
Explainable and Discourse Topic-aware Neural Language Understanding

Yatin Chaudhary^{1,2} Hinrich Schütze² Pankaj Gupta¹

Abstract

Marrying topic models and language models exposes language understanding to a broader source of document-level context beyond sentences via topics. While introducing topical semantics in language models, existing approaches incorporate latent document topic proportions and ignore topical discourse in sentences of the document. This work extends the line of research by additionally introducing an explainable topic representation in language understanding, obtained from a set of key terms correspondingly for each latent topic of the proportion. Moreover, we retain sentence-topic association along with document-topic association by modeling topical discourse for every sentence in the document. We present a novel neural composite language modeling (NCLM) framework that exploits both the latent and explainable topics along with topical discourse at sentence-level in a joint learning framework of topic and language models. Experiments over a range of tasks such as language modeling, word sense disambiguation, document classification, retrieval and text generation demonstrate ability of the proposed model in improving language understanding.

1. Introduction

Topic models (TMs) such as LDA (Blei et al., 2001) facilitate document-level semantic knowledge in the form of topics, explaining the thematic structures hidden in a document collection. In doing so, they learn document-topic association in a generative fashion by counting word-occurrences across documents. Essentially, the generative framework assumes that each document is a mixture of latent topics, i.e., topic-proportions and each latent topic is a unique dis-

tribution over words in vocabulary. Beyond a document representation, topic models also offer interpretability via topics (a set of top key terms). Recently, neural topic models (Gupta et al., 2019b;a; Miao et al., 2016) have been shown to outperform LDA-based models. Thus, we consider neural network based topic models in this work.

Language models (LMs) (Mikolov et al., 2010; Peters et al., 2018) have recently gained success in natural language understanding by predicting the next (target) word in a sequence given its preceding and/or following context(s), accounting for linguistic structures such as word ordering. However, LM are often contextualized by an n-gram window or a sentence, ignoring global semantics in context beyond the sentence boundary especially in modeling documents. To capture long-term semantic dependencies, recent works (Wang et al., 2018; Lau et al., 2017; Dieng et al., 2017) have attempted to introduce document-level semantics in LMs at sentence-level by marrying topic and language models, e.g., augmenting LSTM-based LMs with a latent document-topic proportion (association) obtained from a topic model for the document in which the sentence appears.

Motivation 1: While augmenting LMs with topical semantics, existing approaches incorporate latent document-topic proportions and ignore an explanatory representation for each latent topic of the proportion. Here, the explanatory representation of a topic refers to a vector representation obtained from a set of high-probability terms in its topic-word distribution. For example in Figure 1(a), we run a topic model over a document of three sentences and discover a latent document-topic proportion \hat{h}_d as well as three topics (top-5 key terms) correspondingly explaining each latent topic (T1, T2 and T3) of the proportion. Observe that the context in sent#2 can not resolve the meaning of the word *chip*. However, introducing \hat{h}_d with complementary explainable topics (collections of key terms) provide an abstract (latent) and a fine granularity (explanatory) outlook, respectively. To our knowledge, the scheme of augmenting LMs with both the latent document-topic proportion and explanatory topics remains unexplored.

Contribution 1: Complementing the latent document-topic proportion, we also leverage explanatory topics in augmenting LMs with topical semantics in a neural composite language modeling (NCLM) framework, consisting of a neural

¹Corporate Technology, Machine Intelligence (MIC-DE), Siemens AG, Munich, Germany ²CIS, University of Munich (LMU), Munich, Germany. Correspondence to: Yatin Chaudhary <yatin.chaudhary@drimco.net>.

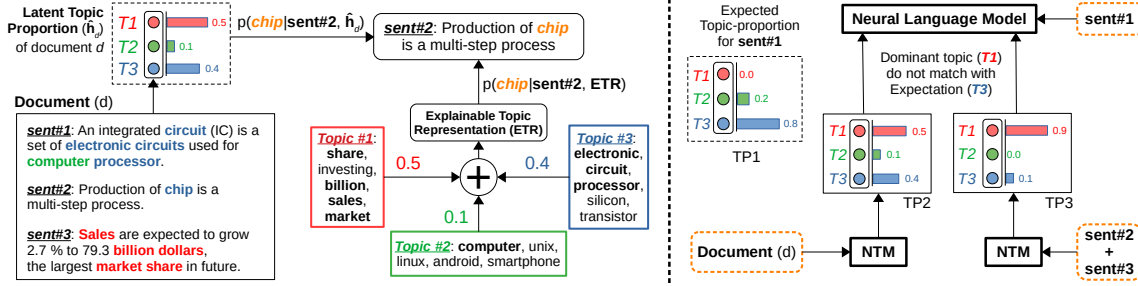


Figure 1. Detailed illustration of: (a) Motivation #1: top 5 key terms of each topic provide a fine-grained outlook of document semantics context than \mathbf{h}_d for prediction of word “chip”; (b) Motivation #2: Negative influence via sentence-level topical discourse mismatch.

topic model (NTM) and a neural language model (NLM).

Motivation 2: A sentence in a document may have a different topical discourse than its neighboring sentences or the document itself. Illustrated in Figure 1 (b), an NTM generates two different document-topic proportions (TP) for input document d and $\text{sent\#2}+\text{sent\#3}$ while modeling sent\#1 in the NLM. Observe that the sent\#1 expects a topic proportion dominated by topic T3 (*electronics*) as in TP1; however NTM generates TP2 or TP3 due to input d or $\text{sent\#2}+\text{sent\#3}$, respectively where both the document-topic proportions are dominated by the topic T1 about *marketing*. Therefore, there is need to deal with such topical discourse mismatch for each sentence in the document.

Contribution 2: In order to retain sentence-level topical semantics, we first extract sentence-topic association, i.e., sentence-level latent topic proportion, for each sentence using NTM and then introduce them in NLM in combination with the document-topic proportion (association).

Contribution 3: We evaluate the proposed NCLM framework over range of tasks such as language modeling, word sense disambiguation, document classification and information retrieval. Experimental results suggest that both the explanatory topics and sentence-topic association help in improving natural language understanding. *Implementation* of NCLM is available at: <https://github.com/YatinChaudhary/NCLM>.

2. Neural Language Model

Language modeling is the task of assigning probability distribution over a sequence of words. Typically, language models (Peters et al., 2018) are applied at the sentence-level. Consider a sentence $s = \{(w_m, y_m) | \forall m=1:M\}$ of length M in document d , where (w_m, y_m) is a tuple containing the indices of input and output words in vocabulary of size V . A LM computes the joint probability $p(s)$ i.e., likelihood of s by a product of conditional probabilities as follows:

$$p(s) = p(y_1, \dots, y_M) = p(y_1) \prod_{m=2}^M p(y_m | y_{1:m-1})$$

where $p(y_m | y_{1:m-1})$ is the probability of word y_m conditioned on preceding context $y_{1:m-1}$. RNN-based LMs capture linguistic properties in their recurrent hidden state $\mathbf{r}_m \in \mathbb{R}^H$ and compute output state $\mathbf{o}_m \in \mathbb{R}^H$ for each y_m :

$$\mathbf{o}_m, \mathbf{r}_m = f(\mathbf{r}_{m-1}, w_m); p(y_m | y_{1:m-1}) = p(y_m | \mathbf{o}_m) \quad (1)$$

