
ConvBERT: Improving BERT with Span-based Dynamic Convolution

Zihang Jiang^{1*†}, Weihao Yu^{1†}, Daquan Zhou¹, Yunpeng Chen², Jiashi Feng¹, Shuicheng Yan²

¹National University of Singapore, ²Yitu Technology

jzihang@u.nus.edu, {weihaoyu6,zhoudaquan21}@gmail.com,

yunpeng.chen@yitu-inc.com, elefjia@nus.edu.sg, shuicheng.yan@yitu-inc.com

Abstract

Pre-trained language models like BERT and its variants have recently achieved impressive performance in various natural language understanding tasks. However, BERT heavily relies on the global self-attention block and thus suffers large memory footprint and computation cost. **Although all its attention heads query on the whole input sequence for generating the attention map from a global perspective, we observe some heads only need to learn local dependencies, which means the existence of computation redundancy.** We therefore propose a **novel span-based dynamic convolution to replace these self-attention heads to directly model local dependencies.** The novel convolution heads, together with the rest self-attention heads, form a new mixed attention block that is more efficient at both **global and local context learning.** We equip BERT with this mixed attention design and build a ConvBERT model. Experiments have shown that ConvBERT significantly outperforms BERT and its variants in various downstream tasks, with lower training costs and fewer model parameters. Remarkably, ConvBERT_{BASE} model achieves 86.4 GLUE score, 0.7 higher than ELECTRA_{BASE}, using less than 1/4 training cost.³

1 Introduction

Language model pre-training has shown great power for improving many natural language processing tasks [63, 52, 51, 33]. Most pre-training models, despite their variety, follow the BERT [14] architecture heavily relying on multi-head self-attention [62] to learn comprehensive representations. It has been found that 1) though the self-attention module in BERT is a highly non-local operator, a large proportion of attention heads indeed learn local dependencies due to the inherent property of natural language [31, 2]; 2) removing some attention heads during fine-tuning on downstream tasks does not degrade the performance [41]. The two findings indicate that heavy computation redundancy exists in the current model design. In this work, we aim to resolve this intrinsic redundancy issue and further improve BERT w.r.t. its efficiency and downstream task performance. We consider such a question: can we reduce the redundancy of attention heads by using a naturally local operation to replace some of them? We notice that convolution has been very successful in extracting local features [35, 32, 55, 23], and thus propose to use convolution layers as a more efficient complement to self-attention for addressing local dependencies in natural language.

Specifically, we propose to integrate convolution into self-attention to form a mixed attention mechanism that combines the advantages of the two operations. Self-attention uses all input tokens to generate attention weights for capturing global dependencies, while we expect to perform local

*Work done during an internship at Yitu Tech.

†Equal contribution.

³Code and pre-trained model will be released at <https://github.com/yitu-opensource/ConvBert>.

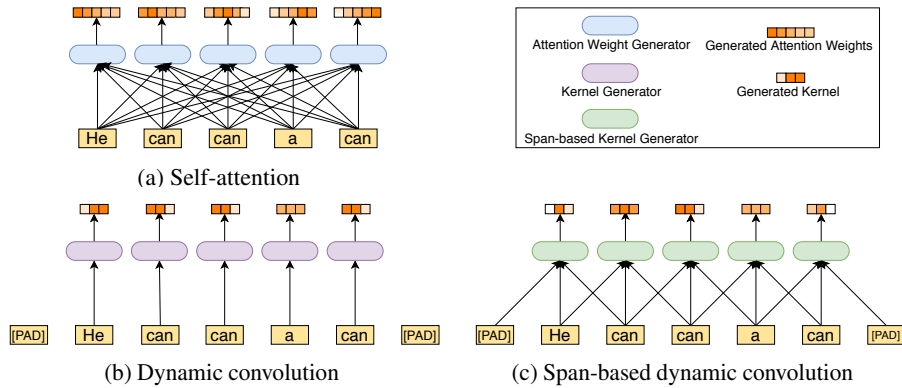


Figure 1: Processes of generating attention weights or convolution kernels. (a) Self-attention: all input tokens are needed to generate attention weights which requires quadratic complexity. (b) Dynamic convolution: **dynamic kernels are generated by taking in one current token only, resulting in generating the same kernels for the same input tokens with different meanings, like “can” token.** (c) Span-based dynamic convolution: kernels are generated by taking in a local span of current token, which better utilizes local dependency and discriminates different meanings of the same token (e.g., if “a” is in front of “can” in the input sentence, “can” is apparently a noun not a verb).

“self-attention”, i.e., taking in a local span of the current token to generate “attention weights” of the span to capture local dependencies. To achieve this, rather than deploying standard convolution with fixed parameters shared for all input tokens, dynamic convolution [66] is a good choice that offers higher flexibility in capturing local dependencies of different tokens. As shown in Fig. 1b, dynamic convolution uses a kernel generator to produce different kernels for different input tokens [66]. **However, such dynamic convolution cannot differentiate the same tokens within different context and generate the same kernels (e.g., the three “can” in Fig. 1b).**

We thus develop the span-based dynamic convolution, a novel convolution that produces more adaptive convolution kernels by receiving an input span instead of only a single token, which enables discrimination of generated kernels for the same tokens within different context. For example, as shown in Fig. 1c, the proposed span-based dynamic convolution produces different kernels for different “can” tokens. With span-based dynamic convolution, we build the mixed attention to improve the conventional self-attention, which brings higher efficiency for pre-training as well as better performance for capturing global and local information.

To further enhance performance and efficiency, we also add the following new architecture design to BERT. First, a bottleneck structure is designed to reduce the number of attention heads by embedding input tokens to a lower-dimensional space for self-attention. This also relieves the redundancy that lies in attention heads and improves efficiency. Second, the feed-forward module in BERT consists of two fully connected linear layers with an activation in between, but the dimensionality of the inner-layer is set much higher (e.g., $4\times$) than that of input and output, which promises good performance but brings large parameter number and computation. Thus we devise a *grouped linear* operator for the feed-forward module, which reduces parameters without hurting representation power. Combining these novelties all together makes our proposed model, termed ConvBERT, small and efficient.

Our contributions are summarized as follows. 1) We propose a new mixed attention to replace the self-attention modules in BERT, which leverages the advantages of convolution to better capture local dependency. To the best of our knowledge, we are the first to explore convolution for enhancing BERT efficiency. 2) We introduce a novel span-based dynamic convolution operation to utilize multiple input tokens to dynamically generate the convolution kernel. 3) Based on the proposed span-based dynamic convolution and mixed attention, we build ConvBERT model. On the GLUE benchmark, ConvBERT_{BASE} achieves 86.4 GLUE score which is 5.5 higher than BERT_{BASE} and 0.7 higher than ELECTRA_{BASE} while requiring less training cost and parameters. 4) ConvBERT also incorporates some new model designs including the bottleneck attention and grouped linear operator that are of independent interest for other NLP model development.

