

Beyond Preserved Accuracy: Evaluating Loyalty and Robustness of BERT Compression

Canwen Xu^{1*}, Wangchunshu Zhou^{2*}, Tao Ge³, Ke Xu⁴, Julian McAuley¹, Furu Wei³

¹ University of California, San Diego ² Stanford University

³ Microsoft Research Asia ⁴ Beihang University

¹{cxu, jmcauley}@ucsd.edu, ²wcszhou@stanford.edu

³{tage, fuwei}@microsoft.com, ⁴kexu@nlsde.buaa.edu.cn

Abstract

Recent studies on compression of pretrained language models (e.g., BERT) usually use preserved accuracy as the metric for evaluation. In this paper, we propose two new metrics, label loyalty and probability loyalty that measure how closely a compressed model (i.e., student) mimics the original model (i.e., teacher). We also explore the effect of compression with regard to robustness under adversarial attacks. We benchmark quantization, pruning, knowledge distillation and progressive module replacing with loyalty and robustness. By combining multiple compression techniques, we provide a practical strategy to achieve better accuracy, loyalty and robustness.¹

1 Introduction

Recently, many large pretrained language models (PLMs, Devlin et al., 2019; Liu et al., 2019; Yang et al., 2019; Shoeybi et al., 2019; Raffel et al., 2020) have been proposed for a variety of Natural Language Processing (NLP) tasks. However, as pointed out in recent studies (Strubell et al., 2019; Schwartz et al., 2020a; Bender et al., 2021), these models suffer from computational inefficiency and high ecological cost. Many attempts have been made to address this problem, including quantization (Zafir et al., 2019; Shen et al., 2020), pruning (Michel et al., 2019; Sanh et al., 2020), knowledge distillation (KD) (Sanh et al., 2019; Sun et al., 2019, 2020; Turc et al., 2019; Jiao et al., 2020; Wang et al., 2020; Zhou et al., 2021) and progressive module replacing (Xu et al., 2020).

BERT (Devlin et al., 2019) is a representative PLM. Many works compressing BERT use preserved accuracy with computational complexity (e.g., speed-up ratio, FLOPS, number of parameters) as metrics to evaluate compression. This

*Equal Contribution. Work done at Microsoft Research Asia.

¹Our code is available at <https://github.com/JetRunner/beyond-preserved-accuracy>.

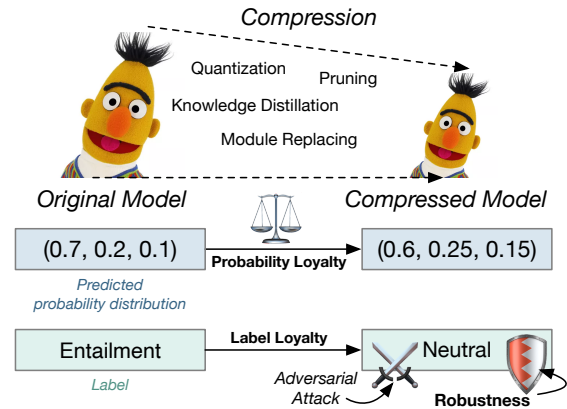


Figure 1: Three metrics to evaluate the compressed models beyond preserved accuracy. For each input, **label and probability loyalty** measure the shift of label and predicted probability distribution, respectively. **Robustness** measures the performance of the compressed model under adversarial attacks.

evaluation scheme is far from perfect: (1) Preserved accuracy cannot reflect how alike the teacher and student² models behave. This can be problematic when applying compression techniques in production (to be detailed in Section 3). (2) Using preserved accuracy to evaluate models compressed with more data or data augmentation (Jiao et al., 2020) can be misleading, since one cannot tell whether the improvement should be attributed to the innovation of the compression technique or addition of data. (3) Model robustness, which is critical for production, is often missing from evaluation, leaving a possible safety risk.

