

# Enriching and Controlling Global Semantics for Text Summarization

Thong Nguyen<sup>1,3\*</sup>, Anh Tuan Luu<sup>2†</sup>, Truc Lu<sup>3</sup>, Tho Quan<sup>3</sup>

<sup>1</sup>VinAI Research, Vietnam

<sup>2</sup>Nanyang Technological University, Singapore

<sup>3</sup>Ho Chi Minh City University of Technology, Vietnam National University Ho Chi Minh City

v.thongnt66@vinai.io, tuancs219@yahoo.com, qttho@hcmut.edu.vn

## Abstract

Recently, Transformer-based models have been proven effective in the abstractive summarization task by creating fluent and informative summaries. Nevertheless, these models still suffer from the short-range dependency problem, causing them to produce summaries that miss the key points of document. In this paper, we attempt to address this issue by introducing a neural topic model empowered with normalizing flow to capture the global semantics of the document, which are then integrated into the summarization model. In addition, to avoid the overwhelming effect of global semantics on contextualized representation, we introduce a mechanism to control the amount of global semantics supplied to the text generation module. Our method outperforms state-of-the-art summarization models on five common text summarization datasets, namely CNN/DailyMail, XSum, Reddit TIFU, arXiv, and PubMed.

## 1 Introduction

Automatic text summarization corresponds to text understanding and text generation processes. In general, there are two main approaches to perform this task. Extractive systems (Liu, 2019; Narayan et al., 2020; Zhang et al., 2019; Jia et al., 2020) highlight salient words or sentences from the source text and form the final summary by concatenating them. On the other hand, abstractive methods (See et al., 2017; Zhang et al., 2020; Zou et al., 2020) switch among generating new words, choosing phrases from the source document, and rephrasing them. Abstractive summarization, which is the focus of this paper, is usually more advanced and closer to human-like interpretation.

Recently, abstractive summarization studies (Lewis et al., 2019; Zhang et al., 2020; Chen and Yang, 2020) are dominated by Transformer-based

\*Work done during internship at VinAI Research

†Corresponding Author

**DOCUMENT:** While Richie Benaud rose from the suburbs to captain Australia, he will be remembered forever for his mastery of commentating. The champion leg spinner turned cricket commentating into an art form, earning him the title of 'the Voice of Cricket.' His commentary was understated, measured and often extremely funny, and were perfectly timed. Scroll down for video. 84-year-old cricket commentator Richie Benaud has passed away after a battle with skin cancer. His sayings from the hundreds of Test and One Day cricket matches he commentated on across the world were often what fans remembered from important moments. His signature one liners soon dropped to a simple word. 'Marvellous...' will forever be linked to the cricket legend. On commentating, Benaud said: 'My mantra is - put your brain into gear and if you can add to what's on the screen then do it, otherwise shut up.' He once described the scene on the field: 'From our broadcast box you can't see any grass at all, it is simply a carpet of humanity.' On captaincy, and he was one of the best Test captains Australia ever had, Benaud was modest: 'The hallmark of a great captain is the ability to win the toss, at the right time.' The former leg spinner turned cricket commentating into an art form, giving him the title 'the Voice of Cricket'. But he cautioned that description with: 'Captaincy is 90 per cent luck and 10 per cent skill. But don't try it without that 10 per cent.' [...]

**GOLD SUMMARY:** Cricket commentator Richie Benaud has passed away after cancer battle. The 84-year-old will be remembered for his mastery of commentating. The former leg spinner earned himself the title of the 'Voice of Cricket'. His trademark line was 'Marvellous'.

**PEGASUS:** The champion leg spinner turned cricket commentating into an art form, earning him the title of 'the Voice of Cricket'. His commentary was understated, measured and often extremely funny, and were perfectly timed.

**Our model:** 84-year-old cricket commentator Richie Benaud has passed away after a battle with skin cancer. The champion leg spinner earned the title of 'the Voice of Cricket'. His commentary was understated, measured and often extremely funny. His trademark word, 'Marvellous...' will forever be linked to the cricket legend.

Table 1: An example of summarization outputs.

architecture (Vaswani et al., 2017). Despite good performance in large scale datasets, Transformer-based summarization models have been proven to have the tendency to favor encoding short-range dependencies (Zhang et al., 2020), i.e., whenever there is one word from the input generated in the summary, the model tends to continue generating the nearby words due to their high attention scores




Figure 1: Self-attention weights of “commentary” in the PEGASUS model

to the previous generated word. As such, if the main content of the document is out of reach from the generated word, the final summary can miss that key information. For example, in Table 1, PEGASUS, a state-of-the-art Transformed-based model, failed to capture one key information of the document, i.e., “84-year-old cricket commentator Richie Benaud has passed away after a battle with skin cancer”. To understand this phenomenon, we visualize the attention scores in the model during the generation process. As shown in Figure 1, when the model generates “commentary”, the main subject of the blue sentence, it tends to point to and generate nearby words such as “his”, “understated”, “funny”, etc. due to their high attention scores while words in the further range such as “Richard”, “Benaud”, “pass”, and “away” receive little weight. Consequently, although PEGASUS generates a grammatically correct summary, the summary lacks the key content which describes the death of “Richie Benaud”.

To avoid missing key points in summarizing, one solution is to furnish the models with global semantics by using probabilistic topic models such as LDA (Narayan et al., 2018), Poisson factor analysis (Wang et al., 2020), or inner hidden states (Liu et al., 2019). Nevertheless, traditional topic models were shown to be inefficient in scalability for large-scale datasets (Hoffman et al., 2013; Rezende and Mohamed, 2015) and have limited capability of describing documents (Ding et al., 2018).

To overcome the above problems, we propose a novel method that integrates neural topic model into summarization architecture. Specifically, we aim to utilize the posterior distribution learned from the neural topic model as an approximation of global semantics of the document and from that, provide a signal for summarization model to have a better understanding of overall document. However, there is one critical question: how can we match the neural topic model’s posterior distribution with the true posterior as it has been proven in improving the performance of variational inference (Rezende and Mohamed, 2015)? To this end, we propose a method to adapt normalizing flow

in the neural topic model to have a better approximation of true distribution and integrate it into the summarization model. Integrating flow mechanism to better approximate the true posterior has been proven to improve performance for variational inference (Rezende and Mohamed, 2015) as well as for downstream tasks such as image synthesis (Kingma et al., 2016), etc. However, to the best of our knowledge, there is no study to investigate the benefit of flow mechanism for the abstractive summarization task.

On the other hand, even though rich global semantics is beneficial, there are recent studies showing that the redundant amount of global semantics may cause harm to hidden representation since it introduces detrimental noise to the model (Tenney et al., 2019; Li et al., 2020). Therefore, we propose a novel contextualized gating mechanism to control the flow of global semantics and maintain important information of the hidden states in the main summarization model.