2 Related work

Language model pre-training Language model pre-training first pre-trains a model on large-scale unlabeled text corpora, and then fine-tunes the model on downstream tasks [12, 46, 34, 69]. It was proposed to learn separate word representations [42, 45]. Later, LSTM based CoVe [40] and ELMo [46] were developed to generate contextualized word representations. Recently, since transformer architecture including multi-head self-attention and feed-forward modules [62] has shown better effectiveness than LSTMs in many NLP tasks, GPT [49] deploys transformer as its backbone for generative pre-training and achieves large performance gain on downstream tasks. To further improve pre-trained models, more effective pre-training objectives have been developed, including Masked Language Modeling and Next Sentence Prediction from BERT [34], Generalized Autoregressive Pretraining from XLNet [69], Span Boundary Objective from SpanBERT [26], and Replaced Token Detection from ELECTRA [9]. Some other works compress the pretrained models by weight pruning [21, 17], weight sharing [34], knowledge distillation [53, 25, 60] and quantization [70, 54]. Our method is orthogonal to the above methods. There are also some works that extend pre-training by incorporating knowledge [73, 47, 37, 68], multiple languages [24, 10, 7], and multiple modalities [39, 58, 5, 59, 8]. However, to the best of our knowledge, since GPT, there is no study on improving pre-trained models w.r.t. backbone architecture design. This work is among the few works that continues the effort on designing better backbone architecture to improve pre-trained model performance and efficiency.

Convolution in NLP models The convolution block has been used in NLP models to encode local information and dependency of the context [71, 29, 28, 16, 72], but not explored in the pre-training field. For instance, 1D convolution is applied to specific sequence-to-sequence learning tasks [18], like machine translation and summarization. The depth-wise separable convolution is deployed in the text encoder and decoder for translation task [27], to reduce parameters and computation cost. A more recent work [66] utilizes the light-weight and dynamic convolution to further enhance the expressive power of convolution. **However, all these models are limited in capability of capturing the whole context of long sentences.** To enhance it, some works [56, 67] combine convolution with transformer in the sequential or multi-branch manner. To the best of our knowledge, our work is the first to explore applying convolution to pre-trained models.

3 Method

We first elaborate how we identify the redundancy of self-attention heads at learning local dependencies by revisiting self-attention and dynamic convolution. Then we explain the novel span-based dynamic convolution that models local dependencies and finally our proposed ConvBERT model built by the mixed attention block.

3.1 Motivation

Self-attention The self-attention block is the basic building block of BERT, which effectively models global dependencies in the input sequence of tokens. As shown in Fig. 3a, given the input $X \in \mathbb{R}^{n \times d}$ where d is the hidden dimension and n is the number of tokens, the self-attention module applies three linear transformations on the inputs X and embeds them to the key K , query Q and value V respectively, where $K, Q, V \in \mathbb{R}^{n \times d}$. Suppose there are H self-attention heads. The key, query and value embeddings are uniformly split into $d_k = d/H$ -dimensional segments. The self-attention module gives outputs in the form:

$$\text{Self-Attn}(Q, K, V) = \text{softmax} \left(\frac{Q^\top K}{\sqrt{d_k}} \right) V. \quad (1)$$

BERT [14] and its variants successfully apply self-attention and achieve high performance. **However, despite non-local essentials of the self-attention operator, some attention heads in BERT indeed learn**

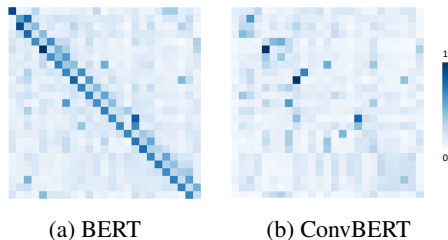


Figure 2: Visualization of average attention map from self-attention in BERT and our ConvBERT. The example is randomly sampled from MRPC development set. The tokenized sentence is "[CLS] he said the foods ##er ##vic ##e pie business doesn't fit the company's long-term growth strategy . [SEP]".

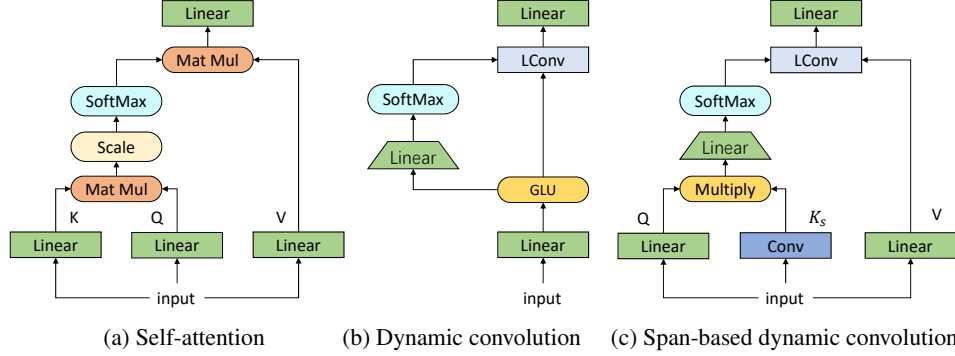


Figure 3: Illustration of self-attention, dynamic convolution [66] and proposed span-based dynamic convolution. Here LConv denotes the light-weight depth-wise convolution that ties all weights along channel dimension but uses a different convolution kernel at each position.

the local dependency of the input sequence due to the inherent property of natural language [31, 2]. As shown in Fig. 2a, the averaged attention map of BERT apparently exhibits diagonal patterns similar to the analysis of [31], meaning a large proportion of attention heads learn local dependency. Since the self-attention computes attention weights between all token pairs as in Eqn. 1, many attention weights beyond the captured local context are unnecessary to compute as they contribute much less compared to the local ones. This leads to unnecessary computation overhead and model redundancy. Motivated by this, we adopt convolution operation to capture the local dependency, considering it is more suitable for learning local dependencies than self-attention.

