
Undivided Attention: Are Intermediate Layers Necessary for BERT?

Sharath Nittur Sridhar

Intel Labs

sharath.nittur.sridhar@intel.com

Anthony Sarah

Intel Labs

anthony.sarah@intel.com

Abstract

In recent times, BERT-based models have been extremely successful in solving a variety of natural language processing (NLP) tasks such as reading comprehension, natural language inference, sentiment analysis, etc. All BERT-based architectures have a self-attention block followed by a block of intermediate layers as the basic building component. However, a strong justification for the inclusion of these intermediate layers remains missing in the literature. In this work we investigate the importance of intermediate layers on the overall network performance of downstream tasks. We show that reducing the number of intermediate layers and modifying the architecture for BERT_{BASE} results in minimal loss in fine-tuning accuracy for downstream tasks while decreasing the number of parameters and training time of the model. Additionally, we use centered kernel alignment and probing linear classifiers to gain insight into our architectural modifications and justify that removal of intermediate layers has little impact on the fine-tuned accuracy.

1 Introduction

Language model pre-training has led to a number of breakthroughs in NLP (Devlin et al. (2018), Brown et al. (2020), Peters et al. (2018), Liu et al. (2019)) and have achieved state-of-the-art results on many non-trivial NLP tasks such as question-answering or natural language inference. In this problem domain, the BERT model, based on the Vaswani et al. (2017) architecture has risen to prominence.

The BERT model architecture has multiple bidirectional Transformer (Vaswani et al. (2017)) encoder blocks stacked together. Training BERT consists of first pre-training on large corpus of unlabeled data followed by fine-tuning on a much smaller, task-specific data set. Pre-training is computationally expensive and often requires several days to complete while fine-tuning can be completed in a much shorter time.

BERT models have significant parameter count and memory footprint, from 110M (BERT_{BASE}) to 340M (BERT_{LARGE}) to 8.3B (Megatron Shoeybi et al. (2019)) parameters. In this work, we demonstrate decreased model size and decreased training time with a negligible change in fine-tuning accuracy. While demonstrating decreased model size and decreased training time is relevant to the field in and of itself, we feel it is also important to provide a deeper understanding of why our architectural modifications are effective. Similar to the work done in Clark et al. (2019) we use probing linear classifiers and centered kernel alignment (Kornblith et al. (2019)) to gain insight into the improvements made by our architectural modifications.

2 Related Work

Approaches such as Sanh et al. (2020), Sun et al. (2020), Jiao et al. (2020), Sun et al. (2019) and Xu et al. (2020) use knowledge distillation techniques to reduce the parameter count in BERT and improve training speed. The authors of DistilBERT (Sanh et al. (2020)) decreased the model size by 40% and improved the training speed by 60%, with minimal loss in accuracy. Similarly, the authors of TinyBERT (Jiao et al. (2020)) employ a two-stage distillation framework, that performs distillation on both the pre-training and task specific learning stage to decrease the model size by 7.5x. BERT-PKD (Patient Knowledge Distillation) (Sun et al. (2019)) learns from multiple layers of the teacher model for incremental knowledge extraction, whereas the BERT-of-Theseus approach (Xu et al. (2020)) progressively substitutes modules of BERT with modules having fewer parameters. However, all knowledge distillation techniques require a pre-trained teacher model, which may not be available in many cases.

Other approaches, such as ALBERT (Lan et al. (2020)), employ a number of techniques, such as decomposing the vocabulary matrix and cross-layer parameter sharing, to achieve similar accuracy to BERT but with 18x fewer parameters and 1.7x faster training. The authors of Press et al. (2020) changed the ordering of the sub-layers to create the *sandwich transformer*. Like us, the authors recognize that the ordering of the sub-layers in BERT networks is not well justified nor necessarily optimal. However, unlike ALBERT, their goal is to improve language modeling performance and not to decrease the network size or complexity. In this work, we share similar goals to ALBERT but employ a different approach that removes blocks of intermediate layers from the network.

3 Architectural Modifications

BERT Devlin et al. (2018) is a stack of multiple Transformer encoder blocks with each encoder block further consisting of a separate multi-head self-attention block followed by an intermediate block (Figure 2). The importance of the self-attention block has been extensively analyzed in Clark et al. (2019), Michel et al. (2019), Tenney et al. (2019) and others.