As illustrated in Figure 1, to measure the resemblance between the student and teacher models, we propose *label loyalty* and *probability loyalty* to target different but important aspects. We also explore the robustness of the compressed models

²Teacher and student are originally concepts in knowledge distillation. In this paper, we will sometimes use teacher and student to refer to the original model and compressed model in other compression methods for simplicity.

by conducting black-box adversarial attacks. We apply representative BERT compression methods of different types to the same teacher model and benchmark their performance in terms of accuracy, speed, loyalty and robustness. We find that methods with a knowledge distillation loss perform well on loyalty and that post-training quantization can drastically improve robustness against adversarial attacks. We use the conclusions drawn from these experiments to combine multiple techniques together and achieve significant improvement in terms of accuracy, loyalty and robustness.

2 BERT Compression

Compressing and accelerating pretrained language models like BERT has been an active field of research. Some initial work employs conventional methods for neural network compression to compress BERT. For example, Q8-BERT (Zafrir et al., 2019) and Q-BERT (Shen et al., 2020) employ weight quantization to reduce the number of bits used to represent a parameter in a BERT model. Pruning methods like Head Prune (Michel et al., 2019) and Movement Pruning (Sanh et al., 2020) remove weights based on their importance to reduce the memory footprint of pretrained models. Another line of research focuses on exploiting the knowledge encoded in a large pretrained model to improve the training of more compact models. For instance, DistilBERT (Sanh et al., 2019) and BERT-PKD (Sun et al., 2019) employ knowledge distillation (Hinton et al., 2015) to train compact BERT models in a task-specific and task-agnostic fashion respectively by mimicking the behavior of large teacher models. Recently, Xu et al. (2020) proposed progressive module replacing, which trains a compact student model by progressively replacing the teacher layers with their more compact substitutes.

3 Metrics Beyond Accuracy

3.1 Loyalty

3.1.1 Label Loyalty

Model compression is a common practice to optimize the efficiency of a model for deployment (Cheng et al., 2017). In real-world settings, training and deployment are often separate (Paley et al., 2020). As such it is desirable to have a metric to measure to what extent the “production model” is different from the “development model”. Moreover, when discussing ethical concerns, previous

studies (Wang et al., 2020; Zhou et al., 2020) ignore the risk that model compression could introduce additional biases. However, a recent work (Hooker et al., 2020) strongly contradicts this assumption. In a nutshell, we would desire the student to behave as closely as possible to the teacher, to make it more predictable and minimize the risk of introducing extra bias. Label loyalty directly reflects the resemblance of the labels predicted between the teacher and student models. It is calculated in the same way as accuracy, but between the student’s prediction and the teacher’s prediction, instead of ground labels:

$$L_l = \text{Accuracy}(pred_t, pred_s) \quad (1)$$

where $pred_t$ and $pred_s$ are the predictions of the teacher and student, respectively.

3.1.2 Probability Loyalty

Except for the label correspondence, we argue that the predicted probability distribution matters as well. In industrial applications, calibration (Guo et al., 2017; Li et al., 2020), which focuses on the meaningfulness of confidence, is an important issue for deployment. Many dynamic inference acceleration methods (Xin et al., 2020b; Schwartz et al., 2020b; Liu et al., 2020; Xin et al., 2020a; Li et al., 2020) use entropy or the maximum value of the predicted probability distribution as the signal for early exiting. Thus, a shift of predicted probability distribution in a compressed model could break the calibration and invalidate calibrated early exiting pipelines.

Kullback–Leibler (KL) divergence is often used to measure how one probability distribution is different from a reference distribution.