Light-weight and dynamic convolution Light-weight convolution [66] can efficiently model the local dependency. Let convolution kernel be denoted as $W \in \mathbb{R}^{d \times k}$. The output of depth-wise convolution at position i and channel c can be formulated as $\text{DWConv}(X, W_{c,:}, i, c) = \sum_{j=1}^k W_{c,j} \cdot X_{(i+j-\lceil \frac{k+1}{2} \rceil),c}$. By tying the weight along the channel dimension, we can simplify the convolution kernel to $W \in \mathbb{R}^k$, giving the following light-weight convolution:

$$\text{LConv}(X, W, i) = \sum_{j=1}^k W_j \cdot X_{(i+j-\lceil \frac{k+1}{2} \rceil)}. \quad (2)$$

This largely reduces the parameter size by d (256 or 768 in practice) compared to the conventional depth-wise convolution. However, after training, **the kernel parameters would be fixed for any input token, not favourable for capturing diversity of the input tokens.** Thus, we further consider dynamic convolution [66] that generates the convolution kernel parameters conditioned on the specific input token. As illustrated in Fig. 3b, **after a linear projection and a Gated Linear Unit (GLU) [13], a dynamic convolution kernel is generated from the current input token, and applied to convolve with the nearby tokens to generate new representation embedding.** Compared to standard convolution kernels that are fixed after training, dynamic convolution can better utilize the input information and generate kernel conditioned on the input token. To be specific, a position dependent kernel $W=f(X_i)$ for position i is used, where f is a linear model with learnable weight $W_f \in \mathbb{R}^{k \times d_k}$ followed by a softmax. We denote the dynamic convolution as

$$\text{DConv}(X, W_f, i) = \text{LConv}(X, \text{softmax}(W_f X_i), i). \quad (3)$$

Compared with self-attention, which is of quadratic computation complexity w.r.t. the input sequence length, the linear complexity dynamic convolution is more efficient and more suitable for modelling local dependency and has shown effectiveness on machine translation, language modeling and abstractive summarization tasks [66]. However, we observe its convolution kernel only depends on a single token of the input, ignoring the local context. It would hurt model performance to use the same kernel for the same token in different context which may have different meanings and relations with its contextual tokens. Thus, we propose the span-based dynamic convolution as below.

3.2 Span-based dynamic convolution

The span-based dynamic convolution first uses a depth-wise separable convolution to gather the information of a span of tokens as shown in Fig. 3c and then dynamically generates convolution kernels. This helps the kernel capture local dependency more effectively by generating local relation

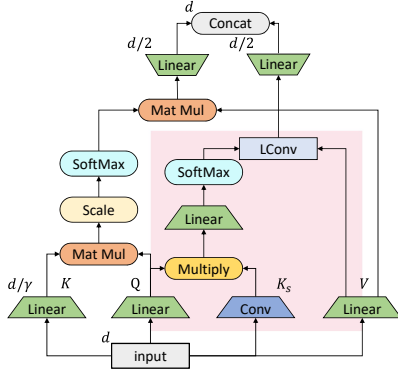


Figure 4: Illustration of mixed-attention block. It is a mixture of self-attention and span-based dynamic convolution (highlighted in pink). They share the same Query but use different Key to generate the attention map and convolution kernel respectively. We reduce the number of attention heads by directly projecting the input to a smaller embedding space to form a bottleneck structure for self-attention and span-based dynamic convolution. Dimensions of the input and output of some blocks are labeled on the left top corner to illustrate the overall framework, where d is the embedding size of the input and γ is the reduction ratio.

of the input token conditioned on its local context instead of a single token. Besides, to make the span-based dynamic convolution compatible with self-attention, we apply linear transformation on the input X to generate query Q and value V , and a depth-wise separable convolution to generate the span-aware key K_s . The result of point-wise multiplication of query Q and span-aware key K_s pairs transformed from input X is then used to generate the dynamic convolution kernel. Specifically, with query and key pair Q, K_s as input, the kernel is generated by

$$f(Q, K_s) = \text{softmax}(W_f(Q \odot K_s)), \quad (4)$$

where \odot denotes point-wise multiplication. As illustrated in Fig. 3c, we call this new operator span-based dynamic convolution. The output can be written as

$$\text{SDConv}(Q, K_s, V; W_f, i) = \text{LConv}(V, \text{softmax}(W_f(Q \odot K_s)), i). \quad (5)$$

Then a linear layer is applied for further process. If not otherwise stated, we always keep the same kernel size for the depth-wise separable convolution and span-based dynamic convolution.

3.3 ConvBERT architecture

With the above span-based dynamic convolutions, we develop the novel mixed attention block and the efficient ConvBERT model.

Mixed attention The mixed attention block integrates the span-based dynamic convolution and self-attention to better model both global and local dependencies with reduced redundancy, as shown in Fig. 4. The self-attention and span-based dynamic convolution share the same query but use different keys as reference to generate attention maps and convolution kernels. Denote $\text{Cat}(\cdot, \cdot)$ as the concatenate operation. We formulate our mixed attention as

$$\text{Mixed-Attn}(K, Q, K_s, V; W_f) = \text{Cat}(\text{Self-Attn}(Q, K, V), \text{SDConv}(Q, K_s, V; W_f)). \quad (6)$$

The final outputs are fed to the feed-forward layer for further process.

Bottleneck design for self-attention Some of the heads are redundant [41] and thus we propose to reduce the number of heads while introducing the span-based dynamic convolution module. We call this bottleneck structure, since the input embeddings are first projected to a lower-dimensional space and then pass through the self-attention module, as shown in Fig. 4. Specifically, in original BERT, the embedding feature of dimension d is projected to query, key and value with the same dimension d in the original transformer architecture by linear transformations. Instead, we project the embedding feature to a smaller space of dimension d/γ for further process, where $\gamma > 1$ is the reduction ratio. Meanwhile, we reduce the number of attention heads by ratio γ . This largely saves computation cost within the self-attention and forces attention heads to yield more compact and useful attention information.

Grouped feed-forward module The large number of parameters in the transformer model actually comes from the feed-forward module. To maintain the representation power while reducing parameters and computation cost, we propose to apply a grouped linear (GL) operator to the feed-forward module in a grouped manner which is defined as follows

$$M = \Pi_{i=0}^g \left[f_{\frac{d}{g}}^i \rightarrow \frac{m}{g} \left(H_{[:, i-1:i \times \frac{d}{g}]} \right) \right], M' = \text{GeLU}(M), H' = \Pi_{i=0}^g \left[f_{\frac{m}{g}}^i \rightarrow \frac{d}{g} \left(M'_{[:, i-1:i \times \frac{m}{g}]} \right) \right], \quad (7)$$

where $H, H' \in \mathbb{R}^{n \times d}$, $M, M' \in \mathbb{R}^{n \times m}$, $f_{d_1 \rightarrow d_2}(\cdot)$ indicates a fully connected layer that transforms dimension d_1 to d_2 , g is the group number and Π means concatenation. This is in line with the original multi-head attention mechanism, where the input features are split into multiple groups on embedding dimension and processed independently, and all processed features are concatenated again on the embedding dimension. This is more efficient than the fully connected layer and costs negligible performance drop.

By stacking the mixed attention and grouped feed-forward modules in an iterative manner, we build our ConvBERT model. As experimentally demonstrated in Sec. 4, ConvBERT is more light-weighted and efficient at capturing both global and local contexts with better performance.

4 Experiment

4.1 Implementation

We evaluate the effectiveness of our proposed architecture on the task of replaced token detection pre-training proposed by [9]. More details about the task and implementation are listed in Appendix.