The contributions of our paper can be summarized as follows:

- We propose a novel architecture which takes the global semantics into consideration when performing abstractive summarization.
- To this end, we introduce a neural topic model which is empowered with normalizing flow to enrich the global semantics and contextualized gating mechanism to better control the effect of global semantics on hidden representations.
- We conduct extensive experiments and outperform other state-of-the-art summarization models on five benchmark datasets, i.e. CNN/DailyMail, XSum, Reddit TIFU, PubMed, and arXiv, while generating summaries which favor human judgments, and producing human-interpretable topics.

## 2 Related Work

### 2.1 Transformer-based Text Summarization

Transformer (Vaswani et al., 2017) and its variants have demonstrated high efficiency in text summarization. (Liu and Lapata, 2019) first use to perform

extractive summarization. (Zhong et al., 2020) propose using Siamese BERT to score among summaries extracted from the source document, exploring the rich semantic space that those summaries are projected onto. (Narayan et al., 2020) combine HiBERT and structured transformers to extract the document incrementally to form the final summary.

For abstractive approaches, (Zhang et al., 2020) develop a pretraining scheme well-suited for abstractive summarization. Other frameworks uniting language understanding and text generation such as BART (Lewis et al., 2019), UniLM (Dong et al., 2019), T5 (Raffel et al., 2019), (Tuan et al., 2020), and MASS (Song et al., 2019) provide further standards for future works. Unified system such as BottomUp (Gehrmann et al., 2018) extracts salient phrases and then generates the summary based upon the extracted content. (Subramanian et al., 2019) further improve with their decoder as a Transformer language model.

## 2.2 Topic-aware Summarization Models

Various works integrate global semantics of topic model into the sequential information. One method is to attend topic vectors with the hidden states, only choosing entries with high document-level representations (Zheng et al., 2020). (Wang et al., 2020) design three modules to incorporate topic information to attentive heads, provide topical embedding, and form document-related representations. Other works integrate topical information into convolutional-based models (Narayan et al., 2018; Wang et al., 2018). Ailem et al. 2019 have their pointer-generator conditioned on both the input document and the latent vector. (Fu et al., 2020) study how to effectively assist deep-learning summarization frameworks with external global information. Arguing that each paragraph in the document possesses a separate subtopic, they propose to merge topic information hierarchically with the dense word embedding.

Unfortunately, there is still limited effort controlling the effect of global semantics on the contextualized representations and enriching the global semantics for summarization performance.

## 3 Methodology

The overall architecture of our approach is given in Figure 2. It comprises of a topic-oriented encoder, a topic-oriented decoder, and a flow-based neural topic model.

Formally, given a document as input, we process it into a sequence of tokens  $X = \{x_i\}$ , and the bag-of-word (BoW) representation  $\mathbf{x}_{bow}$ .  $X$  is taken as the input for the text summarization module, while  $\mathbf{x}_{bow}$  serves as the input for the neural topic model.

### 3.1 Flow-based Neural Topic Model

The architecture of the neural topic model (NTM) takes inspiration from (Miao et al., 2017) based on variational autoencoder (Kingma and Welling, 2013). In this work, we adapt the normalizing flow to the neural topic model to better grasp the global semantic patterns of the document.

**BoW Encoder.** In particular, the input  $\mathbf{x}_{bow}$  is first encoded into a latent variable  $\mathbf{z}$  by a topic encoder. Each input is passed to obtain the prior mean  $\mu$  and prior standard deviation  $\sigma$

$$\pi = f_{MLP}(\mathbf{x}_{bow}), \mu = f_1(\pi), \log \sigma = f_2(\pi) \quad (1)$$

where  $f_{MLP}$  is a non-linear transformation with a tanh activation function;  $f_1$  and  $f_2$  are two linear transformations with bias. To obtain the topic distribution, we draw the latent variable  $\mathbf{z} \sim \mathcal{N}(\mu, \sigma^2)$ .

**Flow.** Different from conventional neural topic model, a flow is applied to map the latent vector to a more complicated distribution. Formally, flow is a chain of transformations  $f_1, f_2, \dots, f_K$  which are all invertible and have the Jacobian easy to compute.

$$\mathbf{z}_K = f_0 \circ f_1 \dots \circ f_K(\mathbf{z}) \quad (2)$$

**BoW Decoder.** Given the new topic vector, the BoW decoder retains the original input  $\mathbf{x}_{bow}$  by generating  $\mathbf{x}'_{bow}$ . We take the following procedure to simulate the reconstruction of  $\mathbf{x}_{bow}$

- Topic mixture  $\theta = \text{softmax}(f_\theta(\mathbf{z}_K))$
- For each word  $w \in \mathbf{x}_{bow}$ , draw  $w \sim \text{softmax}(f_\phi(\theta))$

where  $f_*(\cdot)$  is a ReLU-activated non-linear transformation. The weight matrix of  $f_\phi$  is chosen as the topic-word distribution  $(\phi_1, \phi_2, \dots, \phi_K)$ . We proceed to employ the topic mixture  $\theta$  to guide the text summarization process.

### 3.2 Neural Topic Modeling for Transformer

Text summarization model is passed a source document  $X = \{x_i\}_{i=1}^N$  and its task is to predict the target summary  $Y = \{y_j\}_{j=1}^M$ . In this setting, the document  $D$  has  $N$  tokens and the summary  $S$  has  $M$  tokens ( $M < N$ ).

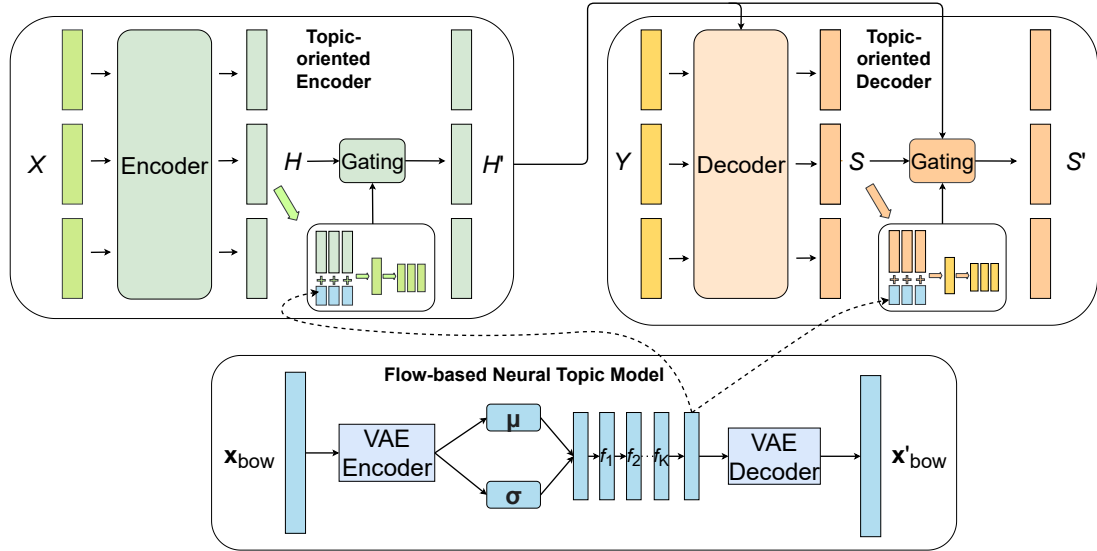


Figure 2: Our overall architecture

Our model inherits the Transformer-based architecture. Particularly, it consists of an encoder and decoder. The encoder learns the context of the source text, and the decoder then predicts the target summary, by learning the context of the generated tokens and attending over encoder hidden states. In our case, we make both the encoder and decoder conditioned on the latent topic yielded by the neural topic model.

