

Topic-Guided Abstractive Multi-Document Summarization

Peng Cui and Le Hu

School of Computer Science and Technology
Harbin Institute of Technology, Harbin, China
{pcui, lhu}@insun.hit.edu.cn

Abstract

A critical point of multi-document summarization (MDS) is to learn the relations among various documents. In this paper, we propose a novel abstractive MDS model, in which we represent multiple documents as a heterogeneous graph, taking semantic nodes of different granularities into account, and then apply a graph-to-sequence framework to generate summaries. Moreover, we employ a neural topic model to jointly discover latent topics that can act as cross-document semantic units to bridge different documents and provide global information to guide the summary generation. Since topic extraction can be viewed as a special type of summarization that “summarizes” texts into a more abstract format, i.e., a topic distribution, we adopt a multi-task learning strategy to jointly train the topic and summarization module, allowing the promotion of each other. Experimental results on the Multi-News dataset demonstrate that our model outperforms previous state-of-the-art MDS models on both Rouge metrics and human evaluation, meanwhile learns high-quality topics.

1 Introduction

Multi-document summarization (MDS) is the task to create a fluent and concise summary for a collection of thematically related documents. Compared to single document summarization, it requires the ability to incorporate the perspective from multiple sources and therefore is arguably more challenging (Lin and Ng, 2019). Broadly, existing studies can be classified into two categories: extractive and abstractive. Extractive approaches directly select important sentences from the input documents, which is usually regarded as a sentence labeling (Nallapati et al., 2016; Zhang et al., 2018; Dong et al., 2018) or sentence ranking task (Narayan et al., 2018). By contrast, abstractive models typically use the natural language generation technology to produce a word-by-word summary. In general, extractive

methods are more efficient and can avoid grammatical errors (Cui et al., 2020), while abstractive methods are more flexible and human-like because they can generate absent words (Lin and Ng, 2019).

Recently, with the development of representation learning for NLP (Vaswani et al., 2017; Devlin et al., 2018) and large-scale datasets (Fabbri et al., 2019), some studies have achieved promising results on abstractive MDS (Liu and Lapata, 2019; Jin et al., 2020). Nevertheless, we found there are two limitations that have not been addressed by previous studies. First, some works simply concatenate multiple documents into a flat sequence and then apply single-document summarization approaches (Liu et al., 2018; Fabbri et al., 2019). However, this paradigm fails to consider the hierarchical document structures, which plays a key role in MDS task (Jin et al., 2020). Also, the concatenation operation inevitably produces a lengthy sequence, and encoding long texts for summarization is a challenge (Cohan et al., 2018).

Second, when dealing with multiple documents, a critical point is to learn the cross-document relations. Some studies address this problem by mining the co-occurrence words or entities (Wang et al., 2020a), which can hardly capture implicit associations due to the diverse language expressions. Some other studies (Jin et al., 2020; Liu and Lapata, 2019) first generate low-dimensional vectors in sentence- or paragraph-level and then build interaction based on these highly compressed representations. These methods inevitably result in the loss of large amounts of fine-grained interaction features and would damage the interpretability of models. Therefore, how to learn the relation across documents effectively remains an open question.

To shed lights on these missing points, this paper proposes a novel abstractive MDS model that marries topic modeling into abstractive summary generation. The motivation is that both tasks aim to distil salient information from massive text and

therefore could provide complementary features for each other. Concretely, we jointly optimize a *neural topic model* (NTM) (Miao et al., 2017; Srivastava and Sutton, 2017) that learns topic distribution of source documents and corpus-level topic representations, and an *abstractive summarizer* that incorporates latent topics to summary generation process. In the encoding process, we represent multiple documents as a heterogeneous graph consisting of word, topic, and document nodes and encode it with a graph neural network to capture the interactions among different semantic units. In the decoding process, we devise a topic-aware decoder that leverages learned topics to guide the summary generation. We train the two modules with a multi-task learning framework, where an *inconsistency loss* is applied to penalize the difference between the topic distribution of source documents and that of generated summaries. It encourages the summarizer to generate a summary that is thematically consistent with its source documents and also helps the two modules learn from each other. In this manner, our model is learned such that better topics can yield better summaries and vice versa.

We conduct throughout experiments on the recently released Multi-News dataset (Fabbri et al., 2019). The results demonstrate the effectiveness and superiority of our model. To sum up, the contributions of this paper are threefolds:

- 1) To the best of our knowledge, we carry out the first systematic study on jointly modeling topic inference and abstractive MDS and demonstrate the positive mutual effect between the two tasks.

- 2) We propose a novel MDS model that joint optimizes a neural topic model and an abstractive summarizer. We propose an inconsistency loss to penalize the disagreement between the two modules and help them learn from each other.