We evaluate ConvBERT with a variety of model sizes. Following [9], for a *small-sized* model, the hidden dimension is 256 while the word embedding dimension is reduced to 128. Like the original transformer architecture, the intermediate layer size of the feed-forward module is 1024 (4 times of the hidden dimension) and we keep the number of layers as 12. For the number of attention heads, we keep it as 4 for the small-sized models. In addition, we also use a *medium-small-sized* model with 384 dimension embedding and 8 attention heads. By inserting span-based dynamic convolution and applying the grouped linear operation, the model can be reduced to a size comparable to a small-sized one while enjoying more representation power. For the *base-sized* model, we adopt the commonly used BERT-base configuration with 768 hidden dimension and 12 layers. For the number of heads, we use 12 for the base-sized model as our baseline. When applying the bottleneck structure (i.e. reducing the dimension of the hidden space for self-attention) we also reduce the head number by a factor $\gamma = 2$ to keep the size of each head stays the same.

During pre-training, the batch size is set to 128 and 256 respectively for the small-sized and base-sized model. An input sequence of length 128 is used to update the model. We show the results of these models after pre-training for 1M updates as well as pre-training longer for 4M updates. More detailed hyper-parameters for pre-training and fine-tuning are listed in the Appendix.

4.2 Pre-training and evaluation

Pre-training dataset The pre-training tasks largely rely on a large corpus of text data. The dataset WikiBooks originally used to train BERT [14] is a combination of English Wikipedia and BooksCorpus [74]. RoBERTa [38], XLNet [69] and ELECTRA [9] further propose to use larger corpora including OpenWebText [50, 20], STORIES [61], CCNEWS [43], ClueWeb [3] and Gigaword [44] to improve overall performance. However, some datasets like BooksCorpus [74] are no longer publicly available. In this paper, unless otherwise stated, we train the models on an open-sourced dataset OpenWebText [20, 50] (32G) to ease reproduction, which is of similar size with the combination of English Wikipedia and BooksCorpus that is used for BERT training. We also show the results of our model trained on the same data as BERT (i.e. WikiBooks) in the Appendix. More details of the corpus information and how to collect and filter the text data can be found in [48].

Evaluation We evaluate our model on the General Language Understanding Evaluation (GLUE) benchmark [63] as well as Question answering task SQuAD [52]. GLUE benchmark includes various tasks which are formatted as single sentence or sentence pair classification. See Appendix for more details of all tasks. SQuAD is a question answering dataset in which each example consists of a context, a question and an answer from the context. The target is to locate the answer with the given context and question. In SQuAD V1.1, the answers are always contained in the context, where in V2.0 some answers are not included in the context.

We measure accuracy for MNLI, QNLI, QQP, RTE, SST, Spearman correlation for STS and Matthews correlation for CoLA. The GLUE score is the average of all 8 tasks. Since there is nearly no

single model submission on SQuAD leaderboard,⁴ we only compare ours with other models on the development set. We report the Exact Match and F1 score on the development set of both v1.1 and v2.0. For fair comparison, unless otherwise stated, we keep the same configuration as in ELECTRA [9]. In addition to BERT and ELECTRA, we also take knowledge distillation based methods including TinyBERT [25], MobileBERT [60] and DistillBERT [53] for comparison. All results are obtained by single task fine-tuning.

4.3 Ablation study

To better investigate each component, we add bottleneck structure, span-based dynamic convolution and grouped linear operation one by one to the original BERT [14] architecture. Additionally, we also increase the hidden dimension to show the performance gain by increasing the parameter size. Detailed configuration and results are shown in Table 1.

Table 1: Detailed configuration of different sized models. GLUE score is the average score on GLUE dev set. Here, BNK represents bottleneck structure; GL denotes grouped linear operator; SDConv is the proposed span-based dynamic convolution. Larger denotes increasing the hidden dimension and head dimension of the model.

Model	Modification	Hidden dim	Head dim	Head	Group	Params	GLUE
BERTSMALL		256	64	4	1	14M	75.1
ELECTRASMALL		256	64	4	1	14M	80.4
ConvBERTSMALL	+BNK	256	64	2	1	12M	80.6
	+BNK, +SDConv	256	64	2	1	14M	81.4
ELECTRAMEDIUM-SMALL		384	48	8	1	26M	82.0
ConvBERTMEDIUM-SMALL	+BNK	384	48	4	1	23M	80.9
	+BNK,+GL	384	48	4	2	14M	81.0
	+BNK,+GL,+Larger	432	54	4	2	17M	81.1
	+BNK,+GL,+SDConv	384	48	4	2	17M	82.1
BERTBASE		768	64	12	1	110M	82.2
ELECTRABASE		768	64	12	1	110M	85.1
ConvBERTBASE	+BNK	768	64	6	1	96M	84.9
	+BNK,+SDConv	768	64	6	1	106M	85.7

Bottleneck structure and grouped linear operation An interesting finding is that, introducing the bottleneck structure and grouped linear operation can reduce the number of parameters and computation cost without hurting the performance too much. It can even bring benefits in the small-sized model setting. This is possibly because the rest of the attention heads are forced to learn more compact representation that generalizes better.

Kernel size Another factor we investigate is the kernel size of the dynamic convolution. We show the results of applying different convolution kernel sizes on the small-sized ConvBERT model in Fig. 5. It can be observed that a larger kernel gives better results as long as the receptive field has not covered the whole input sentence. However when the kernel size is large enough and the receptive field covers all the input tokens, the benefit of using large kernel size diminishes. In later experiments, if not otherwise stated, we set the convolution kernel size as 9 for all dynamic convolution since it gives the best result.

Ways to integrate convolution We here test the different ways to integrate convolution into the self-attention mechanism. As shown in Table 2, directly adding a conventional depth-wise separable convolution parallel to the self-attention module will hurt the performance while inserting dynamic convolution gives little improvement over the baseline BERT architecture w.r.t. the average GLUE score. By further increasing the local dependency with span-based dynamic convolution, the performance can be improved by a notable margin.

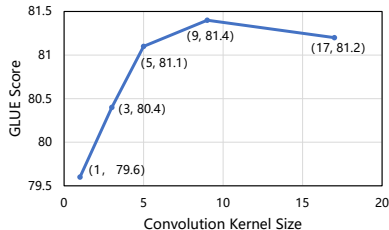


Figure 5: Results on GLUE dev set with different kernel size for span-based dynamic convolution.

⁴<https://rajpurkar.github.io/SQuAD-explorer/>

Table 2: Comparison of ConvBERT with different convolutions. Median results of 9 runs on GLUE dev set are reported.

Model	Convolution	MNLI	QNLI	QQP	RTE	SST-2	MRPC	CoLA	STS-B	Avg.
ELECTRASMALL [9]	-	78.9	87.9	88.3	68.5	88.3	87.4	56.8	86.8	80.4
ConvBERTSMALL	Conventional	79.9	85.3	89.2	63.9	86.9	83.1	53.4	83.9	78.2
	Dynamic	81.1	87.6	90.1	64.3	88.9	86.8	59.3	86.7	80.6
	Span-based Dynamic	81.5	88.5	90.4	67.1	90.1	86.8	59.7	87.7	81.4

4.4 Comparison results

We compare our ConvBERT model with BERT [14] and ELECTRA [9] as well as state-of-the-art methods [25, 53, 60, 50] on GLUE [63] and SQuAD [52] to validate the advantages of our method.

