

End-to-End Segmentation-based News Summarization

Yang Liu, Chenguang Zhu, Michael Zeng

Microsoft Cognitive Services Research Group
{yaliu10, chezhu, nzeng}@microsoft.com

Abstract

In this paper, we bring a new way of digesting news content by introducing the task of segmenting a news article into multiple sections and generating the corresponding summary to each section. We make two contributions towards this new task. First, we create and make available a dataset, SEGNEWS, consisting of 27k news articles with sections and aligned heading-style section summaries. Second, we propose a novel segmentation-based language generation model adapted from pre-trained language models that can jointly segment a document and produce the summary for each section. Experimental results on SEGNEWS demonstrate that our model can outperform several state-of-the-art sequence-to-sequence generation models for this new task.

Introduction

In recent years, automatic summarization has received extensive attention in the natural language processing community, due to its potential for processing redundant information. The evolution of neural network models and availability of large-scale datasets have driven the rapid development of summarization systems.

Despite promising results, there are specific characteristics of the traditional summarization task that impedes it to provide more beneficial ways of digesting long news articles. For instance, current news summarization system only provides one genetic summary of the whole article, and when users want to read in more details, the generated summary is not capable of helping navigate the reading. For example, given a news report, current system will output several highlight summaries (Nallapati, Zhai, and Zhou 2017; Liu and Lapata 2019; Zhang et al. 2020). Under this circumstance, if a user expect to read more details about one highlight, he will still need to browse the whole article to locate related paragraphs. Meanwhile, when processing a long news article, current systems usually truncate the text and only generate a summary based on the partial article (Cheng and Lapata 2016a; Zhang et al. 2020). Although this is reasonable since most important content usually lies in the initial portion, it also makes it difficult for users to quickly access information beyond the truncated portion.

In this paper, we propose a new task of Segmentation-based News Summarization. Given a news article, we aim to identify its potential sections and at the same time, to gen-

erate the corresponding summary for each section. This new task provides a novel alternative to summarizing a news article. We argue that it can lead to a more organized way of understanding long articles and facilitates a more effective style of reading documents.

First, segmenting a news article can provide a structural organisation of the content, which is not only helpful to reading but also benefit many important NLP tasks. For example, Brown et al. (1983) states that this kind of multi-paragraph division is one of the most fundamental tasks in discourse. However, many expository texts, like news articles, instruction manuals, or textbooks consist of long sequences of paragraphs with very little structural demarcation (Hearst 1994), and for these documents a subtopical segmentation can be useful. Second, generating concise text descriptions of each sections further reduces the cognitive burden of reading the article (Florax and Ploetzner 2010). Previous studies (Paice 1990; Hearst 1997) present that subtopic segments with their headings is an effective alternative to traditional summarization tasks.

In this paper, we make two main contributions towards the development of Segmentation-based News Summarization systems.

First, we create and publicize a large-scale benchmark¹, SEGNEWS, for Segmentation-based News Summarization task. Figure 1 shows one example article and its aligned segmentation and summaries from SEGNEWS.

Second, we propose a novel end-to-end approach for this task, which can jointly segment an article while generating the corresponding summaries. These two sub-tasks can learn from each other via a shared encoder. The model is equipped with a segmentation-aware attention mechanism, allowing it to capture segmentation information during summary generation. One important advantage of our framework is that it is a non-invasive adaptation of the Transformer (Vaswani et al. 2017) model, i.e. it does not alter the inner structure of Transformers. And our framework can integrate many pre-trained language generation models, including BART (Lewis et al. 2020), GPT (Radford et al. 2019) and UNILM (Bao et al. 2020). This enables our framework to enjoy a high degree of flexibility and better performance.

We compare the proposed framework with several state-

¹We will share the dataset link after paper acceptance.

