
Poolingformer: Long Document Modeling with Pooling Attention

Hang Zhang^{1,2} Yeyun Gong³ Yelong Shen⁴ Weisheng Li⁵ Jiancheng Lv¹ Nan Duan³ Weizhu Chen⁴


Abstract

In this paper, we introduce a two-level attention schema, Poolingformer, for long document modeling. Its first level uses a smaller sliding window pattern to aggregate information from neighbors. Its second level employs a larger window to increase receptive fields with pooling attention to reduce both computational cost and memory consumption. We first evaluate Poolingformer on two long sequence QA tasks: the monolingual NQ and the multilingual TyDi QA. Experimental results show that Poolingformer sits atop three official leaderboards measured by F1, outperforming previous state-of-the-art models by 1.9 points (79.8 vs. 77.9) on NQ long answer, 1.9 points (79.5 vs. 77.6) on TyDi QA passage answer, and 1.6 points (67.6 vs. 66.0) on TyDi QA minimal answer. We further evaluate Poolingformer on a long sequence summarization task. Experimental results on the arXiv benchmark continue to demonstrate its superior performance.

1. Introduction

Transformer (Vaswani et al., 2017) architecture has been widely used in various natural language processing tasks with impressive results such as Translation (Lewis et al., 2020), Summarization (Qi et al., 2020), Text Classification (He et al., 2020), and Language Modeling (Brown et al., 2020). Self-attention is one of the key components in Transformer, which allows text tokens to interact with each other, and produce contextual representations. Despite the effectiveness of self-attention, its computational and memory complexity increases quadratically with respect to the sequence length. Therefore, most of existing transformer-based pretrained models (Alberti et al., 2019; He et al., 2020; Liu et al., 2019) set the maximum sequence length

¹College of Computer Science, Sichuan University ²During Internship at MSRA ³Microsoft Research Asia ⁴Microsoft Azure AI ⁵University of Science and Technology of China. Correspondence to: Yeyun Gong <yegong@microsoft.com>, Jiancheng Lv <lvjiancheng@scu.edu.cn>.



(a) Single-level local attention (b) Two-level pooling attention

Figure 1. (a): The receptive field of single-level local attention (b): The receptive field of our two-level pooling attention.

to 512 due to either memory or computational constraints, which often leads to a worse performance in long sequence tasks (Kwiatkowski et al., 2019; Cohan et al., 2018).

A lot of works have been proposed to adapt the self-attention layer in transformer to better model long sequence (Miculicich et al., 2018; Liu & Lapata, 2019; Beltagy et al., 2020; Zaheer et al., 2020; Wang et al., 2020b). For example, Longformer (Beltagy et al., 2020) puts forward a combination of both local and global attention patterns to reduce computational cost. Hierarchical transformer (Liu & Lapata, 2019) proposes to split the long document into shorter paragraphs, and apply inter-self-attentions within a paragraph and intra-self-attentions across paragraphs.

Inspired by previous works (Miculicich et al., 2018; Liu & Lapata, 2019; Beltagy et al., 2020; Zaheer et al., 2020; Wang et al., 2020b), we propose Poolingformer in which it revises the full self-attention to a two-level attention schema. The first level adopts a sliding-window attention pattern where each token only attends to its neighbor tokens within the window, as shown in Figure 1 (a). In the second level attention, it increases the receptive fields with a larger window size, followed by a pooling operation on both the key and value vectors in transformer to decrease the number of tokens to be attended. This multi-level design combining both sliding window and pooling can significantly reduce the computational cost and memory consumption while still attain exceptional model performance. Compared with the models with single-level local attention (Beltagy et al., 2020; Zaheer et al., 2020), Poolingformer allows a larger attention

receptive field per token via the benefit from the second-level pooling attention mechanism, as shown in Figure 1 (b). In the meantime, it preserves the sliding-window attention pattern at the first-level to mitigate the information loss due to the pooling operation at the second-level attention. Compared with Hierarchical Transformer (Liu & Lapata, 2019), Poolingformer obviates the need to explicitly split a long document into paragraphs. Thus, it is a more general long-sequence model which can be applied to extremely long text sequence in a cohesive manner. Compared with Transformer (Vaswani et al., 2017), the computational and memory complexity of Poolingformer only increase linearly with respect to sequence length.

In the experiment, we first demonstrate the superior performance of Poolingformer using two QA datasets: the monolingual NQ¹ and multilingual TyDi QA². Experimental results show that Poolingformer has achieved new state-of-the-art results on their official leaderboards. We continue to evaluate Poolingformer on the extremely long summarization task arXiv (Cohan et al., 2018). Experimental results show Poolingformer has set up new state-of-the-art results on this challenging benchmark.

2. Model

In the section, we present the model architecture of Poolingformer. We start with an introduction to the self-attention mechanism in Transformer model in section 2.1 and elaborate the details of Poolingformer self-attention in section 2.2.

2.1. Transformer Self-Attention

Given a sequence of text embeddings denoted as $\mathbf{X} = (\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_n)$, n is the text sequence length and $\mathbf{x}_i \in \mathbb{R}^d$ is the embedding vector of the i -th token. In the transformer model, it produces the query, key, and value vectors for each token by a linear projection of the embeddings, as in Eqn 1.

$$\begin{pmatrix} \mathbf{Q} \\ \mathbf{K} \\ \mathbf{V} \end{pmatrix} = \begin{pmatrix} \mathbf{W}_q \\ \mathbf{W}_k \\ \mathbf{W}_v \end{pmatrix} \mathbf{X} + \begin{pmatrix} \mathbf{b}_q \\ \mathbf{b}_k \\ \mathbf{b}_v \end{pmatrix} \quad (1)$$

where \mathbf{Q} , \mathbf{K} and \mathbf{V} are the query, key and value matrices respectively. Specifically, let \mathbf{q}_i be the i -th column of matrix \mathbf{Q} which indicates the i -th token’s query vector, \mathbf{k}_i and \mathbf{v}_i are defined in the same way.

