

Long-Span Summarization via Local Attention and Content Selection

Potsawee Manakul and Mark J. F. Gales

Department of Engineering, University of Cambridge

pm574@cam.ac.uk, mjfg@eng.cam.ac.uk

Abstract

Transformer-based models have achieved state-of-the-art results in a wide range of natural language processing (NLP) tasks including document summarization. Typically these systems are trained by fine-tuning a large pre-trained model to the target task. One issue with these transformer-based models is that they do not scale well in terms of memory and compute requirements as the input length grows. Thus, for long document summarization, it can be challenging to train or fine-tune these models. In this work, we exploit large pre-trained transformer-based models and address long-span dependencies in abstractive summarization using two methods: local self-attention; and explicit content selection. These approaches are compared on a range of network configurations. Experiments are carried out on standard long-span summarization tasks, including Spotify Podcast, arXiv, and PubMed datasets. We demonstrate that by combining these methods, we can achieve state-of-the-art results on all three tasks in the ROUGE scores. Moreover, without a large-scale GPU card, our approach can achieve comparable or better results than existing approaches.¹

1 Introduction

Transformer-based models (Vaswani et al., 2017) are ubiquitously state-of-art across many natural language processing (NLP) tasks, including summarization. To achieve the best results, the community has trained ever larger transformer models on larger amount of data, and/or more task-specific optimization objectives (Devlin et al., 2019; Raffel et al., 2020; Lewis et al., 2020; Brown et al., 2020). In long document summarization, the input

sequences could be more than an order of magnitude longer than the limits of these transformer models. Although the limits can be extended, training large transformer models on long sequences is expensive and may not be possible on a standard GPU card because of the self-attention mechanism that grows quadratically with sequence length.

To tackle the quadratic characteristic, recent works have modified self-attention mechanism and proposed variants of the transformer such that the quadratic complexity is reduced (Tay et al., 2020b; Kitaev et al., 2020; Child et al., 2019; Beltagy et al., 2020; Ainslie et al., 2020; Zaheer et al., 2020). However, pre-trained weights of the modified models are not readily available. In contrast, standard models such as BERT (Devlin et al., 2019) or BART (Lewis et al., 2020) have been trained on various target tasks, including text summarization (Liu and Lapata, 2019b). This allows practitioners to achieve good performance with less training time. Thus, we are interested in exploiting pre-trained models for long-span summarization tasks.

We study a range of design configurations empirically and theoretically in regards to memory and compute requirements as well as their performance. We propose that long-span dependencies can be handled by two complementary methods. Firstly, inspired by modified self-attention transformers, we exploit standard transformer models by constraining attention mechanism to be local, allowing longer input spans during training. Secondly, because abstractive summarization systems perform content selection implicitly (Nallapati et al., 2016; Lebanoff et al., 2020), to reduce memory and compute requirements an alternative method is to perform content selection explicitly before the abstractive stage. We study content selection during two phases: training time and test time. At training time, we investigate methods to select data for training fixed-span abstractive models. At test

¹Our code is available at <https://github.com/potsawee/longsum0>.

time, we extend existing model-based selection methods, and we propose a multitask content selection method that ranks sentences through extractive labelling based module (Cheng and Lapata, 2016) and attention based module (See et al., 2017). Ultimately, we explore the combined approach, consisting of local self-attention transformer and content selection for long-document summarization.

We conduct our experiments using a number of design configurations on the Spotify open-domain Podcast summarization dataset (Clifton et al., 2020). This dataset is challenging not only because of its long-span nature, but also because transcribed spoken utterances typically have lower information density (Li et al., 2019; Manakul et al., 2020). Furthermore, we carry out experiments on arXiv and PubMed datasets (Cohan et al., 2018) to further demonstrate and verify the effectiveness of our approach as well as making comparisons to existing approaches. We highlight the strengths and weaknesses of our approach in different resources and tasks. The main contributions of this paper are:

- On local self-attention, we show how to exploit a standard transformer model for long-span summarization, and we show good design considerations based on empirical results.
- On content selection, we demonstrate the best selection method at training time, and we propose a multitask content selection (MCS) method outperforming baselines at test time.
- Our work has set new state-of-the-art results on Spotify Podcast, arXiv and PubMed datasets in the ROUGE scores. Furthermore, with a small-scale GPU card, our approach achieves comparable or superior performance to previous state-of-the-art systems.

2 Related Work

Efficient Transformers. Pre-trained transformer models have shown success and become the starting point for various NLP problems such as BERT (Devlin et al., 2019) in contextual representation, GPT2 in text generation (Radford et al., 2019), or BART in seq2seq tasks (Lewis et al., 2020). However, the memory and time requirements for transformer models grow quadratically with the sequence length, and for long-span tasks this quickly leads to GPU running out of memory in training. To mitigate the quadratic nature, a wide range of modified architectures have recently been proposed

(Tay et al., 2021). They reduce the quadratic complexity of the full self-attention mechanism by using fixed attention patterns (Parmar et al., 2018; Dai et al., 2019; Child et al., 2019; Qiu et al., 2020; Ainslie et al., 2020; Zaheer et al., 2020; Beltagy et al., 2020), learnable patterns (Kitaev et al., 2020; Tay et al., 2020a), low-rank matrix approximation (Wang et al., 2020), or kernel method (Choromanski et al., 2021). Alternatively, it has been shown that some attention heads are redundant and can be pruned to reduce model size (Voita et al., 2019; Michel et al., 2019). Knowledge distillation reduces memory and compute by compressing a large model to a smaller one (Hinton et al., 2015; Sanh et al., 2019). In contrast, we focus on the dependencies of long input and target sequences in encoder-decoder architectures, and we exploit publicly available transformer models with summarization weights to long-span summarization tasks.

Long-span Summarization. Efficient transformer architectures have been applied to summarize long documents such as BigBird (Zaheer et al., 2020), and Longformer-Encoder-Decoder (LED) (Beltagy et al., 2020), which has recently been revised parallel to this work.² Hierarchical transformer architectures have been applied to multi-document summarization (Liu and Lapata, 2019a), and extractive news and table-to-text summarization (Zhang et al., 2019; Narayan et al., 2020). Hierarchical attention RNN system has been applied to summarize long articles (Cohan et al., 2018).

Alternatively, earlier methods show that good content selection helps abstractive news summarization systems (Chen and Bansal, 2018; Gehrmann et al., 2018; Hsu et al., 2018). Hybrid systems that select sentences and generate an abstractive summary have been proposed such as extractive system + TLM for scientific articles (Pilault et al., 2020), simple selection + BART for podcasts (Manakul and Gales, 2020; Song et al., 2020), and guided summarization by BERT-based keyword/sentence extraction + BART for news and scientific articles (He et al., 2020; Dou et al., 2021).

Other work includes dividing the source and target into multiple smaller pairs to train abstractive summarizers (Gidiotis and Tsoumakas, 2020). Extractive methods with and without redundancy reduction techniques for long-span summarization have been studied (Xiao and Carenini, 2019, 2020).

²On the self-attention aspect, we believe this system is the most comparable to ours (see comparisons in Sec. 6.2).




Figure 1: Overview of the combined architecture where we highlight different aspects of this work. N_0 is the original document length, N is the input length to the generation system, and M is the summary length.