where function $f(\cdot)$ can be a standard LSTM (Hochreiter & Schmidhuber, 1997) or GRU (Cho et al., 2014) cell and H is the number of hidden units. As illustrated in Figure 2 (c), the NLM component in our proposed NCLM framework is based on LSTM cell, i.e., $f = f^{LSTM}$. Then, the conditional $p(y_m | \mathbf{o}_m)$ is computed using multinomial logistic as:

$$p(y_m | \mathbf{o}_m) = \frac{\exp(\mathbf{o}_m^T \mathbf{U}_{:,y_m} + \mathbf{a}_{y_m})}{\sum_{j=1}^V \exp(\mathbf{o}_m^T \mathbf{U}_{:,j} + \mathbf{a}_j)} \quad (2)$$

where $\mathbf{U} \in \mathbb{R}^{H \times V}$, $\mathbf{a} \in \mathbb{R}^V$ are NLM decoding parameters and V is the vocabulary size. Here, the input w_m and output y_m indices are related as $y_m = w_{m+1}$. Finally, NLM computes log-likelihood \mathcal{L}^{NLM} of s as a training objective and maximizes it:

$$\mathcal{L}^{NLM} = \log p(y_1) \sum_{m=2}^M \log p(y_m | \mathbf{o}_m) \quad (3)$$

3. Neural Composite Language Model

While NLM captures sentence-level (short-range dependencies) linguistic properties, they tend to ignore the document-level (long-range) context across sentence boundaries. Khandelwal et al. (2018) have shown that even by considering multiple preceding sentences as the context to predict the current word, it is often difficult to capture long-term dependencies beyond a distance of 200 words in context. Therefore, a composition of NLM and NTM provides a broader document-level semantic awareness during sequence modeling leveraging document-topic proportion (association) extracted using NTM. The complementary learning leads to an improved language understanding, accounting for both sentence and document-level semantics. See Table 1 for description of notations used.

Table 1. Description of the notations used in this work

Notation	Description	Notation	Description
NTM, NLM	Neural Topic Model, Neural Language Model	V, Z, K	NLM Vocab, NTM Vocab, Number of topics
LTR, ETR	Latent and Explainable Topic Representations	H, D_E	LSTM hidden size, Word embedding size
d, s, y	a document, a sentence in d , a word in s	$\mathbf{V} \in \mathbb{R}^Z$	BoW representation of document d
$d-s, s-y$	d after removing s , s after removing y	$[\mathbf{W}, \mathbf{U}] \in \mathbb{R}^{H \times Z}$	Decoding matrix of NTM, NLM
$\mathcal{N}, \epsilon \in \mathbb{R}^K$	Gaussian distribution, a sample from \mathcal{N}	$\mathbf{E} \in \mathbb{R}^{D_E \times Z}$	Pre-trained word embedding matrix
$[\boldsymbol{\mu}, \boldsymbol{\sigma}] \in \mathbb{R}^K$	mean, variance of approximate \mathcal{N}	$[\mathbf{r}_m, \mathbf{o}_m] \in \mathbb{R}^H$	hidden, output vector of LSTM cell
$[\mathbf{h}_{d-s}, \mathbf{z}_{d-s}^{att}] \in \mathbb{R}^Z$	LTR, ETR representations of document $d-s$	$[\mathbf{o}_d^{LTA}, \mathbf{o}_d^{ETA}, \mathbf{o}_d^{LETA}] \in \mathbb{R}^H$	Topic composition of document d with \mathbf{o}_m
$\mathbf{t} \in \mathbb{R}^{K \times \text{top}N}$	A list of topics with topN words	$[\mathbf{o}_{d,s}^{LTA}, \mathbf{o}_{d,s}^{ETA}, \mathbf{o}_{d,s}^{LETA}] \in \mathbb{R}^H$	Topic composition of sentence s with \mathbf{o}_m

3.1. Neural Topic Model

In this work, NTM (Figure 2 (a)) is based on Neural Variational Document Model (NVDM) proposed by Miao et al. (2016). It is an unsupervised generative model that learns to regenerate an input document \mathbf{V} using a continuous latent topic representation \mathbf{h} which is sampled from a prior Gaussian distribution $p(\mathbf{h})$. NVDM adopts the neural variational inference framework to compute a posterior Gaussian distribution $q(\mathbf{h}|\mathbf{V})$ which approximates the true prior $p(\mathbf{h})$.

Given a document d , consider $\mathbf{V} \in \mathbb{R}^Z$ be its bag-of-words (BoW) representation and $\mathbf{v}_i \in \mathbb{R}^Z$ is the one-hot representation of the i th word of the vocabulary of size Z . The generative process (Algorithm 1: lines #9-18) of NVDM is:

Step 1: Latent topic representation $\mathbf{h} \in \mathbb{R}^K$ is sampled by encoding \mathbf{V} using an MLP encoder f^{MLP} followed by two linear projections l_1 and l_2 as shown in Figure 1(a), where \mathbf{I} is the identity matrix. To elaborate further, for each input \mathbf{V} encoder network generates the parameters mean $\boldsymbol{\mu}(\mathbf{V})$ and deviation $\boldsymbol{\sigma}(\mathbf{V})$ required to parameterize the approximate posterior distribution $q(\mathbf{h}|\mathbf{V})$ in diagonal Gaussian form and samples \mathbf{h} from it (Algorithm 2: lines #13-20).

$$\mathbf{h} \sim q(\mathbf{h}|\mathbf{V}) \equiv \mathcal{N}(\mathbf{h}|\boldsymbol{\mu}(\mathbf{V}), \text{diag}(\boldsymbol{\sigma}^2(\mathbf{V})))$$

Step 2: Conditional word probabilities $p(\mathbf{v}_i|\mathbf{h})$ are computed independently for each word, using multinomial logistic regression with parameters shared across all documents:

$$p(\mathbf{v}_i|\mathbf{h}) = \frac{\exp\{\mathbf{h}^T \mathbf{W}_{:,i} + \mathbf{b}_i\}}{\sum_{j=1}^{|Z|} \exp\{\mathbf{h}^T \mathbf{W}_{:,j} + \mathbf{b}_j\}} \quad (4)$$

where $\mathbf{W} \in \mathbb{R}^{K \times Z}$ & $\mathbf{b} \in \mathbb{R}^Z$ are NTM decoding parameters. The word probabilities $p(\mathbf{v}_i|\mathbf{h})$ are further used to compute document probability $p(\mathbf{V}|\mathbf{h})$ conditioned on \mathbf{h} . By marginalizing $p(\mathbf{V}|\mathbf{h})$ over latent representation \mathbf{h} , we get the likelihood $p(\mathbf{V})$ of document d as shown below.

$$p(\mathbf{V}) = \int_{\mathbf{h} \sim p(\mathbf{h})} p(\mathbf{V}|\mathbf{h}) d\mathbf{h} \quad \text{and} \quad p(\mathbf{V}|\mathbf{h}) = \prod_{i=1}^{N_d} p(\mathbf{v}_i|\mathbf{h})$$

where N_d is the number of words in document d . However, it is intractable to sample all possible configurations of $\mathbf{h} \sim$

$p(\mathbf{h})$. Therefore, NVDM uses neural variational inference framework to compute evidence lower bound \mathcal{L}^{NTM} as:

$$\mathcal{L}^{NTM} = \mathbb{E}_{q(\mathbf{h}|\mathbf{V})} \left[\sum_{i=1}^{N_d} \log p(\mathbf{v}_i|\mathbf{h}) \right] - \text{KLD} \quad (5)$$

Here \mathcal{L}^{NTM} being a lower bound i.e., $\log p(\mathbf{V}) \geq \mathcal{L}^{NTM}$, NVDM maximizes the log-likelihood of documents $\log p(\mathbf{V})$ by maximizing the evidence lower bound itself. The \mathcal{L}^{NTM} can be maximized via back-propagation of gradients w.r.t. model parameters using the samples generated from posterior distribution $q(\mathbf{h}|\mathbf{V})$. NVDM assumes both prior $p(\mathbf{h})$ and posterior $q(\mathbf{h}|\mathbf{V})$ distributions as Gaussian and hence employ KL-Divergence as a regularizer term to conform $q(\mathbf{h}|\mathbf{V})$ to the Gaussian assumption i.e., $\text{KLD} = \text{KL}[q(\mathbf{h}|\mathbf{V})||p(\mathbf{h})]$, mentioned in equation 5.