**Topic-oriented Encoder** We add the special token “CLS” to the beginning of the input. At each iteration, the encoder outputs a localized representation  $H = \{\mathbf{h}_i\}_{i=1}^N$  for each token in the source document  $X$

$$\mathbf{h}_{CLS}, \mathbf{h}_1, \dots, \mathbf{h}_N = \text{Encoder}(x_{CLS}, x_1, \dots, x_N) \quad (3)$$

This explores the relationship among the input tokens (the encoder), or discovering the context each token stays in. We relate the context of each word to the main topic of the document by modulating the  $i$ -th hidden state  $\mathbf{h}_i$

$$\mathbf{h}'_i = g(\mathbf{h}_i, \theta) \quad (4)$$

where  $g$  is a function used to introduce the global semantics to the hidden representations which we will discuss later as the contextualized gating mechanism in section 3.3

**Topic-oriented Decoder** We also make “CLS” the first input of the decoder. The decoder bridges the summary  $Y$  and document  $X$ , creating target hidden states  $S = \{\mathbf{s}_j\}_{j=1}^M$  aligned with the source text. Because of the uni-directionality of the text summarization task, the decoder must work in a

left-to-right fashion

$$\mathbf{s}_j = \text{Decoder}(y_{CLS}, y_1, y_2, \dots, y_{j-1}, \{\mathbf{h}'_i\}_{i=1}^N) \quad (5)$$

Similar to the Encoder, we seek to inject the semantics of the topic model into the output hidden state.

$$\mathbf{s}'_j = g(\{\mathbf{h}'_i\}_{i=1}^N, \mathbf{s}_j, \theta) \quad (6)$$

### 3.3 Contextualized Gating Mechanism

Because a specified amount of semantic meaning, whether it is local or global, has been embedded in the contextualized representations, it is reasonable to only append sufficient information to the calculated hidden states to maximize the efficiency of the topical information. We adapt the gating mechanism of (Cho et al., 2014) to achieve this goal. In our contextualized gating mechanism, we approximate the necessary amount of global semantics based on the obtained hidden states.

**Encoder Gating** For the encoder, we take the hidden representation of “CLS” token to control the amount of additional global information

$$\lambda^E = \text{Sigmoid}(W^E \mathbf{h}_{CLS} + b^E) \quad (7)$$

where  $W^E \in \mathbb{R}^{d \times d}$ , and  $d$  is the dimension of the hidden representation. We form the topic-aware hidden state by merging it with the topic mixture and mapping it onto a topical space

$$\mathbf{u}_i = [\mathbf{h}_i, \theta] \quad (8)$$

$$\mathbf{c}_i = f_{enc\_topic}(\mathbf{u}_i) \quad (9)$$

where  $f_{enc\_topic}$  is a non-linear transformation. The topic-oriented encoder hidden state of every token is the fusion of the topic-aware and the original representation.

$$\mathbf{h}'_i = \lambda^E \mathbf{c}_i + (1 - \lambda^E) \mathbf{h}_i \quad (10)$$

**Decoder Gating** The amount of topic mixture used for the decoder is controlled by both encoder and decoder hidden state

$$\lambda^D = \text{Sigmoid}(W_1^D \mathbf{h}_{CLS} + W_2^D \mathbf{s}_{CLS} + b^D) \quad (11)$$

where  $W_1^D \in \mathbb{R}^{d \times d}$ ,  $W_2^D \in \mathbb{R}^{d \times d}$ . This switching probability is used to modulate the decoder hidden state, which follows the same computation with the encoder gating.

$$\mathbf{t}_j = [\mathbf{s}_j, \theta_{dec}] \quad (12)$$

$$\mathbf{e}_j = f_{dec\_topic}(\mathbf{t}_j) \quad (13)$$

$$\mathbf{s}'_j = \lambda^D \mathbf{e}_j + (1 - \lambda^D) \mathbf{s}_j \quad (14)$$

### 3.4 Training Objective

Our framework favors end-to-end learning of neural topic modeling and text summarization. In this section, we formally define the objective functions for the two modules.

For our neural topic model, the objective function is derived from the evidence lower bound (Blei et al., 2017). We adapt the change of variables in normalizing flow that determine the distribution of the variable at the end of the flow to the loss of neural topic model

$$\begin{aligned} \mathcal{L}_{\text{NTM}} &= \log p(\mathbf{x}, \mathbf{z}) - \log q(\mathbf{z}|\mathbf{x}) \\ &= -\log q(\mathbf{z}_0) + \sum_{i=1}^K \log \left| \det \frac{\partial f_i}{\partial \mathbf{z}_{i-1}} \right| \\ &\quad + \log p(\mathbf{x}|\mathbf{z}_K) + \log p(\mathbf{z}_K) \end{aligned} \quad (15)$$

where  $p(\mathbf{z}_K)$  denotes the prior distribution constructed by the flow;  $p(\mathbf{x}|\mathbf{z}_K)$  stands for the log-likelihood of the document;  $\log q_K(\mathbf{z}_K)$  denotes the approximate posterior distribution. Detailed derivation is available in Appendix.

For text summarization, we minimize the cross-entropy loss

$$\mathcal{L}_{\text{sum}} = - \sum_{j=1}^M \log p(y_j | \{x_i\}_{i=1}^N, y_{i < j}) \quad (16)$$

where  $N$  and  $M$  are the length of the document  $X$  and summary  $Y$ , respectively. The entire framework is trained with the linear combination of two loss functions  $\mathcal{L}_{\text{sum}}$  and  $\mathcal{L}_{\text{NTM}}$

$$\mathcal{L} = \mathcal{L}_{\text{sum}} + \lambda \mathcal{L}_{\text{NTM}} \quad (17)$$

where  $\lambda$  is the hyperparameter balancing the effect of neural topic model on the training process.

## 4 Experimental Setup

### 4.1 Datasets

We evaluate our proposed method on five benchmark datasets: CNN/DailyMail (CNN/DM) (Hermann et al., 2015), XSum (Narayan et al., 2018), Reddit TIFU (Kim et al., 2018), arXiv, and PubMed (Cohan et al., 2018). The datasets possess various styles and varying lengths.

**CNN/DM** is constructed by collecting news articles written by CNN and DailyMail journalists. For each article, its highlights are chosen as the summary. We use the non-anonymized version and follow the conventional training/validation/test split in (Hermann et al., 2015).