- 3) Experimental results on the Multi-News dataset demonstrate that our model achieves the state-of-the-art performance on both Rouge scores and human evaluation, meanwhile learns high-quality topics.

2 Related Work

Multi-document summarization is a challenging subtask of text summarization with a long history. Many previous methods are extractive partly due to the lack of sufficient training data. These methods usually compute sentence salience over graph structures (Mihalcea and Tarau, 2004; Wang et al.,

2020a). Abstractive MDS studies have been fueled by the recent development of large-scale datasets (Fabbri et al., 2019) and representation learning of NLP (Vaswani et al., 2017). Among them, hierarchical networks (Liu and Lapata, 2019) and graph neural networks (Jin et al., 2020) are widely used to capture the cross-document relations. However, most of them build interaction based on word- or paragraph-level representations, which are not flexible or straightforward. In comparison, we propose to model multiple documents more effectively by mining their subtopics.

Datasets for multi-document summarization Recently, Fabbri et al. (2019) released the first large-scale news dataset for MDS. Each article is collected from real-life scenarios and the golden summaries are written by human, which ensures the data quality. Prior to them, some studies tried to construct dataset in automatic manners. For example, Liu et al. (2018) and Liu and Lapata (2019) built datasets based on Wikipedia pages, regarding the first section as the summary and others as different documents. However, modeling the relations among different documents is a different task from modeling that of paragraphs from a same document. Therefore, the generalization ability of models built on such data could be questionable. For this reason, we do not consider such auto-constructed datasets but focus on the Multi-News dataset curated by human.

Topic modeling for text summarization Topic model is widely used for document modeling. Nevertheless, few studies have applied it in summarization task. Previous studies regarded topical distributions as additional features to enrich word or sentence representations (Wei, 2012; Narayan et al., 2018; Wang et al., 2020b). However, these methods use a pipeline process where topic extraction and summary generation are separately performed. In comparison, we adopt a multi-task learning strategy so that the two tasks can learn complementary features from each other. Recently, Cui et al. (2020) has applied NTM to extractive summarization. Though inspired by it, the motivation and proposed method of this study differ from it by a large margin. Cui et al. (2020) use latent topics to preselect salient sentences, while we use them to capture cross-document relations for abstractive MDS. Besides, Cui et al. (2020) solely explores the effect of topic modeling on summarization, while we systematically explore their interplay.

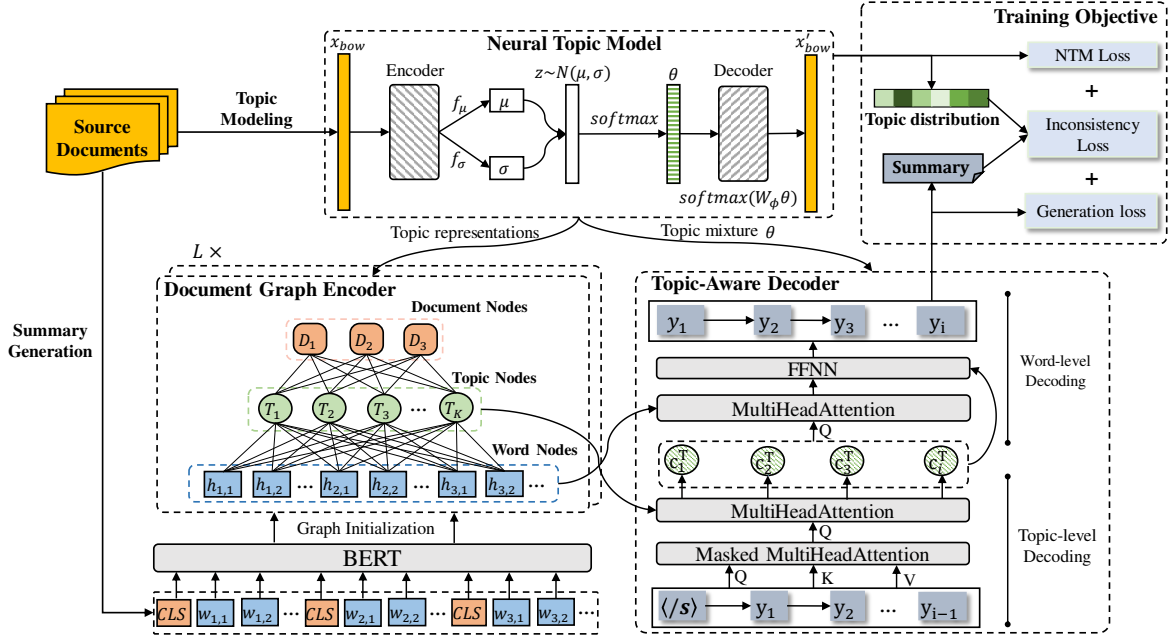


Figure 1: An illustration of the proposed model (*TG-MultiSum*). The summarizer consists of a **document graph encoder** (left bottom) to encode the source documents and a **topic-aware decoder** (bottom right) to generate summary words. A **neural topic model** (top) is applied to provide topical information as guidance signals. The two parts are jointly trained with an *inconsistency loss* to penalize their disagreement.