3.2. Topical Representation Extraction

To exploit document-level semantics while language modeling, we extract topics using NVDM and represent semantics of the extracted topics in the following two forms:

Latent Topic Extraction: We sample latent topic representation $\mathbf{h} \in \mathbb{R}^K$ as shown in Figure 2 (a) and Algorithm 2 (lines #13-20). Essentially, the topic vector \mathbf{h} is an abstract (latent) representation of topic-word distributions for K topics and represents a document-topic proportion (association) as a mixture of K latent topics about the document being modeled. Precisely, each scalar value $h^k \in \mathbb{R}$ denotes the contribution of k th topic in representing a document d by \mathbf{h} . We name \mathbf{h} as *Latent Topic Representation* (LTR) and denote it as \mathbf{h}_d for an input document d .

Explainable Topic Extraction: Beyond the latent topic-proportion, we also extract explainable topics (a fine-granularity description) that can be obtained from high probability key terms of a topic-word distribution corresponding to each latent topic k . In doing so using NTM, we use the decoding weight parameter $\mathbf{W} \in \mathbb{R}^{K \times Z}$, i.e., a topic matrix where each k th row $\mathbf{W}_{k,:} \in \mathbb{R}^Z$ denotes a distribution over vocabulary words for k th topic. As illustrated in Figure 2 (b), we extract key terms for each topic using the utility TOPIC-EXTRACT. Algorithm 2 (lines #1-11) and Algorithm 2 (lines #22-30) describe the mechanism of topic

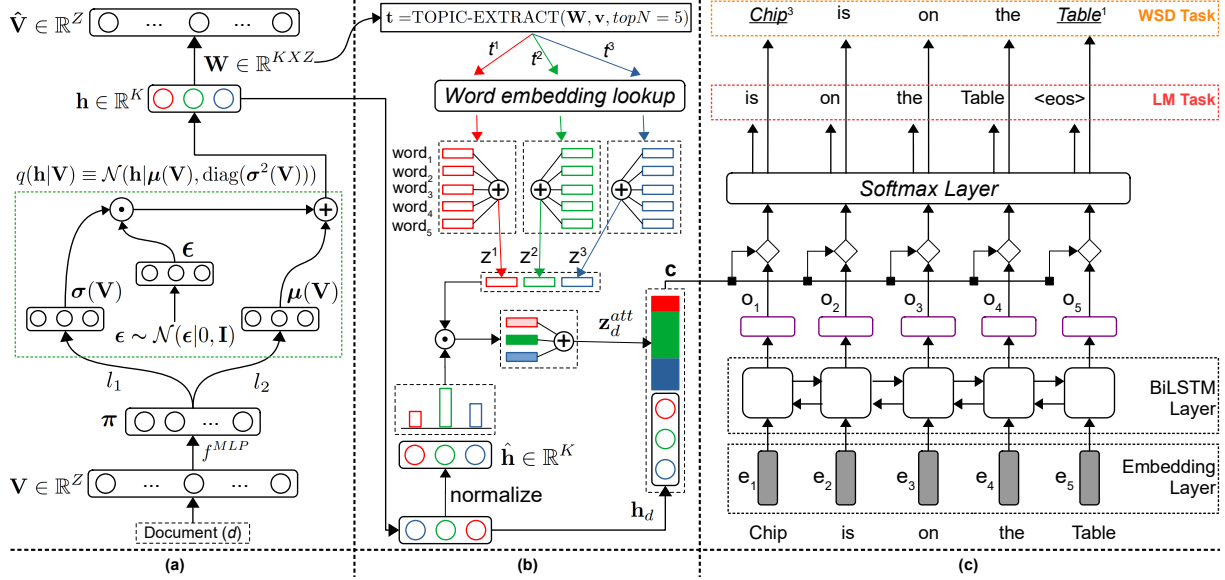


Figure 2. Illustration of our proposed Neural Composite Language Modeling (NCLM) framework: (a) Neural Topic Model (NTM), (b) Latent and Explainable Topic Representation Extraction and (c) Neural Language Model (NLM).

learning and extracting explainable topic using GET-ETR. Observe that the utility TOPIC-EXTRACT filters out the top key-terms not appearing in the document being modeled in order to highlight the contribution of those topical words shared in topic-word distribution and the document itself. Specifically, TOPIC-EXTRACT returns K lists of key terms explaining each latent topic h_k , i.e., $\mathbf{t} = [t^k]_{k=1:K}$ such that t^k had $topN$ key terms for k th topic. We use the mask \mathbf{D} to apply the filter as:

$$\mathbf{t} = \text{row-argmax}[\mathbf{W} \odot \mathbf{D}]_{1:topN}$$

where “row-argmax” is a function which returns indices of $topN$ values from each row of input matrix, \odot is an element-wise hadamard product and $\mathbf{D} \in \mathbb{R}^{K \times Z}$ is an indicator matrix where each column $\mathbf{D}_{:,i} \in \{1^K \text{ if } v_i \neq 0; 0^K \text{ otherwise}\}$. Now for each topic k , we perform embedding lookup using matrix $\mathbf{E} \in \mathbb{R}^{D_E \times Z}$ (pretrained word embeddings (Bojanowski et al., 2017)) for each word index in t^k and then average them to compute the explanatory topic-embedding vector \mathbf{z}^k as shown below:

$$\mathbf{z}^k = \frac{\sum_{j=1}^{topN} \text{emb_lookup}(\mathbf{E}, t_j^k)}{topN}$$

Finally, we perform weighted sum of topic vectors \mathbf{z}^k using document-topic proportion vector $\hat{\mathbf{h}}$ as weights to compute \mathbf{z}^{att} . We name \mathbf{z}^{att} as *Explainable Attentive Topic Representation* (ETR) and denote it as \mathbf{z}_d^{att} for a document d .

$$\mathbf{z}^{att} = \sum_{k=1}^K (\mathbf{z}^k \cdot \hat{h}^k) \quad \text{and} \quad \hat{\mathbf{h}} = \text{softmax}(\mathbf{h})$$