3 Experimental Setup

3.1 Dataset

Spotify Podcast.³ The dataset consists of ASR transcripts with human descriptions as summaries (Clifton et al., 2020). We follow the data processing at TREC2020 (Jones et al., 2020) in removing bad transcript-summary pairs from a total of 105,360+1,027 episodes, resulting in train/valid/test splits of 60,415/2,189/1,027 episodes the same as Manakul and Gales (2020).

arXiv and PubMed. Popular long document summarization datasets consist of academic articles with abstracts as summaries (Cohan et al., 2018) and train/valid/test splits of 203,037/6,436/6,440 for arXiv and 119,924/6,633/6,658 for PubMed.

Dataset	#Doc	Input	90 th %	Target
Podcast	106k	5,727	11,677	61.1
arXiv	216k	8,584	16,108	367
PubMed	133k	3,865	7,234	260

Table 1: Length Statistics (mean & 90th%-ile).

3.2 Models

BART and LoBART. We use the publicly released BART model (Lewis et al., 2020) fine-tuned on CNNDM (Hermann et al., 2015).⁴ Following the local window attention in Sparse Transformer (Child et al., 2019) and Longformer (Beltagy et al., 2020), we modify the self-attention mechanism in the encoder to local self-attention (see Figure 2), and we refer to this local self-attention BART as LoBART. It has the same architecture as BART, e.g. the number of parameters, except that we extend positional embedding beyond 1,024 by copying BART’s positional embedding with flipping to allow a smoother transition. See details in Appendix B.1.

³<https://podcastsdataset.byspotify.com>

⁴<https://huggingface.co/facebook/bart-large-cnn>




Figure 2: Self-Attention Pattern.

Hierarchical RNN. The content selection model is based on a hierarchical encoder-decoder architecture that has been shown effective on meeting and long document summarization (Cohan et al., 2018; Zhao et al., 2019; Li et al., 2019). The model consists of word-level and sentence-level GRUs (Cho et al., 2014). We add a linear layer on top of the sentence-level GRU to perform extractive labelling. The sentence-level attention mechanism and extractive labelling modules form our multitask content selection (MCS). More details in Section 5.2.

We provide the full details about our implementation, model parameters, hyperparameters, optimizer, and training configurations in Appendix B.

4 Longer Span via Local Self-Attention

It has been known that memory and compute complexity of transformers is *quadratic* with the sequence length. However, in encoder-decoder architectures, the exact dependencies on input length N , target length M , and batch size B are less understood. This is particularly important in long-span seq2seq tasks because large memory or compute requirement could make training impractical. Thus, this work studies these dependencies, and shows the trade-off between the size of input span and the size of attention span in local self-attention.

4.1 Memory Analysis and LoBART Design

Firstly, through a regression analysis for an encoder-decoder architecture such as BART, the

memory required in training is:

$$c_1^b + B(c_2^b M + c_3^b N + c_4^b MN + c_5^b M^2 + c_6^b N^2)$$

The term c_1^b depends on only the model size and optimizer, and it is *constant* (theoretical calculation provided in Appendix A). The remaining terms are activation memory associated with the activation outputs cached for backpropagation, and they grow with N , M , and B . Table 2 shows system-independent⁵ regression results for the memory in training BART. It is apparent that as N grows the dominant term is $c_6^b N^2$, which is associated with the encoder self-attention. Thus, this motivates us to modify self-attention only on the encoder side.

Term	c_1^b	$c_2^b M$	$c_3^b N$	$c_4^b MN$	$c_5^b M^2$	$c_6^b N^2$
GiB	6.05	0.23	0.84	0.21	0.02	1.53

Table 2: BART’s Memory Profile ($N=1024$, $M=144$).

By introducing local self-attention of width W , the memory in training LoBART becomes:

$$c_1^l + B(c_2^l M + c_3^l N + c_4^l MN + c_5^l M^2 + c_6^l NW)$$

For large N , the memory is now dominated by $c_6^l NW$. The coefficient $c_6^l \approx 1.72c_6^b$, suggesting that W should be at most $0.58N$ to reduce memory. We provide more details about the exact theoretical calculation for model and optimizer memory as well as time complexity in Appendix A.

The memory for training BART/LoBART in Figure 3 enables us to choose an operating point. Additionally, other *complementary* techniques for reducing memory in training include: (i) gradient-checkpoint where a subset of intermediate values in the computation graph are cached, and the rest are re-computed during backpropagation (Chen et al., 2016), but this requires changes to optimization and leads to longer training time; (ii) half/mixed-precision training (Micikevicius et al., 2018) that would almost halve y-axis in Figure 3, but this requires changes to the model precision and may result in lower performance; (iii) model parallelism with micro-batching (Huang et al., 2019), but this method requires multiple accelerators.

4.2 BART and LoBART

We study the characteristics of the full self-attention in BART by defining the mean attention

⁵system-independent across hardware and machines; albeit implementation-dependent. This analysis is based on widely used PyTorch and Huggingface implementation.




Figure 3: Operating points for $B=1$ and $M=144$. (1) Section 4 studies local attention to reduce quadratic complexity to linear. As W decreases, the gradient of linear complexity decreases. (2) Section 5 studies content selection to move an operating point to the left.

distance in a particular layer and head as follows:

$$D = \frac{1}{N} \sum_{i=1}^N \left(\sum_{j=1}^N \alpha_{i,j} \times |i - j| \right) \quad (1)$$

where $\alpha_{i,j}$ is the attention weight of position i attending to position j ($\sum_{j=1}^N \alpha_{i,j} = 1$). This measure corresponds to the average distance of self-attention. If the attention weight is uniform, $D_U = \frac{N^2-1}{3N}$. For $N = 1024$, $D_U = 341$. In Figure 4, our results show that most layers have a shorter mean distance than D_U , supporting that the information is more localized. The mean distances of differently initialized BART models computed on the podcast data also show that the attention mechanism is learned during pre-training stage as there is little variation after the pre-training stage. As illustrated in Figure 4, the average attention distance D of the BART model is around 250-350 tokens. This suggests the window size W should be designed to be above 700, allowing half local attention window $W/2$ be greater than 250-350 to effectively match BART and to exploit transfer learning more efficiently.

Subsequently, we train different configurations of BART/LoBART models up to our GPU memory limit of 32GiB. The results in Table 3 show that: (i) expanding the model to accommodate longer input spans improve over the baseline BART(1k) as opposed to Manakul and Gales (2020) that trained longer-span models by freezing bottom layers and did not show any improvement over their baseline; (ii) Although LoBART(8k) with $W=512$ can process longer input spans than LoBART(4k) with $W=1024$, it performs worse and we suggest that this is because LoBART(8k)’s window is too small,




Figure 4: The average mean distance across multi-heads for each layer. The average mean distance of the random weight model is slightly lower than D_U as some inputs are shorter than 1,024.

e.g. <700 , to utilize transfer learning efficiently and its effective receptive field is also smaller.

System	W	GiB	R1	R2	RL
BART(1k)	Full	8.9	26.43	9.22	18.35
LoBART(2k)	128	9.6	25.88	8.89	17.87
LoBART(2k)	256	10.2	25.93	8.80	17.82
LoBART(2k)	512	11.6	26.35	8.98	18.19
LoBART(2k)	1024	14.2	26.44	9.26	18.25
BART(2k)	Full	14.5	26.63	9.41	18.65
LoBART(4k)	128	12.8	26.42	9.02	18.12
LoBART(4k)	256	14.1	26.66	9.22	18.33
LoBART(4k)	512	16.7	26.75	9.54	18.54
LoBART(4k)	1024	22.0	27.02	9.57	18.78
LoBART(8k)	128	19.3	26.45	9.04	18.23
LoBART(8k)	256	21.1	26.72	9.30	18.36
LoBART(8k)	512	27.1	26.90	9.47	18.50

Table 3: BART & LoBART memory requirement in training and performance. (nk) denotes maximum input length of $n \times 1024$.

