

Multi-News: a Large-Scale Multi-Document Summarization Dataset and Abstractive Hierarchical Model

Alexander R. Fabbri Irene Li
Tianwei She Suyi Li Dragomir R. Radev

Department of Computer Science
Yale University
{alexander.fabbri,irene.li,tianwei.she,suyi.li,dragomir.radev}@yale.edu

Abstract

Automatic generation of summaries from multiple news articles is a valuable tool as the number of online publications grows rapidly. Single document summarization (SDS) systems have benefited from advances in neural encoder-decoder model thanks to the availability of large datasets. However, multi-document summarization (MDS) of news articles has been limited to datasets of a couple of hundred examples. In this paper, we introduce Multi-News, the first large-scale MDS news dataset. Additionally, we propose an end-to-end model which incorporates a traditional extractive summarization model with a standard SDS model and achieves competitive results on MDS datasets. We benchmark several methods on Multi-News and release our data and code in hope that this work will promote advances in summarization in the multi-document setting¹.

1 Introduction

Summarization is a central problem in Natural Language Processing with increasing applications as the desire to receive content in a concise and easily-understood format increases. Recent advances in neural methods for text summarization have largely been applied in the setting of single-document news summarization and headline generation (Rush et al., 2015; See et al., 2017; Gehrmann et al., 2018). These works take advantage of large datasets such as the Gigaword Corpus (Napoles et al., 2012), the CNN/Daily Mail (CNNDM) dataset (Hermann et al., 2015), the New York Times dataset (NYT, 2008) and the Newsroom corpus (Grusky et al., 2018), which contain on the order of hundreds of thousands to millions of article-summary pairs. However, multi-document summarization (MDS), which aims to

¹<https://github.com/Alex-Fabbri/Multi-News>

Source 1 Meng Wanzhou, Huawei’s chief financial officer and deputy chair, was arrested in Vancouver on 1 December. Details of the arrest have not been released...
Source 2 A Chinese foreign ministry spokesman said on Thursday that Beijing had separately called on the US and Canada to “clarify the reasons for the detention” immediately and “immediately release the detained person”. The spokesman...
Source 3 Canadian officials have arrested Meng Wanzhou, the chief financial officer and deputy chair of the board for the Chinese tech giant Huawei,...Meng was arrested in Vancouver on Saturday and is being sought for extradition by the United States. A bail hearing has been set for Friday...
Summary ...Canadian authorities say she was being sought for extradition to the US, where the company is being investigated for possible violation of sanctions against Iran. Canada’s justice department said Meng was arrested in Vancouver on Dec. 1... China’s embassy in Ottawa released a statement.. “The Chinese side has lodged stern representations with the US and Canadian side, and urged them to immediately correct the wrongdoing” and restore Meng’s freedom, the statement said...

Table 1: An example from our multi-document summarization dataset showing the input documents and their summary. Content found in the summary is color-coded.

output summaries from document clusters on the same topic, has largely been performed on datasets with less than 100 document clusters such as the DUC 2004 (Paul and James, 2004) and TAC 2011 (Owczarzak and Dang, 2011) datasets, and has benefited less from advances in deep learning methods.

Multi-document summarization of news events offers the challenge of outputting a well-organized summary which covers an event comprehensively while simultaneously avoiding redundancy. The input documents may differ in focus and point of view for an event. We present an example of multiple input news documents and their summary in

Figure 1. The three source documents discuss the same event and contain overlaps in content: the fact that *Meng Wanzhou was arrested* is stated explicitly in Source 1 and 3 and indirectly in Source 2. However, some sources contain information not mentioned in the others which should be included in the summary: Source 3 states that (*Wanzhou is being sought for extradition by the US*) while only Source 2 mentioned *the attitude of the Chinese side*.

Recent work in tackling this problem with neural models has attempted to exploit the graph structure among discourse relations in text clusters (Yasunaga et al., 2017) or through an auxiliary text classification task (Cao et al., 2017). Additionally, a couple of recent papers have attempted to adapt neural encoder decoder models trained on single document summarization datasets to MDS (Lebanoff et al., 2018; Baumel et al., 2018; Zhang et al., 2018b).

However, data sparsity has largely been the bottleneck of the development of neural MDS systems. The creation of large-scale multi-document summarization dataset for training has been restricted due to the sparsity and cost of human-written summaries. Liu et al. (2018) trains abstractive sequence-to-sequence models on a large corpus of Wikipedia text with citations and search engine results as input documents. However, no analogous dataset exists in the news domain. To bridge the gap, we introduce Multi-News, the first large-scale MDS news dataset, which contains 56,216 articles-summary pairs. We also propose a hierarchical model for neural abstractive multi-document summarization, which consists of a pointer-generator network (See et al., 2017) and an additional Maximal Marginal Relevance (MMR) (Carbonell and Goldstein, 1998) module that calculates sentence ranking scores based on relevancy and redundancy. We integrate sentence-level MMR scores into the pointer-generator model to adapt the attention weights on a word-level. Our model performs competitively on both our Multi-News dataset and the DUC 2004 dataset on ROUGE scores. We additionally perform human evaluation on several system outputs.

