanding and content generation.

Acknowledgements

Yuan Jin and Wray Buntine were supported by the Australian Research Council under awards DE170100037. Wray Buntine was also sponsored by DARPA under agreement number FA8750-19-2-0501.

References

- Nikolaos Aletras and Mark Stevenson. 2013. [Evaluating topic coherence using distributional semantics](#). In *Proceedings of the 10th International Conference on Computational Semantics (IWCS 2013) – Long Papers*, pages 13–22, Potsdam, Germany. Association for Computational Linguistics.
- Hesam Amoualian, Wei Lu, Eric Gaussier, Georgios Balikas, Massih R. Amini, and Marianne Clausel. 2017. [Topical coherence in LDA-based models through induced segmentation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1799–1809, Vancouver, Canada. Association for Computational Linguistics.
- Georgios Balikas, Massih-Reza Amini, and Marianne Clausel. 2016a. [On a topic model for sentences](#). In *Proceedings of the 39th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '16*, page 921–924, New York, NY, USA. Association for Computing Machinery.
- Georgios Balikas, Hesam Amoualian, Marianne Clausel, Eric Gaussier, and Massih R. Amini. 2016b. [Modeling topic dependencies in semantically coherent text spans with copulas](#). In *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, Osaka, Japan. The COLING 2016 Organizing Committee.
- Federico Bianchi, Silvia Terragni, and Dirk Hovy. 2021. [Pre-training is a hot topic: Contextualized document embeddings improve topic coherence](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 2: Short Papers)*, pages 759–766, Online. Association for Computational Linguistics.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. [Latent Dirichlet allocation](#). *J. Mach. Learn. Res.*, 3:993–1022.
- Gerlof Bouma. 2009. Normalized (pointwise) mutual information in collocation extraction. *Proceedings of the Biennial GSCS Conference 2009*, pages 31–40.
- Samuel R. Bowman, Luke Vilnis, Oriol Vinyals, Andrew Dai, Rafal Jozefowicz, and Samy Bengio. 2016. [Generating sentences from a continuous space](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 10–21, Berlin, Germany. Association for Computational Linguistics.
- Dallas Card, Chenhao Tan, and Noah A. Smith. 2018. [Neural models for documents with metadata](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2031–2040, Melbourne, Australia. Association for Computational Linguistics.
- Yatin Chaudhary, Pankaj Gupta, Khushbu Saxena, Vivek Kulkarni, Thomas Runkler, and Hinrich Schütze. 2020. [TopicBERT for energy efficient document classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1682–1690, Online. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Adji B. Dieng, Francisco J. R. Ruiz, and David M. Blei. 2020. [Topic modeling in embedding spaces](#). *Transactions of the Association for Computational Linguistics*, 8:439–453.
- Ran Ding, Ramesh Nallapati, and Bing Xiang. 2018. [Coherence-aware neural topic modeling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 830–836, Brussels, Belgium. Association for Computational Linguistics.
- Lan Du, Wray Buntine, Huidong Jin, and Changyou Chen. 2012. Sequential latent Dirichlet allocation. *Knowledge and Information Systems*, 31(3):475–503.
- Arthur Gretton, Karsten M. Borgwardt, Malte J. Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. [A kernel two-sample test](#). *Journal of Machine Learning Research*, 13(25):723–773.
- Alexander Miserlis Hoyle, Pranav Goel, and Philip Resnik. 2020. [Improving neural topic models using knowledge distillation](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1752–1771, Online. Association for Computational Linguistics.
- Haixin Jiang, Rui Zhou, Limeng Zhang, Hua Wang, and Yanchun Zhang. 2019. Sentence level topic models for associated topics extraction. *World Wide Web*, 22(6):2545–2560.
- Diederik P. Kingma and Max Welling. 2014. [Auto-encoding variational bayes](#). In *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings*.
- Jey Han Lau, David Newman, and Timothy Baldwin. 2014. [Machine reading tea leaves: Automatically evaluating topic coherence and topic model quality](#). In *Proceedings of the 14th Conference of the European Chapter of the Association for Computational*

- Linguistics*, pages 530–539, Gothenburg, Sweden. Association for Computational Linguistics.
- Tianyi Lin, Zhiyue Hu, and Xin Guo. 2019. Sparse-max and relaxed Wasserstein for topic sparsity. In *Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining, WSDM '19*, page 141–149, New York, NY, USA. Association for Computing Machinery.
- André F. T. Martins and Ramón F. Astudillo. 2016. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *Proceedings of the 33rd International Conference on International Conference on Machine Learning - Volume 48, ICML'16*, page 1614–1623. JMLR.org.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2410–2419. PMLR.
- Yishu Miao, Lei Yu, and Phil Blunsom. 2016. Neural variational inference for text processing. In *Proceedings of the 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 1727–1736, New York, New York, USA. PMLR.
- Ramesh Nallapati, Igor Melnyk, Abhishek Kumar, and Bowen Zhou. 2017. SenGen: Sentence generating neural variational topic model. *CoRR*, arXiv:1708.00308.
- Feng Nan, Ran Ding, Ramesh Nallapati, and Bing Xiang. 2019. Topic modeling with Wasserstein autoencoders. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 6345–6381, Florence, Italy. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. Sentence-BERT: Sentence embeddings using Siamese BERT-networks. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic backpropagation and approximate inference in deep generative models. In *Proceedings of the 31st International Conference on Machine Learning*, volume 32 of *Proceedings of Machine Learning Research*, pages 1278–1286, Beijing, China. PMLR.
- Akash Srivastava and Charles Sutton. 2017. Autoencoding variational inference for topic models. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Runzhi Tian, Yongyi Mao, and Richong Zhang. 2020. Learning VAE-LDA models with rounded reparameterization trick. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, Online. Association for Computational Linguistics.
- He Zhao, Lan Du, Wray Buntine, and Gang Liu. 2019. Leveraging external information in topic modelling. *Knowledge and Information Systems*, 61(2):661–693.
- He Zhao, Dinh Phung, Viet Huynh, Yuan Jin, Lan Du, and Wray Buntine. 2021a. Topic modelling meets deep neural networks: A survey. In *Proceedings of the Thirtieth International Joint Conference on Artificial Intelligence, IJCAI-21*, pages 4713–4720. International Joint Conferences on Artificial Intelligence Organization. Survey Track.
- He Zhao, Dinh Phung, Viet Huynh, Trung Le, and Wray Buntine. 2021b. Neural topic model via optimal transport. In *ICLR*.
- Deyu Zhou, Xuemeng Hu, and Rui Wang. 2020. Neural topic modeling by incorporating document relationship graph. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3790–3796, Online. Association for Computational Linguistics.