5 Longer Span via Content Selection

Some input sequences still exceed LoBART’s longer fixed-span limit. Further extending the input span would lead to a small local attention span, a diminishing improvement, or GPU running out of memory. Alternatively, it has been shown that a better content selection improves abstractive summarization in news (Chen and Bansal, 2018; Gehrmann et al., 2018; Hsu et al., 2018), multi documents (Liu and Lapata, 2019a; Liu et al., 2018), and scientific articles (Pilault et al., 2020). Thus, we propose to tackle the excess length by content selection. Here, we distinguish between two phases of content selection: training time and test time.

5.1 Training-time Content Selection

During training, ground-truth targets are available. We categorize selection methods in this phase into two types: ground-truth based (model-free), which is also referred to as *oracle*; and model-based. Ground-truth based methods cannot be used at test time, while model-based methods can be applied at both phases. Although model-based methods do not rely on ground-truth targets, they have the advantage of matching in training and test phases. Existing oracle methods include using ROUGE-2 recall (Liu et al., 2018) or the average of ROUGE-1,2,L recall (Pilault et al., 2020). We discuss model-based methods in Section 5.2, where we propose the MCS method. Let the subscript (i, j) denote the position of the j -th word in the i -th input sentence, the full input $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_i, \dots, \mathbf{x}_{N_1}\} = \underbrace{[x_{1,1}, x_{1,2}, x_{1,J_1}, \dots]}_{\text{sent 1}}, \underbrace{[x_{i,1}, x_{i,J_i}, \dots]}_{\text{sent } i}, \dots, \underbrace{[x_{N_1,1}, x_{N_1,J_{N_1}}]}_{\text{sent } N_1}$.

Content selection re-ranks, truncates, and sorts \mathbf{X} to get \mathbf{X}^{cs} for training BART/LoBART as follows:

$$\bar{\mathbf{X}} = \{\mathbf{x}_{r_1}, \mathbf{x}_{r_2}, \mathbf{x}_{r_3}, \dots, \mathbf{x}_{r_R}\} \quad (2)$$

$$\mathbf{X}^{\text{cs}} = \text{SortOrig}(\text{TruncateN}(\bar{\mathbf{X}})) \quad (3)$$

where r_i is the index of the sentence of rank i , the TruncateN operation filters $\bar{\mathbf{X}}$ such that the total of number of words is less than N , and SortOrig retains the original sentence order. The following ranking methods are considered:

- Truncation (TRC): $r_k = k$.
- Model-based: Given the score f of model ϕ , $r_k = \{i \in N_1 : f_\phi(i|\mathbf{X}) \text{ is ranked } k\text{-th}\}$
- Oracle (ORC): Given the ground-truth summary \mathbf{y} and similarity measure d , $r_k = \{i \in N_1 : d(\mathbf{x}_i, \mathbf{y}) \text{ is ranked } k\text{-th}\}$

In this work, we use ROUGE-2 recall as the similarity measure d . For the ORC method, first, we retain only sentences with positive d , leading to $R \leq N_1$. We found that the number of sentences with positive d is low at 21.3% of the total number of sentences in average on podcast data. This corresponds to 56% of training instances being shorter than BART input span of 1024.⁶ This no-padding oracle method (ORC_{no-pad}) is highly *aggressive*, potentially preventing the downstream summarizer

⁶We refer to this percentage as %AgORC_{no-pad} (the percentage of inputs aggressively extracted by the oracle method).

from learning complex abstraction. Hence, we propose variants of oracle methods to extend the $\text{ORC}_{\text{no-pad}}$ -selected input to the max input span N :

- $\text{ORC}_{\text{pad-lead}}$: Pad by leading unselected sentences and keep the original sentence order.
- $\text{ORC}_{\text{pad-rand}}$: Pad by random unselected sentences and keep the original sentence order.




Figure 5: The impact of training-time content selection methods on BART(1k) performance.

In Figure 5, since any oracle method is considered cheating at test time, the best performance is obtained by MCS (in blue), and the upper bound performance is obtained by optimal oracle method (in green). The results show that although $\text{ORC}_{\text{no-pad}}$ yields the highest upper bound, the abstractive model in fact does not learn how to perform abstraction. For instance, with TRC or MCS at test time, $\text{ORC}_{\text{no-pad}}$ yields the lowest performance level. The best way to fine-tune the abstractive model shown in Figure 5 is using $\text{ORC}_{\text{pad-rand}}$. Compared to $\text{ORC}_{\text{pad-lead}}$, $\text{ORC}_{\text{pad-rand}}$ is better as it introduces more diversity to the abstractive model. Compared to the model-based method, $\text{ORC}_{\text{pad-rand}}$ is also computationally less expensive.

In addition, Table 5 shows that when there is no content selection at test time (i.e. TRC applied), LoBART(4k) and LoBART(8k) benefit from $\text{ORC}_{\text{pad-rand}}$, whereas BART(1k) does not. This is because in the 1k setting, content selection is more aggressive; as a result, the large mismatch between training and test leads to a poor result. Thus, we suggest that the best content selection during training is $\text{ORC}_{\text{pad-rand}}$ given that content selection will be used at test time, or model’s input span is long.

5.2 Multitask Content Selection (MCS)

To process long input sequences entirely, we consider RNN, whose memory requirement grows lin-

early with the sequence length, and hierarchical architectures which have been shown effective for long seq2seq tasks (Cohan et al., 2018; Li et al., 2019). In this work, the hierarchical RNN model described in Section 3.2 has memory requirement given the target length of 144 during training of $0.83 + B(3.96 \times 10^{-5} + 3.33 \times 10^{-5} N_2) N_1$,⁷ where N_1 is #sentences, and N_2 is the maximum number of words in a sentence, and B is batch size. By setting $N_1=1000$ and $N_2=50$, only 2% of podcast data exceeds this limit, while taking GPU memory to only 2.53GiB for $B=1$. Thus, this shows that this model can cover long sequences.