3 Model

This section describes our model, named as topic-guided multi-document summarization (*TG-MultiSum*). The overall architecture is presented in Figure 1. Given a set of documents $\{D_1, D_2, \dots, D_N\}$, the goal of our model is to generate a word sequence $S = \{y_1, y_2, \dots, y_s\}$ as the summary. Our model consists of three major components: 1) the **neural topic model** aims to learn the topical information of source documents; 2) the **document graph encoder** builds the interaction among different documents and various semantic units. 3) the **topic-aware decoder** generates summary words based on the learned node representations. The entire model is trained in an end-to-end manner. We explain each part below.

3.1 Neural Topic Model

One innovation of this study is that it incorporates topical information into summarization explicitly. Based on the current development of topic modeling, we employ a VAE-based neural topic model proposed in Miao et al. (2017) to discover latent topics. Compared with conventional LDA-style topic models, it can be trained together with neural networks and therefore has better adaptability

(Zeng et al., 2018; Cui et al., 2020).

Similar to LDA, NTM assumes the existence of K underlying topics throughout the corpus. Each document can be represented as a K -dim topic mixture, and each topic can be represented as a distribution over the vocabulary. NTM learns the topics through an encoding–decoding process. Let $x_{bow} \in \mathbb{R}^{|V|}$ denote the bag-of-words term vector of input documents, where $|V|$ is the vocabulary size. We first use an MLP encoder to estimate its exclusive priors σ and μ , which are used to generate the topic distribution through a Gaussian softmax, as shown in the following:

$$\sigma = f_{\sigma}(x_{bow}), \mu = f_{\mu}(x_{bow}), \quad (1)$$

$$z \sim N(\sigma, \mu^2), \theta_x = \text{softmax}(z), \quad (2)$$

where $f_*(\cdot)$ is a neural perceptron with ReLU activation. $z, \theta_x \in \mathbb{R}^K$ are the latent variable and topic distribution of input documents, respectively.

Then, we use a *softmax* layer to reconstruct the input text, i.e., $x'_{bow} = \text{softmax}(W_{\phi}\theta_x)$. In particular, the weight matrix $W_{\phi} \in \mathbb{R}^{|V| \times K}$ can be regarded as the unnormalized topic–word distributions, where $W_{\phi}^{i,j}$ indicates the relevance between the i -th word and j -th topic.

3.2 Document Graph Encoder

Graph Construction Inspired by the assumption of LDA, we view the input documents as a three-layer graph consisting of document, topic, and word units. Formally, let $\mathcal{G} = (V, E)$ denote the constructed graph. The node set V consists of N document nodes $\{v_1^d, v_2^d, \dots, v_N^d\}$, K topic nodes $\{v_1^t, \dots, v_K^t\}$, and M word nodes $\{v_1^w, \dots, v_M^w\}$. The edge set is defined as $E = E_{D,T} \cup E_{T,W}$, where $E_{D,T} = \{e_{1,1}, \dots, e_{N,K}\}$ represents the document–topic edges, and $E_{D,T} = \{e_{1,1}, \dots, e_{K,M}\}$ represents the topic–word edges.

Graph Initialization To capture the contextual information, we use a shared BERT encoder to encode each document independently, which has been proved effective in summarization task. The output states of each word $\mathbf{H}^W = \{h_1^1, \dots, h_{|D_1|}^1, \dots, h_1^N, \dots, h_{|D_N|}^N\}$ are used as the initial word node representations, and those of [CLS] tokens $\mathbf{H}^D = \{h_{CLS}^1, \dots, h_{CLS}^N\}$ are used as the initial document node representations.

As for the topic nodes, we use the weight matrix W_ϕ learned from NTM as raw features and transform it to low-dim topic representations via $\mathbf{H}^T = f_\phi(W_\phi)$, where $f_\phi(\cdot)$ is a Tanh-activated neural perceptron. Each row of $\mathbf{H}^T \in \mathbb{R}^{K \times d}$ is a topic vector with predefined dimension d .

Graph Propagation Given the constructed graph and its initial node representations, we then use a graph neural network to capture the relations among different semantic units. Here we present a single *Document Graph Encoder* (DGE) layer. Multiple DGE layers are stacked in our experiments.