An intermediate block is a four-layer feed-forward network containing a Gaussian error linear unit (GELU) (Hendrycks and Gimpel (2020)) in between two linear layers followed by a dropout layer (see *Intermediate Block* of Figure 1). The intermediate block was added mainly to enrich the representations obtained from the self-attention block. However, the relevance of this block has not been well studied and no strong justification for their inclusion in Transformer-type architectures has been made in the literature.

Our main motivation is to decrease model size and complexity while quantifying and understanding any negative impact of reducing the number of intermediate blocks on model accuracy. To this end, we modify the BERT architecture by removing some of the intermediate blocks within the network. More specifically, an intermediate block will be added only after every n self-attention blocks. If the total number of self-attention blocks in the network is m then the modified network will contain $\lfloor \frac{m}{n} \rfloor$ intermediate blocks. For example, when $n = 1$ the network is unmodified and contains m intermediate blocks. However, when $n = \infty$ the modified network contains no intermediate blocks. See Figure 3 for an example of the modified BERT_{BASE} network with $n = 2$.

Note that n is an architectural hyper-parameter that can be changed to make trade-offs in network size, complexity and accuracy. We experiment with different values of n and analyze its effects with multiple fine-tuning tasks.

4 Results

In this section we present the convergence and accuracy results for $n \in \{1, 2, 3, 4, 6, \infty\}$ where $n = 1$ represents the unmodified network and $n = \infty$ represents the modified network with no intermediate blocks. We chose BERT_{BASE} as our reference model which has 12 self-attention blocks each of which is followed by an intermediate block with a hidden size of 768. We perform all our experiments on BERT_{BASE}, mainly due to limited computational resources and the extremely long times required to train larger models. Table 1 shows the results of fine-tuning on downstream tasks along with the total number of parameters and the relative decrease in size and relative increase in throughput (measured as examples/sec) of the unmodified and modified BERT_{BASE} networks.




Figure 1: Reference unit in the unmodified BERT network. The unmodified BERT_{BASE} network contains 12 of these units arranged sequentially.




Figure 2: First six blocks of the unmodified network architecture. This BERT_{BASE} network will have a total of 12 self-attention blocks and 12 intermediate blocks.

4.1 Pre-Training

All variants were pre-trained from scratch on the English Wikipedia and BookCorpus (Zhu et al. (2015)) data sets. Similar to Devlin et al. (2018), we follow a two-phase pre-training approach. For the first 90% of iterations, the model is trained with a sequence length of 128 and the remaining 10% is trained with a sequence length of 512. To avoid having to search for optimal training hyper-parameters, we chose to use the same values for all variants. An initial learning rate of 10^{-4} was used for all variants including a linear warm-up schedule for the first 1% of training.

4.2 Fine-Tuning

During fine-tuning, all weights of each network investigated are modified through training with a task-specific data set. We evaluate our models on two popular benchmarks: The General Language Understanding Evaluation (GLUE) benchmark Wang et al. (2019) and the Stanford Question Answering Dataset (SQuAD) Rajpurkar et al. (2016). From the GLUE benchmark, we chose MNLI, SST-2, and QQP as our primary tasks since they have sufficiently large data sets and provide stable results across runs. All fine-tuning results are shown on the development sets of SQuAD, MNLI, SST-2 and QQP.

From Table 1 we see that the modified BERT_{BASE} networks perform quite well compared to the unmodified network. Specifically, with $n \in \{2, 3, 4\}$ there is approximately 1-2% loss in accuracy across fine-tuning tasks while simultaneously providing a significant improvement in size and throughput (see Table 1). These results indicate that the n hyper-parameter can be used to make trade-offs in network size, speed and accuracy. For example, if an F1 score on SQuAD of approximately 87% is acceptable then a modified ($n = 3$) network could be used which would be approximately




Figure 3: First six blocks of the modified network architecture with $n = 2$. Note the removal of intermediate blocks between every $n = 2$ self-attention blocks (compare to Figure 2). Using this architecture, the modified BERT_{BASE} network will have only six intermediate blocks instead of the 12 which are found in the unmodified network.