A typical self-attention mechanism computes the inner product between the query and key vectors as the attention scores, and performs weighted intra-aggregation of value vectors to produce contextualized representations. For instance, token

i ’s output vector \mathbf{y}_i is calculated in Eqn. 2.

$$\mathbf{y}_i^T = \text{Softmax}(\alpha \mathbf{q}_i^T \mathbf{K}) \mathbf{V}^T \quad (2)$$

where α is a constant scalar and usually set as: $\alpha = 1/\sqrt{d}$. Therefore, the computation of the full self-attention comes with a $O(n^2)$ memory and computational complexity, which limits its ability for processing extremely long text sequence.

2.2. Poolingformer Self-Attention

Poolingformer revises the full self-attention mechanism to a two-level attention schema: the first level attention adopts the sliding window pattern to let each token only attend to its neighbor tokens within the window; the second level attention increases the receptive fields with a larger window size and performs attention over pooled key and value matrices. We provide an illustration of the Poolingformer self-attention in Figure 2 with more details elaborated in following subsections.

2.2.1. FIRST-LEVEL: SLIDING WINDOW ATTENTION

The first level self-attention sets a sliding window attention pattern to allow each token only attend to its neighbor tokens. For instance, token i ’s neighbor set within the window size w_1 is defined as $\mathcal{N}(i, w_1)$:

$$\mathcal{N}(i, w_1) = \{i - w_1, \dots, i, \dots, i + w_1\} \quad (3)$$

The sub-matrices of \mathbf{K} and \mathbf{V} with corresponding column indexes $\mathcal{N}(i, w_1)$ are denoted by $\mathbf{K}_{\mathcal{N}(i, w_1)}$ and $\mathbf{V}_{\mathcal{N}(i, w_1)}$. According to the sliding window pattern, each \mathbf{q}_i only attends to neighbor set $\mathcal{N}(i, w_1)$. Therefore, token i ’s output of the first-level attention is computed as:

$$\mathbf{y}_i^T = \text{Softmax}(\alpha \mathbf{q}_i^T \mathbf{K}_{\mathcal{N}(i, w_1)}) \mathbf{V}_{\mathcal{N}(i, w_1)}^T \quad (4)$$

Since size of the receptive field is limited to w_1 , it could lead to a worse model performance for long document understanding tasks.

2.2.2. SECOND-LEVEL: POOLING ATTENTION

The second level pooling attention module is built upon the outputs of the first level attention $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_n)$. It first produces new query, key and value matrices from \mathbf{Y} :

$$\begin{pmatrix} \tilde{\mathbf{Q}} \\ \tilde{\mathbf{K}} \\ \tilde{\mathbf{V}} \end{pmatrix} = \begin{pmatrix} \mathbf{W}_{\tilde{q}} \\ \mathbf{W}_{\tilde{k}} \\ \mathbf{W}_{\tilde{v}} \end{pmatrix} \mathbf{Y} + \begin{pmatrix} \mathbf{b}_{\tilde{q}} \\ \mathbf{b}_{\tilde{k}} \\ \mathbf{b}_{\tilde{v}} \end{pmatrix} \quad (5)$$

The query vector of token i and its corresponding key/value matrices are $\tilde{\mathbf{q}}_i$, $\tilde{\mathbf{K}}_{\mathcal{N}(i, w_2)}$, and $\tilde{\mathbf{V}}_{\mathcal{N}(i, w_2)}$ respectively, with a larger window size w_2 . (w_2 can be set to n in the extreme case). Since w_2 could be very large, we apply a pooling

¹<https://ai.google.com/research/NaturalQuestions/dataset>

²<https://ai.google.com/research/tydiqa>




Figure 2. The illustration of the two-level self-attention in PoolingFormer. Left block is the first level sliding window attention; Right block is the second level pooling attention.

layer to compress $\tilde{\mathbf{K}}_{\mathcal{N}(i, w_2)}$ and $\tilde{\mathbf{V}}_{\mathcal{N}(i, w_2)}$ respectively.

$$\bar{\mathbf{K}}_i = \text{Pooling}(\tilde{\mathbf{K}}_{\mathcal{N}(i, w_2)}; \kappa, \xi) \quad (6)$$

$$\bar{\mathbf{V}}_i = \text{Pooling}(\tilde{\mathbf{V}}_{\mathcal{N}(i, w_2)}; \kappa, \xi) \quad (7)$$

where κ and ξ are the pooling kernel size and stride size respectively. $\bar{\mathbf{K}}_i$ and $\bar{\mathbf{V}}_i$ are the compressed key, value matrices, and their size is ξ times smaller than $\tilde{\mathbf{K}}_i$ and $\tilde{\mathbf{V}}_i$.

The output of the second level pooling attention for token i is calculated in Eqn. 8.

$$z_i^T = \text{Softmax}(\alpha \tilde{q}_i^T \bar{\mathbf{K}}_i) \bar{\mathbf{V}}_i^T \quad (8)$$

In addition, we adopt a residual connection between the first level and second level attention modules, such that the final output of the two-level self-attention in Poolingformer is the sum of y_i (as in Eqn. 4) and z_i (as in Eqn. 8).

