

1 Appendix

2 1.1 Datasets

3 1.1.1 GLUE dataset

4 GLUE benchmark introduced by [19] is a collection of nine natural language understanding tasks.
5 The authors hide the labels of testing set and researchers need to submit their predictions to the
6 evaluation server¹ to obtain results on testing sets. We only present results of single-task setting for
7 fair comparison. The GLUE benchmark includes the following datasets.

8 **MNLI** The Multi-Genre Natural Language Inference Corpus [21] is a dataset of sentence pairs
9 with textual entailment annotations. Given a premise sentence and a hypothesis sentence, the task is
10 to predict their relationships including ENTAILMENT, CONTRADICTION and NEUTRAL. The data
11 is from ten distinct genres of written and spoken English.

12 **QNLI** Question Natural Language Inference is a binary sentence pair classification task converted
13 from The Stanford Question Answering Dataset [17], a question-answering dataset. An example of
14 QNLI contains a context sentence and a question, and the task is to determine whether the context
15 sentence contains the answer to the question.

16 **QQP** The Quora Question Pairs dataset [3] is a collection of question pairs from Quora, a com-
17 munity question-answering website, and the task is to determine whether a pair of questions are
18 semantically equivalent.

19 **RTE** The Recognizing Textual Entailment (RTE) dataset is similar to MNLI which only has two
20 classes, i.e., *entailment* and *not entailment*. It is from a series of annual textual entailment challenges
21 including RTE1 [5], RTE2 [10], RTE3 [8], and RTE5 [1].

22 **SST-2** The Stanford Sentiment Treebank [18] is a dataset that consists of sentences from movie
23 reviews and human annotations of their sentiment. GLUE uses the two-way (POSITIVE/NEGATIVE)
24 class split.

25 **MRPC** The Microsoft Research Paraphrase Corpus [7] is a dataset from online news that consists
26 of sentence pairs with human annotations for whether the sentences in the pair are semantically
27 equivalent.

28 **CoLA** The Corpus of Linguistic Acceptability [20] is a binary single-sentence classification dataset
29 containing the examples annotated with whether it is a grammatical English sentence.

30 **SST-B** The Semantic Textual Similarity Benchmark [2] is a collection of sentence pairs human-
31 annotated with a similarity score from 1 to 5, in which models are required to predict the scores.

32 **WNLI** Winograd NLI [13] is a small natural language inference dataset, but as GLUE web page²
33 noted, there are issues with the construction of it. Thus like previous works, GPT [15] and BERT
34 [12] etc., we exclude this dataset for fair comparison.

35 1.1.2 SQuAD dataset

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

¹<https://gluebenchmark.com>

²<https://gluebenchmark.com/faq>

40 1.2 Pre-training details

41 We first give a brief introduction to the replaced token detection task we used for pre-training proposed
42 by [4]. It trains the model in a discriminative way by predicting whether the token in the sequence is
43 replaced. Meanwhile, to generate the sentence with replaced tokens as training example, they propose
44 to use a small-sized generator trained with masked language modelling [6]. The full input sequence
45 is first masked and then feed to the generator to get the prediction of the masked tokens. The target
46 model then serves as a discriminator to distinguish the tokens that are wrongly predicted by the
47 generator. The generator and the discriminator are jointly trained with masked language modelling
48 loss and replaced token detection loss.

49 For the pre-training configuration, we mostly use the same hyper-parameters as ELECTRA [4]. See
50 Table 1 for more details. While using examples of 128 sequence length for pre-training can save a lot
51 of computation, we also find that using examples with longer sequence length can help to improve
52 the performance on downstream task that has longer context. We pre-train our model with input
53 sequence of length 512 for the 10% more updates before fine-tuning it for task with longer context
54 like SQuAD. This helps the positional embedding generalize better to the downstream tasks.

Table 1: 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 [11], 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

55 1.3 Fine-tuning details

56 Following previous work [4, 6], we search for learning rate among {5e-5, 1e-4, 2e-4, 3e-4} and
57 weight decay among {0.01, 0.1}. For the number of training epoch, we search for the best among
58 {10, 3}. All other parameters are kept the same as [4]. See Table 2.

59 1.4 More results

60 We present more results on GLUE dev set with different model sizes and pre-training settings in
61 Table 3. As can be seen, regardless of the pre-training task and dataset size, our method consistently
62 outperform the original BERT [6] architecture.

63 1.5 More examples and analysis of attention map

64 We provide more examples of the attention map in Figure 1. we also compute the diagonal concentra-
65 tion 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 2: Fine-tuning hyper-parameters. The optimizer used here is an Adam optimizer [11], 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 3: 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 [16, 9], and 160G represents the combination of several corpus datasets used by ELECTRA [4] and RoBERTa [14]. * 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	1.4e19	106M	86.0
	RTD	32G	4M	5.6e19	106M	87.7

66 how much local dependency that the attention map captures. The result in Table 4 shows that the
67 attention in BERT concentrates more on the local dependency.

Table 4: Average concentration on MRPC.