Table 1: Fine-tuning results for unmodified and modified BERT_{BASE} networks with different values of n along with the number of parameters, decrease in size and increase in throughput. The F1 and EM scores are shown for SQuAD v1.1, m/mm for MNLI, accuracy for SST-2 and F1 score for QQP. We compare our approach against DistilBERT (Sanh et al. (2020)), BERT-PKD (Sun et al. (2019)) and BERT-of-Theseus (Xu et al. (2020)). For DistilBERT, we show the published scores and parameter count reported in Sanh et al. (2020). Similarly, for the BERT-PKD and BERT-of-Theseus networks, we show the values reported in Xu et al. (2020). Note that certain fine-tuning scores for DistilBERT, BERT-PKD and BERT-of-Theseus are missing in the table as they were not available in the published literature. We do not show the throughput increase for these networks due to the difference in hardware platforms.

Network	SQuAD	MNLI	SST-2	QQP	Parameter Count	Size Decrease	Throughput Increase
$n = 1$	88.43/80.97	82.83/83.38	91.97	87.50	110.10M	1.00x	1.00x
$n = 2$	87.09/79.66	81.64/82.32	89.44	87.38	81.76M	1.35x	1.39x
$n = 3$	86.91/79.28	81.41/82.20	90.13	87.13	72.31M	1.52x	1.59x
$n = 4$	86.63/78.86	81.47/82.45	90.71	86.76	67.59M	1.63x	1.72x
$n = 6$	85.73/77.59	80.65/80.85	88.64	86.42	62.86M	1.75x	1.87x
$n = \infty$	83.27/74.59	79.07/79.85	89.33	85.44	53.41M	2.06x	2.28x
DistilBERT	85.8/77.7	82.2/-	91.3	-	66M	1.66x	-
BERT-PKD	-/-	81.3/-	91.3	-	66M	1.66x	-
BERT-of-Theseus	-/-	82.3/-	91.5	-	66M	1.66x	-

1.5x smaller and 1.6x faster. Other trade-offs can be made for situations where the network is to be deployed onto a system with very limited memory and computational resources. In that case, minimizing network size and computational complexity would be critical and choosing a modified ($n = \infty$) network would decrease memory usage by more than 2x and be approximately 2.3x faster than the unmodified ($n = 1$) network.

In addition, the specific *location* of any one intermediate block in the modified BERT_{BASE} networks does not seem to have a disproportionate affect on performance. For example, with $n = 2$, half of all intermediate blocks are removed from multiple locations within in the network and this results in only a slight accuracy degradation. Thus, the removal of any *one* intermediate block, regardless of its location, would have an even smaller accuracy loss.




Figure 4: F1 scores on SQuAD of modified BERT_{BASE} networks with different values of n along with the corresponding decrease in size and increase in throughput as a relative percentage to the unmodified ($n = 1$) network. F1 score, throughput increase and size as a function of n (left). On the right, the trade-off between accuracy (absolute decrease in F1 score) and efficiency in either throughput (top right) or memory footprint (bottom right). Modified networks with $n \in \{2, 3, 4\}$ result in a significant size decrease and throughput increase while retaining a similar F1 score on SQuAD.

Using the values in Table 1, Figure 4 shows the SQuAD F1 scores of the modified BERT_{BASE} networks along with their corresponding size decrease and throughput increase relative to the unmodified ($n = 1$) network. The contribution of intermediate blocks to the final fine-tuning accuracy is not significant

but their removal can significantly decrease size and increase speed, allowing trade-offs between accuracy and efficiency. For example, the F1 score of the modified ($n = 3$) network decreases by approximately 1.5% but it is more than 34% smaller and 52% faster.

We compare our approach against certain knowledge distillation techniques such as DistilBERT, BERT-PKD and BERT-of-Theseus (see Table 1). The SQuAD F1 and EM scores of our modified $n = 4$ network is approximately 1% higher when compared to DistilBERT, while having a similar parameter count. The modified $n = 4$ network also performs slightly better than BERT-PKD for the MNLI task.

Finally, since we believe that ALBERT would be complementary to our method and could be used to further improve training speed and decrease model size, we did not perform a direct comparison to it.