**XSum** comprises of 226,711 news articles, each of which is linked with a one-sentence summary. Our preprocessing and training/validation/test split is analogous to (Narayan et al., 2018).

**Reddit TIFU** includes 120K informal posts from the online discussion forum Reddit, strictly following the rule of constructing an expressive ‘‘TL;DR’’ summary. In this work, the long subset of the dataset is applied for performance evaluation.

**arXiv, PubMed** are two long-document datasets of scientific publications. For each document, the abstract is chosen to be the summary.

We present the statistics of datasets in Table 2.

Dataset	Train	Val	Test	$l_D$	$l_S$
CNN/DM	287113	13368	11490	781	56
XSum	204045	11332	11334	431	23
Reddit TIFU	33711	4214	4214	433	23
arXiv	203037	6436	6440	4938	220
PubMed	119924	6633	6658	3016	203

Table 2: Description of the evaluation datasets.  $l_D$  and  $l_S$  stand for average length of document and summary

## 4.2 Implementation Details

**Neural Topic Model** Provided a dataset, firstly we pretrain the flow-based topic model so that it is able to obtain the prior context of the downstream documents. We experimented with different choices of the topic number  $T \in \{50, 100, 200\}$  and the number of invertible transformations applied in the flow of neural topic model  $K \in \{1, 4, 16\}$  on CNN/DailyMail dataset.

**Summarization Module** We use the pretrained checkpoint open-sourced by (Zhang et al., 2020), integrate and jointly finetune with the flow-based neural topic model on downstream datasets. Following (Zhang et al., 2020), during test time, our beam search is conducted with a beam size of 8, and top-3 checkpoints are selected based on their evaluation loss in the validation set, and we average their results on the test set. More detailed settings can be found in Appendix.

## 4.3 Comparisons

As baselines, we compare our proposed architecture against a wide variety of previous studies:

- **PTGEN** (See et al., 2017): a pointer-generator baseline that allows switching between generating words from the vocabulary and copying words from the source.
- **PTGEN+Cov** (See et al., 2017): a pointer-generator baseline with coverage mechanism.
- **DRM** (Paulus et al., 2017): a deep reinforced model which handles the coverage mechanism by using intra-attention among decoder tokens.
- **DiscAttn** (Cohan et al., 2018): a Seq2Seq model which targets the long-document summarization.
- **BertSUM** (Liu and Lapata, 2019): a baseline with finetuning strategy is designed based on the discrepancy between the encoder and decoder.
- **ExtSum-LG+RdLoss** (Xiao and Carenini, 2020): a Transformer-based model with a training scheme to explicitly reduce redundancy.
- **MatchSUM** (Zhong et al., 2020): a baseline that makes use of Siamese BERT to score among candidate summaries.
- **BottomUp** (Gehrmann et al., 2018): a baseline uses extractive-abstractive approach: initially extracts salient phrases and performs abstractive summarization based on extracted content.
- **TLM-I+E** (Subramanian et al., 2019): a baseline improved on (Gehrmann et al., 2018) by utilizing Transformer language model as the decoder.

Model	R1	R2	RL
PTGEN	36.44	15.66	33.42
PTGEN + Cov	39.56	17.28	36.38
DRM	41.16	15.75	39.08
BertSUM	43.85	20.34	39.90
MatchSUM	44.41	20.86	40.55
BottomUp	41.22	18.68	38.34
BART	44.16	21.28	40.90
PEGASUS	44.17	21.47	41.11
VHTM	40.57	18.05	37.18
BertSUM + TA	43.06	20.58	39.67
BART + TA	44.47	21.39	41.32
Our Model	<b>44.52</b>	<b>21.95</b>	<b>41.39</b>

Table 3: Results in text summarization on CNN/DailyMail

Model	R1	R2	RL
PTGEN	29.70	9.21	23.24
PTGEN + Cov	28.10	8.02	21.72
BertSUM	38.81	16.50	31.27
MatchSUM	24.86	4.66	18.41
BART	45.14	22.27	37.25
PEGASUS	47.21	24.56	39.25
BertSUM + TA	39.77	17.39	32.39
BART + TA	45.76	22.68	38.03
Our Model	<b>49.57</b>	<b>25.08</b>	<b>41.81</b>

Table 4: Results in text summarization on XSum

- **BART** (Lewis et al., 2019): a baseline which is pretrained with denoising tasks.
- **PEGASUS** (Zhang et al., 2020): a Transformer-based model with pretraining procedure comprised of two tasks: masked sentences prediction and masked language modeling.
- **BertSUM + TA** (Wang et al., 2020): a BertSUM model equipped with the topic assistant inspired by the Poisson Factor Analysis topic model.
- **BART + TA** (Wang et al., 2020): a BART version with a plugged-in topic assistant inspired by the Poisson Factor Analysis topic model.
- **VHTM** (Fu et al., 2020): a baseline which takes hierarchical structure of the source text into account and considers each section as a subtopic.

## 5 Experimental Results

Model	R1	R2	RL
BART	24.19	8.12	21.31
MatchSUM	25.09	6.17	20.13
PEGASUS	26.63	9.01	21.60
Our Model	<b>27.96</b>	<b>9.43</b>	<b>23.08</b>

Table 5: Results in text summarization on Reddit TIFU

Model	R1	R2	RL
PTGEN + Cov	32.06	9.04	25.16
DiscAttn	35.80	11.05	31.80
ExtSum-LG+RdLoss	44.01	17.79	39.09
TLM-I+E	41.62	14.69	38.03
PEGASUS	43.82	16.74	39.15
Our Model	<b>44.53</b>	<b>19.22</b>	<b>40.61</b>

Table 6: Results in text summarization on arXiv dataset

Model	R1	R2	RL
PTGEN + Cov	31.55	8.52	27.38
DiscAttn	38.93	15.37	35.21
MatchSUM	41.21	14.91	20.13
ExtSum-LG+RdLoss	45.30	20.42	40.95
Sent-CLF	42.13	16.27	39.21
PEGASUS	44.29	19.19	40.42
Our Model	<b>45.99</b>	<b>20.49</b>	<b>41.25</b>

Table 7: Results in text summarization on PubMed

## 5.1 Automatic Evaluation

We use the automatic metrics of ROUGE scores (Lin, 2004). In Table 3, 4, 5, 6, and 7, we report the unigram overlap (ROUGE-1), bigram overlap (ROUGE-2) to assess informativeness, and longest common subsequence (ROUGE-L) for the fluency of the generated summary. Our model outperforms prior works on five standard datasets.

For CNN/DailyMail, we achieve an absolute improvement of 0.35 in ROUGE-1, 0.48 in ROUGE-2, and 0.28 in ROUGE-L over PEGASUS. Furthermore, our model obtains better results than the previous topic-aware model BART + TA in ROUGE-2 with 0.6 points. This shows that our methods can generate summaries that include important content in the document.