Model	C (diagonal-concentration)
BERT	0.941
ConvBERT	0.608

68 1.6 Inference speed

69 We test our mixed-attention block and self-attention baseline from base-sized model on Intel CPU
70 (i7-6900K@3.20GHz). The mixed-attention has lower Flops and is much faster than self-attention,
71 as shown in Table 5. On the other hand, our implementation for mixed-attention on GPU and TPU is

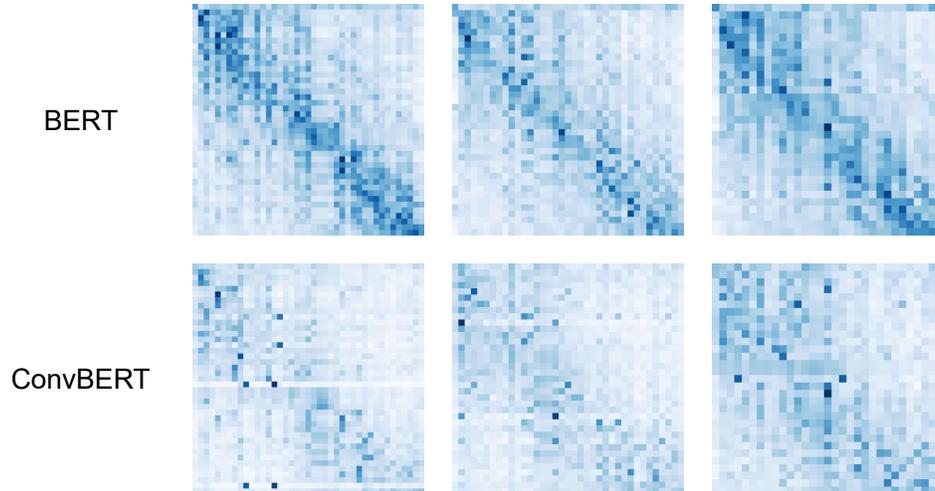


Figure 1: More examples of attention maps.

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

Table 5: Inference speed.

Block	Flops	Speed (ms/sample)
self-attention	26.5G	17.66
mixed-attention	19.3G	12.94

74 References

- 75 [1] Luisa Bentivogli, Peter Clark, Ido Dagan, and Danilo Giampiccolo. The fifth pascal recognizing textual
76 entailment challenge. In *TAC*, 2009.
- 77 [2] Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. SemEval-2017 task 1:
78 Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th*
79 *International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada, August
80 2017. Association for Computational Linguistics.
- 81 [3] Zihan Chen, Hongbo Zhang, Xiaoji Zhang, and Leqi Zhao. Quora question pairs, 2018.
- 82 [4] Kevin Clark, Minh-Thang Luong, Quoc V Le, and Christopher D Manning. Electra: Pre-training text
83 encoders as discriminators rather than generators. In *International Conference on Learning Representations*,
84 2019.
- 85 [5] Ido Dagan, Oren Glickman, and Bernardo Magnini. The pascal recognising textual entailment challenge.
86 In *Machine Learning Challenges Workshop*, pages 177–190. Springer, 2005.
- 87 [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep bidi-
88 rectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North*
89 *American Chapter of the Association for Computational Linguistics: Human Language Technologies*,
90 *Volume 1 (Long and Short Papers)*, pages 4171–4186, 2019.
- 91 [7] William B. Dolan and Chris Brockett. Automatically constructing a corpus of sentential paraphrases. In
92 *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*, 2005.
- 93 [8] Danilo Giampiccolo, Bernardo Magnini, Ido Dagan, and Bill Dolan. The third pascal recognizing
94 textual entailment challenge. In *Proceedings of the ACL-PASCAL workshop on textual entailment and*
95 *paraphrasing*, pages 1–9. Association for Computational Linguistics, 2007.
- 96 [9] Aaron Gokaslan and Vanya Cohen. Openwebtext corpus, 2019.
- 97 [10] R Bar Haim, Ido Dagan, Bill Dolan, Lisa Ferro, Danilo Giampiccolo, Bernardo Magnini, and Idan Szpektor.
98 The second pascal recognising textual entailment challenge. In *Proceedings of the Second PASCAL*
99 *Challenges Workshop on Recognising Textual Entailment*, 2006.
- 100 [11] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *ICLR*, 2015.
- 101 [12] Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Sori-
102 cut. ALBERT: A lite bert for self-supervised learning of language representations. *arXiv preprint*
103 *arXiv:1909.11942*, 2019.
- 104 [13] Hector Levesque, Ernest Davis, and Leora Morgenstern. The winograd schema challenge. In *Thirteenth*
105 *International Conference on the Principles of Knowledge Representation and Reasoning*, 2012.
- 106 [14] Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis,
107 Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining approach. *arXiv*
108 *preprint arXiv:1907.11692*, 2019.
- 109 [15] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving language
110 understanding by generative pre-training. URL [https://s3-us-west-2. amazonaws. com/openai-](https://s3-us-west-2.amazonaws.com/openai-assets/researchcovers/languageunsupervised/language_understanding_paper.pdf)
111 *assets/researchcovers/languageunsupervised/language_understanding_paper. pdf*, 2018.
- 112 [16] Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. Language models
113 are unsupervised multitask learners. *OpenAI Blog*, 1(8):9, 2019.
- 114 [17] Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. SQuAD: 100,000+ questions for
115 machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural*
116 *Language Processing*, pages 2383–2392, Austin, Texas, November 2016. Association for Computational
117 Linguistics.
- 118 [18] Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and
119 Christopher Potts. Recursive deep models for semantic compositionality over a sentiment treebank.
120 In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages
121 1631–1642, 2013.
- 122 [19] Alex Wang, Amapreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. GLUE: A
123 multi-task benchmark and analysis platform for natural language understanding. In *ICLR*, 2019.
- 124 [20] Alex Warstadt, Amanpreet Singh, and Samuel R Bowman. Neural network acceptability judgments.
125 *Transactions of the Association for Computational Linguistics*, 7:625–641, 2019.
- 126 [21] Adina Williams, Nikita Nangia, and Samuel Bowman. A broad-coverage challenge corpus for sentence
127 understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of*
128 *the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*,
129 pages 1112–1122, New Orleans, Louisiana, June 2018. Association for Computational Linguistics.