4.3 Intermediate Layer Removal Prior to Fine-Tuning

Results in 4.2 show that removing intermediate layers have little impact on network accuracy. However, can these intermediate layers be removed *before* the modified network is fine-tuned?

The unmodified BERT_{BASE} network is first pre-trained as described in Section 4 and then the intermediate layers are removed. After removal of the intermediate layers, the now modified BERT_{BASE} network is fine-tuned on different tasks.

Table 2: Results with intermediate blocks removed prior to fine-tuning for modified BERT_{BASE} networks and different values of n . The F1 and EM scores are shown for SQuAD v1.1, m/mm for MNLI, accuracy for SST-2 and F1 score for QQP.

Network	SQuAD	MNLI	SST-2	QQP
$n = 2$	84.20/75.63	79.32/79.93	87.73	83.94
$n = 3$	76.99/67.23	74.36/74.13	83.03	82.59
$n = 4$	72.03/60.80	73.55/73.77	82.68	82.16
$n = 6$	65.45/53.10	72.11/72.23	82.34	82.01
$n = \infty$	56.82/43.66	72.55/73.20	83.02	81.57

Table 2 clearly shows a significant degradation in all fine-tuning tasks for the modified BERT_{BASE} networks. For example, in SQuAD v1.1, the decrease in F1 score for $n = 3$ when the intermediate layers are removed before fine-tuning is 11.44% compared with Table 1 where the decrease is 1.52% for $n = 3$. During pre-training it appears that the intermediate layers are performing important transformations to the representations learned in the self-attention blocks. By removing these layers prior to fine-tuning, the self-attention blocks are unable to compensate during fine-tuning.

4.4 Dependency Parsing

We use the analysis methods presented in Clark et al. (2019) to further justify the removal of intermediate layers. In BERT, each individual attention head is known to specialize to particular aspects of syntax and the model’s overall knowledge about syntax is distributed across multiple attention heads. Attention-based probing classifiers proposed in Clark et al. (2019) measure the overall ability of the model to learn syntactic information. In Clark et al. (2019), the probing classifiers are evaluated on the dependency parsing task for the Penn Tree Bank (PTB) data set Marcus et al. (1993). These classifiers treat the BERT attention outputs as fixed and train only a small number of parameters. Basically, the classifier produces a probability distribution over each word that indicates how likely each other word in the sentence is its syntactic head. The Attn + GloVe probing classifier substantially outperforms other reference baselines in Clark et al. (2019). This suggests BERT’s attention maps have a fairly thorough representation of English syntax.

We use the same Attn + GloVe probing classifier introduced in Clark et al. (2019) for our analysis on intermediate layers. However, we evaluate on the Universal Dependency gold standard for English, the English Web Treebank (EWT) data set Silveira et al. (2014) instead of PTB. Table 3 shows the Unlabeled Attachment Scores (UAS) on EWT for different values of n .

For this particular task it appears that fewer intermediate blocks result in an increase of the UAS score. Therefore, this indicates that removing intermediate blocks has little impact on the model’s ability to learn syntactic information.

Table 3: Attention-based probing classifier results for unmodified and modified BERT_{BASE} networks with different values of n .

Network	UAS Score
Untrained	24.20%
$n = 1$	70.60%
$n = 2$	69.80%
$n = 3$	71.50%
$n = 4$	70.80%
$n = 6$	71.50%
$n = \infty$	73.00%

5 Analysis of Intermediate Layers

In this section, we analyze the similarities of self-attention blocks of unmodified and modified BERT_{BASE} networks using centered kernel alignment (CKA) Kornblith et al. (2019) to better understand the importance of intermediate layers. The CKA similarity metric has useful properties such as invariance to both orthogonal transformations and isotropic scaling which is important when comparing architecturally similar but not identical networks. We follow the procedure in Kornblith et al. (2019) and perform the similarity analysis over multiple pre-training and fine-tuning trials of the unmodified ($n=1$) and modified BERT_{BASE} ($n=2, 3$) networks. A total of 5 trials for each unmodified and modified BERT_{BASE} network are pre-trained and then fine-tuned (see sections 4.1 and 4.2).