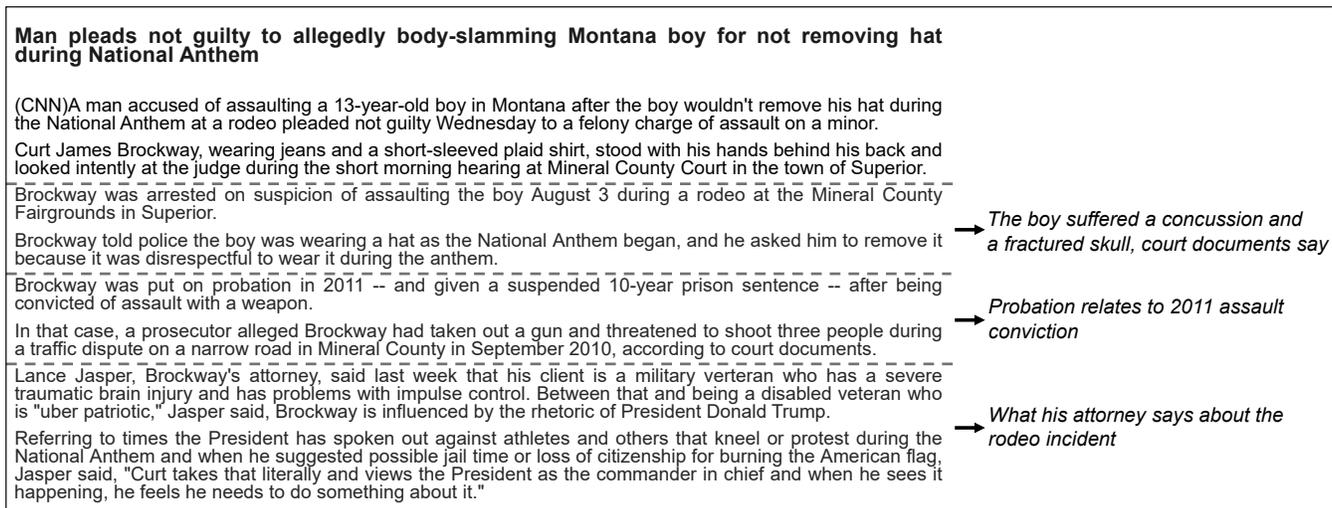


Figure 1: One example from the segmentation-based summarization task SEGNEWS. The news article is taken from a CNN news article and we truncate the article for display. CNN editors have divided this article into several sections and written a heading to section. The goal of this task is to automatically identify sub-topic segments of multiple paragraphs, and generate the heading-style summary for each segment. Dotted lines in the figure indicate segment boundaries. In this article, paragraphs 1,2 are annotated as the first segment, paragraphs 3,4 are annotated as the second segment, paragraphs 5,6 are annotated as the third segment, and paragraphs 7,8 are annotated as the fourth segment. To the right of the article are the heading-style summaries for segments. Since the first segment is usually an overview of the news, we do not assign a summary to it.

of-the-art methods on the SEGNEWS benchmark. Both automatic evaluation and human evaluation demonstrate the superiority of our model.

Related Work

Document Summarization

Document summarization is the task of automatically generating a shorter version text of one or multiple documents while retaining its most important information (Radev, Hovy, and McKeown 2002). The task has received much attention in the natural language processing community due to its potential for various information access applications. Most large-scale summarization datasets are built on news articles. Popular single-document summarization benchmarks include CNN/DM (Nallapati et al. 2016; Cheng and Lapata 2016a), NYT (Durrett, Berg-Kirkpatrick, and Klein 2016) and XSum (Narayan, Cohen, and Lapata 2018).

Document summarization can be classified into different paradigms by different factors (Nenkova and McKeown 2011). And among them, two have consistently attracted attention. *extractive* approaches form summaries by copying and concatenating the most important spans in a document; while in *abstractive* summarization, various text rewriting operations generate summaries using words or phrases that are not in the original text.

Recent approaches to extractive summarization frame the task as a sequence labeling problem by taking advantage of the success of neural network architectures (Bahdanau, Cho, and Bengio 2015). The idea is to predict a label for each sentence specifying whether it should be included in the sum-

mary. Existing systems mostly rely on recurrent neural networks (Hochreiter and Schmidhuber 1997) or Transformer model (Vaswani et al. 2017) to encode the document and obtain a vector representation for each sentence (Nallapati, Zhai, and Zhou 2017; Cheng and Lapata 2016b; Liu, Titov, and Lapata 2019).

In recent years, neural sequence-to-sequence approaches dominate abstractive summarization methods. Rush, Chopra, and Weston (2015) and Nallapati et al. (2016) are among the first to apply the neural encoder-decoder architecture to text summarization. See, Liu, and Manning (2017) enhance this model with a pointer-generator network and a coverage mechanism. Pretrained language models have recently emerged as a key technology for improving abstractive summarization systems. These models first pretrain a language model with self-supervised objectives on large corpora and then fine-tune it on summarization datasets. Liu and Lapata (2019) combine a pretrained encoder based on BERT (Devlin et al.) with a randomly initialized decoder, demonstrating substantial gains on summarization performance. MASS (Song et al. 2019) is an encoder-decoder neural model pretrained with the objective of reconstructing a masked text and can be fine-tuned on summarization tasks. BART (Lewis et al. 2020) is an encoder-decoder Transformer (Vaswani et al. 2017) pretrained by reconstructing a text corrupted with several arbitrary noising functions. Bao et al. (2020) design UNILMv2, a Transformer-based neural network pretrained as a pseudo-masked language model.

# news articles	26,876
# paragraphs	40.31
# sections per article	3.17
# tokens per article	1362.24
# tokens per section summary	4.70

Table 1: Data statistics of the SEGNEWS dataset.

Text Segmentation and Outline Generation

Text segmentation has been widely used in the fields of natural language processing and information extraction. Existing methods for text segmentation fall into two categories: unsupervised and supervised. TextTiling (Hearst 1997) is one of the first unsupervised topic segmentation algorithms. It segments texts in linear time by calculating the similarity between two blocks of words based on the cosine similarity. Choi (2000) introduce a statistical model which can calculate the maximum-probability segmentation of a given text. The TopicTiling (Riedl and Biemann 2012) algorithm is based on TextTiling, which uses the Latent Dirichlet Allocation to find topical changes within documents. LCSeg (Galley et al. 2003) computes lexical chains of documents and segments texts by a score which captures the sharpness of the change in lexical cohesion.