Let u_i^l be the i -th node representation in the l -th layer. The updating process of $u_i^{(i)}$ is denoted as follows:

$$\tilde{u}_i^l = W_1^l \text{Relu}(W_2^l u_i^l + b_1^l) + b_2^l, \quad (3)$$

$$z_{i,j}^l = \text{LeaklyReLU}(W_3^l [f_t(\tilde{u}_i^l); f_t(\tilde{u}_j^l)]), \quad (4)$$

$$\alpha_{i,j}^l = \frac{\exp(z_{i,j}^l)}{\sum_{k \in \mathcal{N}_i} \exp(z_{i,k}^l)}, \quad (5)$$

$$\tilde{u}_i^l = \left\| \sum_{m=1}^M \sum_{j \in \mathcal{N}_i} \tanh(\alpha_{i,j}^{l,m} W_4^m \tilde{u}_j^l) \right\|, \quad (6)$$

where W_* and b_* denote trainable parameters. Eqs.4–6 are the graph attention network (GAT)(Velickovic et al., 2018) which updates each node by aggregating its neighbor nodes \mathcal{N}_* .

Note that the vanilla graph attention network is designed for homogeneous graphs. However, in

our task, word, document, and topic nodes should be considered different semantic units. Therefore, we make a modification in Eq.4 by adding a node-type function $f_t(\cdot)$. It uses exclusive parameters for different node types to project them into a common vector space where the attention score is calculated.

To construct deep networks, we further add a residual connection and layer normalization Ba et al. (2016) to connect adjacent DGE layers.

$$u_i^{l+1} = \text{LayerNorm}(u_i^l + \text{Dropout}(\tilde{u}_i^l)). \quad (7)$$

3.3 Topic-Aware Decoder

To better utilize the guidance effect of latent topics, our decoder, being topic-aware, adopts a two-step decoding process. In each step, it first decodes the current topic, and then generates summary words correspondingly.

Topic-level decoding The topic context c_i^T in i -th step is conditioned on the previous decoded words $\mathcal{Y}_{<i} = \{y_1, y_2, \dots, y_{i-1}\}$ and the topic representations \mathbf{H}^T output from the graph encoder, as shown as follows:

$$u_{i-1} = \text{MHAttn}(e_{i-1}, \mathbf{E}_{<i-1}, \mathbf{E}_{<i-1}), \quad (8)$$

$$c_i^T = \text{MHAttn}(u_i, \mathbf{H}_T, \mathbf{H}_T), \quad (9)$$

where $\text{MHAttn}(Q, K, V)$ denotes the multi-head attention introduced in Vaswani et al. (2017). The first attention layer is used to capture contextual feature of decoded sequence, while the second is to incorporate topical information.

In effect, c_i^T can be viewed as a *topic pointer* that indicates which topics should be discussed in the current step.

Word-level decoding We then use the generated c_i^T to guide the word prediction. Another MHAttn layer is first applied to select relevant parts of source word sequence \mathbf{H}^W with c_i^T as the query, followed by a neural perceptron to inject the current topic focus.

$$v_i = \text{MHAttn}(c_i^T, \mathbf{H}^W, \mathbf{H}^W), \quad (10)$$

$$o_i = \tanh(W_o[v_i; c_i^T] + b_o), \quad (11)$$

where o_i is the final output representation in i -th step of the decoder.

The predicted word distribution over the vocabulary is computed through a softmax layer, i.e., $p_i^g = \text{softmax}(W_g o_{i-1} + b_g)$. To alleviate the out-of-vocabulary (OOV) problem, we employ the

copy mechanism (See et al., 2017) to allow the generator to copy words from source documents. The copy distribution p_i^c is computed as follows.

$$\varepsilon_t = \text{softmax}(\mathbf{H}^W o_t), \quad (12)$$

$$p_i^c = \sum_{i \leq N} \sum_{j \leq |D_i|} \varepsilon_t z_{i,j}, \quad (13)$$

where ε_t is the attention weight of source words, and $z_{i,j}$ is the one-hot indicator vector for word $w_{i,j}$. The final generation distribution is the linear combination of p_i^g and p_i^c , as shown as follows:

$$p_i = \eta_i * p_i^c + (1 - \eta_i) * p_i^g, \quad (14)$$

$$\eta_i = \sigma(W_\eta o_i + b_\eta), \quad (15)$$

where σ indicates sigmoid function and η_i is the copy weight.