3.3. Joint Topic and Language Model

For simplification of notation in further sections, we drop the position index in $(w_m, y_m, \mathbf{o}_m, \mathbf{r}_m)$ from equations {1, 7, 3} and simply refer to them as $(w, y, \mathbf{o}, \mathbf{r})$, since our method is independent of word positions. In this section we describe the *composition* of topical representation $\mathbf{c} \in \{\mathbf{h}_d, \mathbf{z}_d^{att}\}$ with the output vector \mathbf{o} of NLM such that NLM is aware of document-level semantics while language modeling. We denote composition function by $(\mathbf{o} \diamond \mathbf{c})$, where we first concatenate the two complementary representations (\mathbf{o} and \mathbf{c}) and then perform a projection as:

$$\hat{\mathbf{o}} = (\mathbf{o} \diamond \mathbf{c}) = \text{sigmoid}([\mathbf{o}; \mathbf{c}]^T \mathbf{W}^p + \mathbf{b}^p) \quad (6)$$

where $\mathbf{W}^p \in \mathbb{R}^{\hat{H} \times H}$ and $\mathbf{b}^p \in \mathbb{R}^H$ are projection parameters, and $\hat{H} = H + K$. We then compute prediction probability of output word y using equation 7 as:

$$p(y|\mathbf{o}, \mathbf{c}) = \frac{\exp\{\hat{\mathbf{o}}^T \mathbf{U}_{:,y} + \mathbf{a}_y\}}{\sum_{j=1}^V \exp\{\hat{\mathbf{o}}^T \mathbf{U}_{:,j} + \mathbf{a}_j\}} \quad (7)$$

Using this composition scheme, we employ the two representations: LTR and ETR exclusively or in combination while performing composition within NCLM framework. Following are the proposed configurations in NCLM:

Latent Topic-aware NLM: Existing works (Lau et al., 2017; Dieng et al., 2017; Wang et al., 2018) in marrying topic and language models leverage latent document-topic representation \mathbf{h} to incorporate document-level semantics into sequence modeling. Also, modeling in such composite setting can be tricky. To remove the chances of NLM memorizing the next word due to input to NTM, the prior works

Algorithm 1 Computation of combined loss \mathcal{L}

```

1: Input: sentence  $s = \{(w_m, y_m) | \forall m=1:M\}$ 
2: Input:  $\mathbf{V} \in \mathbb{R}^Z$  of document  $d$ -s containing  $N_{d-s}$  words
3: Input: pretrained embedding matrix  $\mathbf{E}$ 
4: Parameters:  $\{\mathbf{W}, \mathbf{U}, \mathbf{b}, \mathbf{a}, f^{MLP}, l_1, l_2, f^{LSTM}\}$ 
5: Hyper-parameters:  $\{\alpha, topN, g\}$ 
6: Initialize:  $p(\mathbf{h}) \equiv \mathcal{N}(\mathbf{h} | 0, \text{diag}(\mathbf{I}))$ 
7: Initialize:  $p(\mathbf{V} | \mathbf{h}) \leftarrow 0; p(s | \mathbf{V}) \leftarrow 0; \mathbf{r}_0 \leftarrow 0$ 
8:
9: Neural Topic Model:
10: Sample Latent Topic Representation (LTR)  $\mathbf{h}$ 
11:  $\mathbf{h}, \mathbf{q}(\mathbf{h} | \mathbf{V}) \leftarrow \text{SAMPLE-h}(f^{MLP}, g, \mathbf{V}, l_1, l_2, \text{sigmoid})$ 
12: Compute KL divergence between true prior  $p(\mathbf{h})$  and  $q(\mathbf{h} | \mathbf{V})$ 
13:  $\text{KLD} \leftarrow \text{KL}[q(\mathbf{h} | \mathbf{V}) || p(\mathbf{h})]$ 
14: for  $i$  from 1 to  $N_{d-s}$  do
15:    $p(\mathbf{v}_i | \mathbf{h}) \leftarrow \frac{\exp\{\mathbf{h}^T \mathbf{W}_{:,i} + \mathbf{b}_i\}}{\sum_{j=1}^Z \exp\{\mathbf{h}^T \mathbf{W}_{:,j} + \mathbf{b}_j\}}$ 
16:    $p(\mathbf{V} | \mathbf{h}) \leftarrow p(\mathbf{V} | \mathbf{h}) \cdot p(\mathbf{v}_i | \mathbf{h})$ 
17: end for
18:  $\mathcal{L}^{NTM} \leftarrow -(\log p(\mathbf{V} | \mathbf{h}) - \text{KLD})$ 
19: if ETA or LETA then
20:   Extract Explainable Topic Representation (ETR)
21:    $\mathbf{z}_{d-s}^{att} \leftarrow \text{GET-ETR}(\mathbf{W}, \mathbf{V}, topN, \mathbf{h}, \mathbf{E})$ 
22: end if
23:
24: Neural Composite Language Model:
25: for  $m$  from 1 to  $M$  do
26:    $\mathbf{o}_m, \mathbf{r}_m \leftarrow f^{LSTM}(\mathbf{r}_{m-1}, w_m)$ 
27:   Composition of NTM and NLM
28:   if LTA then
29:      $\hat{\mathbf{o}}_m \leftarrow (\mathbf{o}_m \diamond \mathbf{h}_{d-s})$ 
30:   else if ETA then
31:      $\hat{\mathbf{o}}_m \leftarrow (\mathbf{o}_m \diamond \mathbf{z}_{d-s}^{att})$ 
32:   else if LETA then
33:      $\hat{\mathbf{o}}_m \leftarrow (\mathbf{o}_m \diamond [\mathbf{h}_{d-s}; \mathbf{z}_{d-s}^{att}])$ 
34:   end if
35:    $p(y_m | \mathbf{o}_m, \mathbf{V}) \leftarrow \frac{\exp\{\hat{\mathbf{o}}_m^T \mathbf{U}_{:,y_m} + \mathbf{a}_{y_m}\}}{\sum_{j=1}^V \exp\{\hat{\mathbf{o}}_m^T \mathbf{U}_{:,j} + \mathbf{a}_j\}}$ 
36:    $p(s | \mathbf{V}) \leftarrow p(s | \mathbf{V}) \cdot p(y_m | \mathbf{o}_m, \mathbf{V})$ 
37: end for
38:  $\mathcal{L}^{NLM} \leftarrow -\log p(s | \mathbf{V})$ 
39:  $\mathcal{L} \leftarrow \alpha \cdot \mathcal{L}^{NTM} + (1 - \alpha) \cdot \mathcal{L}^{NLM}$ 

```

exclude the current sentence from the document before input to NTM. Thus for a given document d and a sentence s on NLM, we compute an LTR vector \mathbf{h}_{d-s} by modeling $d - s$ sentences on NTM. Then, we compose it with output vector \mathbf{o} of NLM to obtain a representation \mathbf{o}_d^{LTA} using equation 6, i.e., $\mathbf{o}_d^{LTA} = (\mathbf{o} \diamond \mathbf{h}_{d-s})$. We name this scheme of composition as *LTA-NLM*, a baseline for our contributions.