Supervised methods have also been proposed for text segmentation. Hsueh, Moore, and Renals (2006) integrate lexical and conversation-based features for topic and sub-topic segmentation. Hernault, Bollegala, and Ishizuka (2010) use CRF to train a discourse segmenter with a set of lexical and syntactic features. Li, Sun, and Joty (2018) propose SEGBOT which uses a neural network model with a bidirectional recurrent neural network together with a pointer network to select text boundaries in the input sequence.

Recently, Zhang et al. (2019) propose Outline Generation task, aiming to identify potential sections of a multi-paragraph document and generate the corresponding section headings as outlines. This task is in form similar to segmentation-based summarization. However, there are two main differences. First, outline generation focused on academic or encyclopaedic documents, where the section headings are extremely short (on average less than two words) and cannot be considered as a summarization task. Second, since outlines care more about briefly describing their corresponding sections, headings in outlines are independently from each other. In segmentation-based summarization task, despite describing the sections, heading-style summaries also devote to navigating the reading, and they are usually related and coherent in content.

The SEGNEWS Benchmark

In order to study and evaluate the Segmentation-based News Summarization task, we build a new benchmark dataset SEGNEWS. We take CNN website as our article source. As illustrated in Figure 1, there are a large part of CNN news articles which are divided by human editors into several sub-topic sections (see Appendix for details). And each section is assigned a heading-style summary also written by

these editors. We collect news articles published from 2017 to 2021, covering multiple CNN news channels, including US Politics, World, Business, Health, Entertainment, Travel and Sports. We filter articles with no sub-topic structures or editor written heading-style summaries. Since the first segment is usually an overview of the news, editors do not assign a summary to it. The resulting dataset contains 26,876 news articles. For each article, it has human annotated segmentation structures and each segment has a human-written heading-style summary.

Table 1 shows the overall statistics of our SEGNEWS benchmark dataset. We can see that the news articles in SEGNEWS contain rich structural information and are much longer (1,362 tokens per article) than traditional news summarization datasets: articles in CNN/DM (Cheng and Lapata 2016b) dataset has an average length of 686.63 tokens and articles in NYT (Sandhaus 2008) dataset has an average length of 800.04 tokens. This is in line with our motivation that segmentation-based summarization can help readers better understand longer articles.

Task Formulation

Given a multi-paragraph article, the segmentation-based summarization task aims to: i) identify sections of the article to unveil its inherent sub-topic structure, where each section consists of neighboring paragraphs with a coherent topic, and ii) generate the heading-style summary for each section to concisely summarize the section. Particularly, in one article, summaries of different sections should be coherent in content and consistent in style.

Formally, let d indicate a document consisting of paragraphs $[p_1, p_2, \dots, p_M]$. The segmentation-based summarization task aims to recognize a sequence of section boundaries $[b_1, b_2, \dots, b_{N-1}]$. These boundaries divide the document into N sections $s_1 = [p_1, \dots, p_{b_1}]$, $s_2 = [p_{b_1+1}, \dots, p_{b_2}]$, \dots , $s_N = [p_{b_{N-1}+1}, \dots, p_M]$. Meanwhile, summarization systems will generate the corresponding section summaries $[y_1, y_2, \dots, y_N]$.

Systems for Segmentation-based News Summarization

In this section, we present two different frameworks to tackle the segmentation-based summarization task. In the *Pipeline* approach, we first apply a segmentation model to identify the potential sections, and then apply a generation model to produce the headings. In the *Joint* approach, one neural model is able to jointly segment an article and produce the summaries. To achieve this, we design a novel segmentation-aware attention mechanism, which allows the model to capture segmentation information when generating summaries. This new attention mechanism can also be considered as a non-invasive adaption for conventional Transformer models. Thus, to take the most advantage of existing pre-trained language models, we propose SEGUNILM and SEGBART which are respectively based on pre-trained UNILM model and BART model. They can be initialized completely from pre-trained models and achieve substantial improvement on segmentation-based summarization task.