Previous model-based methods treat content selection as extractive labelling and create labels heuristically (Pilault et al., 2020), or using encoder-decoder attention mechanism (Manakul and Gales, 2020). To utilize both of these in one framework, we propose a Multitask Content Selection (MCS) method where we train the hierarchical encoder-decoder with attention mechanism and a classification layer on top of the encoder (described in Section 3.2). First, the model is trained on seq2seq abstractive summarization objective:

$$\mathcal{L}_{\text{seq2seq}} = - \sum_{m=1}^M \log P(y_m | \mathbf{y}_{<m}, \mathbf{X}) \quad (4)$$

Second, we create binary labels as follows: for sentence i , the label z_i is 1 if $d(\mathbf{x}_i, \mathbf{y}) > 0$; else z_i is 0, and d is the ROUGE-2 recall measure. The extractive labelling task objective is:

$$\mathcal{L}_{\text{label}} = - \sum_{i=1}^{N_1} (z_i \log \hat{z}_i + (1 - z_i) \log(1 - \hat{z}_i)) \quad (5)$$

$$\hat{z}_i = \text{sigmoid}(\mathbf{W}_{\text{cls}}^T \mathbf{h}_i + \mathbf{b}_{\text{cls}}) \quad (6)$$

where \mathbf{h}_i is the sentence-level encoder output associated with sentence i , and $\mathbf{W}_{\text{cls}}, \mathbf{b}_{\text{cls}}$ are the parameters of the classification layer. Thus, the MCS training loss is defined as follows:

$$\mathcal{L}_{\text{MCS}} = \gamma \mathcal{L}_{\text{label}} + (1 - \gamma) \mathcal{L}_{\text{seq2seq}} \quad (7)$$

At inference stage, there are two modes: (i) standard abstractive summary generation, e.g. via beam search decoding; (ii) ranking input sentences via labelling score and seq2seq attention score. The latter is how we use MCS during inference.⁸ For sentence i , the scores are:

$$\text{score}_{i,(\text{label})} = \hat{z}_i, \text{score}_{i,(\text{seq2seq})} = \sum_{m=1}^M \alpha_{m,i}^s \quad (8)$$

⁷Obtained by least-squares regression with 20 samples.

⁸In practice, we run beam search decoding of width 4, and we obtain the attention score from the top beam.

where $\alpha_{m,i}^s$ is the sentence-level attention weight at decoder step m over input sentence i . Since the scores are on different scales, rather than using the scores defined in Eq. 8, we simply rank the scores, and then normalize the score ranks into the range 0.0 to 1.0. Let nscore denote the normalized ranking score, the MCS inference score is:

$$f_\phi(i|\mathbf{X}) = \text{nscore}_{i,(\text{label})} + \text{nscore}_{i,(\text{seq2seq})} \quad (9)$$

In our preliminary experiments, we vary the amount of selected sentences from the limit of BART/LoBART to a few sentences, and we found that more aggressive selection at test time degrades the performance. Therefore, our MCS selects input sentences up to the limit of BART/LoBART.

By setting $\gamma=0.0$, our method is comparable to the attention-based method in Manakul and Gales (2020). By setting $\gamma=1.0$, our method is similar to the extractive models in Hsu et al. (2018); Pilault et al. (2020). In Table 4, we show that when coupled with BART, MCS yields better summarization performance than both Attn-only and Ext-only baselines. MCS also achieves higher recall rate of sentences with $d(\mathbf{x}_i, \mathbf{y}) > 0$ than the two baselines.

System	%Recall	R1	R2	RL
Attn ($\mathcal{L}_{\text{seq2seq}}$)	38.85	26.90	9.70	18.78
Ext ($\mathcal{L}_{\text{label}}$)	35.26	26.39	8.90	18.03
MCS (\mathcal{L}_{MCS})	40.50	27.28	9.82	19.00

Table 4: The impact of test-time content selection methods on BART(1k) trained using $\text{ORC}_{\text{pad-rand}}$. Optimal $\gamma=0.2$ is tuned between 0.0-1.0 on the validation set.

6 Combined Approach

6.1 Spotify Podcast results

In Table 5, a performance gain is obtained in all settings by adding MCS. By comparing different configurations with MCS, it can be seen that the gain from MCS in LoBART(8k) system is the lowest. This is because the average length is 5,727, meaning that many Podcasts inputs to LoBART(8k) do not benefit from content selection.

CUED-filt, the best single-model system in Manakul and Gales (2020), uses an attention-based content selection at both training and test time, and it is combined with fine-tuned vanilla BART. Our approach outperforms CUED-filt by improved content selection at both training time and test time as

demonstrated by BART(1k)-ORC+MCS. Additionally, local self-attention allows training on longer sequences, and our LoBART(4k)-ORC+MCS system has yielded the best results. Lastly, even though LoBART(8k) requires more resource to train, it does not perform as well as LoBART(4k) due to its smaller attention window, and it also has a lower improvement when adding MCS.

System	CS-trn	CS-tst	R1	R2	RL
CUED-filt*	✓	✓	26.96	9.75	18.90
BART(1k)	✗	✗	26.43	9.22	18.35
BART(1k)	✗	MCS	26.82	9.39	18.57
BART(1k)	ORC	✗	25.54	9.00	17.83
BART(1k)	ORC	MCS	27.28	9.82	19.00
LoBART(4k)	✗	✗	27.02	9.57	18.78
LoBART(4k)	✗	MCS	27.53	9.95	19.08
LoBART(4k)	ORC	✗	27.36	10.04	19.33
LoBART(4k)	ORC	MCS	27.81	10.30	19.61
LoBART(8k)	✗	✗	26.90	9.47	18.50
LoBART(8k)	✗	MCS	27.02	9.52	18.62
LoBART(8k)	ORC	✗	27.16	9.84	19.08
LoBART(8k)	ORC	MCS	27.49	9.98	19.25

Table 5: Podcast Results. The impact of training-time $\text{ORC}_{\text{pad-rand}}$ and test-time MCS. *CUED systems were the top systems by human evaluation at Spotify Challenge 2020; CUED systems use BART with a model-based (trained on $\mathcal{L}_{\text{seq2seq}}$) content selection in both training and test stages.

6.2 ArXiv and PubMed results

To verify the effectiveness of our systems, we re-train BART(1k) and LoBART(4k) on arXiv and PubMed datasets. Our training is different from Ext+TLM (Pilault et al., 2020) where their abstractive models are trained using inputs extracted from top two sentences in ROUGE recall for each target sentence without padding, similar to $\text{ORC}_{\text{no-pad}}$. Although in 1k setting, $\text{ORC}_{\text{no-pad}}$ yields $\% \text{AgORC}_{\text{no-pad}}$ (defined in Section 5.1) of only 2.8% on arXiv (12% on PubMed), in 4k setting this is 39% on arXiv (71% on PubMed). Based on the best configurations on podcast data, we train BART(1k) and LoBART(4k) using TRC or $\text{ORC}_{\text{pad-rand}}$ content selection, and we train the hierarchical model on arXiv/PubMed for MCS.

ArXiv. In Table 6, both BART(1k)+MCS and LoBART(4k)+MCS outperform all existing systems. To better understand the advantages of our approach, the following systems are compared:

	Type	System	arXiv			PubMed		
			R1	R2	RL	R1	R2	RL
Previous Work	Abs	Discourse-Aware (Cohan et al., 2018)	35.80	11.05	31.80	38.93	15.37	35.21
	Mix	Ext+TLM (Pilault et al., 2020)	41.62	14.69	38.03	42.13	16.27	39.21
	Ext	ExtSum-LG+Rd(Xiao and Carenini, 2020)	44.01	17.79	39.09	45.30	20.42	40.95
	Abs	Pegasus (Zhang et al., 2020)	44.21	16.95	38.83	45.97	20.15	41.34
	Abs	DANCER (Gidiotis and Tsoumakos, 2020)	45.01	17.60	40.56	46.34	19.97	42.42
	Abs	BigBird(3k) (Zaheer et al., 2020)	46.63	19.02	41.77	46.32	20.65	42.33
	Abs	LED(4k) (Beltagy et al., 2020)	44.40	17.94	39.76	-	-	-
	Abs	LED(16k) (Beltagy et al., 2020)	46.63	19.62	41.83	-	-	-
	Mix	CTRLsum(BART+BERT) (He et al., 2020)	46.91	18.02	42.14	-	-	-
This Work	Abs	[†] BART(1k)	44.96	17.25	39.76	45.06	18.27	40.84
	Mix	[‡] BART(1k)+MCS	47.68	19.77	42.25	46.49	19.45	42.04
	Abs	[‡] LoBART(4k)	46.59	18.72	41.24	47.47	20.47	43.02
	Mix	[‡] LoBART(4k)+MCS	48.79	20.55	43.31	48.06	20.96	43.56

Table 6: Results on arXiv and PubMed. [†]denotes TRC applied, and [‡]denotes ORC_{pad-rand} applied at training time.