Explainable Topic-aware NLM: Discussed in section 3.2, the ETR explains each latent topic and facilitates a fine-granularity descriptive outlook by a set of key-terms. Complementary to LTR, we use the ETR vector in composition with NLM. In doing so, we first compose ETR representation \mathbf{z}_{d-s}^{att} of d -s sentences in a document d with NLM output vector \mathbf{o} to obtain \mathbf{o}_d^{ETA} using equation 6, i.e., $\mathbf{o}_d^{ETA} = (\mathbf{o} \diamond \mathbf{z}_{d-s}^{att})$. This newly composite vector \mathbf{o}_d^{ETA} encodes fine-grained explainable topical semantics to be

Algorithm 2 Utility functions

```

1: function GET-ETR( $\mathbf{W}, \mathbf{V}, topN, \mathbf{h}, \mathbf{E}$ )
2:   Extract  $topN$  words from each topic belonging to  $d$ 
3:    $\mathbf{t} \leftarrow \text{TOPIC-EXTRACT}(\mathbf{W}, \mathbf{V}, topN)$ 
4:   Embedding lookup and summation to get topic embedding
5:   for  $k$  from 1 to  $K$  do
6:      $\mathbf{z}^k \leftarrow \frac{\sum_{j=1}^{topN} \text{emb\_lookup}(\mathbf{E}, t_j^k)}{topN}$ 
7:   end for
8:   Weighted sum of all topic embeddings
9:    $\mathbf{z}^{att} \leftarrow \sum_{k=1}^K (\mathbf{z}^k \cdot \hat{h}^k); \hat{\mathbf{h}} \leftarrow \text{softmax}(\mathbf{h})$ 
10:  return  $\mathbf{z}^{att}$ 
11: end function
12:
13: function SAMPLE-h( $f, g, \mathbf{V}, l_1, l_2, \text{act}$ )
14:  Sample  $\mathbf{h}$  via gaussian distribution conditioned on  $\mathbf{V}$ 
15:   $\boldsymbol{\pi} \leftarrow \text{act}(f(\mathbf{V})) ; \boldsymbol{\epsilon} \sim \mathcal{N}(\boldsymbol{\epsilon} | 0, \text{diag}(\mathbf{I}))$ 
16:   $\boldsymbol{\mu}(\mathbf{V}) \leftarrow l_1(\boldsymbol{\pi}) ; \boldsymbol{\sigma}(\mathbf{V}) \leftarrow l_2(\boldsymbol{\pi})$ 
17:   $q(\mathbf{h} | \mathbf{V}) \equiv \mathcal{N}(\mathbf{h} | \boldsymbol{\mu}(\mathbf{V}), \text{diag}(\boldsymbol{\sigma}^2(\mathbf{V})))$ 
18:   $\mathbf{h} \leftarrow (\boldsymbol{\mu}(\mathbf{V}) + \boldsymbol{\epsilon} \odot \boldsymbol{\sigma}(\mathbf{V})) \sim q(\mathbf{h} | \mathbf{V})$ 
19:  return  $g(\mathbf{h}), q(\mathbf{h} | \mathbf{V})$ 
20: end function
21:
22: function TOPIC-EXTRACT( $\mathbf{W}, \mathbf{V}, topN$ )
23:  Create mask matrix  $\mathbf{D} \in \mathbb{R}^{K \times Z}$  initialized with 0
24:  for  $i$  from 1 to  $Z$  do
25:    replace all 0 with 1 in column  $\mathbf{D}_{:,i}$  if the count of the  $i$ th word of the vocabulary is non-zero in  $\mathbf{V}$ 
26:  end for
27:  Take hadamard product and find  $topN$  max values
28:   $\mathbf{t} = \text{row-argmax}[\mathbf{W} \odot \mathbf{D}]_{1:topN}$ 
29:  return  $\mathbf{t}$ 
30: end function

```

used in sequence modeling task. We name this composition as *ETA-NLM*. To our knowledge, none of the existing approaches of joint topic and language modeling leverage explainable topics i.e., topic-word distributions into NLMs. The proposed *ETA-NLM* is the first one to exploit it.

Latent and Explainable Topic-aware NLM: We now leverage the two complementary topical representations using the latent \mathbf{h}_{d-s} and explainable \mathbf{z}_{d-s}^{att} vectors jointly. We concatenate them together and compose it with the output vector \mathbf{o} of NLM to obtain \mathbf{o}_d^{LETA} using equation 6, i.e., $\mathbf{o}_d^{LETA} = (\mathbf{o} \diamond [\mathbf{h}_{d-s}; \mathbf{z}_{d-s}^{att}])$. We name this composition as *LETA-NLM* due to latent and explainable topic vectors.

3.4. Sentence-level Topical Discourse

Discussed in section 1 and illustrated in Figure 1(b), there is a need for sentence-level topics in order to avoid dominant topic mismatch. Thus, we retain sentence-level topical discourse (*SDT*) by incorporating sentence-level topic association (latent and/or explainable) while modeling the sentence on NLM. To avoid memorization of current word being predicted y , we remove it from sentence s i.e., $s-y$ is input to NTM to compute its topic-proportion. Given the latent and explainable representations, we first extract

sentence-level LTR \mathbf{h}_{s-y} and ETR \mathbf{z}_{s-y}^{att} vectors and then concatenate these with the corresponding document-level LTR and/or ETR vectors before composing them with NLM. Following are the additional compositions for every sentence s in a document d :

$$LTA-NLM +SDT: \mathbf{o}_{d,s}^{LTA} = (\mathbf{o} \diamond [\mathbf{h}_{d-s}; \mathbf{h}_{s-y}])$$

$$ETA-NLM +SDT: \mathbf{o}_{d,s}^{ETA} = (\mathbf{o} \diamond [\mathbf{z}_{d-s}^{att}; \mathbf{z}_{s-y}^{att}])$$

$$LETA-NLM +SDT: \mathbf{o}_{d,s}^{LETA} = (\mathbf{o} \diamond [\mathbf{h}_{d-s}; \mathbf{h}_{s-y}; \mathbf{z}_{d-s}^{att}; \mathbf{z}_{s-y}^{att}])$$

Similarly, these composed output vectors are used to assign probability to the output word y using equation 7.

To summarize, we have presented six different configurations of our proposed NCLM framework based on the composition of different latent or explainable representations as well as document-topic and sentence-topic associations:

$$p(y|\mathbf{o}, \mathbf{c}) = \frac{\exp\{\hat{\mathbf{o}}^T \mathbf{U}_{:,y} + \mathbf{a}_y\}}{\sum_{j=1}^V \exp\{\hat{\mathbf{o}}^T \mathbf{U}_{:,j} + \mathbf{a}_j\}}$$

where, $\hat{\mathbf{o}} \in \{\mathbf{o}_d^{LTA}, \mathbf{o}_d^{ETA}, \mathbf{o}_d^{LETA}, \mathbf{o}_{d,s}^{LTA}, \mathbf{o}_{d,s}^{ETA}, \mathbf{o}_{d,s}^{LETA}\}$

3.5. Training Objective

Training of the joint topic and language model is performed by maximizing the joint log-likelihood objective \mathcal{L} which is a linear combination of the log-likelihood of document d via NTM and sentence s via NLM i.e., $\mathcal{L} = \alpha \cdot \mathcal{L}^{NTM} + (1 - \alpha) \cdot \mathcal{L}^{NLM}$ where, $\alpha \in [0, 1]$ is a hyper-parameter, maintaining a balance between NTM and NLM during joint training by updating model parameters at different scales.

3.6. Computational Complexity of NCLM Framework

NTM component complexity: Computational complexities of extracting latent and explainable topic representations of document d are described below.