3.4 Joint Learning with Inconsistency Loss

Since text summarization and topic modeling both aim to distill salient information from input documents, we jointly train the two modules to help them learn complementary information from each other. The loss function of our model consists of three parts. The summary generation loss \mathcal{L}_{gen} is defined as the negative log-likelihood of ground-truth words, i.e.,

$$\mathcal{L}_{gen} = \sum_{c \in \mathcal{C}} \sum_{w \in y^c} \log p(w). \quad (16)$$

The NTM loss \mathcal{L}_{NTM} is based on the evidence lower bound, i.e.,

$$\mathcal{L}_{NTM} = \text{KL}(p(z) \parallel q(z|x)) - \mathbb{E}_{q(z)}[p(x|z)], \quad (17)$$

where the first term is the KL divergence, and the second term indicates the construction loss. $p(\cdot)$ and $q(\cdot)$ are the encoder and decoder networks described in § 3.1, respectively.

We also devise an inconsistency loss \mathcal{L}_{inc} to penalize the disagreement between the topic distribution of generated summary and that of source documents, as shown as follows:

$$\mathcal{L}_{inc} = \text{KL}(\theta_x \parallel \sum_{1 \leq i \leq L} \theta_{dec}^i), \quad (18)$$

where θ_x is the document topic mixture learned from NTM (Eq.2), and θ_{dec}^i is the topic distribution of i -th decoding step computed in Eq.9.

The final loss is the linear combination of the three parts, i.e., $L = \mathcal{L}_{gen} + \gamma * \mathcal{L}_{NTM} + \tau * \mathcal{L}_{inc}$, where γ and τ are hyperparameters.

# of source	Freq	Prop	# of source	Freq	Prop
2	23,894	53.1%	7	382	0.8%
3	12,707	28.3%	8	209	0.5%
4	5,022	11.2%	9	89	0.2%
5	1,873	4.2%	10	33	0.1%
6	763	1.7%			

Table 1: The distribution of the number of source documents in the Multi-News dataset.

4 Experimental Setup

4.1 Dataset

We conduct experiments on the recently constructed dataset Multi-News (Fabbri et al., 2019). The standard split contains 44972/5622/5622 instances for training, validation, and test. Each instance consists of a set of news articles paired with a human-written summary. The average summary length and article cluster length are 264 and 2103, respectively. In Table 1, we present the distribution of the number of source articles per summary. As shown, nearly half of the summaries are paired with at least three source articles, which highly demands the ability to process multi-source information. The average input length (2103) also brings difficulty for the encoder network. These characteristics make the dataset a good challenge for the MDS task.

4.2 State-of-the-art Baselines

We compare our model with the state-of-the-art extractive and abstractive models. The abstractive baselines are as follows.

PGN (See et al., 2017), pointer-generator network extends the standard seq2seq framework with copy and coverage mechanism.

Hi-MAP (Fabbri et al., 2019) extends PGN into a hierarchical structure and integrates a MMR module to minimize redundancy.

CopyTransformer (Gehrmann et al., 2018) randomly chooses one of the attention heads of Transformer as the copy distribution.

MGSum-abs (Jin et al., 2020) is a state-of-the-art abstractive MDS model. It designs an interaction network to integrate information from different granularities.

We also compare with the following extractive baselines:

HiBERT (Zhang et al., 2019) modifies the standard BERT to a hierarchical structure. We migrate it to MDS by concatenating the input documents.

model	R-1	R-2	R-SU
<i>Non-Neural methods</i>			
Lead-3 [†]	39.41	11.77	14.51
LexRank [†]	38.27	12.70	13.20
TextRank [†]	38.44	13.10	13.50
MMR [†]	38.77	11.98	12.91
<i>Neural-based Extractive Models</i>			
Hi-BERT	43.86	14.62	18.34
MGSUM-ext	44.75	15.75	19.30
HeterGraphSum	46.05	16.35	17.81
MatchSum	46.20	16.51	20.05
<i>Abstractive Models</i>			
PGN	41.85	12.91	16.46
Hi-Map	43.47	14.89	17.41
Copy Transformer	43.57	14.03	17.37
Hi-Transformer	43.85	15.60	18.80
MGSUM-abs	46.00	16.81	20.09
<i>Ours</i>			
TG-MultiSum _{pip}	46.04	16.43	19.82
TG-MultiSum	47.10	17.55	20.73

Table 2: Rouge F1 score of different models. We also report the results ([†]) of several non-neural methods cited from Fabbri et al. (2019).

Hi-Transformer (Liu and Lapata, 2019) adds additional attention heads to the Transformer to share the information across documents.

HeterGrapSum (Wang et al., 2020a) uses a heterogeneous graph neural network to encode word, sentence, and document nodes.

MatchSum (Zhong et al., 2020) regards content selection as a text matching problem. It has reported the state-of-the-art results on Multi-News dataset.

MGSUM-ext (Jin et al., 2020) is the extractive version of MGSUM-abs.