Our contributions are as follows: We introduce the first large-scale multi-document summarization datasets in the news domain. We propose an end-to-end method to incorporate MMR into pointer-generator networks. Finally, we bench-

mark various methods on our dataset to lay the foundations for future work on large-scale MDS.

2 Related Work

Traditional non-neural approaches to multi-document summarization have been both extractive (Carbonell and Goldstein, 1998; Radev et al., 2000; Erkan and Radev, 2004; Mihalcea and Tarau, 2004; Haghighi and Vanderwende, 2009) as well as abstractive (McKeown and Radev, 1995; Radev and McKeown, 1998; Barzilay et al., 1999; Ganesan et al., 2010). Recently, neural methods have shown great promise in text summarization, although largely in the single-document setting, with both extractive (Nallapati et al., 2016a; Cheng and Lapata, 2016; Narayan et al., 2018b) and abstractive methods (Chopra et al., 2016; Nallapati et al., 2016b; See et al., 2017; Paulus et al., 2017; Cohan et al., 2018; Çelikyilmaz et al., 2018; Gehrmann et al., 2018)

In addition to the multi-document methods described above which address data sparsity, recent work has attempted unsupervised and weakly supervised methods in non-news domains (Chu and Liu, 2018; Angelidis and Lapata, 2018). The methods most related to this work are SDS adapted for MDS data. Zhang et al. (2018a) adopts a hierarchical encoding framework trained on SDS data to MDS data by adding an additional document-level encoding. Baumel et al. (2018) incorporates query relevance into standard sequence-to-sequence models. Lebanoff et al. (2018) adapts encoder-decoder models trained on single-document datasets to the MDS case by introducing an external MMR module which does not require training on the MDS dataset. In our work, we incorporate the MMR module directly into our model, learning weights for the similarity functions simultaneously with the rest of the model.

3 Multi-News Dataset

Our dataset, which we call Multi-News, consists of news articles and human-written summaries of these articles from the site newser.com. Each summary is professionally written by editors and includes links to the original articles cited. We will release stable Wayback-archived links, and scripts to reproduce the dataset from these links. Our dataset is notably the first large-scale dataset for MDS on news articles. Our dataset also comes

# of source	Frequency	# of source	Frequency
2	23,894	7	382
3	12,707	8	209
4	5,022	9	89
5	1,873	10	33
6	763		

Table 2: The number of source articles per example, by frequency, in our dataset.

from a diverse set of news sources; over 1,500 sites appear as source documents 5 times or greater, as opposed to previous news datasets (DUC comes from 2 sources, CNNDM comes from CNN and Daily Mail respectively, and even the Newsroom dataset (Grusky et al., 2018) covers only 38 news sources). A total of 20 editors contribute to 85% of the total summaries on newser.com. Thus we believe that this dataset allows for the summarization of diverse source documents and summaries.

3.1 Statistics and Analysis

The number of collected Wayback links for summaries and their corresponding cited articles totals over 250,000. We only include examples with between 2 and 10 source documents per summary, as our goal is MDS, and the number of examples with more than 10 sources was minimal. The number of source articles per summary present, after downloading and processing the text to obtain the original article text, varies across the dataset, as shown in Table 2. We believe this setting reflects real-world situations; often for a new or specialized event there may be only a few news articles. Nonetheless, we would like to summarize these events in addition to others with greater news coverage.

We split our dataset into training (80%, 44,972), validation (10%, 5,622), and test (10%, 5,622) sets. Table 3 compares Multi-News to other news datasets used in experiments below. We choose to compare Multi-News with DUC data from 2003 and 2004 and TAC 2011 data, which are typically used in multi-document settings. Additionally, we compare to the single-document CNNDM dataset, as this has been recently used in work which adapts SDS to MDS (Lebanoff et al., 2018). The number of examples in our Multi-News dataset is two orders of magnitude larger than previous MDS news data. The total number of words in the concatenated inputs is shorter than other MDS datasets, as those consist of 10 input documents, but larger than SDS datasets, as expected. Our

summaries are notably longer than in other works, about 260 words on average. While compressing information into a shorter text is the goal of summarization, our dataset tests the ability of abstractive models to generate fluent text concise in meaning while also coherent in the entirety of its generally longer output, which we consider an interesting challenge.