Results on GLUE We evaluate the performance of all methods on different downstream tasks on development set and test set of GLUE. See Table 3. Due to space limit, we only show detailed results on test set and defer those on development set to Appendix. As can be seen from Table 3, our small and base-sized models outperform other baseline models of similar size while requiring much less pre-training cost. For example, compared with the strong baseline ELECTRABASE, our ConvBERTBASE achieves better performance with less than 1/4 training cost. Note that TinyBERT [25] and MobileBERT [60] rely on a pre-trained large teacher network. Our model is actually a new backbone and thus is orthogonal to these compression techniques. Due to computation resource limitations, we remain the comparison of large-sized models for future work.

Table 3: Comparison of models with similar size on GLUE test set. Pre-training computation cost is also reported. † denotes knowledge distillation based methods relying on large pre-trained teacher models.

Model	Train FLOPs	Params	MNLI	QNLI	QQP	RTE	SST-2	MRPC	CoLA	STS-B	Avg.
TinyBERT† [25]	6.4e19+ (49x+)	15M	84.6	90.4	89.1	70.0	93.1	82.6	51.1	83.7	80.6
MobileBERT† [60]	6.4e19+ (49x+)	25M	84.3	91.6	88.3	70.4	92.6	84.5	51.1	84.8	81.0
ELECTRASMALL [9]	1.4e18 (1.1x)	14M	79.7	87.7	88.0	60.8	89.1	83.7	54.6	80.3	78.0
train longer [9]	3.3e19 (25x)	14M	81.6	88.3	88.0	63.6	91.1	84.9	55.6	84.6	79.7
GPT [49]	4.0e19 (31x)	117M	82.1	88.1	88.5	56.0	91.3	75.7	45.4	80.0	75.9
BERTBASE [14]	6.4e19 (49x)	110M	84.6	90.5	89.2	66.4	93.5	84.8	52.1	85.8	80.9
ELECTRABASE [9]	6.4e19 (49x)	110M	85.8	92.7	89.1	73.1	93.4	86.7	59.7	87.7	83.5
train longer [9]	3.3e20 (254x)	110M	88.5	93.1	89.5	75.2	96.0	88.1	64.6	90.2	85.7
ConvBERTSMALL	1.3e18 (1x)	14M	81.5	88.5	88.0	62.2	89.2	83.3	54.8	83.4	78.9
train longer for 4M updates	5.2e18 (4x)	14M	82.1	88.5	88.4	65.1	91.2	85.1	56.7	83.8	80.1
ConvBERTMEDIUM-SMALL	1.5e18 (1.2x)	17M	82.1	88.7	88.4	65.3	89.2	84.6	56.4	82.9	79.7
train longer for 4M updates	6.0e18 (4.6x)	17M	82.9	89.2	88.6	66.0	92.3	85.7	59.1	85.3	81.1
ConvBERTBASE	1.9e19 (15x)	106M	85.3	92.4	89.6	74.7	95.0	88.2	66.0	88.2	84.9
train longer for 4M updates	7.6e19 (59x)	106M	88.3	93.2	90.0	77.9	95.7	88.3	67.8	89.7	86.4

Results on SQuAD We also evaluate our model on question answering task benchmark SQuAD [52]. Table 4 shows the results of our proposed ConvBERT as well as other methods [53, 25, 60, 14, 9] with similar model size. For small model size, our ConvBERTSMALL and ConvBERTMEDIUM-SMALL outperform the baseline ELECTRASMALL and achieve results comparable to BERTBASE. The results of MobileBERT are much higher since they use knowledge distillation and search for model architecture and hyper-parameters based on the development set of SQuAD. With much less pre-training cost, our base-sized model outperforms all other models of similar size.

Table 4: Comparison of models with similar model size on SQuAD dev set. Pre-training computation cost is also reported. * denotes results obtained by running official code [9]. † denotes knowledge distillation based methods.

Model	Train FLOPs	Params	SQuAD v1.1		SQuAD v2.0	
			EM	F1	EM	F1
DistillBERT [†] [53]	6.4e19+ (49x+)	52M	71.8	81.2	60.6	64.1
TinyBERT [†] [25]	6.4e19+ (49x+)	15M	72.7	82.1	65.3	68.8
MobileBERT [†] [60]	6.4e19+ (49x+)	25M	83.4	90.3	77.6	80.2
ELECTRASMALL [9]	1.4e18 (1.1x)	14M	65.6*	73.8*	62.8*	65.2*
train longer [9]	3.3e19 (25x)	14M	75.8	85.2*	70.1	73.2*
BERTBASE [14]	6.4e19 (49x)	110M	80.7	88.4	74.2	77.1
ELECTRABASE [9]	6.4e19 (49x)	110M	84.5	90.8	80.5	83.3
ConvBERTSMALL	1.3e18 (1x)	14M	77.1	84.1	68.7	70.5
train longer	5.2e18 (4x)	14M	78.8	85.5	71.2	73.6
ConvBERTMEDIUM-SMALL	1.5e18 (1.2x)	17M	79.4	86.0	71.7	74.3
train longer	6.0e18 (4.6x)	17M	81.5	88.1	74.3	76.8
ConvBERTBASE	1.9e19 (15x)	106M	84.7	90.9	80.6	83.1

5 Conclusion

We present a novel span-based dynamic convolution operator and integrate it into the self-attention mechanism to form our mixed attention block for language pre-training. We also devise a bottleneck structure applied to the self-attention module and a grouped linear operation for the feed-forward module. Experiment results show that our ConvBERT model integrating above novelties achieves consistent performance improvements while costing much less pre-training computation.

Broader impact

Positive impact The pre-training scheme has been widely deployed in the natural language processing field. It proposes to train a large model by self-supervised learning on large corpus at first and then fine-tune the model on downstream tasks quickly. Such a pre-training scheme has produced a series of powerful language models and BERT is one of the most popular one. In this work, we developed a new pre-training based language understanding model, ConvBERT. It offers smaller model size, lower training cost and better performance, compared with the BERT model. ConvBERT has multiple positive impacts. In contrary to the trend of further increasing model complexity for better performance, ConvBERT turns to making the model more efficient and saving the training cost. It will benefit the applications where the computation resource is limited. In terms of the methodology, it looks into the model backbone designs, instead of using distillation-alike algorithms that still require training a large teacher model beforehand, to make the model more efficient. We encourage researchers to build NLP models based on ConvBERT for tasks we can expect to be particularly beneficial, such as text-based counselling.

Negative impact Compared with BERT, ConvBERT is more efficient and saves the training cost, which can be used to detect and understand personal text posts on social platforms and brings privacy threat.