1. **Latent topic extraction:** complexity of computing the latent topic representation (LTR) vector $\mathbf{h}_d \in \mathbb{R}^K$, via matrix projection on the encoder side, is given as $\mathcal{O}(KZ)$, where K is number of topics.
2. **Explainable topic extraction:** complexity of computing explainable attentive topic representation (ETR) vector \mathbf{z}_d^{att} is given as $\mathcal{O}(KZ + K(Z \log Z + D_{EtopN}))$, where (a) $\mathcal{O}(KZ)$ is the complexity of computing mask matrix \mathbf{D} and taking hadamard product with the topic matrix \mathbf{W} , (b) $\mathcal{O}(Z \log Z)$ is the complexity of sorting k th row $\mathbf{W}_{k,:}$ of topic matrix $\mathbf{W}^{K \times Z}$, and (c) $\mathcal{O}(D_{EtopN})$ is the complexity of extracting and adding pre-trained word embeddings, via embedding matrix \mathbf{E} , for $topN$ key terms of k th topic. Therefore, at asymptotic limits, computational complexity for \mathbf{z}_d^{att} becomes $\mathcal{O}(K(Z \log Z + D_{EtopN}))$.

NLM component complexity: NLM component of our NCLM framework is implemented as Recurrent Neural Network using LSTM cell(s). For each input word w_m of sentence $s = \{(w_m, y_m) | \forall m=1:M\}$ in document d , the complexity of our NLM component can be sub-divided into three parts:

1. **Hidden state computation:** complexity of computing output state \mathbf{o}_m , via LSTM cell, is given as $\mathcal{O}(H(H + HH_i))$, where H is the number of hidden units of LSTM cell and H_i is the size of input word embedding.
2. **Topic composition:** complexity of composition (\diamond), via concatenation and projection, of topic representation $\mathbf{c} \in \{\mathbf{h}_d, \mathbf{z}_d^{att}, [\mathbf{h}_d; \mathbf{z}_d^{att}]\}$ with output state \mathbf{o}_m of NLM is given as $\mathcal{O}((H + H_c)H)$, where $H_c \in \{K, D_E, (K + D_E)\}$ is the size of topic representation \mathbf{c} .
3. **Word prediction:** complexity of output word prediction via multinomial logistic regression over NLM vocabulary is given as $\mathcal{O}(HV)$, where V is NLM vocabulary size.

Therefore, combined computational complexity for all M words in sentence s is given as

$$\mathcal{O}_s = \mathcal{O}(MH(H + H_i + H_c + V))$$

Composite model complexity: Based on the type of topic composition, our proposed models have different computational complexities as mentioned below:

$$LTA-NLM: \mathcal{O}_s + \mathcal{O}(KZ)$$

$$ETA-NLM: \mathcal{O}_s + \mathcal{O}(K(Z \log Z + D_{EtopN}))$$

$$LETA-NLM: \mathcal{O}_s + \mathcal{O}(K(Z \log Z + D_{EtopN}))$$

Sentence-level topical discourse: For each output word y_m in sentence s , we additionally compute latent/explainable topical representation of sentence s after removing y_m i.e., $s-y_m$, to maintain topical discourse across sentences in document d . Therefore, M additional LTR/ETR vectors are computed for all words in sentence s and hence computational complexities are significantly increased by the factor of M as shown below:

$$LTA-NLM +SDT: \mathcal{O}_s + \mathcal{O}(MKZ)$$

$$ETA-NLM +SDT: \mathcal{O}_s + \mathcal{O}(MK(Z \log Z + D_{EtopN}))$$

$$LETA-NLM +SDT: \mathcal{O}_s + \mathcal{O}(MK(Z \log Z + D_{EtopN}))$$

4. Experiments and Results

To demonstrate the positive influence of composing LTR and ETR representations in neural language modeling, we

Table 2. Language Modeling Perplexity scores on three datasets under two different settings of NLM i.e., S \rightarrow small-NLM and L \rightarrow large-NLM. Here, (+) \rightarrow augment this feature to the model in previous row, NLM \rightarrow *LSTM-LM*, (*) \rightarrow scores taken from Wang et al. (2018). Dotted line separates our baseline models from the previous works (*). Here, **bold** values indicate best performing proposed model in comparison to *LSTM-LM* baseline, and **GAIN(%)** indicates the improvement in performance of the same.

	MODEL	APNEWS		IMDB		BNC	
		S	L	S	L	S	L
BASELINES	<i>LDA+LSTM*</i>	54.83	50.17	69.62	62.78	96.38	87.28
	<i>LCLM*</i>	54.18	50.63	67.78	67.86	87.47	80.68
	<i>TopicRNN*</i>	54.12	50.01	66.45	60.14	93.55	84.12
	<i>TDLM*</i>	52.65	48.21	63.82	58.59	86.43	80.58
	<i>TCNLM*</i>	52.59	47.74	62.59	56.12	86.21	80.12
	<i>LSTM-LM</i>	64.95	59.28	72.31	65.54	106.82	98.78
	<i>LTA-NLM</i>	55.48	49.61	68.21	61.49	98.31	89.36
PROPOSED	+ <i>SDT</i>	48.23	42.85	63.81	58.90	90.36	80.30
	<i>ETA-NLM</i>	49.34	43.19	59.20	51.40	95.62	87.22
	+ <i>SDT</i>	48.75	43.50	57.83	50.51	95.64	88.37
	<i>LETA-NLM</i>	48.33	43.17	58.10	52.35	94.78	86.73
	+ <i>SDT</i>	42.98	39.41	56.65	51.05	88.30	81.12
	GAIN(%)	33.8	33.5	21.6	22.1	17.3	18.7

perform quantitative and qualitative evaluation of our proposed models on five NLP tasks.

4.1. Evaluation: Language Modeling

We present experimental results of language modeling using our proposed models on APNEWS, IMDB and BNC datasets (Lau et al., 2017). For NLM, we tokenize sentences and documents into words, lowercase all words and remove those words which occur less than 10 times. For NTM, we additionally remove stopwords, word occurring less than 100 times and top 0.1% most frequent words. We use standard language model perplexity as the evaluation measure for our proposed models. For data statistics and time complexity of experiments refer supplementary.

Experimental setup: We follow Wang et al. (2018) for our experimental setup. See supplementary for detailed hyperparameter settings. Sentence s being modeled at NLM side is removed from document d at NTM side. We use two settings of NLM component: (1) small-NLM (1-layer, 600 hidden units), and (2) large-NLM (2-layer, 900 hidden units). We fix the NLM sequence length to 30 and bigger sentences are split into multiple sequences of length less than 30. We initialize the input word embeddings for NLM with 300-dimensional pretrained embeddings extracted from word2vec (Mikolov et al., 2013) model trained on Google News. We perform an ablation study to get the best setting of hyperparameters α and $topN$ (see supplementary).

Table 3. Top 5 words of two randomly selected topics extracted using *ETA-NLM* model for APNEWS, IMDB and BNC datasets.

APNEWS		IMDB		BNC	
army	legal	comedy	disney	music	art
soldiers	jurors	jokes	daffy	album	art
infantry	trial	unfunny	cindrella	guitar	paintings
brigade	execution	satire	alladin	band	painting
veterans	jury	sandler	looney	music	museum
battalion	verdict	streep	bambi	pop	gallery

Baselines: We compare our proposed models with seven baseline models: (i) *LDA-LSTM*: concatenating pretrained LDA topic-proportion vector with *LSTM-LM*; (ii) *LCLM* (Wang & Cho, 2016); (iii) *TopicRNN* (Dieng et al., 2017); (iv) *TDLM* (Lau et al., 2017); (v) *TCNLM* (Wang et al., 2018); (vi) *LSTM-LM*: NLM component of our proposed models; and (vii) *LTA-NLM*: our baseline model.