CTRLsum versus our BART(1k) baseline; LED and BigBird versus our LoBART(4k) system.

CTRLsum extends BART by conditioning it with extracted keywords \mathbf{v} using a BERT-based model, e.g. $p(\mathbf{y}|\mathbf{X}, \mathbf{v})$. Their BERT-based model uses sliding window allowing it to extract \mathbf{v} in long sequences, but their BART is still limited to the first 1,024 tokens. As a result, it performs better than BART(1k), but worse than BART(1k)+MCS.

LoBART(4k) has a similar architecture to LED(4k) without the global attention pattern for special tokens. Instead, our LoBART(4k) benefits from knowledge transferred from CNNDM and the ORC_{pad-rand} training-time content selection, which yields a larger gain when MCS is applied, i.e. the system trained with truncated data has a smaller gain when MCS is applied. Transfer learning comparison and additional results on the impact of ORC_{pad-rand} are provided in Appendix C.


Compared to BigBird, LoBART(4k) has a longer input span, e.g. 3,072 vs. 4,096. However, BigBird benefits from utilizing more recent summarization specific pre-training Pegasus (Zhang et al., 2020) which is better than our transfer learning. BigBird incorporates a global attention pattern similar to LED, and it also has a random attention pattern. Hence, LoBART without MCS performs worse.

Ultimately, we show that adding MCS to either BART(1k) or LoBART(4k) yields a significant improvement, resulting in state-of-the-art results in both settings. Moreover, although the gain from adding MCS is comparable to the gain observed in extending LED(4k) to LED(16k), the content selection method adds less training cost.


PubMed. Similarly, LoBART(4k)+MCS achieves state-of-the-art results shown in Table 6. In contrast to the arXiv results, BART(1k)+MCS does not outperform LoBART(4k) nor BigBird, and the gain from MCS is not as high in both 1k and 4k settings.

6.3 Local Attention v.s. MCS.

Local attention yields better performance on PubMed, while MCS yields better performance on arXiv. To understand this discrepancy, a fine-grained analysis is conducted.



(a) arXiv (Len:Avg=8,584, 90th%=16,108)



(b) PubMed (Len:Avg=3,865, 90th%=7,234)

Figure 6: ROUGE-1 score relative to that of BART(1k) system evaluated on different partitions by length.

In Figure 6, we partition the test sets by input lengths, and we evaluate the performance improvement in each partition with respect to the BART(1k) baseline.⁹ The results illustrate that as the input length N increases:

- The improvement of systems *with* MCS increases and subsequently plateaus out.
- The improvement of systems *without* MCS decreases once the input exceeds the length limit but then plateaus, suggesting that fixed-span systems without content selection perform worse once the maximum fixed-span is reached. For instance, below 4,000 input words, LoBART(4k) without MCS performs better than BART(1k)+MCS on both datasets.

Therefore, our MCS method is more effective on arXiv compared to PubMed because the average length of PubMed documents is more than twice shorter than the average length of arXiv documents.

7 Conclusion

We study two methods for long-span summarization tasks. First, on local self-attention transformers, we present the design considerations for local self-attention BART, and we investigate the feasibility and performance of different network configurations. Second, on content selection, we distinguish between training time and test time methods, and we provide a good practice for both phases. At training time, we show that the oracle method with random sentences padded (ORC_{pad-rand}) yields the best results. At test time, we propose multitask content selection (MCS) that shows an improvement over baselines. We demonstrate that content selection is essential, in particular for longer documents such as the articles in the arXiv dataset. Our BART(1k)+MCS outperforms the current best systems on Podcast and arXiv datasets, and this system does not require a large-scale accelerator in training. Ultimately, by combining local self-attention technique with MCS, our LoBART(4k)+MCS system has set new state-of-the-art results in terms of ROUGE scores in all three long-span summarization tasks. Future work will focus on training our LoBART+MCS system in an end-to-end fashion.

⁹For arXiv/PubMed, each test set consists of over 6,000 instances, while Podcast test set has only 1,027 instances. The same analysis is conducted on Podcast, but the results are noisy due to the smaller size of its test set (see Appendix C).

Acknowledgments

This paper reports on research supported by ALTA institute, Cambridge Assessment English, University of Cambridge, and Cambridge International & St John’s College Scholarship. Thanks to Yiting Lu, Qingyun Dou, Xixin Wu, Raf Czlonka, and Kate Knill for interesting discussions and computing resource support. Thanks to the anonymous reviewers for their helpful comments.

References

- Joshua Ainslie, Santiago Ontanon, Chris Alberti, Václav Cvicek, Zachary Fisher, Philip Pham, Anirudh Ravula, Sumit Sanghai, Qifan Wang, and Li Yang. 2020. [ETC: Encoding long and structured inputs in transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 268–284, Online. Association for Computational Linguistics.
- Iz Beltagy, Matthew E Peters, and Arman Cohan. 2020. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Tianqi Chen, Bing Xu, Chiyuan Zhang, and Carlos Guestrin. 2016. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*.
- Yen-Chun Chen and Mohit Bansal. 2018. [Fast abstractive summarization with reinforce-selected sentence rewriting](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 675–686, Melbourne, Australia. Association for Computational Linguistics.
- 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 (Volume 1: Long Papers)*, pages 484–494, Berlin, Germany. Association for Computational Linguistics.