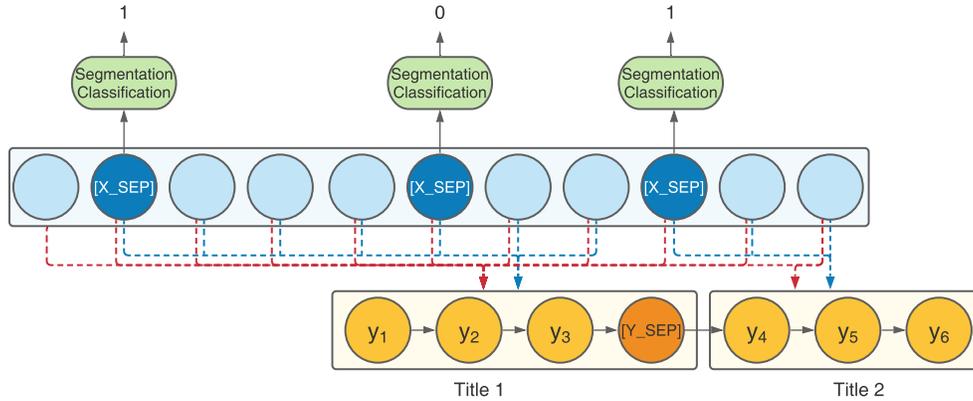


Figure 2: The overall framework of SEGTRANS model. The blue circles indicate input source text, where dark blue circles indicate special symbols representing paragraph boundaries. The yellow circles indicate output target text, where orange circles indicate special symbols representing title boundaries. Dotted blue lines indicate attention heads with segmentation-aware attention mechanism and red dotted lines indicate attention heads with original full attention mechanism.

Pipeline Approach

Segmentation model We formulate the section identification process as a sequence labeling task. We insert a special symbol $[X_SEP]$ at the boundary of paragraph p_i and p_{i+1} , and then concatenate all paragraphs into a single text input. A neural encoder is then applied to encode this input. Define u_i as the output vector of $[X_SEP]$ after paragraph p_i . We then apply a binary classification layer over u_i to obtain $y_i \in \{0, 1\}$. $y_i = 0$ indicates paragraph p_i and p_{i+1} are in one segmentation, and $y_i = 1$ indicates p_{i+1} should be the start of a new segment.

Generation model We then generate an aligned heading-style summary for each identified section s_j . The generation of each heading is independent. Here, we can choose existing extractive or abstractive summarization methods.

- TOPICRANK (Bougouin, Boudin, and Daille 2013) is an extractive method for keyphrase extraction which represents a document as a complete graph depending on topical representations. We use the top ranked phrase as the summary for input section;
- SEQ2SEQ represents the sequence-to-sequence neural model, which is usually used in abstractive summarization. It first encodes the concatenated text of all paragraphs within this section, and then decodes the heading in an auto-regressive manner. In experiments, we try both non-pretrained Transformer model and pretrained UNILM and BART models as SEQ2SEQ models.

Joint Approach

Instead of relying on a pipeline framework, we can also tackle the segmentation-based summarization task with a single encoder-decoder neural model. This brings two main advantages. First, the encoders for segmentation and generation can be shared, benefiting both tasks as a multi-task

learner. Second, we can decode all summaries in an auto-regressive manner. In this way, when the decoder generates the l -th heading, it will be exposed to the 1st to $(l - 1)$ -th generated headings. This is considerably helpful since in a news article, many headings are highly related and coherent in their content.

We use Transformer (Vaswani et al. 2017) as base model for the encoder and decoder. Formally, the encoder maps a sequence of tokens in the source document $\mathbf{x} = [x_1, \dots, x_n]$ into a sequence of continuous representations $\mathbf{t} = [t_1, \dots, t_n]$. Then a segmentation classification layer is applied over output vectors of paragraph boundaries to identify correct segments $B = [b_1, b_2, \dots, b_{N-1}]$ for the input article. The decoder then generates the tokens of the target text $\mathbf{y} = (y_1, \dots, y_m)$ auto-regressively based on the conditional probability: $p(y_1, \dots, y_m | x_1, \dots, x_n, B)$. As the decoder produces summaries for all sections in one pass, we add a special symbol $[Y_SEP]$ between summaries from neighboring sections to indicate their boundaries.

However, in this vanilla sequence-to-sequence model, during inference, the decoder is not aware of the segmentation results and can only implicitly use this information when decoding the summaries. Thus, to better jointly learn segmentation and generation tasks, we propose SEGTRANS model, which is equipped with Segmentation-aware Attention mechanism.

Segmentation-aware attention The multi-head decoder-to-encoder attention in a Transformer decoder defines that for a head $z \in \{1, \dots, n_{head}\}$ at each layer, the model calculates attention probabilities a_{ij}^z against each source token x_j when generating the i -th token y_i .

$$q_i^z = W_q^z Y_i, \quad (1)$$

$$k_j^z = W_k^z X_j, \quad (2)$$

$$a_{ij}^z = \frac{\exp(q_i^{zT} k_j^z)}{\sum_{o=1}^n \exp(q_i^{zT} k_o^z)}, \quad (3)$$

<i>Vanilla Seq2Seq</i>		R1		R2		RL	
TRANS		8.66		1.51		8.16	
UNILM		19.22		7.18		16.99	
<i>Pipeline</i>		With Gold Segments			With Predicted Segments		
Segmentor	Generator	R1	R2	RL	R1	R2	RL
Transformer	Transformer	8.69	1.83	9.09	–	–	–
Transformer	TopicRank	5.09	1.14	6.28	–	–	–
BART	BART	21.42	7.76	19.28	16.01	5.27	14.37
UNILM	UNILM	21.76	8.22	19.75	16.27	5.45	14.65
<i>Joint</i>		R1	R2	RL	R1	R2	RL
SEGTRANS		8.94	1.85	9.35	–	–	–
SEGBART		21.49	8.29	19.52	16.36	5.14	14.96
SEGUNILM		22.17	8.86	20.17	17.59	6.20	15.90