3.2 Diversity

We report the percentage of n-grams in the gold summaries which do not appear in the input documents as a measure of how abstractive our summaries are in Table 4. As the table shows, the smaller MDS datasets tend to be more abstractive, but Multi-News is comparable and similar to the abstractiveness of SDS datasets. Grusky et al. (2018) additionally define three measures of the extractive nature of a dataset, which we use here for a comparison. We extend these notions to the multi-document setting by concatenating the source documents and treating them as a single input. Extractive fragment coverage is the percentage of words in the summary that are from the source article, measuring the extent to which a summary is derivative of a text:

$$COVERAGE(A,S) = \frac{1}{|S|} \sum_{f \in F(A,S)} |f| \quad (1)$$

where A is the article, S the summary, and $F(A, S)$ the set of all token sequences identified as extractive in a greedy manner; if there is a sequence of source tokens that is a prefix of the remainder of the summary, that is marked as extractive. Similarly, density is defined as the average length of the extractive fragment to which each summary word belongs:

$$DENSITY(A,S) = \frac{1}{|S|} \sum_{f \in F(A,S)} |f|^2 \quad (2)$$

Finally, compression ratio is defined as the word ratio between the articles and its summaries:

$$COMPRESSION(A,S) = \frac{|A|}{|S|} \quad (3)$$

These numbers are plotted using kernel density estimation in Figure 1. As explained above, our summaries are larger on average, which corresponds to a lower compression rate. The variability along the x-axis (fragment coverage), suggests

Dataset	# pairs	# words (doc)	# sents (docs)	# words (summary)	# sents (summary)	vocab size
Multi-News	44,972/5,622/5,622	2,103.49	82.73	263.66	9.97	666,515
DUC03+04	320	4,636.24	173.15	109.58	2.88	19,734
TAC 2011	176	4,695.70	188.43	99.70	1.00	24,672
CNNNDM	287,227/13,368/11,490	810.57	39.78	56.20	3.68	717,951

Table 3: Comparison of our Multi-News dataset to other MDS datasets as well as an SDS dataset used as training data for MDS (CNNNDM). Training, validation and testing size splits (article(s) to summary) are provided when applicable. Statistics for multi-document inputs are calculated on the concatenation of all input sources.

% novel n-grams	Multi-News	DUC03+04	TAC11	CNNNDM
uni-grams	17.76	27.74	16.65	19.50
bi-grams	57.10	72.87	61.18	56.88
tri-grams	75.71	90.61	83.34	74.41
4-grams	82.30	96.18	92.04	82.83

Table 4: Percentage of n-grams in summaries which do not appear in the input documents, a measure of the abstractiveness, in relevant datasets.




Figure 1: Density estimation of extractive diversity scores as explained in Section 3.2. Large variability along the y-axis suggests variation in the average length of source sequences present in the summary, while the x axis shows variability in the average length of the extractive fragments to which summary words belong.

variability in the percentage of copied words, with the DUC data varying the most. In terms of y-axis (fragment density), our dataset shows variability in the average length of copied sequence, suggesting varying styles of word sequence arrangement. Our dataset exhibits extractive characteristics similar to the CNNNDM dataset.

3.3 Other Datasets

As discussed above, large scale datasets for multi-document news summarization are lacking. There have been several attempts to create MDS datasets in other domains. Zopf (2018) introduce a multi-lingual MDS dataset based on English and German Wikipedia articles as summaries to create a

set of about 7,000 examples. Liu et al. (2018) use Wikipedia as well, creating a dataset of over two million examples. That paper uses Wikipedia references as input documents but largely relies on Google search to increase topic coverage. We, however, are focused on the news domain, and the source articles in our dataset are specifically cited by the corresponding summaries. Related work has also focused on opinion summarization in the multi-document setting; Angelidis and Lapata (2018) introduces a dataset of 600 Amazon product reviews.

4 Preliminaries

We introduce several common methods for summarization.

4.1 Pointer-generator Network

The pointer-generator network (See et al., 2017) is a commonly-used encoder-decoder summarization model with attention (Bahdanau et al., 2014) which combines copying words from source documents and outputting words from a vocabulary. The encoder converts each token w_i in the document into the hidden state h_i . At each decoding step t , the decoder has a hidden state d_t . An attention distribution a^t is calculated as in (Bahdanau et al., 2014) and is used to get the context vector h_t^* , which is a weighted sum of the encoder hidden states, representing the semantic meaning of the related document content for this decoding time step:

$$\begin{aligned}
 e_i^t &= v^T \tanh(W_h h_i + W_s d_t + b_{attn}) \\
 a^t &= \text{softmax}(e^t) \\
 h_t^* &= \sum_i a_i^t h_i^t
 \end{aligned} \tag{4}$$

The context vector h_t^* and the decoder hidden state d_t are then passed to two linear layers to produce the vocabulary distribution P_{vocab} . For each word, there is also a copy probability P_{copy} . It is the sum

of the attention weights over all the word occurrences:

$$\begin{aligned} P_{vocab} &= \text{softmax}(V'(V[d_t, h_t^*] + b) + b') \\ P_{copy} &= \sum_{i:w_i=w} d_i^t \end{aligned} \quad (5)$$

The pointer-generator network has a soft switch p_{gen} , which indicates whether to generate a word from vocabulary by sampling from P_{vocab} , or to copy a word from the source sequence by sampling from the copy probability P_{copy} .

$$p_{gen} = \sigma(w_{h^*}^T h_t^* + w_d^T d_t + w_x^T x_t + b_{ptr}) \quad (6)$$

where x_t is the decoder input. The final probability distribution is a weighted sum of the vocabulary distribution and copy probability:

$$P(w) = p_{gen}P_{vocab}(w) + (1 - p_{gen})P_{copy}(w) \quad (7)$$

4.2 Transformer

The Transformer model replaces recurrent layers with self-attention in an encoder-decoder framework and has achieved state-of-the-art results in machine translation (Vaswani et al., 2017) and language modeling (Baeviski and Auli, 2019; Dai et al., 2019). The Transformer has also been successfully applied to SDS (Gehrmann et al., 2018). More specifically, for each word during encoding, the multi-head self-attention sub-layer allows the encoder to directly attend to all other words in a sentence in one step. Decoding contains the typical encoder-decoder attention mechanisms as well as self-attention to all previous generated output. The Transformer motivates the elimination of recurrence to allow more direct interaction among words in a sequence.