Pooling: we explore a few different pooling operations to compute $\bar{\mathbf{K}}_i$ and $\bar{\mathbf{V}}_i$ in our empirical studies, including the mean pooling, the max pooling and the convolution pooling (Wu et al., 2019). For a more comprehensive study, we introduce two trainable pooling mechanisms: the lightweight dynamic convolution pooling (LDConv) (Wu et al., 2019) and its variant mean-LDConv:

The input matrix $\mathbf{V} = (v_1, \dots, v_m)$ is first chunked into a list of segments: $((v_1, \dots, v_\kappa), (v_{1+\xi}, \dots, v_{1+\xi+\kappa}), \dots)$ in the pooling according to the kernel size κ and stride size ξ . The LDConv then maps each segment, i.e. (v_1, \dots, v_κ) , into a single vector for information compression in Eqn. 9

$$\text{LDConv}(v_1, \dots, v_\kappa) = \sum_{i=1}^{\kappa} \delta_i \cdot v_i \quad (9)$$

where $(\delta_1, \dots, \delta_\kappa)$ are called dynamic weights, computed by the context of $v_{i=\lceil \frac{1+\kappa}{2} \rceil}$,

$$(\delta_1, \dots, \delta_\kappa)^T = \text{Softmax}(\mathbf{W}_p v_i) \quad (10)$$

$\mathbf{W}_p \in \mathbb{R}^{\kappa \times d}$ is a learnable weight matrix. In the mean-LDConv, the dynamic weights $(\delta_1, \dots, \delta_\kappa)$ are computed by the mean of the context $\bar{v} = \frac{1}{\kappa} \sum_{i=1}^{\kappa} v_i$,

$$(\delta_1, \dots, \delta_\kappa)^T = \text{Softmax}(\mathbf{W}_p \bar{v}) \quad (11)$$

A detailed comparison on different pooling approaches is presented in section 3.4.

2.2.3. TASK SPECIFIC GLOBAL ATTENTION

In some specific long document modeling tasks, i.e., Question Answering, the question tokens are important to all the document tokens. Therefore, we follow Longformer (Beltagy et al., 2020) to append the indexes of query tokens into a global set $\mathcal{G} = \{q_1, \dots, q_l\}$ and allow all the tokens to attend to both the tokens in the global set and the tokens within its sliding window.

We integrate the global tokens into the first-level attention module in Poolingformer. The receptive field for each token i (not in the global set) is the union of $\mathcal{N}(i, w_1)$ and \mathcal{G} . For the tokens in the global set, the receptive field is the entire text sequence.

$$\mathcal{N}_{\mathcal{G}}(i, w_1) = \begin{cases} \mathcal{N}(i, w_1) \cup \mathcal{G} & i \notin \mathcal{G} \\ [1, \dots, n] & i \in \mathcal{G} \end{cases} \quad (12)$$

The output for token i of the first-level attention in Eqn. 4 is

Table 1. Computational complexity of several related models.

Model	Complexity
Transformer (Vaswani et al., 2017)	$\mathcal{O}(n^2)$
Reformer (Kitaev et al., 2020)	$\mathcal{O}(n \log n)$
Cluster-Former (Wang et al., 2020b)	$\mathcal{O}(n \log n)$
Longformer (Beltagy et al., 2020)	$\mathcal{O}(n)$
BigBird (Zaheer et al., 2020)	$\mathcal{O}(n)$
Poolingformer	$\mathcal{O}(n)$

revised accordingly in Eqn. 13

$$\mathbf{y}_i^T = \text{Softmax}(\alpha \mathbf{q}_i^T \mathbf{K}_{\mathcal{N}_{\mathcal{G}}(i, w_1)}) \mathbf{V}_{\mathcal{N}_{\mathcal{G}}(i, w_1)}^T \quad (13)$$

2.2.4. COMPLEXITY ANALYSIS

In this section, we simply analyze the complexity of Poolingformer. The computational complexity of the first-level sliding window attention is $\mathcal{O}(w_1 n)$. Considering w_1 is a constant and usually much smaller than n , the computational complexity can be simplified as $\mathcal{O}(n)$. The computational complexity of the second-level pooling attention is $\mathcal{O}(n \frac{w_2}{\xi})$, in which w_2 and ξ are two hyper-parameters. Compared with n , we usually configure the ratio $\frac{w_2}{\xi}$ to be a relatively small constant. Therefore, the complexity of the second-level pooling attention is $\mathcal{O}(n)$. In summary, the overall complexity of Poolingformer is $\mathcal{O}(n)$, we list the computational complexity of different long document modeling methods in Table 1 for comparison.

3. Experiments

3.1. Datasets

We evaluate Poolingformer on two long document tasks: Question Answering and Summarization. For QA, we report the results on the monolingual Natural Question (NQ) and the multilingual TyDi QA. For long document summarization, we report the results on the arXiv dataset (Cohan et al., 2018).

Natural Questions: This dataset collected real questions in Google’s search engine. Each question is paired with a Wikipedia page. Given a question and a document, NQ requires the model to find (1) an answer span (short answer) and (2) a paragraph that contains the information required to answer the question (long answer). If the question can not be answered from the given document, the model is asked to return NULL ANSWER. NQ provides a blind test set consisting of 7,842 examples, whose labels are hidden to us. Any submission to the public leaderboard will be evaluated on this hidden dataset. The leaderboard system will produce the rank of the submission according to the F1 metric.

TyDi QA: TyDi QA is a multilingual question answering dataset consisting of 11 typologically diverse languages with 200K human-annotated question-answer pairs. Similar to NQ, each question is paired with a Wikipedia article. The model need to make two predictions: (1) index of the passage that answers the question (Passage Selection Task) (2) minimal span that completely answers the question (Minimal Answer Span Task). TyDi QA also provides a blind test set and maintains a leaderboard like NQ with the same evaluation metrics.

arXiv: arXiv (Cohan et al., 2018) is a long document summarization dataset collected from scientific repositories—arxiv.org. The dataset contains about 215k long Scientific papers and uses the paper abstract as the summary. About the length of the document, the mean, median and 90th percentile are about 5k, 6.1k and 16.5k, respectively. Following previous work, We use ROUGE-1, ROUGE-2, and ROUGE-L as automatic evaluation metrics.