$$D_{\text{KL}}(P\|Q) = \sum_{x \in \mathcal{X}} P(x) \log \left(\frac{P(x)}{Q(x)} \right) \quad (2)$$

where \mathcal{X} is the probability space; P and Q are predicted probability distributions of the teacher and students, respectively. Here, we use its variant, the Jensen–Shannon (JS) divergence, since it is symmetric and always has a finite value which is desirable for a distance-like metric:

$$D_{\text{JS}}(P\|Q) = \frac{1}{2} D_{\text{KL}}(P\|M) + \frac{1}{2} D_{\text{KL}}(Q\|M) \quad (3)$$

where $M = \frac{1}{2}(P + Q)$. Finally, the probability loyalty between P and Q is defined as:

$$L_p(P\|Q) = 1 - \sqrt{D_{\text{JS}}(P\|Q)} \quad (4)$$

Method	# Layer	Speed -up \uparrow	MNLI m/mm \uparrow	Loyalty		Adversarial Attack	
				Label \uparrow	Probability \uparrow	Acc \uparrow	#Query \uparrow
Teacher (Devlin et al., 2019)	12	1.0 \times	84.5 / 83.3	100	100	8.1(\pm 0.1)	89.6(\pm 0.1)
Truncate & Finetune	6	2.0 \times	81.1 / 80.0	87.7(\pm 0.2)	84.9(\pm 0.7)	4.4(\pm 0.1)	78.0(\pm 0.1)
Pure KD	6	2.0 \times	81.1 / 80.8	89.2(\pm 0.1)	89.5(\pm 0.2)	6.2(\pm 0.1)	80.1(\pm 0.2)
Q8-PTQ (Zafir et al., 2019)	12	1.8 \times [‡]	80.7 / 80.4	89.6(\pm 0.5)	80.8(\pm 0.4)	40.2(\pm 0.1)	91.6(\pm 0.1)
Q8-QAT [†] (Zafir et al., 2019)	12	1.8 \times [‡]	83.4 / 82.4	89.7(\pm 0.2)	88.2(\pm 0.3)	6.8(\pm 0.2)	82.7(\pm 0.2)
Head Prune (Michel et al., 2019)	12	1.2 \times	80.9 / 80.6	87.8(\pm 0.1)	85.5(\pm 0.6)	9.1(\pm 0.1)	90.5(\pm 0.2)
DistilBERT [†] (Sanh et al., 2019)	6	2.0 \times	82.4 / 81.4	88.9(\pm 0.2)	88.4(\pm 0.4)	5.9(\pm 0.1)	80.8(\pm 0.2)
TinyBERT [†] (Jiao et al., 2020)	6	2.0 \times	82.7 / 82.7	88.9(\pm 0.1)	88.4(\pm 0.7)	6.7(\pm 0.1)	82.1(\pm 0.2)
BERT-PKD (Sun et al., 2019)	6	2.0 \times	81.3 / 81.1	88.9(\pm 0.1)	89.0(\pm 0.2)	6.4(\pm 0.2)	81.9(\pm 0.2)
BERT-of-Theseus (Xu et al., 2020)	6	2.0 \times	81.8 / 80.7	88.1(\pm 0.2)	82.5(\pm 0.3)	8.3(\pm 0.2)	89.7(\pm 0.2)

Table 1: Accuracy, loyalty and robustness of compressed models on the test set of MNLI (3 runs). Accuracy scores are from the GLUE (Wang et al., 2019) test server. [†]These models are not initialized from (a part of) the finetuned BERT teacher. [‡]The speed-up ratio of quantization is benchmarked on CPU. [†]Higher is better.

where $L_p \in [0, 1]$; higher L_p represents higher resemblance. Note that Equation 2 is also known as the KD loss (Hinton et al., 2015), thus KD-based methods will naturally have an advantage in terms of probability loyalty.

3.2 Robustness

Deep Learning models have been shown to be vulnerable to adversarial examples that are slightly altered with perturbations often indistinguishable to humans (Kurakin et al., 2017). Previous work (Su et al., 2018) found that small convolutional neural networks (CNN) are more vulnerable to adversarial attacks compared to bigger ones. Likewise, we intend to investigate how BERT models perform and the effect of different types of compression in terms of robustness. We use an off-the-shelf adversarial attack method, TextFooler (Jin et al., 2020), which demonstrates state-of-the-art performance on attacking BERT. TextFooler conducts black-box attacks by querying the BERT model with the adversarial input where words are perturbed based on their part-of-speech role. We select two metrics from (Jin et al., 2020), after-attack accuracy and query number, to evaluate a model’s robustness. After-attack accuracy represents the remaining accuracy after the adversarial attack. Query number represents how many queries with perturbed input have been made to complete the attack.