All BERT_{BASE} networks denoted as `Unmodified` or `Modified` are first fully pre-trained while networks denoted `Untrained` are not pre-trained and have their weights initialized randomly. However, *all* networks are fine-tuned using SQuAD v1.1. This training procedure is used so that the similarity analysis is applied to only fine-tuned networks.

After fine-tuning, each BERT_{BASE} network is used to encode the same 500 sentences from the WebText Gokaslan et al. (2019) data set to produce output activations from each of the self-attention blocks. These output activations are then used as the input to linear-kernel CKA.

Table 4: CKA similarities for multiple pairs of unmodified and modified BERT_{BASE} networks fine-tuned using SQuAD. Average linear-kernel CKA similarities along with standard deviation are computed over all 12 self-attention blocks using 5 trials of each network.

Network Pair	Average CKA Similarity
Untrained vs. $n = 1$	0.17 ± 0.06
Untrained vs. $n = 2$	0.22 ± 0.09
Untrained vs. $n = 3$	0.21 ± 0.10
$n = 1$ vs. $n = 1$	0.90 ± 0.08
$n = 1$ vs. $n = 2$	0.72 ± 0.15
$n = 1$ vs. $n = 3$	0.72 ± 0.13

Table 4 shows the average similarities over all 12 self-attention blocks for different pairs of untrained, unmodified and modified BERT_{BASE} networks. Similarities between unmodified and modified networks are much higher than similarities with untrained networks. This high similarity translates to a minimal loss in fine-tuning accuracy.

6 Conclusion

In this work we proposed a modification to the BERT architecture focusing on reducing the number of intermediate layers in the network. With the modified BERT_{BASE} network we show that the network complexity can be significantly decreased while preserving accuracy on fine-tuning tasks. In addition, our analysis with CKA and probing linear classifiers justifies the removal of intermediate layers.

For future work, we plan to apply our modifications to larger models like BERT_{LARGE} and GPT to show that the architectural changes can be used more widely. We also intend to integrate other network pruning techniques (e.g., ALBERT) with our modified architecture to further decrease size and increase throughput.

References

- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language models are few-shot learners.
- Kevin Clark, Urvashi Khandelwal, Omer Levy, and Christopher D. Manning. 2019. What does bert look at? an analysis of bert’s attention.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*.
- Aaron Gokaslan, Vanya Cohen, Ellie Pavlick, and Stefanie Tellex. 2019. Openwebtext corpus. [urlhttp://Skylion007.github.io/OpenWebTextCorpus](http://Skylion007.github.io/OpenWebTextCorpus).
- Dan Hendrycks and Kevin Gimpel. 2020. Gaussian error linear units (gelus).
- Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2020. Tinybert: Distilling bert for natural language understanding.
- Simon Kornblith, Mohammad Norouzi, Honglak Lee, and Geoffrey Hinton. 2019. Similarity of neural network representations revisited.
- Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2020. Albert: A lite bert for self-supervised learning of language representations.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. Building a large annotated corpus of english: The penn treebank. *Comput. Linguist.*, 19(2):313–330.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one?
- Matthew E. Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. Deep contextualized word representations.
- Ofir Press, Noah A. Smith, and Omer Levy. 2020. Improving transformer models by reordering their sublayers.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. Squad: 100,000+ questions for machine comprehension of text. *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using gpu model parallelism. *arXiv preprint arXiv:1909.08053*.
- Natalia Silveira, Timothy Dozat, Marie-Catherine De Marneffe, Samuel R Bowman, Miriam Connor, John Bauer, and Christopher D Manning. 2014. A gold standard dependency corpus for english. In *LREC*, pages 2897–2904. Citeseer.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. Patient knowledge distillation for bert model compression.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic bert for resource-limited devices.

- Ian Tenney, Dipanjan Das, and Ellie Pavlick. 2019. Bert rediscovers the classical nlp pipeline.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bowman. 2019. Glue: A multi-task benchmark and analysis platform for natural language understanding.
- Canwen Xu, Wangchunshu Zhou, Tao Ge, Furu Wei, and Ming Zhou. 2020. Bert-of-theseus: Compressing bert by progressive module replacing.
- Yukun Zhu, Ryan Kiros, Richard Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books.