Table 2: ROUGE F1 results on SEGNEWS test set. R1 and R2 are shorthands for ROUGE scores of unigram and bigram overlap; RL is the ROUGE score of longest common subsequence. In pipeline approach, we try combinations of different segmentators and generators. Due to their failure on segmentation, non-pretrained models have very low ROUGE scores with predicted segments, and we do not compare them in the table.

where $Y_i, X_j \in \mathbb{R}^d$ are the layer’s input vectors corresponding to the token y_i and x_j , respectively. $W_q^z, W_k^z \in \mathbb{R}^{d_{head} * d}$ are learnable weights. n is the number of tokens in source input.

However, in segmentation-based summarization, when generating the heading for the i -th section, the decoder should focus more on the input tokens belonging to that section. Thus, we propose the segmentation-aware attention as follows.

We select a subset \hat{z} of decoder heads to apply a segmentation mask to enforce that these heads only attend to the corresponding section. For a head in \hat{z} , Eq. 3 is modified to:

$$a_{ij}^z = \frac{\exp(q_i^{zT} k_j^z) \text{seg}(y_i, x_j)}{\sum_{o=1}^n \exp(q_i^{zT} k_o^z) \text{seg}(y_i, x_j)} \quad (4)$$

where $\text{seg}(y_i, x_j)$ is a indicator function. It equals 1 if and only if y_i and x_j both belong to the same section, and 0 otherwise. In this manner, parts of the heads in multi-head attention are able to dynamically capture segmentation information, while the other heads still model global features of the entire input article.

We illustrate a detailed example of our framework with segmentation-aware attention in Figure 2. We first encode the source text, and apply a segmentation classification layer over output vectors of paragraph boundaries. For this example input, the model classifies the first and the third paragraph boundaries to be segmentation points. Then the decoder will apply a segmentation-aware multi-head attention over the source outputs. It generates the summary for the first identified section with parts of the attention heads over only the first and the second paragraphs. After generating the first heading ending symbol [Y_SEP], the decoder changes the segmentation-aware attention to the third paragraph for generating the summary for the second section.

The final loss \mathcal{L} for training SEGTRANS is the summation of the segmentation loss (binary classification loss) \mathcal{L}_{seg} and

generation loss (negative likelihood loss) \mathcal{L}_{gen} :

$$\mathcal{L} = \mathcal{L}_{\text{seg}} + \mathcal{L}_{\text{gen}} \quad (5)$$

One advantage of our framework is that it is a non-invasive adaptation of the Transformer model, i.e. it does not alter the inner structure of Transformers. This is important since this adaptation can be applied to many popular pretrained language generation models (e.g. MASS, BART, GPT and UNILM), offering our framework a high degree of flexibility and better performance. In this paper, we also augment pre-trained UNILM model and BART model with this mechanism and propose SEGUNILM and SEGBART to further boost their performance.

Experiments

In this section, we conduct experiments on SEGNEWS dataset by comparing our proposed model with several strong baselines.

Experimental Settings

In pre-processing, all the words in news articles and headings are transformed to lower case and tokenized with word-piece tokenizer from BERT (Devlin et al.). In data splitting, we guarantee the headings of articles in the test set have low bigram overlap with articles in the training set. We obtain a splitting of 21,748 articles in training set, 2,688 in validation set and 2,444 in test set.

We experiment under both non-pretrained and pretrained settings. In non-pretrained setting, we use a 6-layer Transformer encoder-decoder model (SEGTRANS) with 512 hidden size and 2,048 feed-forward filter size. In pretrained setting, we propose SEGUNILM and SEGBART which adopts the base version of UNILMv2 (Bao et al. 2020) and the large version of BART (Lewis et al. 2020) as the pretrained model. UNILMv2 is a Transformer-based neural network with 12 Transformer layers and 12 attention heads. It is pretrained as a pseudo-masked language model on a large corpus. BART

is a Transformer-based neural encode-decoder model with 12 layers and 16 attention heads, pretrained via a denoising auto-encoder loss. Label smoothing is used with smoothing factor 0.1. For segmentation-aware attention, we choose the best c (number of segmentation-aware heads) by experiments on the validation set, and $c = 9$ for SEGUNILM and $c = 13$ for SEGBART provide the best performance.

During all decoding we use beam search (size 5), and tune α for the length penalty (Wu et al. 2016) between 0.6 and 1 on the validation set. To guarantee the number of generated headings can match the number of predicted source segments, we take a trick of only generating the end-of-generation token (EOS) when these two numbers match.

We compare the proposed joint models with two sets of strong baselines. The first set of baselines are vanilla sequence-to-sequence models. These models take raw article as input and output the concatenated headings. The second set are pipeline models. As described, these systems first use a segmentor to divide the article into several sections, and then apply a generator to produce summary for each section.