3.2. Implementation Details

Question Answer:

For NQ and TyDi QA, We split documents into multiple spans with a sliding window approach (Alberti et al., 2019). The size and stride of the sliding window are set to 4,096 and 1,568, respectively. Each instance is formed by a start placeholder, a question, and a document span. The question and the document span are separated by a special placeholder. Since many instances contain no answer, the number of negative instances and positive instances is imbalanced. We follow Liu et al. (2020) to sub-sample negative instances during training. The ratio of the sub-sampling set to 0.5. Similar to Alberti et al. (2019), we use token features to predict the short answer (Minimal Answer Span for TyDi QA). During inference, the distance between the start position and the end position is limited to 30 tokens. To predict Long Answer (Passage Selection for TyDi QA), we generate paragraph representations by applying a mean pooling to the tokens within the same paragraph. The answer type is predicted by the document representation which is the mean of all the paragraph representations.

We use RoBERTa-large (Liu et al., 2019) for NQ and XLM-RoBERTa (Conneau et al., 2020) for TyDi QA to initialize our models for training Poolingformer. Both models contain 24 Transformer encoder Layers. Since the maximum length of our model is several times that of the pretrained model, we follow Beltagy et al. (2020) to loop copying the position embedding of pretrained model to initialize our model. From the 15th to the 20th layer of our models, we apply two-level pooling self-attention, with other layers adopting the sliding window self-attention. The reason why we only utilize the two-level pooling attention in part of the layers is to avoid catastrophic forgetting of the prior knowledge in

Table 2. Results on the dev set and the blind test set of NQ. We report the evaluation results of the precision (P), the recall (R), and the F1 score for both long-answer (LA) and short-answer (SA) tasks.

	NQ LA Dev			NQ LA Test			NQ SA Dev			NQ SA Test		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
DocumentQA (Clark & Gardner, 2018)	47.5	44.7	46.1	48.9	43.3	45.7	38.6	33.2	35.7	40.6	31.0	35.1
DecAtt (Parikh et al., 2016) + DocReader (Chen et al., 2017)	52.7	57.0	54.8	54.3	55.7	55.0	34.3	28.9	31.4	31.9	31.1	31.5
BERTjoint (Alberti et al., 2019)	61.3	68.4	64.7	64.1	68.3	66.2	59.5	47.3	52.7	63.8	44.0	52.1
RikiNet (Liu et al., 2020)	74.3	76.4	75.3	-	-	-	61.4	57.3	59.3	-	-	-
-Ensemble model	73.3	78.7	75.9	78.1	74.2	76.1	66.6	56.4	61.1	67.6	56.1	61.3
ReflectionNet (Wang et al., 2020c)	79.4	72.7	75.9	-	-	-	69.3	55.0	61.3	-	-	-
-Ensemble model	78.2	75.9	77.0	76.8	77.6	77.2	67.9	59.4	63.4	70.4	58.8	64.1
Sparse Transformer (Child et al., 2019)	-	-	74.5	-	-	-	-	-	56.1	-	-	-
Reformer (Kitaev et al., 2020)	-	-	75.5	-	-	-	-	-	56.4	-	-	-
BigBird-ETC (Zaheer et al., 2020)	-	-	-	77.5	78.1	77.8	-	-	-	63.7	53.4	57.9
Cluster-Former (Wang et al., 2020b)	-	-	76.5	-	-	-	-	-	57.1	-	-	-
-Ensemble model	-	-	-	78.5	77.5	78.0	-	-	-	62.1	59.8	60.9
Poolingformer	77.7	77.3	77.5	-	-	-	62.3	55.3	58.6	-	-	-
-Ensemble model	-	-	-	78.5	81.2	79.8	-	-	-	70.4	54.8	61.6

the initialization model. Since question tokens are very important in the QA tasks, we treat question tokens the global tokens, as described in 2.2.3. The window sizes of the first-level and second-level is set to 128 and 512, respectively. The pooling kernel size, stride size are set to 5, 4. We use Adam optimizer (Kingma & Ba, 2015) with linear learning rate decay. The batch size, the training epoch, the learning rate, and the learning rate warmup proportion are set to 64, 2, 2×10^{-5} and 0.1 respectively. For the NQ leaderboard, the model we submitted is an ensemble of three models using different hyper-parameters. For the TyDi QA leaderboard, we use a single model for the submission.

Summarization:

For the arXiv dataset, we use the Encoder-Decoder framework following previous works (Zhang et al., 2020; Gidiotis & Tsoumakas, 2020). Pretrained model BART (Lewis et al., 2020) is used to initialize our model which consists of 12 encoder and 12 decoder layers. We expand the position embedding of encoder using the same method as QA. We apply the Poolingformer structure on the encoder side and keep the decoder structure unchanged. For encoder, the 6th to 11th layers adopt our two-level pooling self-attention and others adopt the single-level sliding window self-attention. Besides, we set the first token in encoder to global token as described in 2.2.3. The sizes of the first-level and second-level window are set to 128 and 512 respectively. The pooling kernel size and stride size are set to 5 and 4 respectively. We use Adam optimizer (Kingma & Ba, 2015) with linear learning rate decay. The batch size, the training epoch, the learning rate, and the learning rate warmup step are set to 128, 10, 2×10^{-4} , 1000, respectively. During inference, the beam size and length penalty are set to 5, 2 respectively.

For all experiments, we use 8 NVIDIA Tesla V100 GPUs.

All the experiments are conducted on Huggingface Transformers (Wolf et al., 2020) and Fairseq (Ott et al., 2019). We utilize Gradient Checkpointing (Chen et al., 2016), Apex³, and Gradient Accumulation to save GPU memory.