Results: Language modeling perplexity scores are presented in Table 2. All topic composition models outperform *LSTM-LM* baseline which demonstrate the advantage of composing document topical semantics in NLM. Based on the results, here are three key observations: (i) *ETA-NLM* always performs better than *LTA-NLM* resulting in 11% (49.34 vs 55.48), 13.2% (59.20 vs 68.21) and 2.7% (95.62 vs 98.31) improvement for APNEWS, IMDB and BNC datasets respectively under small-NLM configuration. This behaviour asserts that ETR vector effectively captures fine-grained document topical semantics compared to LTR vector; (ii) *LETA-NLM* outperforms both *LTA-NLM* & *ETA-NLM* by exploiting complementary semantics of ETR and LTR vectors; (iii) composing sentence-level topic representations i.e., +*SDT*, further boost the performance by maintaining sentence-level topical discourse and hence *LETA-NLM* +*SDT* improve upon *LTA-NLM* model by 22% (42.98 vs 55.48), 17% (56.65 vs 68.21) and 10% (88.30 vs 98.31) for APNEWS, IMDB and BNC datasets respectively under small-NLM configuration.

4.2. Evaluation: Topic Modeling

Typically, topic models are evaluated using perplexity measure. However, in the context of our NCLM framework where we compose topical semantics in language modeling, we investigate quality of topics generated by our composite models. We follow Wang et al. (2018) to infer topic coherence of top 5/10/15/20 topic words for each topic using pairwise NPMI scores and average them to get an average coherence score. We use the same experimental setup and hyperparameter settings as described in section 4.1.

Baselines: Following Wang et al. (2018), we compare average coherence scores of our proposed models with the following baselines: (i) LDA (Blei et al., 2001); (ii)




Figure 3. Topic coherence score comparison of our proposed models and multiple baselines. (*) \rightarrow taken from Wang et al. (2018)

Table 4. Text Classification accuracy scores on three datasets. *CNN* \rightarrow model proposed by Kim (2014), *+Topic* \rightarrow augment topic feature in the above model, and **bold** values indicate best models.

	MODEL	20NS	R21578	IMDB
BASELINES	<i>CNN-Rand</i>	.721	.690	.888
	<i>+Topic</i>	.724	.699	.891
	<i>CNN-LSTM</i>	.745	.750	.899
PROPOSED	<i>CNN-LTA</i>	.753	.759	.907
	<i>CNN-ETA</i>	.775	.763	.903
	<i>CNN-LETA</i>	.770	.750	.908

TDLM (Lau et al., 2017); (iii) TopicRNN (Dieng et al., 2017); (iv) TCNLM (Wang et al., 2018).

Results: Average topic coherence scores are presented as bar-plot in Figure 3. Based on the plot, here are two key observations: (i) all of our proposed models outperform every baseline by a significant margin; (ii) however, there is no discernible pattern in the topic coherence scores of our proposed models, hence, an improvement in language modeling performance does not correspondingly improve topic coherence. For a qualitative overview, Table 3 shows 2 randomly chosen topics for each dataset. See *supplementary* for examples of topic-aware sentence generation.

4.3. Evaluation: Text Classification

We evaluate the quality of representations learned by our proposed models via document classification. We use three labeled datasets: 20Newsgroups (20NS), Reuters (R21578) and IMDB movie reviews (IMDB) (See *supplementary for data statistics*). Based on the scores in Table 2, we employ our best performing composite language models as static feature extractors. For each document d , we extract: (1) output state \mathbf{o}_m for each input word x_m via NLM

component ; and (2) document topic representation vector $\mathbf{c} \in \{\mathbf{h}_d, \mathbf{z}_d^{att}, [\mathbf{h}_d; \mathbf{z}_d^{att}]\}$ via NTM component based on the model configuration. We then concatenate \mathbf{c} with each \mathbf{o}_m and use them as inputs to train a CNN based text classifier proposed by Kim (2014). For IMDB movie reviews dataset, we use our best models trained on unlabeled IMDB dataset as feature extractor. However, as 20NS and R21578 are news-domain datasets, we employ best models trained on APNEWS because of its bigger corpus size than BNC.

Baselines: Using *LTA-NLM*, *ETA-NLM* and *LETA-NLM* as feature extractors, we propose *CNN-LTA*, *CNN-ETA* and *CNN-LETA* respectively. We compare these models with: (1) *CNN-Rand*: randomly initialized CNN text classifier (Kim, 2014) without embedding update; (2) *+Topic*: additionally concatenate LTR vector \mathbf{h}_d with each word embedding input in *CNN-Rand*; and (3) *CNN-LSTM*: use \mathbf{o}_m extracted using pre-trained *LSTM-LM* as input to CNN classifier.

Results: Document classification results are presented in Table 4. Based on the results, there are two notable key findings: (1) *CNN-Rand* performed worst among all models, but incorporating document topical semantics i.e., *+Topic*, provided a boost in classification scores for 20NS (.724 vs .721), R21578 (.699 vs .690) and IMDB (.891 vs .888) datasets which shows the advantage of composing document topic representations during language modeling; (2) however, the best performance comes from *CNN-ETA* for 20NS (.775 vs .745) & R21578 (.763 vs .750) datasets and *CNN-LETA* for IMDB (.908 vs .899). This shows that *ETA-NLM* and *LETA-NLM* learn better representations than *LTA-NLM* and *LSTM-LM* and suggests that ETR vector effectively captures fine-grained document semantics than LTR vector.

4.4. Evaluation: Information Retrieval

We further evaluate the quality of learned representations via document retrieval task. We show retrieval performance on three datasets: 20Newsgroups (20NS), Reuters (R21578) and AGnews. Following Gupta et al. (2019a), we treat all test documents as queries and retrieve a fraction of training documents closest to each query using cosine similarity measure. Then, we compute precision for each query as the fraction of all retrieved documents with same label as query and average over precision scores of all queries to get an final precision score. Similar to text classification, for each document d of length N_d , we extract the final output state \mathbf{o}_{N_d} of the NLM component and concatenate it with $\mathbf{c} \in \{\mathbf{h}_d, \mathbf{z}_d^{att}, [\mathbf{h}_d; \mathbf{z}_d^{att}]\}$ extracted via NTM component to get a composite representation. We then compute cosine similarity of each query-document pair using this composite representation. We employ our proposed models pre-trained on APNEWS dataset as feature extractors. We compute precision scores for top-5 and top-10 retrieved documents for each dataset.

Table 5. Information Retrieval Evaluation: Average precision scores on three labeled datasets for top-5 (P@5) & top-10 (P@10) retrievals. **Bold** value indicates best score in each column.

MODEL	20NS		R21578		AGNEWS	
	P@5	P@10	P@5	P@10	P@5	P@10
<i>NTM</i>	.198	.190	.581	.567	.607	.600
<i>LTA-NLM</i>	.264	.217	.585	.558	.682	.666
<i>ETA-NLM</i>	.287	.242	.590	.562	.683	.665
<i>LETA-NLM</i>	.281	.236	.615	.589	.694	.675

Table 6. WSD evaluation results using F1 scores (micro). Among proposed models, **Bold** values indicates best model compared to *BiLSTM-LM*.