4.3 MMR

Maximal Marginal Relevance (MMR) is an approach for combining query-relevance with information-novelty in the context of summarization (Carbonell and Goldstein, 1998). MMR produces a ranked list of the candidate sentences based on the relevance and redundancy to the query, which can be used to extract sentences. The score is calculated as follows:

$$\begin{aligned} \text{MMR} &= \underset{D_i \in R \setminus S}{\text{argmax}} \left[\lambda \text{Sim}_1(D_i, Q) \right. \\ &\quad \left. - (1 - \lambda) \max_{D_j \in S} \text{Sim}_2(D_i, D_j) \right] \end{aligned} \quad (8)$$




Figure 2: Our Hierarchical MMR-Attention Pointer-generator (Hi-MAP) model incorporates sentence-level representations and hidden-state-based MMR on top of a standard pointer-generator network.

where R is the collection of all candidate sentences, Q is the query, S is the set of sentences that have been selected, and $R \setminus S$ is set of the un-selected ones. In general, each time we want to select a sentence, we have a ranking score for all the candidates that considers relevance and redundancy. A recent work (Lebanoff et al., 2018) applied MMR for multi-document summarization by creating an external module and a supervised regression model for sentence importance. Our proposed method, however, incorporates MMR with the pointer-generator network in an end-to-end manner that learns parameters for similarity and redundancy.

5 Hi-MAP Model

In this section, we provide the details of our Hierarchical MMR-Attention Pointer-generator (Hi-MAP) model for multi-document neural abstractive summarization. We expand the existing pointer-generator network model into a hierarchical network, which allows us to calculate sentence-level MMR scores. Our model consists of a pointer-generator network and an integrated MMR module, as shown in Figure 2.

5.1 Sentence representations

To expand our model into a hierarchical one, we compute sentence representations on both the encoder and decoder. The input is a collection of sentences $D = [s_1, s_2, \dots, s_n]$ from all the source documents, where a given sentence $s_i = [w_{k-m}, w_{k-m+1}, \dots, w_k]$ is made up of input word tokens. Word tokens from the whole document are treated as a single sequential input to a Bi-LSTM encoder as in the original encoder of the pointer-

generator network from See et al. (2017) (see bottom of Figure 2). For each time step, the output of an input word token w_l is h_l^w (we use superscript w to indicate word-level LSTM cells, s for sentence-level).

To obtain a representation for each sentence s_i , we take the encoder output of the last token for that sentence. If that token has an index of k in the whole document D , then the sentence representation is marked as $h_{s_i}^w = h_k^w$. The word-level sentence embeddings of the document $h_D^w = [h_{s_1}^w, h_{s_2}^w, \dots, h_{s_n}^w]$ will be a sequence which is fed into a sentence-level LSTM network. Thus, for each input sentence $h_{s_i}^w$, we obtain an output hidden state $h_{s_i}^s$. We then get the final sentence-level embeddings $h_D^s = [h_1^s, h_2^s, \dots, h_n^s]$ (we omit the subscript for sentences s). To obtain a summary representation, we simply treat the current decoded summary as a single sentence and take the output of the last step of the decoder: s_{sum} . We plan to investigate alternative methods for input and output sentence embeddings, such as separate LSTMs for each sentence, in future work.

5.2 MMR-Attention

Now, we have all the sentence-level representation from both the articles and summary, and then we apply MMR to compute a ranking on the candidate sentences h_D^s . Intuitively, incorporating MMR will help determine salient sentences from the input at the current decoding step based on relevancy and redundancy.

We follow Section 4.3 to compute MMR scores. Here, however, our query document is represented by the summary vector s_{sum} , and we want to rank the candidates in h_D^s . The MMR score for an input sentence i is then defined as:

$$\text{MMR}_i = \lambda \text{Sim}_1(h_i^s, s_{sum}) - (1 - \lambda) \max_{s_j \in D, j \neq i} \text{Sim}_2(h_i^s, h_j^s) \quad (9)$$

We then add a softmax function to normalize all the MMR scores of these candidates as a probability distribution.

$$\overline{\text{MMR}}_i = \frac{\exp(\text{MMR}_i)}{\sum_i \exp(\text{MMR}_i)} \quad (10)$$

Now we define the similarity function between each candidate sentence h_i^s and summary sentence s_{sum} to be:

$$\text{Sim}_1 = h_i^{sT} W_{\text{Sim}} s_{sum} \quad (11)$$

where W_{Sim} is a learned parameter used to transform s_{sum} and h_i^s into a common feature space.