In segmentation-based summarization, summarization systems require segmentation results. We set two settings of segmentation. For the first setting, we provide golden segments to the models to evaluate their performance of generating the summaries when given the correct segments. For the second setting, we require the models to first segment the article and then generate summaries for the predicted segments.

Evaluation Metrics

Evaluation metrics for summarization performance are ROUGE (Lin 2004) F1 scores of the generated headings against the gold headings. We report unigram and bigram overlap (ROUGE-1 and ROUGE-2) as a means of assessing informativeness and the longest common subsequence (ROUGE-L) as a means of assessing fluency.

We use standard metrics Pk (Beeferman, Berger, and Lafferty 1999) and $WinDiff$ (Pevzner and Hearst 2002) to evaluate segmentation results. Lower scores of these two metrics indicate that the predicted segmentation is closer to the ground truth. A EVEN baseline is included for comparison where it segments the whole article evenly.

Results

Table 2 describes our summarization results on the SEGNEWS dataset. The first vertical block includes the results of vanilla sequence-to-sequence models. TRANS is the non-pretrained Transformer encoder-decoder model. UNILM and BART are two pretrained baseline models. The second vertical block contains the results of pipeline models. We present the combinations of different segmentation models and generation models. For segmentor, we experiment non-pretrained Transformer model and pretrained BART and UNILM models. For generator, we also include TOPICRANK, which is a classical extractive summarization method.

The last vertical block includes the results of our joint models: SEGBART, SEGTRANS and SEGUNILM.

Models	R1	R2	RL
SEGUNILM	22.17	8.86	20.17
($c=12$)	22.14	8.81	20.09
($c=8$)	22.13	8.84	20.10
($c=4$)	21.39	7.99	19.23
($c=0$)	19.85	7.74	17.62
(w/o seg loss)	22.06	8.66	20.02

Table 3: Ablation study results on SEGNEWS. We compare multiple variants of SEGUNILM. c indicates the number of decoder heads modified into segmentation-aware attention. By default, SEGUNILM uses $c = 9$ to achieve the best performance. We also present a SEGUNILM model without (w/o) segmentation classification loss, and it is trained solely by generation loss.

Model	WD	PK
EVEN	0.469	0.450
Transformer	0.563	0.462
BART	0.484	0.411
UNILM	0.479	0.391
SEGBART	0.471	0.405
SEGUNILM	0.462	0.380

Table 4: Experimental results on document segmentation task. WD indicates $WinDiff$ metric.

They respectively rely on non-pretrained Transformer and pretrained BART and UNILM as backbone models. Segmentation-aware attention mechanism is used to augment these jointly trained systems.

We can see vanilla sequence-to-sequence models with no segmentation information input perform poorly on this task. End-to-end SEGUNILM model achieves the best performance among all systems. SEGUNILM outperforms the best pipeline system under both settings when gold segments or predicted segments are provided. This indicates SEGUNILM has better overall performance and will be more useful when applied as practical applications. It also shows higher summarization results than vanilla UNILM model, confirming the effectiveness of segmentation-aware attention mechanism. SEGBART and SEGTRANS also show similar superiority over their pipeline versions. Examples of system output are shown in Table 5.

Table 3 summarizes ablation studies aiming to assess the contribution of individual components of SEGUNILM. We first modify SEGUNILM by varying c , the number of heads of segmentation-aware attention. We can see the best results of ROUGE are achieved when $c = 9$. With more or less heads modified as segmentation-aware attention heads, the summarization performance show a clear trend of decreasing. Also, as shown in the last column, when segmentation layer and segmentation loss are removed, we observe a sharp decrease on ROUGE scores. The results prove that both segmentation-aware attention and joint training provide improvement to the summarization results.

Table 4 describes the results on news segmentation task. SEGUNILM achieves the lowest WD and PK scores, reveal-

Title: One JFK conspiracy theory that could be true	
GOLD	1. LBJ had it done; 2. The military industrial complex did it; 3. The mob did it; 4. Oswald acted alone as part of an unknown conspiracy; 5. The CIA did it
Pipeline UNILM	those kennedys will never embarrass me again; did kennedy want to withdraw us troops from vietnam ?; 3. different mobs; other conspirators ?; would america be ok with that ?
SEGBART	they thought he was a crook; he was going to pull american troops out of vietnam; the mob did this; there were others, but who were they?; the russians ordered the killing
SEGUNILM	1. those kennedys will never embarrass me again; 2. he said he'd pull troops out of vietnam; 3. mob members claim they were witnesses to the alleged shootings; 4. there were more people who knew where oswald was; 5. the cia didn t release any of the good stuff
Title: What is the diversity visa lottery?	
GOLD	What is it?; How does it work?; How did it get started?; Has Congress tried to change the program?; What is Trump proposing?
Pipeline UNILM	what is a green card?; how the program works; history of the visa program; schumer helped replace the program; create a point system ...
SEGBART	what is the diversity visa program?; how does the program work?; who created the program?; who has sought reform?; what are the next steps?
SEGUNILM	what is the diversity visa program?; what are the requirements for the visas?; how did it start?; was the "diversity visa" created by the gang of eight?; is there any debate over reform?
Title: This man is tasked with finding out who failed Larry Nassar's victims	
GOLD	Seeking justice; A very youthful 68-year-old; A model independent prosecutor
Pipeline UNILM	searching for truth; he couldn't stay retired; he didn't have an agenda
SEGBART	searching for the truth; working with juveniles; no stone unturned
SEGUNILM	searching for the truth; he's has to do something; he doesn't have an agenda

Table 5: GOLD reference summaries and automatic summaries produced by pipeline UNILM, SEGBART and SEGUNILM on the SEGNEWS datasets. Semicolons indicate the boundaries of headings.