On the XSum dataset, which is more abstractive than CNN/DailyMail (Bommasani and Cardie, 2020), our gain is even more pronounced. Compared with BART + TA, we achieve 3.8 absolute improvement in ROUGE-1, 2.4 in ROUGE-2, and 3.8 in ROUGE-L.

For Reddit TIFU, in which most of the source texts and the target summaries are informal, our model outperforms PEGASUS by 1.3 in ROUGE-1, 0.4 in ROUGE-2, and 1.5 in ROUGE-L. These results show that global semantics is capable of helping the model generate better target summaries.

For arXiv and PubMed dataset, we also achieve improvement over the baseline PEGASUS, which is designed specifically for abstractive text summarization. In particular, for arXiv dataset, we gain an increase of 0.71 in ROUGE-1, 2.48 in ROUGE-2, and 1.46 in ROUGE-L. For PubMed dataset, the

increase is 1.7 in ROUGE-1, 1.3 in ROUGE-2, and 0.83 in ROUGE-L.

## 5.2 Human Evaluation

Model	Preference Score	QA score
BART	-0.286	24.59
PEGASUS	-0.257	26.53
Our Model	0.250	46.94
Gold Summary	<b>0.536</b>	<b>93.88</b>

Table 8: Human evaluation

Since the automatic metric does not fully reveal the true quality of the model, we conduct a human evaluation for further assessment. To achieve that goal, we design two tests in order to elicit human judgements in two ways.

In the first experiment, we presented summaries of PEGASUS (Zhang et al., 2020), BART (Lewis et al., 2019), our model, and the gold summary, then asked four professional English speakers to rate the summaries from worst to best in terms of informativeness, faithfulness, topic coherence, and fluency. We randomly sampled 100 summaries from 100 documents of CNN/DailyMail test set. The score of a system is equal to the percentage of times it was selected as the best minus the percentage of times it was chosen as the worst.

In the second experiment, we applied the question answering (QA) paradigm. For each document, we create two independent questions which emphasizes the key information of the text. Participants would read and answer those questions as best as they could. The score for one system is the percentage of questions the participants answer correctly.

Ten professional English speakers were asked to participate in two assessments. The results in table 9 show that our generated summaries favor human judgements, and are more likely to maintain the important content in the original text than other systems’ summaries.

The Fleiss’ Kappa scores with overall agreement percentages of the first and second human evaluation experiments were denoted in Table 9. As shown in the Table, the measures demonstrate a good inter-agreement among the annotators.

Test	Fleiss’ Kappa	Overall Agreement
Preference	0.61	70.45%
QA	0.77	81.13%

Table 9: Fleiss’ Kappa and Overall Agreement percentage of each human evaluation test. Higher score indicates better agreement.

### 5.3 Flow-based neural topic model with other Transformer-based model

Model	R1	R2	RL
BART	44.16	21.28	40.90
BART + Flow-based NTM + Gating	<b>44.89</b>	<b>21.74</b>	<b>41.48</b>

Table 10: Results when applying flow-based neural topic model and contextualized gating for BART on CNN/DailyMail dataset (Lewis et al., 2019)

Model	R1	R2	RL
BART	45.14	22.27	37.25
BART + Flow-based NTM + Gating	<b>46.86</b>	<b>23.74</b>	<b>38.49</b>

Table 11: Results when applying flow-based neural topic model and contextualized gating for BART on XSum dataset (Lewis et al., 2019)

Model	R1	R2	RL
BART	43.92	16.36	39.16
BART + Flow-based NTM + Gating	<b>47.78</b>	<b>18.28</b>	<b>41.47</b>

Table 12: Results when applying flow-based neural topic model and contextualized gating for BART on arXiv dataset (Lewis et al., 2019)

To study the effect of our topic-oriented module on other abstractive Transformer-based model, we integrate our flow-based neural topic model and contextualized gating into BART (Lewis et al., 2019). In particular, we continue to finetune on CNN/DailyMail, XSum, and arXiv dataset, given the pretrained checkpoint. As can be seen in Table 10, 11, 12, our topic-oriented module is able to improve the performance, showing general effectiveness on other Transformer-based architecture.

### 5.4 Analysis on Neural Topic Model and Traditional Topic Model

To substantiate our hypothesis that neural topic model does enhance the summarization performance in large-scale datasets, we have conducted experiments to combine the Transformer-based summarization module with traditional topic model, i.e. Latent Dirichlet Allocation (LDA) and Poisson Factor Analysis (PFA) on CNN/DailyMail and PubMed. We denoted the results in Table 13 and Table 14. As it can be seen, neural topic models, particularly our proposed model, significantly outperform approaches of traditional topic models on abstractive summarization.

Model	R1	R2	RL
PEGASUS + LDA + Gating	44.17	21.47	41.11
PEGASUS + PFA + Gating	44.18	21.53	41.14
PEGASUS + VAE + Gating	44.33	21.71	41.27
Our Model	<b>44.52</b>	<b>21.95</b>	<b>41.39</b>

Table 13: Results when adapting various topic models on CNN/DailyMail dataset

Model	R1	R2	RL
PEGASUS + LDA + Gating	44.36	19.24	40.53
PEGASUS + PFA + Gating	44.41	19.19	40.55
PEGASUS + VAE + Gating	45.46	19.84	40.89
Our Model	<b>45.99</b>	<b>20.49</b>	<b>41.25</b>

Table 14: Results when adapting various topic models on PubMed dataset

### 5.5 Latent Topic Analysis

Datasets	CNN/DailyMail	XSum
LDA	35.03	26.22
LSA	41.64	27.69
VAE-based NTM	52.60	52.94
Our Model	<b>53.25</b>	<b>53.09</b>

Table 15:  $C_V$  topic coherence score on benchmark datasets. Higher scores mean more coherent topics

It is inherent that latent vector is useful for text summarization, as shown in section 5.1. Here we study whether jointly training with summarization module helps the topic model produce human-interpretable topics.

**Coherence Score Comparison** We decide to evaluate the topic models with the automatic  $C_V$  measure. Following (Zeng et al., 2018; Wang et al., 2020), we pick the top 10 words from each topic and average  $C_V$  scores of all topics. The results are reported on two summarization datasets, CNN/DailyMail and XSum. To conduct the comparisons, we take LDA and LSA as probabilistic baselines, as they are notable and well-known for human interpretability. For both baselines, we execute 1000 iterations to assure convergence. As Table 16 shows, our model outperforms traditional topic models, which implies that jointly training neural topic model and text summarization creates human-understandable topics.

**Sample Topics** To further assess the quality of the topics learned by our system, we continue to extract some sample words (Table 6) indicating the context around “liverpool chelsea” discovered by the model trained on CNN/DailyMail dataset. As can be realized, the topics pertaining to probabilis-



tic topic models such as LSA and LDA contain some mixed topic words. Conversely, our neural topic models trained with the text summarization module produce the topic which looks more coherent. In particular, our words refer to the context which involves the teams competing in the football championship of England, such as “arsenal”, “tottenham”, etc. and related factors, for instance, “balls”, “prize”, “winning”, etc.