4 Experiments

4.1 Dataset and Baselines

We use the Multi-Genre Natural Language Inference dataset (MNLI) (Williams et al., 2018) for

evaluation. The accuracy scores are from the GLUE benchmark (Wang et al., 2019) test server. We select representative methods of different types of compression methods: Head Prune (Michel et al., 2019) for pruning; Post-training Quantization (PTQ) and Quantization-aware Training (QAT) (Zafir et al., 2019) for quantization; DistilBERT (Sanh et al., 2019) and TinyBERT (Jiao et al., 2020) for pretraining-phase knowledge distillation; BERT-PKD (Sun et al., 2019) for downstream knowledge distillation; and BERT-of-Theseus (Xu et al., 2020) for module replacing. Following Sun et al. (2019); Xu et al. (2020), we truncate the first (bottom) 6 layers and then finetune it as a baseline for 6-layer models. Additionally, we directly optimize the KL divergence (i.e., pure KD loss) to set an upper bound for probability loyalty.

4.2 Training Details

Our implementation is based on Hugging Face Transformers (Wolf et al., 2020). We first finetune a BERT-base model to be the teacher for KD and the source model for quantization and pruning. The learning rate is set to 3×10^{-5} and the batch size is 64 with 1,000 warm-up steps. For quantization and pruning, the source model is the same finetuned teacher. For downstream KD and BERT-of-Theseus, we initialize the model by truncating the first (bottom) 6 layers of the finetuned teacher, following the original papers (Sun et al., 2019; Xu et al., 2020). QAT uses pretrained BERT-base for initialization. For pretraining distillation, we directly finetune compressed 6-layer DistilBERT and TinyBERT checkpoints to report results. The pruned

ing percentage for Head Prune is 45%. The hyperparameters of BERT-PKD are from the original implementation. The detailed hyperparameters for each method can be found in Appendix A.

4.3 Experimental Results

We show experimental results in Table 1. First, we find that post-training quantization can drastically improve model robustness. A possible explanation is that the regularization effect of post-training quantization (Paupamah et al., 2020; Wu and Flierl, 2020) helps improve the robustness of the model (Werpachowski et al., 2019; Ma et al., 2020). A similar but smaller effect can be found from pruning. However, as shown in Table 2, if we finetune the low-precision or pruned model again, the model would re-overfit the data and yield even lower robustness than the original model. Second, KD-based models maintains good label loyalty and probability loyalty due to their optimization objectives. Interestingly, compared to *Pure KD* where we directly optimize the KL divergence, DistilBERT, TinyBERT and BERT-PKD trade some loyalty in exchange for accuracy. Compared to DistilBERT, TinyBERT digs up higher accuracy by introducing layer-to-layer distillation, with their loyalty remains identical. Also, we do not observe a significant difference between pretraining KD and downstream KD in terms of both loyalty and robustness ($p > 0.1$). Notably, BERT-of-Theseus has a significantly lower loyalty, suggesting the mechanism behind it is different from KD. We also provide some results on SST-2 (Socher et al., 2013) in Appendix B.

5 Combining the Bag of Tricks

As we described in Section 4.3, we discover that post-training quantization (PTQ) can improve the robustness of a model while knowledge distillation (KD) loss benefits the loyalty of a compressed model. Thus, by combining multiple compression techniques, we expect to achieve a higher speed-up ratio with improved accuracy, loyalty and robustness.