- Rewon Child, Scott Gray, Alec Radford, and Ilya Sutskever. 2019. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. [Learning phrase representations using RNN encoder–decoder for statistical machine translation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar. Association for Computational Linguistics.
- Krzysztof Marcin Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, David Benjamin Belanger, Lucy J Colwell, and Adrian Weller. 2021. [Rethinking attention with performers](#). In *International Conference on Learning Representations*.
- Ann Clifton, Sravana Reddy, Yongze Yu, Aasish Pappu, Rezvaneh Rezapour, Hamed Bonab, Maria Eskevich, Gareth Jones, Jussi Karlgren, Ben Carterette, and Rosie Jones. 2020. [100,000 podcasts: A spoken English document corpus](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 5903–5917, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- 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)*, pages 615–621, New Orleans, Louisiana. Association for Computational Linguistics.
- Zihang Dai, Zhilin Yang, Yiming Yang, Jaime Carbonell, Quoc Le, and Ruslan Salakhutdinov. 2019. [Transformer-XL: Attentive language models beyond a fixed-length context](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2978–2988, Florence, Italy. Association for Computational Linguistics.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [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, Minneapolis, Minnesota. Association for Computational Linguistics.
- Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang, and Graham Neubig. 2021. [GSum: A general framework for guided neural abstractive summarization](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4830–4842, Online. Association for Computational Linguistics.
- Sebastian Gehrmann, Yuntian Deng, and Alexander Rush. 2018. [Bottom-up abstractive summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4098–4109, Brussels, Belgium. Association for Computational Linguistics.
- Alexios Gidiotis and Grigorios Tsoumakas. 2020. [A divide-and-conquer approach to the summarization of long documents](#). *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:3029–3040.
- Junxian He, Wojciech Kryściński, Bryan McCann, Nazneen Rajani, and Caiming Xiong. 2020. CTRL-sum: Towards generic controllable text summarization. *arXiv preprint arXiv:2012.04281*.
- Karl Moritz Hermann, Tomáš Kociský, 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.
- Geoffrey Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). In *NIPS Deep Learning and Representation Learning Workshop*.
- Wan-Ting Hsu, Chieh-Kai Lin, Ming-Ying Lee, Kerui Min, Jing Tang, and Min Sun. 2018. [A unified model for extractive and abstractive summarization using inconsistency loss](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 132–141, Melbourne, Australia. Association for Computational Linguistics.
- Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Xu Chen, HyoukJoong Lee, Jiquan Ngiam, Quoc V. Le, Yonghui Wu, and Zhifeng Chen. 2019. [Gpipe: Efficient training of giant neural networks using pipeline parallelism](#). In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 103–112.
- Rosie Jones, Ben Carterette, Ann Clifton, Maria Eskevich, Gareth J. F. Jones, Jussi Karlgren, Aasish Pappu, Sravana Reddy, and Yongze Yu. 2020. Trec 2020 podcasts track overview. In *The 29th Text Retrieval Conference (TREC) notebook*.
- Diederik P. Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *3rd International Conference on Learning Representations*,

- ICLR 2015, San Diego, CA, USA, May 7-9, 2015, *Conference Track Proceedings*.
- Nikita Kitaev, Lukasz Kaiser, and Anselm Levskaya. 2020. [Reformer: The efficient transformer](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Logan Lebanoff, Franck Dernoncourt, Doo Soon Kim, Walter Chang, and Fei Liu. 2020. [A cascade approach to neural abstractive summarization with content selection and fusion](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 529–535, Suzhou, China. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 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*, pages 7871–7880, Online. Association for Computational Linguistics.
- Manling Li, Lingyu Zhang, Heng Ji, and Richard J. Radke. 2019. [Keep meeting summaries on topic: Abstractive multi-modal meeting summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 2190–2196, Florence, Italy. Association for Computational Linguistics.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Peter J. Liu, Mohammad Saleh, Etienne Pot, Ben Goodrich, Ryan Sepassi, Lukasz Kaiser, and Noam Shazeer. 2018. [Generating wikipedia by summarizing long sequences](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Yang Liu and Mirella Lapata. 2019a. [Hierarchical transformers for multi-document summarization](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5070–5081, Florence, Italy. Association for Computational Linguistics.
- Yang Liu and Mirella Lapata. 2019b. [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)*, pages 3730–3740, Hong Kong, China. Association for Computational Linguistics.
- Potsawee Manakul and Mark Gales. 2020. [CUED_speech at TREC 2020 podcast summarisation track](#). *arXiv preprint arXiv:2012.02535*.
- Potsawee Manakul, Mark J.F. Gales, and Linlin Wang. 2020. [Abstractive Spoken Document Summarization Using Hierarchical Model with Multi-Stage Attention Diversity Optimization](#). In *Proc. Interspeech 2020*, pages 4248–4252.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. [Are sixteen heads really better than one?](#) In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 14014–14024.
- Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory F. Damos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, and Hao Wu. 2018. [Mixed precision training](#). In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Conference Track Proceedings*. OpenReview.net.
- Ramesh Nallapati, Bowen Zhou, Cicero dos Santos, Çağlar Gulçehre, and Bing Xiang. 2016. [Abstractive text summarization using sequence-to-sequence RNNs and beyond](#). In *Proceedings of The 20th SIGNLL Conference on Computational Natural Language Learning*, pages 280–290, Berlin, Germany. Association for Computational Linguistics.
- Shashi Narayan, Joshua Maynez, Jakub Adamek, Daniele Pighin, Blaz Bratanić, and Ryan McDonald. 2020. [Stepwise extractive summarization and planning with structured transformers](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4143–4159, Online. Association for Computational Linguistics.
- Niki Parmar, Ashish Vaswani, Jakob Uszkoreit, Lukasz Kaiser, Noam Shazeer, Alexander Ku, and Dustin Tran. 2018. [Image transformer](#). In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholm, Sweden, July 10-15, 2018*, volume 80 of *Proceedings of Machine Learning Research*, pages 4052–4061. PMLR.
- Jonathan Pilault, Raymond Li, Sandeep Subramanian, and Chris Pal. 2020. [On extractive and abstractive neural document summarization with transformer language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 9308–9319, Online. Association for Computational Linguistics.
- Jiezhong Qiu, Hao Ma, Omer Levy, Wen-tau Yih, Sinong Wang, and Jie Tang. 2020. [Blockwise self-attention for long document understanding](#). In *Findings of the Association for Computational Linguistics*.

- tics: *EMNLP 2020*, pages 2555–2565, Online. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of Machine Learning Research*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- 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)*, pages 1073–1083, Vancouver, Canada. Association for Computational Linguistics.
- Kaiqiang Song, Chen Li, Xiaoyang Wang, Dong Yu, and Fei Liu. 2020. Automatic summarization of open-domain podcast episodes. *arXiv preprint arXiv:2011.04132*.
- Yi Tay, Dara Bahri, Liu Yang, Donald Metzler, and Da-Cheng Juan. 2020a. Sparse sinkhorn attention. In *Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event*, volume 119 of *Proceedings of Machine Learning Research*, pages 9438–9447. PMLR.
- Yi Tay, Mostafa Dehghani, Samira Abnar, Yikang Shen, Dara Bahri, Philip Pham, Jinfeng Rao, Liu Yang, Sebastian Ruder, and Donald Metzler. 2021. Long range arena : A benchmark for efficient transformers. In *International Conference on Learning Representations*.
- Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. 2020b. Efficient transformers: A survey. *arXiv preprint arXiv:2009.06732*.
- 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 *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 5998–6008.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5797–5808, Florence, Italy. Association for Computational Linguistics.
- Sinong Wang, Belinda Li, Madian Khabsa, Han Fang, and Hao Ma. 2020. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. Transformers: State-of-the-art natural language processing. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Wen Xiao and Giuseppe Carenini. 2019. Extractive summarization of long documents by combining global and local context. 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)*, pages 3011–3021, Hong Kong, China. Association for Computational Linguistics.
- Wen Xiao and Giuseppe Carenini. 2020. Systematically exploring redundancy reduction in summarizing long documents. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 516–528, Suzhou, China. Association for Computational Linguistics.
- Manzil Zaheer, Guru Guruganesh, Kumar Avinava Dubey, Joshua Ainslie, Chris Alberti, Santiago Ontanon, Philip Pham, Anirudh Ravula, Qifan Wang, Li Yang, et al. 2020. Big bird: Transformers for longer sequences. *Advances in Neural Information Processing Systems*, 33.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter Liu. 2020. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *International Conference on Machine Learning*, pages 11328–11339. PMLR.
- Xingxing Zhang, Furu Wei, and Ming Zhou. 2019. HiBERT: Document level pre-training of hierarchical bidirectional transformers for document summarization. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5059–5069, Florence, Italy. Association for Computational Linguistics.
- Zhou Zhao, Haojie Pan, Changjie Fan, Yan Liu, Linlin Li, and Min Yang. 2019. Abstractive meeting summarization via hierarchical adaptive segmental network learning. In *The World Wide Web Conference, WWW 2019, San Francisco, CA, USA, May 13-17, 2019*, pages 3455–3461. ACM.