4.3 Implementation Details

We choose “bert-base-uncased” as our pre-trained BERT. For the NTM, we set the topic number $K=50$ and prune the vocabulary to 50,000. For the graph encoder, we set its layer number to 3. The dimension size of nodes representations is set to 768. For the decoder, we set the head of attention number to 6. γ and τ is set to 0.8 and 0.3 to balance different losses. We train our model for up to 1000 epoch with a small batch size of 8. The experiments are based on 2 NVIDIA V100 cards. During the training, an early stop strategy is applied when the loss on validation set does not decrease for three consecutive epochs. We select the hyperparameters with grid search based on the Rouge-2 score on the validation set. In the summary generation, we adopt the beam search strategy with a search size of 5. We report the average results on 3 runs.




Figure 2: Relationship between number of source documents (x-axis) and model performance \hat{R} (y-axis), which is the mean of R-1, R-2, and R-SU.

5 Results and Analysis

5.1 Automatic Evaluation

Overall Performance Table 2 presents the performance (Rouge) of our model against recently released methods on Multi-News. For ours, TG-MultiSum represents our jointly trained model, while TG-MultiSum_{pip} is a pipeline version that separately trains the NTM and the summarizer. We use it to verify whether joint topic inference can bring positive effect on summarization.

As shown, our pipeline version shows competitive results against strong baselines and our full model achieves state-of-the-art performance on the Multi-News dataset, indicating topical information is an effective feature for summarization. Compared to TG-MultiSum_{pip}, our full model achieves 1.06/1.12/0.91 improvements on R-1, R-2, and R-SU. This proves that joint topic inference is effective for abstractive MDS. We also observe that several graph-based models, such as MGSUM and HeterGraphSum, achieve promising results compared to the “flat” models, such as PGN, and Copy-Transformer, implying that graph structure is an effective way to model multiple documents for MDS task. Among the non-neural methods, Lead-3 serves as a simple but effective method. This is because that news articles tend to present key points in the beginning.

Results on varying document numbers We also investigate how the source document number influences the model performance. To this end, we first divide the test set of Multi-News into differ-




Figure 3: Human evaluation results of different abstractive models. Golden lines represent the scores of reference summary

ent intervals based on the number of source documents and discard those with less than 100 examples. Then, we take Hi-Map² as the baseline and compare the results on different parts.

As shown in Figure 2, the Rouge declines with the increasing of document number, indicating that summarizing multiple documents is more challenging. Nevertheless, our two models show better robustness than Hi-Map on increasing document numbers. And joint training can further enhance this ability. Such observation verifies our assumption that latent topics can act as relay nodes to help capture cross-document relations for MDS.

5.2 Human Evaluation

To evaluate the linguistic quality of generated summaries in better granularity, we conduct a human evaluation for the abstractive models based on three aspects: (1) **Fluency** measures whether the summary is coherent and grammatically correct. (2) **Informativeness** focuses on whether the summary covers the salient information of original documents. (3) **Non-redundancy** reflects whether the summary avoids repeated expressions. We sample 100 instances from the test set and generate summaries using different models. Then, we employ five graduates to rate the generated summaries.

As shown in Figure 3, our model beats all baselines in three indicators, especially in informativeness, implying that latent topics are indicative features for capturing salient information. Surprisingly, our model also shows promising improvement in non-redundancy score. This positive effect is probably attributed to the topic context c_*^T (Eq.9) learned in the decoder. It can adaptively decide the current topic focus based on previous decoded words and therefore avoid generating repetitive contents of the same topic.

We also present the human ratings of reference

²We obtain similar results from other abstractive baselines.

Ablated Models	R-1	R-2	R-SU
Our full model	47.10	17.55	20.73
w/o \mathcal{L}_{inc}	46.69	17.04	20.15
w/o topic nodes	46.48	17.11	20.08
w/o topic pointer	46.32	16.97	19.73
w/o DGE	46.08	16.19	19.72
w/o NTM	45.83	16.02	19.45
w/o BERT	45.67	16.13	19.27

Table 3: Performance of different ablated variants against our full model) compared with our full model.

summaries (golden lines). As can be seen, despite the promising improvements of our model, there is a large gap between the quality of model-generated summaries and reference summaries, implying that abstractive MDS remains a challenge.

5.3 Ablation Study

To analyze the relative contributions of different components to the model performance, we compare our full model with the following ablated variants: (1) w/o \mathcal{L}_{inc} removes the inconsistency loss (Eq.18). (2) w/o topic nodes builds the document graph solely with word and document nodes. (3) w/o topic pointer removes the topic pointer (Eq.9) in the decoder. (4) w/o DGE removes the document graph encoder described in § 3.3. (5) w/o NTM removes the NTM module described in § 3.1. For compensation, we use a pre-trained LDA to provide word-topic matrix. (6) w/o BERT removes the BERT encoder and initialize word and document nodes with trainable embeddings.