3.3. Main Results

3.3.1. GOOGLE NQ RESULTS

The results of both dev set and test set on NQ are shown in Table 2. The top block of the table shows the results of several approaches with input length of 512. The first three rows of the top block show the results of three multi-passage baseline models presented in the original NQ paper (Kwiatkowski et al., 2019). The fourth and fifth rows show two previous state-of-the-art models. RikiNet (Liu et al., 2020) adds Dynamic Paragraph Dual-Attention (DPDA) reader and multi-level cascaded answer predictor on top of the pretrained models. ReflectionNet (Wang et al., 2020c) is a two-phase model with an answer verification mechanism. These two models are proposed for NQ task and it is not easy to extend them to other tasks. The middle block lists the results of well-known and strong baselines designed for long documents, including Sparse Transformer (Child et al., 2019), Reformer (Kitaev et al., 2020), Cluster-Former (Wang et al., 2020b), BigBird (Zaheer et al., 2020). The first three rows are borrowed from the Cluster-Former paper (Wang et al., 2020b). The bottom block shows the results from Poolingformer. It is clear that Poolingformer has a significant improvement over previous methods consistently in both dev set and test set. It is worth noting that PoolingFormer achieves the best result in LA task in both single model and ensemble model. For example, in the hidden LA test set, its improvement over

³<https://github.com/NVIDIA/apex>

Table 3. Performance comparisons on the dev set and the blind test set of TyDi QA. We report the results using precision (P), recall (R), and F1 score for both the Passage Answer and the Minimal Answer tasks.

	Passage Answer Dev			Passage Answer Test			Minimal Answer Dev			Minimal Answer Test		
	P	R	F1	P	R	F1	P	R	F1	P	R	F1
Tydiqa-baseline (Clark et al., 2020)	63.1	57.0	59.1	62.3	67.1	64.4	41.3	35.3	50.5	56.4	50.1	52.7
mBERT-mnlp	-	-	-	63.8	60.4	61.7	-	-	-	61.5	47.3	53.2
GAAMA (XLM-R)-with ARES system	-	-	-	73.6	72.1	72.6	-	-	-	70.8	62.2	66.1
BERT with language-clustered vocab (Chung et al., 2020)	-	-	78.0	77.4	78.0	77.7	-	-	65.4	67.2	60.2	63.4
Poolingformer	79.5	78.7	79.1	80.4	78.8	79.5	74.4	63.2	68.5	73.5	63.3	67.7
Lesser Human	84.4	74.5	79.9	-	-	-	70.8	62.4	70.1	-	-	-

the previous state of the art is 1.8%(79.8 vs. 78.0). We treat this significant improvement, since NQ is an extremely competitive leaderboard and these scores are produced by a hidden dataset from the official NQ organizer.

3.3.2. TYDI QA RESULTS

In Table 3, we compare Poolingformer with Tydiqa-baseline (Clark et al., 2020) and previous state-of-the-art models. Tydiqa-baseline utilizes mBert (Alberti et al., 2019) which is a multilingual extended version of Bert. Chung et al. (2020) improve multilingual models with language-clustered vocabularies. We show that Poolingformer achieves a significant improvement over previous state-of-the-art models, improved from 77.7 to 79.5 in the Passage Answer task and 66.1 to 67.7 in the Minimal Answer tasks. The bottom block is a lesser estimate of human performance from Clark et al. (2020). Poolingformer further narrows the gap between machine and human performance. Without the ensemble approach, the gap between Poolingformer and human performance is only 0.4% and 2.3%. At the time of our submission (25 Jan. 2021), Poolingformer achieves the new state-of-the-art result on both LA (F1 79.5) and SA (F1 67.7) on the TyDi QA leaderboard. All of these results demonstrate that Poolingformer is simultaneously shining in multilingual comprehension tasks.

Table 4. The results on the arXiv test set. We report the results of ROUGE-1 (R-1), ROUGE-2 (R-2), and ROUGE-L (R-L).

Model	R-1	R-2	R-L
Sent-PTR (Pilault et al., 2020)	42.32	15.63	38.06
Extr-Abst-TLM (Pilault et al., 2020)	41.62	14.69	38.03
PEGASUS (Zhang et al., 2020)	44.21	16.95	38.83
Dancer (Gidiotis & Tsoumakas, 2020)	45.01	17.60	40.56
BigBird (Zaheer et al., 2020)	46.63	19.02	41.77
LED _{4k} (Beltagy et al., 2020)	44.40	17.94	39.76
LED _{16k} (Beltagy et al., 2020)	46.63	19.62	41.83
Poolingformer _{4k}	47.86	19.54	42.35
Poolingformer _{16k}	48.47	20.23	42.69

3.3.3. SUMMARIZATION RESULTS

The result on the arXiv test set is shown in table 4. The top block presents previous state-of-the-art methods with shorter input sequences. Sent-PTR (Pilault et al., 2020) is an extractive model that uses hierarchical LSTMs and a sentence pointer to select key sentences as the summary. Extr-Abst-TLM (Pilault et al., 2020) is a two-phase model that generates summaries based on sentences selected by an extractive model. PEGASUS (Zhang et al., 2020) is a large pretrained model specifically for abstractive summarization with an input length up to 1,024. Dancer (Gidiotis & Tsoumakas, 2020) breaks a long document into multiple sections to produce partial summaries for different sections and then produces a final complete summary based on the partial summaries.

The middle and bottom blocks show the results of several long document modeling methods for longer input sequence. Both BigBird and LED (Longformer-Encoder-Decoder) use both the slide local attention and the global attention mechanism to encode long documents. BigBird (Zaheer et al., 2020) initializes and continuously pretrain the model with PEGASUS which is dedicated to the summarization task. Following LED (Beltagy et al., 2020), Poolingformer is initialized from BART (Lewis et al., 2020) without the continuous pretraining process. We evaluate the performance of Poolingformer with input lengths of both 4K and 16K. One can clearly see that Poolingformer with input length of 16k greatly outperforms previous state-of-the-art models. Even if the input length is reduced to 4k, Poolingformer can still achieve the best on ROUGE-1 and ROUGE-L.