MODEL	DEV		TEST			
	SE07	SE13	SE15	SE2	SE3	ALL
<i>MFS</i>	54.5	63.8	67.1	65.6	66.0	65.5
<i>BiLSTM-LM</i>	55.0	53.9	60.8	63.6	60.8	59.8
<i>LTA-NLM</i>	56.0	54.8	60.9	64.8	62.2	60.7
<i>ETA-NLM</i>	55.4	54.7	60.6	64.8	62.3	60.7
<i>LETA-NLM</i>	55.6	54.7	61.1	64.7	62.1	60.7

Baselines: Document retrieval is used to evaluate applicability of topic models. Therefore, we compare the retrieval performance using composite representations of our best performing *LTA-NLM*, *ETA-NLM* and *LETA-NLM* models (using table 2) with baseline performance using document LTR vector \mathbf{h}_d extracted via pre-trained NTM component of our proposed composite language model.

Results: Document retrieval results are presented in Table 5. It is worth noting that: (1) all of our proposed composite models performed much better than the NTM component itself, and (2) *ETA-NLM* and *LETA-NLM* models performed much better than *LTA-NLM* which reconfirms that ETR vector is more descriptive than LTR vector and support NLM in encoding long-term semantic dependencies. (3) as compared to *LTA-NLM*, *ETA-NLM* performed best for 20NS (.376 vs .355), while *LETA-NLM* performed best for R21578 (.664 vs .629) and AGnews (.694 vs .682).

4.5. Evaluation: Word Sense Disambiguation

Word sense disambiguation (WSD) deals with correct prediction of appropriate semantic meaning (sense) of a word given its surrounding context. Similar to language modeling, a word can have semantic dependencies across sentence boundaries. Therefore, we show the applicability of our NCLM framework which boosts correct sense prediction by exploiting document-level topical knowledge to capture long-range semantic dependencies. We focus on English all-words WSD task where, the aim is to simultaneously predict correct sense for each word in a given sentence. We use evaluation framework proposed by Navigli et al. (2017)

for training and evaluation.

Experimental setup and Baselines: Following Raganato et al. (2017), we use 1-layer bidirectional-LSTM cell with 100 hidden units in the NLM component of our proposed models *LTA-NLM*, *ETA-NLM* and *LETA-NLM*. In the absence of next word prediction task, we use full document context on NTM side. Our models consider all the words in a sentence as input, and learn to predict the correct sense via multinomial logistic regression over a vocabulary of all unique senses present in training data. Models are trained using a learning rate of 1e-3 & batch size of 32 and predictions are evaluated using micro F1 score. We compare evaluation performance of our model with the following baselines: (1) *MFS*: most frequent sense extracted from WordNet (Miller, 1995); and (2) *BiLSTM-LM*: language model using 1-layer bidirectional-LSTM cell with 100 hidden units.

Results: WSD F1 scores are presented in Table 6. Observe that by averaging F1 scores over all test datasets, our proposed models outperform *BiLSTM-LM* (60.7 vs 59.8) which again confirms the advantage of document-level semantic knowledge in resolving sense ambiguities via composition.

5. Conclusion

We have presented a neural composite language modeling framework that leverages both the latent and explainable topic representations by composing a neural language model and a neural topic model. Moreover, we have introduced sentence-topic association along with document-topic association to retain sentence-level topical discourse. Experimental results on several language understanding tasks have supported our multi-fold contributions.

Acknowledgments

This research was supported by Bundeswirtschaftsministerium (bmwi.de), grant 01MD19003E (PLASS (plass.io)) at Siemens AG - CT Machine Intelligence, Munich Germany.

References

- Blei, D. M., Ng, A. Y., and Jordan, M. I. Latent dirichlet allocation. In *Advances in Neural Information Processing Systems 14 [Neural Information Processing Systems: Natural and Synthetic, NIPS 2001, December 3-8, 2001, Vancouver, British Columbia, Canada]*, pp. 601–608, 2001.
- Bojanowski, P., Grave, E., Joulin, A., and Mikolov, T. Enriching word vectors with subword information. *Transactions of the Association for Computational Linguistics*, 5: 135–146, 2017. ISSN 2307-387X.
- Cho, K., van Merriënboer, B., Gülçehre, Ç., Bahdanau, D.,

- Bougares, F., Schwenk, H., and Bengio, Y. Learning phrase representations using RNN encoder-decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pp. 1724–1734, 2014.
- Dieng, A. B., Wang, C., Gao, J., and Paisley, J. W. Topicrnn: A recurrent neural network with long-range semantic dependency. In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*, 2017.
- Gupta, P., Chaudhary, Y., Buettner, F., and Schütze, H. Document informed neural autoregressive topic models with distributional prior. In *The Thirty-Third AAAI Conference on Artificial Intelligence, AAAI 2019, The Thirty-First Innovative Applications of Artificial Intelligence Conference, IAAI 2019, The Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2019, Honolulu, Hawaii, USA, January 27 - February 1, 2019*, pp. 6505–6512, 2019a.
- Gupta, P., Chaudhary, Y., Buettner, F., and Schütze, H. Text-tovec: Deep contextualized neural autoregressive topic models of language with distributed compositional prior. In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*, 2019b.
- Hochreiter, S. and Schmidhuber, J. Long short-term memory. *Neural Computation*, 9(8):1735–1780, 1997.
- Khandelwal, U., He, H., Qi, P., and Jurafsky, D. Sharp nearby, fuzzy far away: How neural language models use context. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 1: Long Papers*, pp. 284–294, 2018.
- Kim, Y. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*, pp. 1746–1751, 2014.
- Lau, J. H., Baldwin, T., and Cohn, T. Topically driven neural language model. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 1: Long Papers*, pp. 355–365, 2017.
- Miao, Y., Yu, L., and Blunsom, P. Neural variational inference for text processing. In *Proceedings of the 33rd International Conference on Machine Learning, ICML 2016, New York City, NY, USA, June 19-24, 2016*, pp. 1727–1736, 2016.
- Mikolov, T., Karafiát, M., Burget, L., Cernocký, J., and Khudanpur, S. Recurrent neural network based language model. In *INTERSPEECH 2010, 11th Annual Conference of the International Speech Communication Association, Makuhari, Chiba, Japan, September 26-30, 2010*, pp. 1045–1048, 2010.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. Efficient estimation of word representations in vector space. In *1st International Conference on Learning Representations, ICLR 2013, Scottsdale, Arizona, USA, May 2-4, 2013, Workshop Track Proceedings*, 2013.
- Miller, G. A. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, 1995.
- Navigli, R., Camacho-Collados, J., and Raganato, A. Word sense disambiguation: A unified evaluation framework and empirical comparison. In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics, EACL 2017, Valencia, Spain, April 3-7, 2017, Volume 1: Long Papers*, pp. 99–110, 2017.
- Peters, M. E., Neumann, M., Iyyer, M., Gardner, M., Clark, C., Lee, K., and Zettlemoyer, L. Deep contextualized word representations. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pp. 2227–2237, 2018.
- Raganato, A., Bovi, C. D., and Navigli, R. Neural sequence learning models for word sense disambiguation. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing, EMNLP 2017, Copenhagen, Denmark, September 9-11, 2017*, pp. 1156–1167, 2017.
- Wang, T. and Cho, K. Larger-context language modelling with recurrent neural network. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*, 2016.
- Wang, W., Gan, Z., Wang, W., Shen, D., Huang, J., Ping, W., Satheesh, S., and Carin, L. Topic compositional neural language model. In *International Conference on Artificial Intelligence and Statistics, AISTATS 2018, 9-11 April 2018, Playa Blanca, Lanzarote, Canary Islands, Spain*, pp. 356–365, 2018.