Acknowledgement

We would like to thank the anonymous reviewers for their insightful comments and suggestions; Jiashi Feng was partially supported by AISG-100E-2019-035, MOE2017-T2-2-151, NUS-ECRA-FY17-P08 and CRP20-2017-0006. Weihao Yu and Zihang Jiang would like to thank TFRC program for the support of computational resources.

References

- [1] Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. The fifth pascal recognizing textual entailment challenge. In *TAC*, 2009.
- [2] Gino Brunner, Yang Liu, Damian Pascual, Oliver Richter, Massimiliano Ciaramita, and Roger Wattenhofer. On identifiability in transformers. In *International Conference on Learning Representations*, 2020.
- [3] Jamie Callan, Mark Hoy, Changkuk Yoo, and Le Zhao. Clueweb09 data set, 2009.
- [4] Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada, August 2017. Association for Computational Linguistics.
- [5] Yen-Chun Chen, Linjie Li, Licheng Yu, Ahmed El Kholy, Faisal Ahmed, Zhe Gan, Yu Cheng, and Jingjing Liu. Uniter: Learning universal image-text representations. *arXiv preprint arXiv:1909.11740*, 2019.
- [6] Zihan Chen, Hongbo Zhang, Xiaoji Zhang, and Leqi Zhao. Quora question pairs, 2018.
- [7] Zewen Chi, Li Dong, Furu Wei, Wenhui Wang, Xian-Ling Mao, and Heyan Huang. Cross-lingual natural language generation via pre-training. *arXiv preprint arXiv:1909.10481*, 2019.
- [8] Yung-Sung Chuang, Chi-Liang Liu, and Hung-Yi Lee. Speechbert: Cross-modal pre-trained language model for end-to-end spoken question answering. *arXiv preprint arXiv:1910.11559*, 2019.
- [9] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text encoders as discriminators rather than generators. In *International Conference on Learning Representations*, 2019.
- [10] Alexis Conneau and Guillaume Lample. Cross-lingual language model pretraining. In *Advances in Neural Information Processing Systems*, pages 7057–7067, 2019.
- [11] Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pages 177–190. Springer, 2005.
- [12] Andrew M Dai and Quoc V Le. Semi-supervised sequence learning. In *Advances in neural information processing systems*, pages 3079–3087, 2015.
- [13] Yann N Dauphin, Angela Fan, Michael Auli, and David Grangier. Language modeling with gated convolutional networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 933–941. JMLR. org, 2017.
- [14] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 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, 2019.
- [15] William B. Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*, 2005.
- [16] Cícero dos Santos and Maíra Gatti. Deep convolutional neural networks for sentiment analysis of short texts. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 69–78, Dublin, Ireland, August 2014. Dublin City University and Association for Computational Linguistics.
- [17] Angela Fan, Edouard Grave, and Armand Joulin. Reducing transformer depth on demand with structured dropout. In *International Conference on Learning Representations*, 2020.
- [18] Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N Dauphin. Convolutional sequence to sequence learning. *arXiv preprint arXiv:1705.03122*, 2017.
- [19] Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. The third pascal recognizing textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and paraphrasing*, pages 1–9. Association for Computational Linguistics, 2007.
- [20] Aaron Gokaslan and Vanya Cohen. Openwebtext corpus, 2019.
- [21] Mitchell A Gordon, Kevin Duh, and Nicholas Andrews. Compressing bert: Studying the effects of weight pruning on transfer learning. *arXiv preprint arXiv:2002.08307*, 2020.

- [22] R Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor. The second pascal recognising textual entailment challenge. In *Proceedings of the Second PASCAL Challenges Workshop on Recognising Textual Entailment*, 2006.
- [23] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [24] Haoyang Huang, Yaobo Liang, Nan Duan, Ming Gong, Linjun Shou, Daxin Jiang, and Ming Zhou. Unicoder: A universal language encoder by pre-training with multiple cross-lingual tasks. 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 2485–2494, 2019.
- [25] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019.
- [26] Mandar Joshi, Danqi Chen, Yinhan Liu, Daniel S. Weld, Luke Zettlemoyer, and Omer Levy. Spanbert: Improving pre-training by representing and predicting spans. *Trans. Assoc. Comput. Linguistics*, 8:64–77, 2020.
- [27] Lukasz Kaiser, Aidan N Gomez, and Francois Chollet. Depthwise separable convolutions for neural machine translation. *arXiv preprint arXiv:1706.03059*, 2017.
- [28] Nal Kalchbrenner, Edward Grefenstette, and Phil Blunsom. A convolutional neural network for modelling sentences. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 655–665, Baltimore, Maryland, June 2014. Association for Computational Linguistics.
- [29] Yoon Kim. Convolutional neural networks for sentence classification. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1746–1751, Doha, Qatar, October 2014. Association for Computational Linguistics.
- [30] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- [31] Olga Kovaleva, Alexey Romanov, Anna Rogers, and Anna Rumshisky. Revealing the dark secrets of bert. 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 4356–4365, 2019.
- [32] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *NeurIPS*, 2012.
- [33] Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. RACE: Large-scale ReAding comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark, September 2017. Association for Computational Linguistics.
- [34] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. ALBERT: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*, 2019.
- [35] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [36] Hector Levesque, Ernest Davis, and Leora Morgenstern. The winograd schema challenge. In *Thirteenth International Conference on the Principles of Knowledge Representation and Reasoning*, 2012.
- [37] Weijie Liu, Peng Zhou, Zhe Zhao, Zhiruo Wang, Qi Ju, Haotang Deng, and Ping Wang. K-bert: Enabling language representation with knowledge graph. In *AAAI*, pages 2901–2908, 2020.
- [38] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*, 2019.
- [39] Jiasen Lu, Dhruv Batra, Devi Parikh, and Stefan Lee. Vilbert: Pretraining task-agnostic visiolinguistic representations for vision-and-language tasks. In *Advances in Neural Information Processing Systems*, pages 13–23, 2019.
- [40] Bryan McCann, James Bradbury, Caiming Xiong, and Richard Socher. Learned in translation: Contextualized word vectors. In *Advances in Neural Information Processing Systems*, pages 6294–6305, 2017.