In addition, Poolingformer achieves a better computational complexity than models with single-level local attention. On this dataset, LED (Beltagy et al., 2020) sets the local attention one-side window size to 512 to increase the model’s receptive fields. That means the complexity of LED is $\mathcal{O}(1024 \times n)$. With the same receptive field, the complexity of Poolingformer’s two-level attention is $\mathcal{O}((256 + 1024/4) \times n)$, which accounts for only half complexity of LED. In other words, Poolingformer can greatly outperform LED in both accuracy and complexity. This

Table 5. Ablation study of Poolingformer_{base} with different window lengths on NQ dev set. w_1 : the size of the first level window. w_2 : the size of the second level window. C : the compression rate of the second level window controlled by adjusting the kernel size and stride size of the pooling.

Setting	w_1	w_2	C	LA F1	SA F1
RoBERTa _{base}	-	-	-	63.8	43.2
Poolingformer _{base}	128	-	-	66.3	43.1
Poolingformer _{base}	256	-	-	67.4	43.4
Poolingformer _{base}	512	-	-	66.1	42.6
Poolingformer _{base}	128	256	4	67.9	45.0
Poolingformer _{base}	128	512	4	68.7	45.2
Poolingformer _{base}	128	2,048	8	66.9	42.6
Poolingformer _{base}	128	2,048	16	67.0	44.4

demonstrates the effectiveness of the two-level pooling attention schema from both dimensions.

3.4. Ablation Study

For the sake of saving computational resources, we conduct all the ablation studies using one-fifth of the NQ training set using the base-size model. This model is initialized from RoBERTa-base. The 6^{th} – 8^{th} layers of the model adopt two-level pooling self-attention, and other layers adopt sliding window self-attention. The sizes of the first-level, second-level window, the pooling kernel size and stride size are set to 128, 512, 5 and 4 respectively.

Performance improvements of long document modeling: The top block of Table 5 shows a simple setting without the pooling attention. We first explored the advantages of long context modeling. Following previous work, we evaluate the RoBERTa model with the input length of 512. We observe that other models supporting longer input length consistently produce better results than RoBERTa on the LA task.

Useful but redundant information from distant tokens: From the second to the fourth rows in Table 5, we remove the second-level window and explore the relationship between the size of the first-level window and task performance. We may expect the performance becomes better for a larger window size. But the results show that it achieves the best performance when the sliding window size is set to 256. We conjecture that the reason for the poor performance of 512 windows size is that the self-attention mechanism is difficult to deal with remote token due to redundancy noise. In the bottom two rows of the Table 5, the second-level window size is set to the entire input sequence. We compress the sequence length by C times by adjusting the kernel and stride size in pooling attention. Each token will attend to

the tokens in the first-level window and tokens compressed from the entire sequence. From the results, we can see this approach does not work very well. We think that for every distant token, there may be too little useful information to compute attention. With these findings, we designed a two-level pooling attention mechanism to perform coarse-grained compression on farther tokens. For the tokens that are very far away, we will discard them directly. In this way, tokens can pay more attention to key information and reduce computation and memory consumed.

Table 6. Ablation study of pooling and fusion approaches.

Setting	LA F1	SA F1
Poolingformer _{base} (Without 2nd level window)	66.3	43.1
Poolingformer _{base} (MEAN)	68.5	43.7
Poolingformer _{base} (MAX)	68.6	45.3
Poolingformer _{base} (LDConv)	68.7	45.2
Poolingformer _{base} (MeanLDConv)	67.7	44.1
Poolingformer _{base} (LDConv, <i>Mix</i>)	67.5	44.6
Poolingformer _{base} (LDConv, <i>Weight Sharing</i>)	67.2	44.2

Impact of different pooling and fusion approaches: In the experiment, we have explored four different pooling methods while keeping other settings unchanged. The results are shown in Table 6. MEAN and MAX represent Mean pooling and Max pooling, respectively. LDConv refers to stride lightweight and dynamic convolution (Wu et al., 2019), as we discussed in Eqn. 9. Mean-LDConv is an variant of LDConv, refers to the weighted sum of token embeddings within the pooling window, where the weight is dynamically generated using the mean and linear layer. The detail of LDConv and Mean-LDConv is given in section 2.2.2. As presented in Table 6, LDConv and MAX are slightly better than others. We defer a more comprehensive study of different pooling approach in future work.

We explore another two different settings of poolingformer: *Mix* and *Weight Sharing*. In *Mix*, the second-level pooling attention module is built upon the input embeddings instead of the output of the first-level attention. To be more clear, it replaces Y with X in Eqn. 5 in *Mix* setting. From Table 6, we can see that the Poolingformer in the *Mix* setting performs worse on NQ tasks, which illustrates the effectiveness of stacking two level attentions in Poolingformer. In *Weight sharing*, the first level and second level share the linear mapping matrices W_q , W_k , and W_v in Eqn.1 and Eqn.5. From Table 6, we observe that the default setting produces better performance.

Impact of Poolingformer layer number: As shown in Table 7, an appropriate number of Poolingformer layers can greatly improve the model performance, up to 2.4 points and 2.1 points in terms of LA F1 and SA F1, respectively. This continues to demonstrate the value of the Poolingformer

Table 7. Ablation study of the number of Poolingformer layer.