To combine KD with other methods, we replace the original cross-entropy loss in quantization-aware training and module replacing with the knowledge distillation loss (Hinton et al., 2015) as in Equation 2. For pruning, we perform knowledge distillation on the pruned model. We also apply the temperature re-scaling trick from (Hinton

Method	Speed	MNLI	L-L	P-L	AA	# Q
Teacher	1.0×	84.5 / 83.3	100	100	8.1	89.6
Head Prune	1.2×	80.9 / 80.6	87.8	85.5	9.1	90.5
+Finetune	1.2×	83.2 / 81.9	89.1	85.5	7.2	83.2
+KD	1.2×	84.2 / 83.0	93.3	93.0	8.3	90.5
+KD+PTQ	2.2×	80.8 / 80.4	89.6	86.3	38.4	90.9
Q8-QAT	1.8×	83.4 / 82.4	89.7	88.2	6.8	82.7
Q8-PTQ	1.8×	80.7 / 80.4	89.6	80.8	40.2	91.6
+Finetune	1.8×	82.9 / 81.9	89.7	84.8	7.1	84.5
+KD	1.8×	84.1 / 83.5	94.0	93.9	7.5	86.1
BERT-PKD	2.0×	81.3 / 81.1	88.9	89.0	6.4	81.9
Theseus	2.0×	81.8 / 80.7	88.1	82.5	8.3	89.7
+KD	2.0×	82.6 / 81.7	91.2	91.4	8.0	88.7
+KD+PTQ	3.6×	80.2 / 79.9	89.5	80.3	36.5	91.3

Table 2: Accuracy and loyalty for combining multiple compression techniques on the test set of MNLI. **L-L**: label loyalty; **P-L**: probability loyalty; **AA**: after-attack accuracy; **# Q**: Query number for adversarial attack. The number of layers for each group is consistent with Table 1.

et al., 2015) with a fixed temperature of 10. As shown in Table 2, the knowledge distillation loss effectively improves the accuracy and loyalty of pruning, quantization and module replacing.

Furthermore, we post-quantize the KD-enhanced models after they are trained. Shown in Table 2, by adding post-training quantization, the speed and robustness can both be boosted. Notably, the order to apply PTQ and KD does matter. PTQ→KD has high accuracy and loyalty but poor robustness while KD→PTQ remains a good robustness with a lower accuracy performance. To summarize, we recommend the following compression strategy: (1) conduct pruning or module replacing with a KD loss; (2) for speed-sensitive and robustness-sensitive applications, apply post-training quantization afterwards.

6 Conclusion

In this paper, we propose label and probability loyalty to measure the correspondence of label and predicted probability distribution between compressed and original models. In addition to loyalty, we investigate the robustness of different compression techniques under adversarial attacks. These metrics reveal that post-training quantization and knowledge distillation can drastically improve robustness and loyalty, respectively. By combining multiple compression methods, we can further improve speed, accuracy, loyalty and robustness for various applications. Our metrics help mitigate the gap between model training and deployment, shed

light upon comprehensive evaluation for compression of pretrained language models, and call for the invention of new compression techniques.

Ethical Concerns

We include a discussion about the possible ethical risks of a compressed model in Section 3.1. Although our paper is an attempt to mitigate the risk of introducing extra biases to compression, we would like to point out that our metrics do not directly indicate the bias level in the compressed model. That is to say, additional measures should still be taken to evaluate and debias both the teacher and student models.

Acknowledgments

We would like to thank all anonymous reviewers for their insightful comments. Tao Ge is the corresponding author.