For the second term of Equation 9, instead of choosing the maximum score from all candidates except for h_i^s , which is intended to find the candidate most similar to h_i^s , we choose to apply a self-attention model on h_i^s and all the other candidates $h_j^s \in h_D^s$. We then choose the largest weight as the final score:

$$\begin{aligned} v_{ij} &= \tanh\left(h_j^{sT} W_{self} h_i^s\right) \\ \beta_{ij} &= \frac{\exp(v_{ij})}{\sum_j \exp(v_{ij})} \\ \text{score}_i &= \max_j(\beta_{i,j}) \end{aligned} \quad (12)$$

Note that W_{self} is also a trainable parameter. Eventually, the MMR score from Equation 9 becomes:

$$\text{MMR}_i = \lambda \text{Sim}_1(h_i^s, s_{sum}) - (1 - \lambda) \text{score}_i \quad (13)$$

5.3 MMR-attention Pointer-generator

After we calculate $\overline{\text{MMR}}_i$ for each sentence representation h_i^s , we use these scores to update the word-level attention weights for the pointer-generator model shown by the blue arrows in Figure 2. Since $\overline{\text{MMR}}_i$ is a sentence weight for h_i^s , each token in the sentence will have the same value of $\overline{\text{MMR}}_i$. The new attention for each input token from Equation 4 becomes:

$$\overline{a^t} = a^t \overline{\text{MMR}}_i \quad (14)$$

6 Experiments

In this section we describe additional methods we compare with and present our assumptions and experimental process.

6.1 Baseline and Extractive Methods

First We concatenate the first sentence of each article in a document cluster as the system summary. For our dataset, *First-k* means the first k sentences from each source article will be concatenated as the summary. Due to the difference in gold summary length, we only use First-1 for DUC, as others would exceed the average summary length.

LexRank Initially proposed by (Erkan and Radev, 2004), LexRank is a graph-based method for computing relative importance in extractive summarization.

TextRank Introduced by (Mihalcea and Tarau, 2004), TextRank is a graph-based ranking model. Sentence importance scores are computed based on eigenvector centrality within a global graph from the corpus.

MMR In addition to incorporating MMR in our pointer generator network, we use this original method as an extractive summarization baseline. When testing on DUC data, we set these extractive methods to give an output of 100 tokens and 300 tokens for Multi-News data.

6.2 Neural Abstractive Methods

PG-Original, PG-MMR These are the original pointer-generator network models reported by (Lebanoff et al., 2018).

PG-BRNN The PG-BRNN model is a pointer-generator implementation from OpenNMT². As in the original paper (See et al., 2017), we use a 1-layer bi-LSTM as encoder, with 128-dimensional word-embeddings and 256-dimensional hidden states for each direction. The decoder is a 512-dimensional single-layer LSTM. We include this for reference in addition to PG-Original, as our Hi-MAP code builds upon this implementation.

CopyTransformer Instead of using an LSTM, the CopyTransformer model used in Gehrmann et al. (2018) uses a 4-layer Transformer of 512 dimensions for encoder and decoder. One of the attention heads is chosen randomly as the copy distribution. This model and the PG-BRNN are run without the bottom-up masked attention for inference from Gehrmann et al. (2018) as we did not find a large improvement when reproducing the model on this data.

6.3 Experimental Setting

Following the setting from (Lebanoff et al., 2018), we report ROUGE (Lin, 2004) scores, which measure the overlap of unigrams (R-1), bigrams (R-2) and skip bigrams with a max distance of four words (R-SU). For the neural abstractive models, we truncate input articles to 500 tokens in the following way: for each example with S source input documents, we take the first 500/ S tokens from each source document. As some source documents may be shorter, we iteratively determine the number of tokens to take from each document until the 500 token quota is reached. Hav-

²<https://github.com/OpenNMT/OpenNMT-py/blob/master/docs/source/Summarization.md>

ing determined the number of tokens per source document to use, we concatenate the truncated source documents into a single mega-document. This effectively reduces MDS to SDS on longer documents, a commonly-used assumption for recent neural MDS papers (Cao et al., 2017; Liu et al., 2018; Lebanoff et al., 2018). We chose 500 as our truncation size as related MDS work did not find significant improvement when increasing input length from 500 to 1000 tokens (Liu et al., 2018). We simply introduce a special token between source documents to aid our models in detecting document-to-document relationships and leave direct modeling of this relationship, as well as modeling longer input sequences, to future work. We hope that the dataset we introduce will promote such work. For our Hi-MAP model, we applied a 1-layer bidirectional LSTM network, with the hidden state dimension 256 in each direction. The sentence representation dimension is also 256. We set the $\lambda = 0.5$ to calculate the MMR value in Equation 9.