From Table 3, We can obtain the following observations: (1) The removal of topic nodes and topic pointer both lead to performance drops, indicating that latent topics are effective features for both encoding and decoding process. (2) The document graph encoder plays a necessary role in our model since it can aggregate information from different granularities and documents. (3) NTM serves as a better topic learner than LDA in our experiments, and the inconsistency loss demonstrates its effectiveness. We conjecture that NTM can adaptively learn topics that are suitable for summarization under a multi-task setting where \mathcal{L}_{inc} is applied, while the topics learned by LDA keep unchanged as external features (Cui et al., 2020). (4) The performance declines dramatically when removing BERT. This shows that BERT can provide necessary contextual information to better initialize the graph. Similar results have been observed in GNN-based extractive summarization (Wang et al., 2020a).

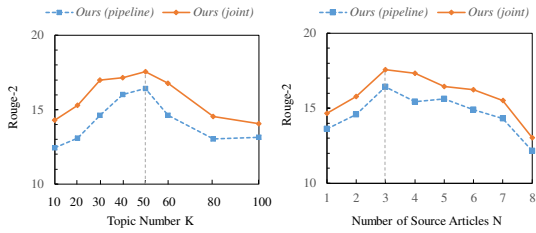


Figure 4: Impact of topic number (left) and graph iteration number (right) on model performance (Rouge-2).

5.4 Impact of Hyperparameters

We further conduct experiments on the validation set of Multi-News to probe into the impact of two important parameters, i.e., the topic number K and the graph iteration number L . The results are presented in Figure 4.

Impact of topic number As can be seen, with a particular range where K is relatively small, more topics produce better performance. However, such increasing trend will reach a saturation when K exceeds a threshold (50 in our experiments). It worth noting that the joint model consistently beat the pipeline model under all different K , implying that joint training can help the NTM adaptively adjust learned topics for better summarization.

Impact of graph iteration number Figure 4 (right) presents the relationship between graph iteration number with model performance. We can see that the two curves show a similar trend. In particular, the performance is dramatically boosted when L goes from 0 to 3. However, such increasing trend is not always monotonous, and a larger L will damage the performance. A possible reason lies in that deep networks could lead to overfit, although we add a residual connection between adjacent layers.

5.5 Topic Quality Analysis

We have shown the effect of latent topics on MDS task. In this subsection, we conduct experiments whether summary generation can in turn help in producing better topics.

We refer NTM_{sum} as our jointly trained topic model and consider three baselines for comparisons. (1) **LDA** (Blei et al., 2003) is a widely used topic model based on Bayesian graphical models. (2) **BTM** (Yan et al., 2013) is an enhanced topic model for short text modeling. (3) **GSM** (Miao et al., 2017) is the model used in our method. Different with NTM_{sum} , it is separately trained on VAE loss. We use it to show the effect of joint summary generation on topic modeling.

Models	C_v	Sample Topics
LDA	0.442	sport NBA <u>green</u> champion watch deal guard brand <u>speak</u> commercial
BTM	0.431	balls sport fight basketball <u>year</u> <u>violence</u> foul superbowl fail <u>crazy</u>
GSM	0.370	player sport <u>eye</u> football national <u>word</u> halftime <u>answer</u> playing <u>day</u>
NTM_{sum}	0.496	sport quarterback scores NBA show play reporter winner Olympic medal

Table 4: Coherence score C_v and inferred topic (*sport*) of different topic models. Off-topic words are underlined and in red.

The three comparison models are all trained on the Multi-News dataset. We run 1,000 Gibbs sampling for LDA and BTM to ensure the convergence. For GSM, we use the same settings described in § 4.3 to make the results comparable.

Topic Coherence Following previous studies (Zeng et al., 2018; Wang et al., 2019), we use the coherence score C_v (Röder et al., 2015) to quantitatively evaluate inferred topics, which has been proved highly consistent with human evaluation. We can see from Table 4 that the separately trained GSM performs rather poorly compared with two traditional models. However, the performance is significantly improved when it is jointly trained with the summarizer. This proves that a joint summarization task can in turn help in topic modeling because a summary usually reflects the major topics of its source document(s).

Sample Topics To obtain a more intuitive comparison of the topic quality learned by different models, we present top 10 representative words of the topic "*sport*" inferred by different models. As can be seen from Table 4, there are mixed off-topic words in three baselines. Besides, compared with them, our inferred topic looks more coherent. For example, it includes less half-related words, such as "*commercial*" (LDA), "*fail*" (BTM), and "*national*" (GSM).