Model	Quality	Fluency
Pipeline UNILM	1.93	2.62
SEGUNILM	2.17	2.59
Gold	2.44	2.79

Table 6: Human evaluation results based on summary quality and fluency.

ing its ability to identify the structure of a news article. Compared with UNILM model without the segmentation-aware attention, SEGUNILM shows clear superiority on both metrics. The same trend is also observed in BART related models.

Human Evaluation

In addition to automatic evaluation, we also assess system performance by eliciting human judgments on 20 randomly selected test instances. The evaluation study assess the overall quality and fluency of the summaries by asking participants to rate them. We present the news article to evaluators along with system generated heading-style summaries, and we ask evaluators to read the complete article, and give scores based on summary quality and fluency respectively. Participants can have three scores (1-low quality/fluency, 2-median quality/fluency, 3-high quality/fluency).

Gold summaries, outputs from pipeline UNILM and SEGUNILM models are compared in evaluation. We invite three evaluators with linguist background to conduct the human evaluation. The averaged results are shown in Table 4.

Overall, we observe pipeline UNILM and SEGUNILM perform similarly on fluency, but SEGUNILM shows its superiority on summary quality. Gold summaries are marginally better than automatic generated summaries.

Conclusion

In this work, we proposed a new task, segmentation-based news summarization. It aims to segment a news article into multiple sections and generate the corresponding summary to each section. This new task provides a novel alternative to digesting a news article. We built a new benchmark dataset SEGNEWS to study and evaluate the task. Furthermore, we designed a segmentation-aware attention mechanism, which allows neural decoder to capture segmentation information in the source texts. The new attention mechanism is a non-invasive adaption of Transformer models and can be integrated with many pretrained language generation models. We jointly train the model for generating summaries and recognizing news segments. Experimental results on SEGNEWS demonstrate that our framework produces better segmentation-based summaries than competitive systems.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2015. Neural machine translation by jointly learning to align and translate. In *Proceedings of the 3rd International Conference on Learning Representations*. San Diego, California.
- Bao, H.; Dong, L.; Wei, F.; Wang, W.; Yang, N.; Liu, X.; Wang, Y.; Piao, S.; Gao, J.; Zhou, M.; and Hon, H.-W. 2020. UniLMv2: Pseudo-masked language models for unified language model pre-training. *arXiv preprint arXiv:2002.12804*.
- Beeferman, D.; Berger, A.; and Lafferty, J. 1999. Statistical Models for Text Segmentation. *Machine Learning*, 1(34): 177–210.
- Bougouin, A.; Boudin, F.; and Daille, B. 2013. Topicrank: Graph-based topic ranking for keyphrase extraction. In *International joint conference on natural language processing (IJCNLP)*, 543–551.
- Brown, G.; Brown, G. D.; Brown, G. R.; Gillian, B.; and Yule, G. 1983. *Discourse analysis*. Cambridge university press.
- Cheng, J.; and Lapata, M. 2016a. Neural Summarization by Extracting Sentences and Words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 484–494. Berlin, Germany.
- Cheng, J.; and Lapata, M. 2016b. Neural summarization by extracting sentences and words. In *Proceedings of the ACL Conference*.
- Choi, F. Y. Y. 2000. Advances in domain independent linear text segmentation. In *1st Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. ??? BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the NAACL Conference*.
- Durrett, G.; Berg-Kirkpatrick, T.; and Klein, D. 2016. Learning-based single-document summarization with compression and anaphoricity constraints. In *Proceedings of the ACL Conference*.
- Florax, M.; and Ploetzner, R. 2010. What contributes to the split-attention effect? The role of text segmentation, picture labelling, and spatial proximity. *Learning and instruction*, 20(3): 216–224.
- Galley, M.; McKeown, K.; Fosler-Lussier, E.; and Jing, H. 2003. Discourse segmentation of multi-party conversation. In *Proceedings of the 41st Annual Meeting of the Association for Computational Linguistics*, 562–569.
- Hearst, M. A. 1994. Multi-Paragraph Segmentation Expository Text. In *32nd Annual Meeting of the Association for Computational Linguistics*, 9–16. Las Cruces, New Mexico, USA: Association for Computational Linguistics.
- Hearst, M. A. 1997. Text Tiling: Segmenting text into multi-paragraph subtopic passages. *Computational linguistics*, 23(1): 33–64.
- Hernault, H.; Bollegala, D.; and Ishizuka, M. 2010. A sequential model for discourse segmentation. In *International Conference on Intelligent Text Processing and Computational Linguistics*, 315–326. Springer.
- Hochreiter, S.; and Schmidhuber, J. 1997. Long short-term memory. *Neural computation*, 9(8): 1735–1780.
- Hsueh, P.-Y.; Moore, J. D.; and Renals, S. 2006. Automatic segmentation of multiparty dialogue. In *11th Conference of the European Chapter of the Association for Computational Linguistics*.
- Lewis, M.; Liu, Y.; Goyal, N.; Ghazvininejad, M.; Mohamed, A.; Levy, O.; Stoyanov, V.; and Zettlemoyer, L. 2020. BART: Denoising Sequence-to-Sequence Pre-training for Natural Language Generation, Translation, and Comprehension. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Li, J.; Sun, A.; and Joty, S. R. 2018. SegBot: A Generic Neural Text Segmentation Model with Pointer Network. In *IJCAI*, 4166–4172.
- Lin, C.-Y. 2004. ROUGE: A Package for Automatic Evaluation of Summaries. In Marie-Francine Moens, S. S., ed., *Text Summarization Branches Out: Proceedings of the ACL-04 Workshop*, 74–81. Barcelona, Spain: Association for Computational Linguistics.
- Liu, Y.; and Lapata, M. 2019. Text Summarization with Pretrained Encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 3730–3740. Hong Kong, China.
- Liu, Y.; Titov, I.; and Lapata, M. 2019. Single Document Summarization as Tree Induction. In *Proceedings of the NAACL Conference*.
- Nallapati, R.; Zhai, F.; and Zhou, B. 2017. SummaRuNNer: A Recurrent Neural Network Based Sequence Model for Extractive Summarization of Documents. In *Proceedings of the AAAI Conference*.
- Nallapati, R.; Zhou, B.; dos Santos, C.; Gulcehre, C.; and Xiang, B. 2016. Abstractive Text Summarization using Sequence-to-sequence RNNs and Beyond. In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, 280–290. Berlin, Germany: Association for Computational Linguistics.
- Narayan, S.; Cohen, S. B.; and Lapata, M. 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*, 1797–1807. Brussels, Belgium.
- Nenkova, A.; and McKeown, K. 2011. Automatic Summarization. *Foundations and Trends in Information Retrieval*, 5(2–3): 103–233.
- Paice, C. D. 1990. Constructing literature abstracts by computer: techniques and prospects. *Information Processing & Management*, 26(1): 171–186.
- Pevzner, L.; and Hearst, M. A. 2002. A Critique and Improvement of an Evaluation Metric for Text Segmentation. *Computational Linguistics*, 28(1): 19–36.