Method	R-1	R-2	R-SU
First	30.77	8.27	7.35
LexRank (Erkan and Radev, 2004)	35.56	7.87	11.86
TextRank (Mihalcea and Tarau, 2004)	33.16	6.13	10.16
MMR (Carbonell and Goldstein, 1998)	30.14	4.55	8.16
PG-Original (Lebanoff et al., 2018)	31.43	6.03	10.01
PG-MMR (Lebanoff et al., 2018)	36.42	9.36	13.23
PG-BRNN (Gehrmann et al., 2018)	29.47	6.77	7.56
CopyTransformer (Gehrmann et al., 2018)	28.54	6.38	7.22
Hi-MAP (Our Model)	35.78	8.90	11.43

Table 5: ROUGE scores on the DUC 2004 dataset for models trained on CNNDM data, as in Lebanoff et al. (2018).³

Method	R-1	R-2	R-SU
First-1	26.83	7.25	6.46
First-2	35.99	10.17	12.06
First-3	39.41	11.77	14.51
LexRank (Erkan and Radev, 2004)	38.27	12.70	13.20
TextRank (Mihalcea and Tarau, 2004)	38.44	13.10	13.50
MMR (Carbonell and Goldstein, 1998)	38.77	11.98	12.91
PG-Original (Lebanoff et al., 2018)	41.85	12.91	16.46
PG-MMR (Lebanoff et al., 2018)	40.55	12.36	15.87
PG-BRNN (Gehrmann et al., 2018)	42.80	14.19	16.75
CopyTransformer (Gehrmann et al., 2018)	43.57	14.03	17.37
Hi-MAP (Our Model)	43.47	14.89	17.41

Table 6: ROUGE scores for models trained and tested on the Multi-News dataset.

³As our focus was on deep methods for MDS, we only tested several non-neural baselines. However, other classical methods deserve more attention, for which we refer the reader to Hong et al. (2014) and leave the implementation of these methods on Multi-News for future work.

Method	Informativeness	Fluency	Non-Redundancy
PG-MMR	95	70	45
Hi-MAP	85	75	100
CopyTransformer	99	100	107
Human	150	150	149

Table 7: Number of times a system was chosen as best in pairwise comparisons according to informativeness, fluency and non-redundancy.

7 Analysis and Discussion

In Table 5 and Table 6 we report ROUGE scores on DUC 2004 and Multi-News datasets respectively. We use DUC 2004, as results on this dataset are reported in Lebanoff et al. (2018), although this dataset is not the focus of this work. For results on DUC 2004, models were trained on the CNNDM dataset, as in Lebanoff et al. (2018). PG-BRNN and CopyTransformer models, which were pretrained by OpenNMT on CNNDM, were applied to DUC without additional training, analogous to PG-Original. We also experimented with training on Multi-News and testing on DUC data, but we did not see significant improvements. We attribute the generally low performance of pointer-generator, CopyTransformer and Hi-MAP to domain differences between DUC and CNNDM as well as DUC and Multi-News. These domain differences are evident in the statistics and extractive metrics discussed in Section 3.

Additionally, for both DUC and Multi-News testing, we experimented with using the output of 500 tokens from extractive methods (LexRank, TextRank and MMR) as input to the abstractive model. However, this did not improve results. We believe this is because our truncated input mirrors the First-3 baseline, which outperforms these three extractive methods and thus may provide more information as input to the abstractive model.

Our model outperforms PG-MMR when trained and tested on the Multi-News dataset. We see much-improved model performances when trained and tested on in-domain Multi-News data. The Transformer performs best in terms of R-1 while Hi-MAP outperforms it on R-2 and R-SU. Also, we notice a drop in performance between PG-original, and PG-MMR (which takes the pretrained PG-original and applies MMR on top of the model). Our PG-MMR results correspond to *PG-MMR w Cosine* reported in Lebanoff et al. (2018). We trained their sentence regression model on Multi-News data and leave the investigation of transferring regression models from SDS to Multi-News for future work.

In addition to automatic evaluation, we performed human evaluation to compare the summaries produced. We used Best-Worst Scaling (Louveire and Woodworth, 1991; Louviere et al., 2015), which has shown to be more reliable than rating scales (Kiritchenko and Mohammad, 2017) and has been used to evaluate summaries (Narayan et al., 2018a; Angelidis and Lapata, 2018). Annotators were presented with the same input that the systems saw at testing time; input documents were truncated, and we separated input documents by visible spaces in our annotator interface. We chose three native English speakers as annotators. They were presented with input documents, and summaries generated by two out of four systems, and were asked to determine which summary was better and which was worse in terms of *informativeness* (is the meaning in the input text preserved in the summary?), *fluency* (is the summary written in well-formed and grammatical English?) and *non-redundancy* (does the summary avoid repeating information?). We randomly selected 50 documents from the Multi-News test set and compared all possible combinations of two out of four systems. We chose to compare PG-MMR, CopyTransformer, Hi-MAP and gold summaries. The order of summaries was randomized per example.

The results of our pairwise human-annotated comparison are shown in Table 7. Human-written summaries were easily marked as better than other systems, which, while expected, shows that there is much room for improvement in producing readable, informative summaries. We performed pairwise comparison of the models over the three metrics combined, using a one-way ANOVA with Tukey HSD tests and p value of 0.05. Overall, statistically significant differences were found between human summaries score and all other systems, CopyTransformer and the other two models, and our Hi-MAP model compared to PG-MMR. Our Hi-MAP model performs comparably to PG-MMR on informativeness and fluency but much better in terms of non-redundancy. We believe that the incorporation of learned parameters for similarity and redundancy reduces redundancy in our output summaries. In future work, we would like to incorporate MMR into Transformer models to benefit from their fluent summaries.