LDA	father liverpool son chelsea called group night child west cup
LSA	chelsea beat half winner jose mourinho table place happy lake
VAE-based NTM	liverpool salmon manchester england everton newcastle bale premiersip fabregas clasico
Our Model	liverpool cup leagues chelsea balls night tot- tenham prize winning arsenal

Table 16: Top 10 words for the topic related to “liverpool chelsea”. Red words highlight non-topic words.

## 5.6 Ablation Study

Model	R1	R2	RL
Our Model (with Flow-based NTM and Gating)	<b>49.57</b>	<b>25.08</b>	<b>41.81</b>
- with VAE-based NTM and Gat- ing	48.13	23.91	40.68
- with Flow-based NTM	46.83	23.89	39.78
- with VAE-based NTM	46.30	23.59	39.05

Table 17: Ablation Study on XSum test set

In this section, we proceed to study the impact that (1) The integration of normalizing flow and (2) The contextualized gating mechanism have on the text summarization performance.

### Impact of the contextualized gating mechanism

It can be seen that plainly incorporating the global semantics into the model makes the performance improvement drop strongly. As shown in table 17, the ROUGE-1 score’s decreases more than 2 points compared with models we apply contextualized gating. We hypothesize that in numerous cases, the effect of global semantics overwhelm the benefits of contextualized representations.

**Impact of integrating normalizing flow** In this ablation, we eliminate the normalizing flow from the neural topic modeling. As shown in Table 17, without the normalizing flow, the improvement that the latent vector brings is downgraded, nearly 0.4 of ROUGE-1 for using contextualized gating and 0.53 of ROUGE-1 in non-gating case . We hypothesize that the plain neural topic model does not give

a sufficiently expressive global semantics as the neural topic model using normalizing flow.

## 5.7 Case Studies

Table 1 shows a case study on the summarization results of PEGASUS and our models. While PEGASUS model misses the key information related to the death of “Richie Benauld”, our model successfully include it in the final summarization. It shows the effectiveness of our model in capturing key information in the document, thanks to the contribution of neural topic model and gating mechanism. Remarkably, our model is also able to rephrase “signature one liners” as “trademark word” when describing *Richie Benauld*’s famous quote, not just copying the words in the original document. More case studies can also be found in Appendix.

## 6 Conclusion

In this paper, we propose a method to utilize global semantics for text summarization task. In particular, we aim to fit the global semantics to expressively describe the documents. Moreover, we find that maintaining the information in the original contextualized representations is also beneficial for the summarization performance. We outperform other state-of-the-art models on five benchmark datasets.

## References

- Melissa Ailem, Bowen Zhang, and Fei Sha. 2019. Topic augmented generator for abstractive summarization. *arXiv preprint arXiv:1908.07026*.
- David M Blei, Alp Kucukelbir, and Jon D McAuliffe. 2017. Variational inference: A review for statisticians. *Journal of the American statistical Association*, 112(518):859–877.
- Rishi Bommasani and Claire Cardie. 2020. Intrinsic evaluation of summarization datasets. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8075–8096.
- Jiaao Chen and Diyi Yang. 2020. Multi-view sequence-to-sequence models with conversational structure for abstractive dialogue summarization. *arXiv preprint arXiv:2010.01672*.
- Kyunghyun Cho, Bart Van Merriënboer, Dzmitry Bahdanau, and Yoshua Bengio. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.

- Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A discourse-aware attention model for abstractive summarization of long documents. *arXiv preprint arXiv:1804.05685*.
- Ran Ding, Ramesh Nallapati, and Bing Xiang. 2018. Coherence-aware neural topic modeling. *arXiv preprint arXiv:1809.02687*.
- Li Dong, Nan Yang, Wenhui Wang, Furu Wei, Xiaodong Liu, Yu Wang, Jianfeng Gao, Ming Zhou, and Hsiao-Wuen Hon. 2019. Unified language model pre-training for natural language understanding and generation. *arXiv preprint arXiv:1905.03197*.
- Xiyan Fu, Jun Wang, Jinghan Zhang, Jinmao Wei, and Zhenglu Yang. 2020. Document summarization with vhtm: Variational hierarchical topic-aware mechanism. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 7740–7747.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M Rush. 2018. Bottom-up abstractive summarization. *arXiv preprint arXiv:1808.10792*.
- Karl Moritz Hermann, Tomáš Kočiský, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching machines to read and comprehend. *arXiv preprint arXiv:1506.03340*.
- Matthew D Hoffman, David M Blei, Chong Wang, and John Paisley. 2013. Stochastic variational inference. *Journal of Machine Learning Research*, 14(5).
- Ruipeng Jia, Yanan Cao, Hengzhu Tang, Fang Fang, Cong Cao, and Shi Wang. 2020. Neural extractive summarization with hierarchical attentive heterogeneous graph network. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3622–3631.
- Byeongchang Kim, Hyunwoo Kim, and Gunhee Kim. 2018. Abstractive summarization of reddit posts with multi-level memory networks. *arXiv preprint arXiv:1811.00783*.
- Diederik P Kingma and Max Welling. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Durk P Kingma, Tim Salimans, Rafal Jozefowicz, Xi Chen, Ilya Sutskever, and Max Welling. 2016. Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29:4743–4751.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Xianming Li, Zongxi Li, Yingbin Zhao, Haoran Xie, and Qing Li. 2020. Incorporating effective global information via adaptive gate attention for text classification. *arXiv preprint arXiv:2002.09673*.
- Chin-Yew Lin. 2004. Rouge: A package for automatic evaluation of summaries. In *Text summarization branches out*, pages 74–81.
- Yang Liu. 2019. Fine-tune bert for extractive summarization. *arXiv preprint arXiv:1903.10318*.
- Yang Liu and Mirella Lapata. 2019. Text summarization with pretrained encoders. *arXiv preprint arXiv:1908.08345*.
- Zhengyuan Liu, Angela Ng, Sheldon Lee, Ai Ti Aw, and Nancy F Chen. 2019. Topic-aware pointer-generator networks for summarizing spoken conversations. In *2019 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 814–821. IEEE.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *International Conference on Machine Learning*, pages 2410–2419. PMLR.
- Shashi Narayan, Shay B Cohen, and Mirella Lapata. 2018. Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization. *arXiv preprint arXiv:1808.08745*.
- Shashi Narayan, Joshua Maynez, Jakub Adamek, Daniele Pighin, Blaž Bratanič, and Ryan McDonald. 2020. Stepwise extractive summarization and planning with structured transformers. *arXiv preprint arXiv:2010.02744*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A deep reinforced model for abstractive summarization. *arXiv preprint arXiv:1705.04304*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2019. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv preprint arXiv:1910.10683*.
- Danilo Rezende and Shakir Mohamed. 2015. Variational inference with normalizing flows. In *International Conference on Machine Learning*, pages 1530–1538. PMLR.
- Abigail See, Peter J Liu, and Christopher D Manning. 2017. Get to the point: Summarization with pointer-generator networks. *arXiv preprint arXiv:1704.04368*.
- Kaitao Song, Xu Tan, Tao Qin, Jianfeng Lu, and Tie-Yan Liu. 2019. Mass: Masked sequence to sequence pre-training for language generation. *arXiv preprint arXiv:1905.02450*.