- [41] Paul Michel, Omer Levy, and Graham Neubig. Are sixteen heads really better than one? In *Advances in Neural Information Processing Systems*, pages 14014–14024, 2019.
- [42] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. In *Advances in neural information processing systems*, pages 3111–3119, 2013.
- [43] Sebastian Nagel. Cc-news. <http://web.archive.org/save/http://commoncrawl.org/2016/10/news-dataset-available>, 2016.
- [44] Robert Parker, David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. English gigaword fifth edition, linguistic data consortium. *Google Scholar*, 2011.
- [45] Jeffrey Pennington, Richard Socher, and Christopher Manning. Glove: Global vectors for word representation. In *EMNLP*, 2014.
- [46] Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. Deep contextualized word representations. In *Proc. of NAACL*, 2018.
- [47] Matthew E Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A Smith. Knowledge enhanced contextual word representations. 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 43–54, 2019.
- [48] Xipeng Qiu, Tianxiang Sun, Yige Xu, Yunfan Shao, Ning Dai, and Xuanjing Huang. Pre-trained models for natural language processing: A survey. *arXiv preprint arXiv:2003.08271*, 2020.
- [49] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language understanding by generative pre-training. URL https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf, 2018.
- [50] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.
- [51] Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don’t know: Unanswerable questions for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 784–789, 2018.
- [52] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas, November 2016. Association for Computational Linguistics.
- [53] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [54] Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Q-bert: Hessian based ultra low precision quantization of bert. *arXiv preprint arXiv:1909.05840*, 2019.
- [55] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [56] David So, Quoc Le, and Chen Liang. The evolved transformer. In *International Conference on Machine Learning*, pages 5877–5886, 2019.
- [57] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642, 2013.
- [58] Weijie Su, Xizhou Zhu, Yue Cao, Bin Li, Lewei Lu, Furu Wei, and Jifeng Dai. Vi-bert: Pre-training of generic visual-linguistic representations. In *International Conference on Learning Representations*, 2020.
- [59] Chen Sun, Austin Myers, Carl Vondrick, Kevin Murphy, and Cordelia Schmid. Videobert: A joint model for video and language representation learning. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 7464–7473, 2019.

- [60] Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. Mobilebert: a compact task-agnostic bert for resource-limited devices. *arXiv preprint arXiv:2004.02984*, 2020.
- [61] Trieu H Trinh and Quoc V Le. A simple method for commonsense reasoning. *arXiv preprint arXiv:1806.02847*, 2018.
- [62] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in neural information processing systems*, pages 5998–6008, 2017.
- [63] Alex Wang, Amapreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. GLUE: A multi-task benchmark and analysis platform for natural language understanding. In *ICLR*, 2019.
- [64] Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641, 2019.
- [65] Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.
- [66] Felix Wu, Angela Fan, Alexei Baevski, Yann N Dauphin, and Michael Auli. Pay less attention with lightweight and dynamic convolutions. *arXiv preprint arXiv:1901.10430*, 2019.
- [67] Zhanghao Wu*, Zhijian Liu*, Ji Lin, Yujun Lin, and Song Han. Lite transformer with long-short range attention. In *International Conference on Learning Representations*, 2020.
- [68] Wenhan Xiong, Jingfei Du, William Yang Wang, and Veselin Stoyanov. Pretrained encyclopedia: Weakly supervised knowledge-pretrained language model. In *International Conference on Learning Representations*, 2020.
- [69] Zhilin Yang, Zihang Dai, Yiming Yang, Jaime Carbonell, Ruslan Salakhutdinov, and Quoc V Le. XLNet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*, 2019.
- [70] Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. Q8bert: Quantized 8bit bert. *arXiv preprint arXiv:1910.06188*, 2019.
- [71] Daojian Zeng, Kang Liu, Siwei Lai, Guangyou Zhou, and Jun Zhao. Relation classification via convolutional deep neural network. In *Proceedings of COLING 2014, the 25th International Conference on Computational Linguistics: Technical Papers*, pages 2335–2344, Dublin, Ireland, August 2014. Dublin City University and Association for Computational Linguistics.
- [72] Xiang Zhang, Junbo Zhao, and Yann LeCun. Character-level convolutional networks for text classification. In *Advances in neural information processing systems*, pages 649–657, 2015.
- [73] Zhengyan Zhang, Xu Han, Zhiyuan Liu, Xin Jiang, Maosong Sun, and Qun Liu. Ernie: Enhanced language representation with informative entities. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 1441–1451, 2019.
- [74] Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *Proceedings of the IEEE international conference on computer vision*, pages 19–27, 2015.

6 Appendix

6.1 Datasets

6.1.1 GLUE dataset

GLUE benchmark introduced by [63] is a collection of nine natural language understanding tasks. The authors hide the labels of testing set and researchers need to submit their predictions to the evaluation server⁵ to obtain results on testing sets. We only present results of single-task setting for fair comparison. The GLUE benchmark includes the following datasets.

MNLI The Multi-Genre Natural Language Inference Corpus [65] is a dataset of sentence pairs with textual entailment annotations. Given a premise sentence and a hypothesis sentence, the task is to predict their relationships including ENTAILMENT, CONTRADICTION and NEUTRAL. The data is from ten distinct genres of written and spoken English.

QNLI Question Natural Language Inference is a binary sentence pair classification task converted from The Stanford Question Answering Dataset [52], a question-answering dataset. An example of QNLI contains a context sentence and a question, and the task is to determine whether the context sentence contains the answer to the question.

QQP The Quora Question Pairs dataset [6] is a collection of question pairs from Quora, a community question-answering website, and the task is to determine whether a pair of questions are semantically equivalent.

RTE The Recognizing Textual Entailment (RTE) dataset is similar to MNLI which only has two classes, i.e., *entailment* and *not entailment*. It is from a series of annual textual entailment challenges including RTE1 [11], RTE2 [22], RTE3 [19], and RTE5 [1].

SST-2 The Stanford Sentiment Treebank [57] is a dataset that consists of sentences from movie reviews and human annotations of their sentiment. GLUE uses the two-way (POSITIVE/NEGATIVE) class split.

MRPC The Microsoft Research Paraphrase Corpus [15] is a dataset from online news that consists of sentence pairs with human annotations for whether the sentences in the pair are semantically equivalent.

CoLA The Corpus of Linguistic Acceptability [64] is a binary single-sentence classification dataset containing the examples annotated with whether it is a grammatical English sentence.

SST-B The Semantic Textual Similarity Benchmark [4] is a collection of sentence pairs human-annotated with a similarity score from 1 to 5, in which models are required to predict the scores.

WNLI Winograd NLI [36] is a small natural language inference dataset, but as GLUE web page⁶ noted, there are issues with the construction of it. Thus like previous works, GPT [49] and BERT [34] etc., we exclude this dataset for fair comparison.

6.1.2 SQuAD dataset

The Stanford Question Answering Dataset (**SQuAD v1.1**), a question answering (reading comprehension) dataset which consists of more than 100K questions. The answer to each question is a span of text from the corresponding context passage, meaning that every question can be answered. Then the following version **SQuAD v2.0** combines the existing data with over 50K unanswerable questions.

⁵<https://gluebenchmark.com>

⁶<https://gluebenchmark.com/faq>

6.2 Pre-training details

We first give a brief introduction to the replaced token detection task we used for pre-training proposed by [9]. It trains the model in a discriminative way by predicting whether the token in the sequence is replaced. Meanwhile, to generate the sentence with replaced tokens as training example, they propose to use a small-sized generator trained with masked language modelling [14]. The full input sequence is first masked and then feed to the generator to get the prediction of the masked tokens. The target model then serves as a discriminator to distinguish the tokens that are wrongly predicted by the generator. The generator and the discriminator are jointly trained with masked language modelling loss and replaced token detection loss.