6 Conclusion and Future Work

This study proposes a novel abstractive MDS model that integrates a joint NTM to discover latent topics. Experimental results demonstrate that our model not only achieves the-state-of-the-art results on summarization but also produce high-quality topics. Further discussions show that topic inference and summary generation can promote each other. In the future, we will explore how to apply latent topics in controllable summarization.

References

- Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. 2016. Layer normalization. *arXiv: Machine Learning*.
- David M. Blei, Andrew Y. Ng, and Michael I. Jordan. 2003. Latent dirichlet allocation. *Journal of Machine Learning Research*, 3:993–1022.
- 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. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, volume 2, pages 615–621.
- Peng Cui, Le Hu, and Yuanchao Liu. 2020. Enhancing extractive text summarization with topic-aware graph neural networks. In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5360–5371.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina N. Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186.
- Yue Dong, Yikang Shen, Eric Crawford, Herke van Hoof, and Jackie Chi Kit Cheung. 2018. Bandit-sum: Extractive summarization as a contextual bandit. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3739–3748.
- Alexander R. Fabbri, Irene Li, Tianwei She, Suyi Li, and Dragomir R. Radev. 2019. Multi-news: A large-scale multi-document summarization dataset and abstractive hierarchical model. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1074–1084.
- Sebastian Gehrmann, Yuntian Deng, and Alexander M. Rush. 2018. Bottom-up abstractive summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4098–4109.
- Hanqi Jin, Tianming Wang, and Xiaojun Wan. 2020. Multi-granularity interaction network for extractive and abstractive multi-document summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6244–6254.
- Hui Lin and Vincent Ng. 2019. Abstractive summarization: A survey of the state of the art. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9815–9822.
- Peter J. Liu, Mohammad Ahmad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating wikipedia by summarizing long sequences. In *International Conference on Learning Representations*.
- Yang Liu and Mirella Lapata. 2019. Hierarchical transformers for multi-document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5070–5081.
- Yishu Miao, Edward Grefenstette, and Phil Blunsom. 2017. Discovering discrete latent topics with neural variational inference. In *the 34th International Conference on Machine Learning (ICML)*, pages 2410–2419.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing order into text. In *Proc. 2004 Conference on Empirical Methods in Natural Language Processing, Barcelona, Spain, July*, pages 404–411.
- Ramesh Nallapati, Feifei Zhai, and Bowen Zhou. 2016. Summarunner: A recurrent neural network based sequence model for extractive summarization of documents. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 31, pages 3075–3081.
- 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. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1797–1807.
- Michael Röder, Andreas Both, and Alexander Hinneburg. 2015. Exploring the space of topic coherence measures. In *Proceedings of the Eighth ACM International Conference on Web Search and Data Mining*, pages 399–408.
- Abigail See, Peter J. Liu, and Christopher D. Manning. 2017. Get to the point: Summarization with pointer-generator networks. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, volume 1, pages 1073–1083.
- Akash Srivastava and Charles Sutton. 2017. Autoencoding variational inference for topic models. In *5th International Conference on Learning Representations*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, volume 30, pages 5998–6008.
- Petar Velickovic, Guillem Cucurull, Arantxa Casanova, Adriana Romero, Pietro Liò, and Yoshua Bengio. 2018. Graph attention networks. In *International Conference on Learning Representations*.

- Danqing Wang, Pengfei Liu, Yining Zheng, Xipeng Qiu, and Xuanjing Huang. 2020a. Heterogeneous graph neural networks for extractive document summarization. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 6209–6219.
- Yue Wang, Jing Li, Hou Pong Chan, Irwin King, Michael R. Lyu, and Shuming Shi. 2019. Topic-aware neural keyphrase generation for social media language. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2516–2526.
- Zhengjue Wang, Zhibin Duan, Hao Zhang, Chaojie Wang, Long Tian, Bo Chen, and Mingyuan Zhou. 2020b. 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.
- Yang Wei. 2012. Document summarization method based on heterogeneous graph. In *2012 9th International Conference on Fuzzy Systems and Knowledge Discovery*, pages 1285–1289.
- Xiaohui Yan, Jiafeng Guo, Yanyan Lan, and Xueqi Cheng. 2013. A biterm topic model for short texts. In *Proceedings of the 22nd international conference on World Wide Web*, pages 1445–1456.
- Jichuan Zeng, Jing Li, Yan Song, Cuiyun Gao, Michael R. Lyu, and Irwin King. 2018. Topic memory networks for short text classification. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 3120–3131.
- Xingxing Zhang, Mirella Lapata, Furu Wei, and Ming Zhou. 2018. Neural latent extractive document summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 779–784.
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. HiberT: Document level pre-training of hierarchical bidirectional transformers for document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5059–5069.
- M. Zhong, P. Liu, Y. Chen, D. Wang, X. Qiu, and X. Huang. 2020. Extractive summarization as text matching.