A Detailed Memory & Time Analysis

Our memory analysis is system-independent, albeit implementation-dependent. We carry out the experiments using PyTorch version 1.2.0. We use `pytorch_memlab`¹⁰ to compute GPU memory during forward and backward passes. Our notation is: input length N , target length M , local self-attention width W , and batch size B .

A.1 BART Memory

We collect 30 samples, spanning $N \in [64, 3000]$ and $M \in [36, 576]$ using batch size of 1. Our least-squared regression of the memory equation $\text{memory} = c_1^b + B(c_2^b M + c_3^b N + c_4^b MN + c_5^b M^2 + c_6^b N^2)$ yields $R^2 = 1$, $\text{RMSE} = 0.026$, and the coefficients are: $c_1^b = 6.054$, $c_2^b = 1.594 \times 10^{-3}$, $c_3^b = 8.192 \times 10^{-4}$, $c_4^b = 1.418 \times 10^{-6}$, $c_5^b = 1.077 \times 10^{-6}$, $c_6^b = 1.456 \times 10^{-6}$.

Model and Optimizer

The *constant* term $c_1^b = 6.054$ GiB is independent of batch size, system, or implementation (given the same floating-point precision). This term comprises model and optimizer memory as follows (in 32-bit floating point, 1 variable takes 4 bytes):

1. Model Parameter: BART has 406,290,432 parameters, yielding $406290432 \times 4 = 1.625 \times 10^9$ bytes = 1.51 GiB.
2. Model Gradient: Each parameter has one corresponding gradient variable, e.g. `.grad` in PyTorch. Thus, this also occupies 1.51 GiB.
3. Optimizer: Adam optimizer (Kingma and Ba, 2015) stores first moment and second moment for each and every model parameters, hence, taking 3.02 GiB.

Activation

The terms corresponding to c_2^b, \dots, c_6^b are associated with activation buffers cached for computing gradients in backpropagation. These terms grow linearly with batch size. The dominant term $c_6^b N^2 B$ grows quadratically with the input length N , motivating encoder’s local self-attention design.

Chen et al. (2016) proposes a method to save the activation memory by only caching buffers of a subset of layers, and re-computing the rest dynamically during backpropagation. This results in repeated computations and more training time.

¹⁰https://github.com/Stonesjtu/pytorch_memlab

A.2 LoBART Memory

We collect 36 samples, spanning $N \in [512, 4096]$, $M \in [100, 400]$, and $W \in [32, 512]$ using batch size of 1. Our least-squared regression of the memory equation $\text{memory} = c_1^l + B(c_2^l M + c_3^l N + c_4^l MN + c_5^l M^2 + c_6^l NW)$ yields $\text{RMSE} = 0.010$, and the coefficients are: $c_1^l = 6.104$, $c_2^l = 1.443 \times 10^{-3}$, $c_3^l = 1.032 \times 10^{-3}$, $c_4^l = 1.487 \times 10^{-6}$, $c_5^l = 1.277 \times 10^{-6}$, $c_6^l = 2.503 \times 10^{-6}$. The model and optimizer memory is similar to the analysis for BART. The activation memory is now dominated by $c_6^l NW \times B$, where $c_6^l = 1.72c_6^b$. Thus, we highlight that once $W > 0.58N$, LoBART no longer reduces memory. Note that we also tried incorporating the terms N^2 and W in the least-squared regression analysis, but their resulting coefficients are small, making both terms negligible. This is expected as quadratic self-attention is replaced by local attention of width W , and the width W only determines the receptive field of each and every position in N , resulting in the NW term.

A.3 Time: BART & LoBART

Unlike memory, time requirement is both system and implementation dependent. In this analysis, we show the results on our infrastructure consisting of a 32 GiB V100 GPU and 32-core Intel Xeon 4215R CPU (3.20GHz). We compute the time required for 50 forward and backward passes in 12 settings for each model configuration. Similar to the memory analysis, we perform least-squared regression where the results are shown in Figure 7. It can be seen that although LoBART reduces memory requirement, when it comes to time requirement, LoBART is only comparable to BART. This is due to the implementation of local self-attention that involves more processes such as chunking.




Figure 7: Quadratic time complexity of BART and Linear time complexity of LoBART for $M=144$ and $B=1$.

B Implementation Details

B.1 Models

BART & LoBART.

We use publicly released BART-large.¹¹ For LoBART, our local self-attention is based on Hugging-Face’s implementation (Wolf et al., 2020).¹² The number of parameters in BART is 406M.

The positional embedding of LoBART beyond 1,024 is created by copying BART’s positional embedding with flipping to allow a smoother transition as shown in Figure 8, and the number of parameters in LoBART(nk) is $406M + 50,264 \times (n-1) \times 1,024$.




Figure 8: LoBART positional embedding is initialized by copying and flipping BART’s positional embedding.

Hierarchical RNN.

The encoder consists of word-level and sentence-level bidirectional GRUs. The word-level GRU takes embedding vector $e_{i,j}$ of word i in sentence j , and outputs forward representation $h_{i,j}^{(f)}$ and backward representation $h_{i,j}^{(b)}$. The sentence-level GRU takes concatenated vector $[h_{N_j,j}^{(f)}; h_{1,j}^{(b)}]$, and outputs sentence representation h_j . The decoder consists of a unidirectional GRU. Each of the encoder GRUs has 2 layers with a dropout layer ($p=0.1$), and the decoder GRU has 1 layer. There are word-level and sentence-level attention mechanisms connecting the encoder and decoder. The classification head is a single-layer feedforward layer. The dimension of embedding space is 256, and the hidden size is 512. The number of parameters is 52M.

B.2 Training & Inference Hyperparameters

We process data using the same byte-pair-encoding tokenizer as the BART-large tokenizer, and we use NLTK tokenizer for sentence splitting. We use 32-bit precision training. We stop training when the loss on the validation set stop improving for 3 times. For example, the training steps are approximately:

¹¹<https://huggingface.co/facebook/bart-large-cnn>

¹²<https://huggingface.co/transformers/>

180k for Podcast; 240k for arXiv; 160k for PubMed. We report the validation performance when training is stopped in Table 10. Adam optimizer is used for all experiments with learning rate:

$$lr = 0.002 \times \min(\text{step}^{-0.5}, \text{step} \times \text{warmup}^{-1.5})$$

Parameter	Podcast	arXiv/PubMed
max. tgt len M	144	400
dropout	0.1	0.1
batch size	1	1
gradient accum.	2	2
warmup	10,000	20,000
valid step	20,000	20,000
loss		cross entropy
compute (BART)	1 × GTX TITAN X (12GiB)	
compute (LoBART)	1 × V100 (32GiB)	

Table 7: BART/LoBART Training Hyperparameters.