8 Conclusion

In this paper we introduce Multi-News, the first large-scale multi-document news summarization

dataset. We hope that this dataset will promote work in multi-document summarization similar to the progress seen in the single-document case. Additionally, we introduce an end-to-end model which incorporates MMR into a pointer-generator network, which performs competitively compared to previous multi-document summarization models. We also benchmark methods on our dataset. In the future we plan to explore interactions among documents beyond concatenation and experiment with summarizing longer input documents.

References

2008. [The New York Times Annotated Corpus](#).

Stefanos Angelidis and Mirella Lapata. 2018. Summarizing Opinions: Aspect Extraction Meets Sentiment Prediction and They Are Both Weakly Supervised. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 3675–3686.

Alexei Baevski and Michael Auli. 2019. Adaptive input representations for neural language modeling. In *International Conference on Learning Representations*.

Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *arXiv preprint arXiv:1409.0473*.

Regina Barzilay, Kathleen R. McKeown, and Michael Elhadad. 1999. Information Fusion in the Context of Multi-Document Summarization. In *27th Annual Meeting of the Association for Computational Linguistics, University of Maryland, College Park, Maryland, USA, 20-26 June 1999*.

Tal Baumel, Matan Eyal, and Michael Elhadad. 2018. Query Focused Abstractive Summarization: Incorporating Query Relevance, Multi-Document Coverage, and Summary Length Constraints into seq2seq Models. *CoRR*, abs/1801.07704.

Ziqiang Cao, Wenjie Li, Sujian Li, and Furu Wei. 2017. Improving Multi-Document Summarization via Text Classification. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence, February 4-9, 2017, San Francisco, California, USA.*, pages 3053–3059.

Jaime Carbonell and Jade Goldstein. 1998. The Use of MMR, Diversity-Based Reranking for Reordering Documents and Producing Summaries. In *Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval*, pages 335–336. ACM.

Aslı Çelikyılmaz, Antoine Bosselut, Xiaodong He, and Yejin Choi. 2018. Deep Communicating Agents for Abstractive Summarization. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1662–1675.

Jianpeng Cheng and Mirella Lapata. 2016. Neural Summarization by Extracting Sentences and Words. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics, ACL 2016, August 7-12, 2016, Berlin, Germany, Volume 1: Long Papers*.

Sumit Chopra, Michael Auli, and Alexander M. Rush. 2016. Abstractive Sentence Summarization with Attentive Recurrent Neural Networks. In *NAACL HLT 2016, The 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, San Diego California, USA, June 12-17, 2016*, pages 93–98.

Eric Chu and Peter J. Liu. 2018. Unsupervised Neural Multi-Document Abstractive Summarization. *CoRR*, abs/1810.05739.

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, NAACL-HLT, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 2 (Short Papers)*, pages 615–621.

Zihang Dai, Zhilin Yang, Yiming Yang, William W. Cohen, Jaime Carbonell, Quoc V. Le, and Ruslan Salakhutdinov. 2019. Transformer-XL: Language modeling with longer-term dependency.

Günes Erkan and Dragomir R Radev. 2004. Lexrank: Graph-Based Lexical Centrality as Saliency in Text Summarization. *Journal of artificial intelligence research*, 22:457–479.

Kavita Ganesan, ChengXiang Zhai, and Jiawei Han. 2010. Opinosis: A Graph Based Approach to Abstractive Summarization of Highly Redundant Opinions. In *COLING 2010, 23rd International Conference on Computational Linguistics, Proceedings of the Conference, 23-27 August 2010, Beijing, China*, pages 340–348.

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, Brussels, Belgium, October 31 - November 4, 2018*, pages 4098–4109.