- Radev, D. R.; Hovy, E.; and McKeown, K. 2002. Introduction to the Special Issue on Summarization. *Computational Linguistics*, 28(4): 399–408.
- Radford, A.; Wu, J.; Child, R.; Luan, D.; Amodei, D.; and Sutskever, I. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8): 9.
- Riedl, M.; and Biemann, C. 2012. Topictiling: a text segmentation algorithm based on lda. In *Proceedings of ACL 2012 Student Research Workshop*, 37–42.
- Rush, A. M.; Chopra, S.; and Weston, J. 2015. A Neural Attention Model for Abstractive Sentence Summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, 379–389. Lisbon, Portugal.
- Sandhaus, E. 2008. The New York Times Annotated Corpus. *Linguistic Data Consortium, Philadelphia*, 6(12).
- See, A.; Liu, P. J.; and Manning, C. D. 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)*, 1073–1083. Vancouver, Canada.
- Song, K.; Tan, X.; Qin, T.; Lu, J.; and Liu, T.-Y. 2019. MASS: Masked sequence to sequence pre-training for language generation. In *Proceedings of International Conference on Machine Learning*, 5926–5936. Long Beach, California.
- Vaswani, A.; Shazeer, N.; Parmar, N.; Uszkoreit, J.; Jones, L.; Gomez, A. N.; Kaiser, L. u.; and Polosukhin, I. 2017. Attention is All you Need. In Guyon, I.; Luxburg, U. V.; Bengio, S.; Wallach, H.; Fergus, R.; Vishwanathan, S.; and Garnett, R., eds., *Advances in Neural Information Processing Systems 30*, 5998–6008.
- Wu, Y.; Schuster, M.; Chen, Z.; Le, Q. V.; Norouzi, M.; Macherey, W.; Krikun, M.; Cao, Y.; Gao, Q.; Macherey, K.; et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. In *arXiv preprint arXiv:1609.08144*.
- Zhang, J.; Zhao, Y.; Saleh, M.; and Liu, P. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, 11328–11339. PMLR.
- Zhang, R.; Guo, J.; Fan, Y.; Lan, Y.; and Cheng, X. 2019. Outline generation: Understanding the inherent content structure of documents. In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 745–754.