Parameter	Podcast	arXiv/PubMed
max. src #sent	1000	640
max. src #words-in-sent	50	120
max. tgt len M	144	400
dropout	0.1	0.1
batch size	2	2
gradient accum.	1	1
warmup	20,000	20,000
valid step	20,000	20,000
loss*		$\mathcal{L}_{\text{seq2seq}}$ & \mathcal{L}_{ext}
compute	1 × GTX TITAN X (12GiB)	

Table 8: RNN Training Hyperparameters. *Both loss functions are cross entropy based.

Parameter	Value
beam width	4
length penalty	2.0
min length	56
max length*	144 & 400
no repeat trigram size	3

Table 9: Inference Hyperparameters. *144 for Podcast, and 400 for arXiv/PubMed.

B.3 Evaluation

Our ROUGE (Lin, 2004) scoring tool is pyrouge, which is a wrapper for perl script.¹³

¹³<https://pypi.org/project/pyrouge/>

System	Attn-Width	CS-train	Podcast	arXiv	PubMed
BART(1k)	Full	Truncate	2.767	2.179	1.867
LoBART(4k)	1024	Truncate	2.680	1.878	1.530
*LoBART(4k)	1024	ORC _{pad-rand}	2.647	1.721	1.474

Table 10: Performance measured by the average cross-entropy on validation set. *Best system on the test set.

System	CS-train	CS-test	arXiv			PubMed		
			R1	R2	RL	R1	R2	RL
BART(1k)	✗	✗	44.96	17.25	39.76	45.06	18.27	40.84
BART(1k)	✗	MCS	46.11	18.79	40.83	46.46	19.54	41.91
BART(1k)	ORC	✗	42.03	15.62	37.15	43.20	17.02	39.19
BART(1k)	ORC	MCS	47.68	19.77	42.25	46.49	19.45	42.04
LoBART(4k)	✗	✗	46.90	18.88	41.50	47.40	20.43	42.95
LoBART(4k)	✗	MCS	48.05	20.11	42.58	47.76	20.76	43.27
LoBART(4k)	ORC	✗	46.59	18.72	41.24	47.47	20.47	43.02
LoBART(4k)	ORC	MCS	48.79	20.55	43.31	48.06	20.96	43.56

Table 11: Extended results on arXiv and PubMed (in Table 6). ORC is ORC_{pad-rand} training-time content selection.

System	Initialization	R1	R2	RL
BART(1k)	Random	14.61	0.82	11.54
	BART-large	25.82	9.07	17.99
	BART-large-CNNNDM	26.43	9.22	18.35

Table 12: Podcast results. The impact of transfer learning. Truncation is applied at both training and test stages.

System	CS-train	Initialization	R1	R2	RL
LED(4k)	Truncate	*BART-large	44.40	17.94	39.76
LoBART(4k)	Truncate	BART-large	46.17	17.96	40.74
	Truncate	BART-large-CNNNDM	46.90	18.88	41.50
	ORC _{pad-rand}	BART-large	45.25	17.40	39.96
	ORC _{pad-rand}	BART-large-CNNNDM	46.59	18.72	41.24

Table 13: arXiv results. The impact of transfer learning on initializing LoBART. At test time, there is *no* content selection. *To our understanding, LED-large was initialized from BART-large as described in [Beltagy et al. \(2020\)](#).

C Additional Results

Losses on Validation Sets

In Table 10, we show the standard cross entropy losses on validation sets of our BART/LoBART.

BART and LoBART on arXiv/PubMed

In Table 11, we provide configurations in addition to Table 6. These results (as well as Podcast results in Table 5) show that: in all settings, applying MCS at test time yields a performance gain; and with ORC applied at training, a larger gain is observed.

Transfer Learning from CNN/DailyMail

In Table 12, we show the impact of transfer learning on fine-tuning BART to Podcast. In Table 13, LED(4k) should be very close to LoBART(4k)-TRC-BART-large, we believe that the performance difference is due to the stochastic nature of training.

Nevertheless, our experiments are carried out using the same training setting, e.g. hyperparameters, optimizer, etc. Thus, based on the results, we believe that there is an observable improvement due to transfer learning from CNNNDM.

Fine-grained analysis on Podcast test set




Figure 9: ROUGE-1 score relative to that of BART(1k) on Spotify Podcast (Len: Avg=5,727, 90th=11,677).

Reference Summary: This week, Irwin and I discuss the iconic designs of the Rolex Submariner and the Porsche 911. Remaining subjectively unchanged through the years. We talk about the subtle changes over the years for this special car and watch and what are the similarities from their history. If you love cars and watches you'll enjoy the show. Grail Watch and Grail Porsche: Irwin: 1959 Rolex Submariner 5512 x 1967 Porsche 911s Anthony: 2020 Moser Streamliner Chronograph x 1988 Porsche 911 959.

LoBART(4k)+MCS: In this episode we talk to @grandobsessionsirwin about his passion for Porsche and vintage Rolex. We talk about the Submariner and 911 Submariners, the design evolution over the decades and the future of Porsche in the world of watchmaking. We also chat about the evolution of the Rolex submariner from the 1950s to the present day.

Table 14: Spotify Podcast test set sample #31.

Reference Summary: we present data from our investigation of the anomalous orange - colored afterglow that was seen in the gammev chameleon afterglow search (chase) . these data includes information about the broad band color of the observed glow , the relationship between the glow and the temperature of the apparatus , and other data taken prior to and during the science operations of chase . while differing in several details , the generic properties of the afterglow from chase are similar to luminescence seen in some vacuum compounds . contamination from this , or similar , luminescent signatures will likely impact the design of implementation of future experiments involving single photon detectors and high intensity light sources in a cryogenic environment .

LoBART(4k)+MCS: the gammev chameleon afterglow search (chase) experiment at the fermilab tevatron reported the discovery of an anomalous afterglows in its apparatus after shining a high - power pulsed laser into the bore of a cryogenic vacuum chamber immersed in a magnetic field . we present all of our data that pertains materially to the characterization of the “ orange glow “ signal . we do not claim any specific explanation of the source or cause of the orange glow , though the dependence upon temperature suggests strongly that the effect is due to some chemical or material property that is excited by the input laser . the data and discussion presented here may be useful for the design of future experiments that use high intensity light sources in conjunction with single photon detectors in cryogenic environments .

Table 15: arXiv test set sample #315.

Reference Summary: the survey of how canadian intensive care units (icu) prevent and diagnose venous thromboembolism (vte) presented in this issue of critical care illustrates considerable variability . lack of optimal patient care reflects how vte is rated in icu . the discussion should no longer focus on the incidence of thrombosis , but rather on its prevention . unfractionated heparin remains the most commonly used agent to prevent vte , despite the recognized efficacy and safety of low - molecular - weight heparins (lmwhs) in the icu setting . in addition , too few icu directors consider the use of mechanical prophylactic measures , such as graded elastic stockings and venous foot pump . the present situation calls for large randomized controlled trials in either medical or surgical icu patients , and for new education programmes in order to modify the care of icu patients with regard to vte .

LoBART(4k)+MCS: deep vein thrombosis (dvt) remains an underestimated problem in intensive care unit (icu) patients , despite the findings of many randomized controlled trials performed in the field of dvt prophylaxis after surgery during the past few decades . the canadian survey reported in the present issue of critical care provides a useful snapshot of daily clinical practice in canada with regard to dvt prevention in icu patients . it strongly suggests that studies dedicated to this topic should be performed in order to develop useful recommendations . furthermore , a great effort should be made to educate physicians regarding both dvt screening and pharmacological aspects .

Table 16: PubMed test set sample #3150.