- Max Grusky, Mor Naaman, and Yoav Artzi. 2018. Newsroom: A Dataset of 1.3 Million Summaries with Diverse Extractive Strategies. *CoRR*, abs/1804.11283.
- Aria Haghighi and Lucy Vanderwende. 2009. Exploring Content Models for Multi-Document Summarization. In *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, May 31 - June 5, 2009, Boulder, Colorado, USA*, pages 362–370.
- Karl Moritz Hermann, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom. 2015. Teaching Machines to Read and Comprehend. In *Advances in Neural Information Processing Systems 28: Annual Conference on Neural Information Processing Systems 2015, December 7-12, 2015, Montreal, Quebec, Canada*, pages 1693–1701.
- Kai Hong, John M. Conroy, Benoıt Favre, Alex Kulesza, Hui Lin, and Ani Nenkova. 2014. A repository of state of the art and competitive baseline summaries for generic news summarization. In *Proceedings of the Ninth International Conference on Language Resources and Evaluation, LREC 2014, Reykjavik, Iceland, May 26-31, 2014.*, pages 1608–1616.
- Svetlana Kiritchenko and Saif Mohammad. 2017. Best-Worst Scaling More Reliable than Rating Scales: A Case Study on Sentiment Intensity Annotation. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics, ACL 2017, Vancouver, Canada, July 30 - August 4, Volume 2: Short Papers*, pages 465–470.
- Logan Lebanoff, Kaiqiang Song, and Fei Liu. 2018. Adapting the Neural Encoder-Decoder Framework from Single to Multi-Document summarization. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 4131–4141.
- Chin-Yew Lin. 2004. Rouge: A Package for Automatic Evaluation of Summaries. *Text Summarization Branches Out*.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. Generating Wikipedia by Summarizing Long Sequences. *CoRR*, abs/1801.10198.
- Jordan Louviere, Terry Flynn, and A. A. J. Marley. 2015. *Best-Worst Scaling: Theory, Methods and Applications*.
- Jordan J Louviere and George G Woodworth. 1991. *Best-Worst Scaling: A Model for the Largest Difference Judgments*.
- Kathleen R. McKeown and Dragomir R. Radev. 1995. Generating summaries of multiple news articles. In *Proceedings, ACM Conference on Research and Development in Information Retrieval SIGIR’95*, pages 74–82, Seattle, Washington.
- Rada Mihalcea and Paul Tarau. 2004. Textrank: Bringing Order into Text. In *Proceedings of the 2004 conference on empirical methods in natural language processing*.
- Ramesh Nallapati, Bowen Zhou, and Mingbo Ma. 2016a. Classify or Select: Neural Architectures for Extractive Document Summarization. *CoRR*, abs/1611.04244.
- Ramesh Nallapati, Bowen Zhou, Cıcer Nogueira dos Santos, Caglar Gulcehre, and Bing Xiang. 2016b. Abstractive Text Summarization Using Sequence-to-Sequence RNNs and Beyond. In *Proceedings of the 20th SIGNLL Conference on Computational Natural Language Learning, CoNLL 2016, Berlin, Germany, August 11-12, 2016*, pages 280–290.
- Courtney Napoles, Matthew R. Gormley, and Benjamin Van Durme. 2012. Annotated Gigaword. In *Proceedings of the Joint Workshop on Automatic Knowledge Base Construction and Web-scale Knowledge Extraction, AKBC-WEKEX@NAACL-HLT 2012, Montreal, Canada, June 7-8, 2012*, pages 95–100.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018a. 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. Association for Computational Linguistics.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018b. Ranking Sentences for Extractive Summarization with Reinforcement Learning. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, NAACL-HLT 2018, New Orleans, Louisiana, USA, June 1-6, 2018, Volume 1 (Long Papers)*, pages 1747–1759.
- Karolina Owczarzak and Hoa Trang Dang. 2011. Overview of the TAC 2011 Summarization Track: Guided Task and AESOP Task.
- Over Paul and Yen James. 2004. An Introduction to DUC-2004. In *Proceedings of the 4th Document Understanding Conference (DUC 2004)*.
- Romain Paulus, Caiming Xiong, and Richard Socher. 2017. A Deep Reinforced Model for Abstractive Summarization. *CoRR*, abs/1705.04304.
- Dragomir R. Radev, Hongyan Jing, and Malgorzata Budzikowska. 2000. Centroid-Based Summarization of Multiple Documents: Sentence Extraction utility-based evaluation, and user studies. *CoRR*, cs.CL/0005020.

- Dragomir R. Radev and Kathleen R. McKeown. 1998. Generating Natural Language Summaries from Multiple On-Line Sources. *Computational Linguistics*, 24(3):469–500.
- Alexander M. Rush, Sumit Chopra, and Jason Weston. 2015. A Neural Attention Model for Abstractive Sentence Summarization. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing, EMNLP 2015, Lisbon, Portugal, September 17-21, 2015*, pages 379–389.
- 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.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems*, pages 5998–6008.
- Michihiro Yasunaga, Rui Zhang, Kshitijh Meelu, Ayush Pareek, Krishnan Srinivasan, and Dragomir R. Radev. 2017. Graph-Based Neural Multi-Document Summarization. In *Proceedings of CoNLL-2017*. Association for Computational Linguistics.
- Jianmin Zhang, Jiwei Tan, and Xiaojun Wan. 2018a. Adapting Neural Single-Document Summarization Model for Abstractive Multi-Document Summarization: A Pilot Study. In *Proceedings of the 11th International Conference on Natural Language Generation, Tilburg University, The Netherlands, November 5-8, 2018*, pages 381–390.
- Jianmin Zhang, Jiwei Tan, and Xiaojun Wan. 2018b. Towards a Neural Network Approach to Abstractive Multi-Document Summarization. *CoRR*, abs/1804.09010.
- Markus Zopf. 2018. Auto-hmds: Automatic Construction of a Large Heterogeneous Multilingual Multi-Document Summarization Corpus. In *Proceedings of the Eleventh International Conference on Language Resources and Evaluation, LREC 2018, Miyazaki, Japan, May 7-12, 2018*.