- Sandeep Subramanian, Raymond Li, Jonathan Pila, and Christopher Pal. 2019. On extractive and abstractive neural document summarization with transformer language models. *arXiv preprint arXiv:1909.03186*.
- Ian Tenney, Patrick Xia, Berlin Chen, Alex Wang, Adam Poliak, R Thomas McCoy, Najoung Kim, Benjamin Van Durme, Samuel R Bowman, Dipanjan Das, et al. 2019. What do you learn from context? probing for sentence structure in contextualized word representations. *arXiv preprint arXiv:1905.06316*.
- Luu Anh Tuan, Darsh Shah, and Regina Barzilay. 2020. Capturing greater context for question generation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 9065–9072.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- Li Wang, Junlin Yao, Yunzhe Tao, Li Zhong, Wei Liu, and Qiang Du. 2018. A reinforced topic-aware convolutional sequence-to-sequence model for abstractive text summarization. *arXiv preprint arXiv:1805.03616*.
- Zhengjue Wang, Zhibin Duan, Hao Zhang, Chaojie Wang, Long Tian, Bo Chen, and Mingyuan Zhou. 2020. Friendly topic assistant for transformer based abstractive summarization. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 485–497.
- Wen Xiao and Giuseppe Carenini. 2020. Systematically exploring redundancy reduction in summarizing long documents. *arXiv preprint arXiv:2012.00052*.
- Jichuan Zeng, Jing Li, Yan Song, Cuiyun Gao, Michael R Lyu, and Irwin King. 2018. Topic memory networks for short text classification. *arXiv preprint arXiv:1809.03664*.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. HiBERT: Document level pre-training of hierarchical bidirectional transformers for document summarization. *arXiv preprint arXiv:1905.06566*.
- Chujie Zheng, Kunpeng Zhang, Harry Jiannan Wang, and Ling Fan. 2020. Topic-aware abstractive text summarization. *arXiv preprint arXiv:2010.10323*.
- Ming Zhong, Pengfei Liu, Yiran Chen, Danqing Wang, Xipeng Qiu, and Xuanjing Huang. 2020. Extractive summarization as text matching. *arXiv preprint arXiv:2004.08795*.
- Yanyan Zou, Xingxing Zhang, Wei Lu, Furu Wei, and Ming Zhou. 2020. Pre-training for abstractive document summarization by reinstating source text. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 3646–3660.

## A Summary examples

We present some summary examples in this section

<p><b>DOCUMENT:</b> New York (CNN) New York state authorities have issued a health alert following a dramatic spike in hospital visits for synthetic marijuana-related emergencies. Gov. Andrew Cuomo said Friday that more than 160 patients in nine days have been rushed to hospitals across the state for adverse reactions to synthetic cannabinoid, known as "spice" or "K2." "Spice" and other similar synthetic drugs are often marketed as legal plant material coated with chemicals that are supposed to mimic the effects of marijuana, according to a statement from the governor's office. "Since the exact compounds contained in synthetic cannabinoid products change so frequently, it's often impossible for users to know exactly what they are putting in their body," acting New York State Health Commissioner Dr. Howard Zucker said. Symptoms after use have a wide range of severity, from confusion, drowsiness and headaches to increased heart rate, seizures and loss of consciousness, according to the New York State Department of Health. Synthetic marijuana is popular among teens because it is marketed as incense or natural products to "mask its true purpose", the health department statement said. "Young people may be fooled into thinking that these substances are safe because they are sold over the counter or are in colorful packaging, but they are not made for human consumption," New York Alcohol and Substance Abuse Service s Commissioner Arlene Gonzalez Sanchez said. "They are dangerous and can have significant, long-term effects on the brain." The recent surge is not isolated in New York; other states across the country have noticed similar trends. Alabama Department of Public Health issued a statement last week acknowledging a rise of synthetic marijuana usage and said there had been 98 overdoses suspected to be linked with "spice" in the previous month. Mobile County alone has seen seven cases in 2015 so far, more than the entire previous year, the statement said. Mississippi health officials are also concerned that synthetic marijuana is on the rise. [...]</p>
<p><b>GOLD SUMMARY:</b> New York reports 160 hospitalizations related to synthetic marijuana . Gov. Andrew Cuomo issued a health alert .</p>
<p><b>PEGASUS:</b> Synthetic marijuana is popular among teens because it is marketed as incense or natural products to "mask its true purpose". "Spice" and other similar synthetic drugs are often marketed as legal plant material coated with chemicals that are supposed to mimic the effects of marijuana. Symptoms after use have a wide range of severity, from confusion, drowsiness and headaches to increased heart rate, seizures and loss of consciousness.</p>
<p><b>BART:</b> "Spice" and other similar synthetic drugs are often marketed as legal plant material. Symptoms after use have a wide range of severity, from confusion, drowsiness and headaches to increased heart rate, seizures and loss of consciousness. "Young people may be fooled into thinking that these substances are safe because they are sold over the counter or are in colorful packaging," official says.</p>
<p><b>Our model:</b> New York Gov. Andrew Cuomo says 160 people have been hospitalized in nine days. Symptoms after use include confusion, drowsiness, headaches, increased heart rate, seizures and loss of consciousness. Health officials are concerned that synthetic marijuana is on the rise.</p>

Table 18: A generated summary example .