References

- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? In *FAccT*.
- Yu Cheng, Duo Wang, Pan Zhou, and Tao Zhang. 2017. A survey of model compression and acceleration for deep neural networks. *arXiv preprint arXiv:1710.09282*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: pre-training of deep bidirectional transformers for language understanding. In *NAACL-HLT*, pages 4171–4186. Association for Computational Linguistics.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *ICML*, volume 70 of *Proceedings of Machine Learning Research*, pages 1321–1330. PMLR.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*.
- Sara Hooker, Nyalleng Moorosi, Gregory Clark, Samy Bengio, and Emily Denton. 2020. Characterising bias in compressed models. *arXiv preprint arXiv:2010.03058*.
- 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. In *EMNLP (Findings)*, pages 4163–4174. Association for Computational Linguistics.
- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is BERT really robust? natural language attack on text classification and entailment. In *AAAI*.
- Alexey Kurakin, Ian J. Goodfellow, and Samy Bengio. 2017. Adversarial examples in the physical world. In *ICLR (Workshop)*.
- Lei Li, Yankai Lin, Shuhuai Ren, Deli Chen, Xuancheng Ren, Peng Li, Jie Zhou, and Xu Sun. 2020. Accelerating pre-trained language models via calibrated cascade. *arXiv preprint arXiv:2012.14682*.
- Weijie Liu, Peng Zhou, Zhiruo Wang, Zhe Zhao, Haotang Deng, and Qi Ju. 2020. Fastbert: a self-distilling BERT with adaptive inference time. In *ACL*, pages 6035–6044. Association for Computational Linguistics.
- 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. *arXiv preprint arXiv:1907.11692*.
- Avery Ma, Fartash Faghri, and Amir-massoud Farahmand. 2020. Adversarial robustness through regularization: A second-order approach. *arXiv preprint arXiv:2004.01832*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? In *NeurIPS*, pages 14014–14024.
- Andrei Paleyev, Raoul-Gabriel Urma, and Neil D Lawrence. 2020. Challenges in deploying machine learning: a survey of case studies. *arXiv preprint arXiv:2011.09926*.
- Kimessha Paupamah, Steven James, and Richard Klein. 2020. Quantisation and pruning for neural network compression and regularisation. *arXiv preprint arXiv:2001.04850*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21:140:1–140:67.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.
- Victor Sanh, Thomas Wolf, and Alexander M. Rush. 2020. Movement pruning: Adaptive sparsity by fine-tuning. In *NeurIPS*.
- Roy Schwartz, Jesse Dodge, Noah A. Smith, and Oren Etzioni. 2020a. Green AI. *Commun. ACM*, 63(12):54–63.

- Roy Schwartz, Gabriel Stanovsky, Swabha Swayamdipta, Jesse Dodge, and Noah A. Smith. 2020b. The right tool for the job: Matching model and instance complexities. In *ACL*, pages 6640–6651. Association for Computational Linguistics.
- Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W. Mahoney, and Kurt Keutzer. 2020. Q-BERT: hessian based ultra low precision quantization of BERT. In *AAAI*, pages 8815–8821. AAAI Press.
- Mohammad Shoeybi, Mostofa Patwary, Raul Puri, Patrick LeGresley, Jared Casper, and Bryan Catanzaro. 2019. Megatron-lm: Training multi-billion parameter language models using model parallelism. *arXiv preprint arXiv:1909.08053*.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Y. Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*.
- Emma Strubell, Ananya Ganesh, and Andrew McCallum. 2019. Energy and policy considerations for deep learning in NLP. In *ACL*, pages 3645–3650. Association for Computational Linguistics.
- Dong Su, Huan Zhang, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, and Yupeng Gao. 2018. Is robustness the cost of accuracy? - A comprehensive study on the robustness of 18 deep image classification models. In *ECCV*, volume 11216 of *Lecture Notes in Computer Science*, pages 644–661. Springer.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. Patient knowledge distillation for BERT model compression. In *EMNLP-IJCNLP*, pages 4322–4331. Association for Computational Linguistics.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic BERT for resource-limited devices. In *ACL*, pages 2158–2170. Association for Computational Linguistics.
- Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: The impact of student initialization on knowledge distillation. *arXiv preprint arXiv:1908.08962*, 13.
- 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. In *ICLR*. OpenReview.net.
- Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. 2020. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. In *NeurIPS*.
- Roman Werpachowski, András György, and Csaba Szepesvári. 2019. Detecting overfitting via adversarial examples. In *NeurIPS*, pages 7856–7866.
- Adina Williams, Nikita Nangia, and Samuel R. Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *NAACL-HLT*, pages 1112–1122. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander M. Rush. 2020. Transformers: State-of-the-art natural language processing. In *EMNLP (Demos)*, pages 38–45. Association for Computational Linguistics.
- Hanwei Wu and Markus Flierl. 2020. Vector quantization-based regularization for autoencoders. In *AAAI*, pages 6380–6387. AAAI Press.
- Ji Xin, Rodrigo Nogueira, Yaoliang Yu, and Jimmy Lin. 2020a. Early exiting bert for efficient document ranking. In *Proceedings of SustaiNLP: Workshop on Simple and Efficient Natural Language Processing*, pages 83–88.
- Ji Xin, Raphael Tang, Jaejun Lee, Yaoliang Yu, and Jimmy Lin. 2020b. Deebert: Dynamic early exiting for accelerating BERT inference. In *ACL*, pages 2246–2251. Association for Computational Linguistics.
- Canwen Xu, Wangchunshu Zhou, Tao Ge, Furu Wei, and Ming Zhou. 2020. Bert-of-theseus: Compressing BERT by progressive module replacing. In *EMNLP*, pages 7859–7869. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Yiming Yang, Jaime G. Carbonell, Ruslan Salakhutdinov, and Quoc V. Le. 2019. Xlnet: Generalized autoregressive pretraining for language understanding. In *NeurIPS*, pages 5754–5764.
- Ofir Zafrir, Guy Boudoukh, Peter Izsak, and Moshe Wasserblat. 2019. Q8bert: Quantized 8bit bert. *arXiv preprint arXiv:1910.06188*.
- Wangchunshu Zhou, Canwen Xu, Tao Ge, Julian J. McAuley, Ke Xu, and Furu Wei. 2020. BERT loses patience: Fast and robust inference with early exit. In *NeurIPS*.
- Wangchunshu Zhou, Canwen Xu, and Julian McAuley. 2021. Meta learning for knowledge distillation. *arXiv preprint arXiv:2106.04570*.