Setting	LA F1	SA F1
Poolingformer _{base} (0 layers)	66.3	43.1
Poolingformer _{base} (1 layers)	68.0	44.5
Poolingformer _{base} (3 layers)	68.7	45.2
Poolingformer _{base} (6 layers)	67.5	43.7
Poolingformer _{base} (all layers)	65.0	41.5

layers. On the other hand, additional Poolingformer layers do not always lead to a better the performance. We observe some performance degradation when all the layers are replaced with the Poolingformer layers. Although the Poolingformer layer can effectively make use of distant information, it is still not fully compatible with the existing pretrained models. This may lead to some catastrophic forgetting of the information in the pretrained models. It is actually a trade-off between distant information and the prior knowledge of the pretrained models. Our experience shared with us that the best results often happen when the number of Poolingformer layer is one fourth of the total number of layers.

4. Related Work

The core limitation of Transformer in long document modeling is the computational complexity since the self-attention mechanism can grow quadratically to the sequence length. There are two widely adopted approaches to mitigate this problem. One is to use kernel functions, random projection, and others to approximate or eliminate the dot product in self-attention. Synthesizer (Tay et al., 2020) directly uses trainable parameters to generate attention weights, avoiding the dot-product interactions. Performer (Choromanski et al., 2020) and Linear Transformer (Katharopoulos et al., 2020) view the attention mechanism through kernelization and design different kernel functions to approximate the attention matrix. Compared with the original attention, these methods can reduce the complexity to linearity. But the performance of these methods comes with no theoretical guarantee. Moreover, it is difficult to make them compatible with existing pretrained models.

Another method is sparse attention, which focuses on making each token attend to less but more important context. Generally, the most important context is the local context. One simple way is the blockwise pattern (Qiu et al., 2020), which cuts the input sequence into multiple fixed chunks, and each token only attends to its neighbors within the same chunk. Furthermore, BP-Transformer (Ye et al., 2019) uses the binary partitioning tree to hierarchically block the sequence, and each token receives information from different

blocks according to distance. Another approach is the sliding window attention pattern that each token can attend to the neighbors in a sliding window. Longformer (Beltagy et al., 2020), BigBird (Zaheer et al., 2020) use this attention pattern to capture local information, and use global attention to capture global information which is similar to Star Transformer (Guo et al., 2019). Moreover, Sparse Transformer (Child et al., 2019) and Longformer (Beltagy et al., 2020) propose dilated window attention pattern which is similar to dilated convolution (Yu & Koltun, 2016). This pattern works well in autoregressive language modeling, but it is also not compatible with existing pretrained models. Linformer (Wang et al., 2020a) assumes that the attention mechanism matrix is low-rank, and utilizes linear mapping to compress sentence sequences. Another related work is Memory Compressed Attention (Liu et al., 2018), which adopts stride convolution to compress sentence information in the decoder and its computational complexity does not increase linearly with length. Cluster-Former (Wang et al., 2020b), Reformer (Kitaev et al., 2020), and Routing Transformer (Roy et al., 2020) utilize locally sensitive hashing and clustering methods to assign tokens with high similarity into buckets. Each token only attends to the tokens within its bucket.

5. Conclusion

We introduce Poolingformer, a two-level attention model for long sequence modeling with linear complexity. In the first level attention, it uses a smaller sliding window pattern to aggregate information from neighbor tokens. In the second level attention, it increases the receptive fields with a larger window size, followed by a pooling operation on both the key and value vectors to reduce the computational cost. Poolingformer achieves new state-of-the-art performance on long-document QA tasks and shows superior performance on long-document summarization task. For future work, we will continue to explore continuous improvement of Poolingformer from the following perspectives: 1) Theoretical analysis of the proposed multi-level attention in contrast to the classical single-level self-attention. 2) Extend Poolingformer to other types of long sequence data, such as image and music.

6. Acknowledgement

This work is supported by the National Key R&D Program of China under contract No. 2017YFB1002201, the National Natural Science Fund for Distinguished Young Scholar (Grant No. 61625204), and partially supported by the Key Program of National Science Foundation of China (Grant No. 61836006). We would like to thank Dayiheng Liu, Weizhen Qi for helpful discussions.