<p><b>DOCUMENT:</b> (CNN)Panic. Tears. Fear. All those feelings and more permeated cities, villages and camps around Nepal on Saturday, after a massive 7.8 magnitude earthquake struck around midday. Hours later, after a wave of relentless aftershocks, many people still were too scared to go back inside any buildings. Others crowded around rubble, including men and women racing to rescue those trapped. And then there are the hundreds already confirmed dead, not to mention the hundreds more who suffered injuries. Below are some accounts from witnesses in the mountainous Asian nation, in their own words. Fast Facts: Earthquakes . Anderson, an American who was in Nepal for trekking and meditation, was in his hotel room when the quake struck. "I went outside five minutes after the major tremors stopped. I went to a parking lot nearby for one hour or so, then walked down the main road," he said. He took a series of photos on the main road between Thamel and Durbar Squares, that he shared via CNN iReport. Kumar posted a photo of people in his neighborhood sheltering in a makeshift tent after the quake. He sent updates via Twitter about what he was seeing in the Lalitpur District of Kathmandu. "It's getting dark, no power and no water supply in Lalitpur area, but people are helping each other with food and other items . "Almost everyone staying outside home...Hard time for small kids and older people . "People are very worried and are planning to stay out on the street overnight, but they lack sufficient food and water." Joshi is a UNICEF communication officer who was on the ground at the time of the quake. "The shake was like nothing I have experienced in my 57 years. It was strong and it shook for a long time." Old monuments and temples fell, Joshi wrote of his experience. There were fears that other buildings would collapse. "When I went out in the evening, I saw many people preparing to camp out in the main open parade ground in the middle of the street. Relatives were crying in the main government hospital where the dead were being lined up in front of the hospital building. "My family is traumatised. We are 5 generations living under one roof – from a 100 year old grandmother to my 16 month old granddaughter. Strong aftershocks are keeping most of us up!" "Some of the historical sites are completely devastated. "Most of the people – a lot of the people – are walking through the city. They're confused and scared. A lot of people are crying. "They're out with their pets and their families and a lot of locals are volunteering in rescue operations. "In several parts of Kathmandu, a lot of people seem trapped under the rubble. Locals are trying to rescue these people because they can still hear them." Are you in Nepal or have loved ones affected? Please share with us if you are in a safe place. "We are scared and waiting for the tremors to end. We are all sitting outside because there is more news of another quake. "There is no power and families are listening to the FM radio inside their cars. News of multiple building collapses. "I've seen many cracked walls and roads and buildings. "The Dharahara was packed with people a while ago. There are police everywhere trying to move rubble to make space on the roads for ambulances. Everyone is very scared. " "I see many cracked buildings and people are panicked and all running down to the streets. "The main landmark in Kathmandu is a spire, Dharahara, and it has fallen down, it is about 140 feet high in the center city. "Another aftershock is hitting now, it is really strong. "Airplanes are circling now overhead and helicopters are flying and not clear if the airport is open. We hear it is damaged." How are earthquakes measured? "Many historic buildings have collapsed in the city. "In all my years I have never seen such a big earthquake here. [...]</p>
<p><b>GOLD SUMMARY:</b> Massive 7.8 magnitude earthquake has struck Nepal near its capital, Kathmandu . As the death toll rises, witnesses describe devastation and panic .</p>
<p><b>PEGASUS:</b> After hours of relentless aftershocks, many people still were too scared to go back inside any buildings. Below are some accounts from witnesses in the mountainous Asian nation, in their own words. Old monuments and temples fell, and there were fears that other buildings would collapse.</p>
<p><b>BART:</b> Hundreds of people are dead and injured in Nepal on Saturday. Aftershocks and aftershocks have been felt in Kathmandu, Nepal's capital city. "The shake was like nothing I have experienced in my 57 years. It was strong and it shook for a long time"</p>
<p><b>Our model:</b> "The shake was like nothing I have experienced in my 57 years," a UNICEF worker says. The 7.8-magnitude earthquake struck around midday Saturday. Many people were too scared to go back inside any buildings. Many buildings in Kathmandu collapsed, including a 140-foot spire.'</p>

Table 19: Another generated summary example

## B Loss of Flow-based Neural Topic Model

We have the loss of neural topic model, called evidence lower bound

$$\mathcal{L}_{\text{NTM}} = \log p(\mathbf{x}, \mathbf{z}) - \log q(\mathbf{z}|\mathbf{x}) \quad (18)$$

Let  $f_1, f_2, \dots, f_K$  be a chain of invertible transformations which have Jacobian easy-to-compute. We have change of variable formula for transforming  $\mathbf{z}_i$  to  $\mathbf{z}_{i+1}$

$$q(\mathbf{z}_{i+1}) = q(\mathbf{z}_i) \left| \det \frac{\partial f_{i+1}}{\partial \mathbf{z}_i} \right|^{-1} \quad (19)$$

Sequentially applying for  $K$  transformations, we have

$$q(\mathbf{z}_K) = q(\mathbf{z}_0) \prod_{i=1}^K \left| \det \frac{\partial f_i}{\partial \mathbf{z}_{i-1}} \right|^{-1} \quad (20)$$

or equivalently,

$$\log q(\mathbf{z}_K) = \log q(\mathbf{z}_0) - \sum_{i=1}^K \left| \det \frac{\partial f_i}{\partial \mathbf{z}_{i-1}} \right| \quad (21)$$

Plugging the formula (21) to equation (18), we obtain

$$\begin{aligned} \mathcal{L}_{\text{NTM}} &= \log p(\mathbf{x}, \mathbf{z}) - \log q(\mathbf{z}|\mathbf{x}) = \log p(\mathbf{x}, \mathbf{z}_K) - \log q(\mathbf{z}_K|\mathbf{x}) \\ &= -\log q(\mathbf{z}_0) + \sum_{i=1}^K \log \left| \det \frac{\partial f_i}{\partial \mathbf{z}_{i-1}} \right| + \log p(\mathbf{x}|\mathbf{z}_K) + \log p(\mathbf{z}_K) \end{aligned} \quad (22)$$

We reach the neural topic model component in our training objective.

## C Implementation Details

**Flow-based Neural Topic Model** Following (Wang et al., 2020), we preprocess to remove stopwords in the BoW representations. We experiment with different number of topics in 50, 100, 150, 200 and the number of invertible transformations in flow-based neural topic model (flow length) on CNN/DailyMail dataset. The results (in the format of R1/R2/RL scores) are shown in Table 20. It can be seen that a simple posterior distribution in neural topic model is not sufficient to describe the large-scale dataset, while a highly complex one can negatively affect the performance slightly. Similarly, it is necessary to set an adequate number of topics. We proceed to use flow length of 4 and topic number of 100 for other datasets.

Topic Num./Flow length	1	4	16
50	44.19/21.49/41.28	44.48/21.54/41.36	44.23/21.51/41.29
100	44.30/21.78/41.37	<b>44.52/21.95/41.39</b>	44.40/21.87/41.38
150	44.25/21.70/41.34	44.44/21.86/41.27	44.30/21.80/41.21
200	44.24/21.61/41.23	44.35/21.75/41.22	44.24/21.69/41.20

Table 20: Comparisons on the number of topics and flow length on CNN/DailyMail dataset

We pretrain our versions of flow-based neural topic model on five downstream datasets CNN/DailyMail, XSum, Reddit TIFU, arXiv, and PubMed with batch size  $\{256, 256, 256, 320, 320\}$ , respectively. All versions are trained with Adadelta optimizer with a learning rate of 0.01.

**Topic-oriented Transformer-based summarization model.** We do not change any settings from original papers of PEGASUS (Zhang et al., 2020) and BART (Lewis et al., 2019). In particular, we finetune all models on 16 Nvidia GeForce A100 GPUs with batch size 256, Adam optimizer of learning rate  $1e - 4$ . For the objective function in Equation 17, we experimented  $\lambda$  with the choices of  $\{0.5, 0.6, 0.75, 0.9\}$  and found that  $\lambda = 0.75$  gives the best performance for all datasets. Models are evaluated and saved checkpoints every one epoch. During training, we keep track three best validated checkpoints in terms of evaluation loss on the validation set. Eventually, for decoding, we run beam search with beam size of 8 and note the best result out of three validated checkpoints.