A Hyperparameter Settings

Teacher Learning rate: $3e-5$; Batch size: 64; Warm-up steps: 1000.

Truncate & Finetune Learning rate: $3e-5$; Batch size: 128.

Pure KD Learning rate: $3e-5$; Batch size: 64; Warm-up steps: 1000; α : 1; Temperature: 10.

Q8-QAT Learning rate: $2e-5$; Batch size: 128.

DistilBERT Learning rate: $3e-5$; Batch size: 64; Warm-up steps: 1000.

TinyBERT Learning rate: $3e-5$; Batch size: 64; Warm-up steps: 1000.

BERT-PKD Learning rate: $3e-5$; Batch size: 64; Warm-up steps: 1000; α : 0.7; β : 500; Temperature: 10.

BERT-of-Theseus Learning rate: $3e-5$; Batch size: 128; Warm-up steps: 8000; k : 0.00002; b : 0.5.

Head Prune + Finetune Learning rate: $2e-5$; Batch size: 128.

Head Prune + KD Learning rate: $2e-5$; Batch size: 128; Temperature: 10.

Q8 + FT Learning rate: $2e-5$; Batch size: 128.

Q8 + KD Learning rate: $2e-5$; Batch size: 128; Temperature: 10.

BERT-of-Theseus + KD Learning rate: $3e-5$; Batch size: 64; Warm-up steps: 1000; Temperature: 10.

B Experimental Results on SST-2

Method	Speed	SST-2	L-L	P-L	AA	# Q
Teacher	1.0×	92.0	100	100	7.5	81.2
KD	2.0×	91.5	93.8	92.4	7.2	80.3
Head Prune	1.3×	90.4	89.5	88.2	8.1	81.0
Q8-QAT	1.8×	91.4	91.8	90.3	6.6	81.4
Q8-PTQ	1.8×	90.1	91.3	88.9	26.5	86.4

Table 3: Accuracy and loyalty of some compression techniques on the test set of SST-2. **L-L**: label loyalty; **P-L**: probability loyalty; **AA**: after-attack accuracy; **# Q**: Query number for adversarial attack. The number of layers for each group is consistent with Table 1.