References

- Alberti, C., Lee, K., and Collins, M. A bert baseline for the natural questions. *arXiv preprint arXiv:1901.08634*, 2019.
- Beltagy, I., Peters, M. E., and Cohan, A. Longformer: The long-document transformer. *arXiv preprint arXiv:2004.05150*, 2020.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., Agarwal, S., Herbert-Voss, A., Krueger, G., Henighan, T., Child, R., Ramesh, A., Ziegler, D. M., Wu, J., Winter, C., Hesse, C., Chen, M., Sigler, E., Litwin, M., Gray, S., Chess, B., Clark, J., Berner, C., McCandlish, S., Radford, A., Sutskever, I., and Amodei, D. Language models are few-shot learners. In *NIPS*, 2020.
- Chen, D., Fisch, A., Weston, J., and Bordes, A. Reading wikipedia to answer open-domain questions. In *ACL*, 2017.
- Chen, T., Xu, B., Zhang, C., and Guestrin, C. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016.
- Child, R., Gray, S., Radford, A., and Sutskever, I. Generating long sequences with sparse transformers. *arXiv preprint arXiv:1904.10509*, 2019.
- Choromanski, K., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020.
- Chung, H. W., Garrette, D., Tan, K. C., and Riesa, J. Improving multilingual models with language-clustered vocabularies. In *EMNLP*, 2020.
- Clark, C. and Gardner, M. Simple and effective multi-paragraph reading comprehension. In *ACL*, 2018.
- Clark, J. H., Choi, E., Collins, M., Garrette, D., Kwiatkowski, T., Nikolaev, V., and Palomaki, J. Tydi qa: A benchmark for information-seeking question answering in typologically diverse languages. *Transactions of the Association for Computational Linguistics*, 8:454–470, 2020.
- Cohan, A., Dernoncourt, F., Kim, D. S., Bui, T., Kim, S., Chang, W., and Goharian, N. A discourse-aware attention model for abstractive summarization of long documents. In *NAACL*, 2018.
- Conneau, A., Khandelwal, K., Goyal, N., Chaudhary, V., Wenzek, G., Guzmán, F., Grave, E., Ott, M., Zettlemoyer, L., and Stoyanov, V. Unsupervised cross-lingual representation learning at scale. In *ACL*, 2020.
- Gidiotis, A. and Tsoumakas, G. A divide-and-conquer approach to the summarization of long documents. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:3029–3040, 2020.
- Guo, Q., Qiu, X., Liu, P., Shao, Y., Xue, X., and Zhang, Z. Star-transformer. In *NAACL*, 2019.
- He, P., Liu, X., Gao, J., and Chen, W. Deberta: Decoding-enhanced bert with disentangled attention, 2020.
- Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are rns: Fast autoregressive transformers with linear attention. In *International Conference on Machine Learning*, pp. 5156–5165. PMLR, 2020.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- Kitaev, N., Kaiser, L., and Levskaya, A. Reformer: The efficient transformer. In *ICLR*, 2020.
- Kwiatkowski, T., Palomaki, J., Redfield, O., Collins, M., Parikh, A., Alberti, C., Epstein, D., Polosukhin, I., Devlin, J., Lee, K., et al. Natural questions: A benchmark for question answering research. *Transactions of the Association for Computational Linguistics*, 7:453–466, 2019.
- Lewis, M., Liu, Y., Goyal, N., Ghazvininejad, M., Mohamed, A., Levy, O., Stoyanov, V., and Zettlemoyer, L. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. In *ACL*, 2020.
- Liu, D., Gong, Y., Fu, J., Yan, Y., Chen, J., Jiang, D., Lv, J., and Duan, N. Rikinet: Reading wikipedia pages for natural question answering. In *ACL*, 2020.
- Liu, P. J., Saleh, M., Pot, E., Goodrich, B., Sepassi, R., Kaiser, L., and Shazeer, N. Generating wikipedia by summarizing long sequences. In *ICLR*, 2018.
- Liu, Y. and Lapata, M. Hierarchical transformers for multi-document summarization. In *ACL*, 2019.
- Liu, Y., Ott, M., Goyal, N., Du, J., Joshi, M., Chen, D., Levy, O., Lewis, M., Zettlemoyer, L., and Stoyanov, V. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- Miculicich, L., Ram, D., Pappas, N., and Henderson, J. Document-level neural machine translation with hierarchical attention networks. In *EMNLP*, 2018.
- Ott, M., Edunov, S., Baevski, A., Fan, A., Gross, S., Ng, N., Grangier, D., and Auli, M. fairseq: A fast, extensible toolkit for sequence modeling. In *Proceedings of NAACL-HLT 2019: Demonstrations*, 2019.

- Parikh, A., Täckström, O., Das, D., and Uszkoreit, J. A decomposable attention model for natural language inference. In *EMNLP*, 2016.
- Pilault, J., Li, R., Subramanian, S., and Pal, C. On extractive and abstractive neural document summarization with transformer language models. In *EMNLP*, pp. 9308–9319, 2020.
- Qi, W., Yan, Y., Gong, Y., Liu, D., Duan, N., Chen, J., Zhang, R., and Zhou, M. Prophetnet: Predicting future n-gram for sequence-to-sequence pre-training. In *EMNLP: Findings*, pp. 2401–2410, 2020.
- Qiu, J., Ma, H., Levy, O., Yih, W.-t., Wang, S., and Tang, J. Blockwise self-attention for long document understanding. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: Findings*, pp. 2555–2565, 2020.
- Roy, A., Saffar, M., Vaswani, A., and Grangier, D. Efficient content-based sparse attention with routing transformers. *arXiv preprint arXiv:2003.05997*, 2020.
- Tay, Y., Bahri, D., Metzler, D., Juan, D.-C., Zhao, Z., and Zheng, C. Synthesizer: Rethinking self-attention in transformer models. *arXiv preprint arXiv:2005.00743*, 2020.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Wang, S., Li, B., Khabsa, M., Fang, H., and Ma, H. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020a.
- Wang, S., Zhou, L., Gan, Z., Chen, Y.-C., Fang, Y., Sun, S., Cheng, Y., and Liu, J. Cluster-former: Clustering-based sparse transformer for long-range dependency encoding. *arXiv preprint arXiv:2009.06097*, 2020b.
- Wang, X., Shou, L., Gong, M., Duan, N., and Jiang, D. No answer is better than wrong answer: A reflection model for document level machine reading comprehension. In *EMNLP: Findings*, 2020c.
- Wolf, T., Debut, L., Sanh, V., Chaumond, J., Delangue, C., Moi, A., Cistac, P., Rault, T., Louf, R., Funtowicz, M., Davison, J., Shleifer, S., von Platen, P., Ma, C., Jernite, Y., Plu, J., Xu, C., Scao, T. L., Gugger, S., Drame, M., Lhoest, Q., and Rush, A. M. Transformers: State-of-the-art natural language processing. In *EMNLP*, 2020.
- Ye, Z., Guo, Q., Gan, Q., Qiu, X., and Zhang, Z. Bp-transformer: Modelling long-range context via binary partitioning. *arXiv preprint arXiv:1911.04070*, 2019.
- Yu, F. and Koltun, V. Multi-scale context aggregation by dilated convolutions. In Bengio, Y. and LeCun, Y. (eds.), *ICLR*, 2016.
- Zaheer, M., Guruganesh, G., Dubey, K. A., Ainslie, J., Alberti, C., Ontañón, S., Pham, P., Ravula, A., Wang, Q., Yang, L., and Ahmed, A. Big bird: Transformers for longer sequences. In *NeurIPS*, 2020.
- Zhang, J., Zhao, Y., Saleh, M., and Liu, P. Pegasus: Pre-training with extracted gap-sentences for abstractive summarization. In *ICML*, 2020.