For the pre-training configuration, we mostly use the same hyper-parameters as ELECTRA [9]. See Table 5 for more details. While using examples of 128 sequence length for pre-training can save a lot of computation, we also find that using examples with longer sequence length can help to improve the performance on downstream task that has longer context. We pre-train our model with input sequence of length 512 for the 10% more updates before fine-tuning it for task with longer context like SQuAD. This helps the positional embedding generalize better to the downstream tasks.

Table 5: Pre-training hyper-parameters. Generator size here is the multiplier for hidden size, feed-forward inner hidden size and attention heads to compute configuration for generator. The optimizer used here is an Adam optimizer [30], and details of the optimizer are listed in the table.

Hyper-parameter	Small	Medium-small	Base
Layer	12	12	12
Hidden dim	256	384	768
Word Embedding dim	128	128	768
feed-forward inner hidden size	1024	1536	3072
Generator size	1/4	1/4	1/3
Attention heads	2	4	6
Attention head size	64	48	64
Learning rate	3e-4	5e-4	2e-4
Learning rate decay	Linear	Linear	Linear
Warmup steps	10k	10k	10k
Adam ϵ	1e-6	1e-6	1e-6
Adam β_1	0.9	0.9	0.9
Adam β_2	0.999	0.999	0.999
Dropout	0.1	0.1	0.1
Batch size	128	128	256
Input sequence length	128	128	128

6.3 Fine-tuning details

Following previous work [9, 14], we search for learning rate among {5e-5, 1e-4, 2e-4, 3e-4} and weight decay among {0.01, 0.1}. For the number of training epoch, we search for the best among {10, 3}. All other parameters are kept the same as [9]. See Table 6.

6.4 More results

We present more results on GLUE dev set with different model sizes and pre-training settings in Table 7. As can be seen, regardless of the pre-training task and dataset size, our method consistently outperform the original BERT [14] architecture.

6.5 More examples and analysis of attention map

We provide more examples of the attention map in Figure 6. we also compute the diagonal concentration for the attention map M as quantitative metric. It is define as $C = \frac{\sum_{|i-j|\leq 4} M_{i,j}}{\sum_{|i-j|>4} M_{i,j}}$. This indicates

Table 6: Fine-tuning hyper-parameters. The optimizer used here is an Adam optimizer [30], and details of the optimizer are listed in the table.

Hyper-parameter	Value
Adam ϵ	1e-6
Adam β_1	0.9
Adam β_2	0.999
Layer-wise LR decay	0.8
Learning rate decay	Linear
Warmup fraction	0.1
Dropout	0.1
Batch size	32

Table 7: Comparison of our proposed ConvBERT architecture with the transformer based BERT architecture in different sizes and different pre-training settings. GLUE score represents the average score of 8 tasks on GLUE development set. MLM represents masked language modelling and RTD represents replaced token detection. The 16G WikiBooks dataset is the combination of EnWiki and BOOKCORPUS, 32G represents the OpenWebText dataset proposed by [50, 20], and 160G represents the combination of several corpus datasets used by ELECTRA [9] and RoBERTa [38]. * denotes the results from ELECTRA and + denotes the result from RoBERTa.

Model	Pre-train task	Training data	update	Train FLOPs	Params	GLUE
BERTSMALL	MLM	16G	1.45M	1.4e18	14M	75.1*
	RTD	16G	1M	1.4e18	14M	79.7*
	RTD	32G	1M	1.4e18	14M	80.3*
	RTD	160G	4M	3.3e19	14M	81.1*
ConvBERTSMALL	MLM	16G	1.45M	1.3e18	14M	75.9
	RTD	16G	1M	1.3e18	14M	80.6
	RTD	32G	1M	1.3e18	14M	81.4
	RTD	32G	4M	5.2e18	14M	81.8
ConvBERTSMALL-PLUS	RTD	32G	1M	1.5e18	17M	82.1
	RTD	32G	4M	6.0e18	17M	82.8
BERTBASE	MLM	16G	1M	6.4e19	110M	82.2*
	MLM	160G	500k	1.0e21	125M	86.4+
	RTD	16G	766k	6.4e19	110M	85.1*
	RTD	160G	4M	3.3e20	110M	87.5*
ConvBERTBASE	RTD	32G	1M	1.9e19	106M	86.0
	RTD	32G	4M	7.6e19	106M	87.7

how much local dependency that the attention map captures. The result in Table 8 shows that the attention in BERT concentrates more on the local dependency.

Table 8: Average concentration on MRPC.

Model	C (diagonal-concentration)
BERT	0.941
ConvBERT	0.608

6.6 Inference speed

We test our mixed-attention block and self-attention baseline from base-sized model on Intel CPU (i7-6900K@3.20GHz). The mixed-attention has lower Flops and is faster than self-attention, as shown in Table 9. On the other hand, our implementation for mixed-attention on GPU and TPU is

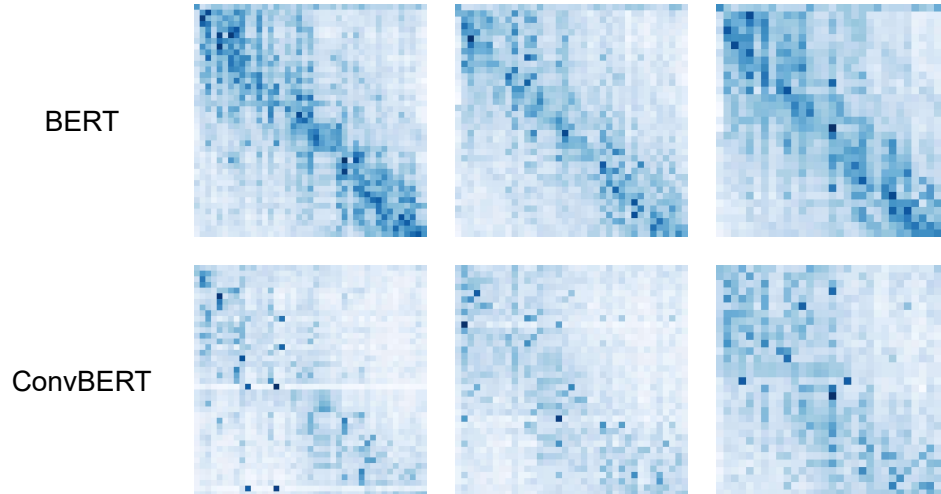


Figure 6: More examples of attention maps.

not well optimized for the efficiency yet. Thus its acceleration may not be obvious when the input sequence length is short. We will work on further improvement on the low-level implementation.

Table 9: Inference speed.

Block	Flops	Speed (ms/sample)
self-attention	2.6e9	17.66
mixed-attention	1.9e9	12.94