

Robustness Challenges in Model Distillation and Pruning for Natural Language Understanding

Mengnan Du^{1*}, Subhabrata Mukherjee², Yu Cheng², Milad Shokouhi²,
Xia Hu³, Ahmed Hassan Awadallah²

¹New Jersey Institute of Technology ²Microsoft Research ³Rice University

mengnan.du@njit.edu, xia.hu@rice.edu

{submukhe, Yu.Cheng, milads, hassanam}@microsoft.com

Abstract

Recent work has focused on compressing pre-trained language models (PLMs) like BERT where the major focus has been to improve the in-distribution performance for downstream tasks. However, very few of these studies have analyzed the impact of compression on the generalizability and robustness of compressed models for out-of-distribution (OOD) data. Towards this end, we study two popular model compression techniques including knowledge distillation and pruning and show that the compressed models are significantly less robust than their PLM counterparts on OOD test sets although they obtain similar performance on in-distribution development sets for a task. Further analysis indicates that the compressed models overfit on the shortcut samples and generalize poorly on the hard ones. We further leverage this observation to develop a regularization strategy for robust model compression based on sample uncertainty. Experimental results on several natural language understanding tasks demonstrate that our bias mitigation framework improves the OOD generalization of the compressed models, while not sacrificing the in-distribution task performance.

1 Introduction

Large pretrained language models (PLMs) (e.g., BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), GPT-3 (Brown et al., 2020)) have obtained state-of-the-art performance in several Natural Language Understanding (NLU) tasks. However, recent studies (Niven and Kao, 2019; Du et al., 2021; Mudrakarta et al., 2018) indicate that PLMs heavily rely on shortcut learning/spurious correlations, rather than acquiring higher level language understanding and semantic reasoning in several NLU tasks. Specifically, these models often exploit dataset biases and artifacts, e.g., lexical bias and

overlap bias, as shortcuts for prediction. Due to the independent and identically distributed (IID) split of training, development, and test sets, these models that learn spurious decision rules from training data can perform well on in-distribution data (Du et al., 2022). Nevertheless, the shortcut learning behavior will result in models that have poor generalization performance on out-of-distribution (OOD) data, raising concerns about their robustness.

On the other hand, it is difficult to use these large PLMs models in real-world applications with latency and capacity constraints, e.g., on edge devices and mobile phones. Thus, model compression emerges as one of the techniques to reduce model size, speed up inference, and save energy without significant performance drop for downstream tasks. State-of-the-art model compression techniques such as knowledge distillation (Sanh et al., 2019; Sun et al., 2019) and pruning (Sanh et al., 2020) primarily focus on evaluating compressed model performance in in-distribution test data. However, in-distribution testing is insufficient to capture the generalizability of PLMs (D’Amour et al., 2020). In contrast to existing work that is geared towards general-purpose PLMs (Niven and Kao, 2019; Du et al., 2021; Mudrakarta et al., 2018), this work aims to study the impact of compression on the shortcut learning and OOD generalization ability of compressed models.

Towards this end, we conduct comprehensive experiments to evaluate the OOD robustness of compressed models, with BERT as the base encoder. **We focus primarily on two popular model compression techniques in the form of pruning and knowledge distillation (Sanh et al., 2019; Wang et al., 2020).** For pruning, we consider two popular techniques including iterative magnitude pruning (Sanh et al., 2020) and structured pruning (Prasanna et al., 2020; Liang et al., 2021). Specifically, we explore the following research questions: *Are distilled and pruned models as ro-*

*Most of the work was completed while the first author was an intern at Microsoft Research during summer 2021.

bust as their PLM counterparts for downstream NLU tasks? What is the impact of varying the level of compression on OOD generalization and bias of compressed models? We evaluate the performance of several compressed models obtained using the above techniques on both standard in-distribution development sets and OOD test sets for downstream NLU tasks. Experimental analysis indicates that distilled and pruned models are consistently less robust than their PLM counterparts. Further analysis of the poor generalization performance of compressed models reveals some interesting observations. For instance, we observe that the compressed models overfit on the easy / shortcut samples and generalize poorly on the hard ones. This motivates our second research question: *How to regularize model compression techniques to generalize across samples with varying difficulty?* This brings some interesting challenges since we do not know which samples are easy or hard a priori.

Based on the above observations, we propose a bias mitigation framework to improve the OOD robustness of compressed models, termed as RMC (Robust Model Compression). First, we leverage the uncertainty of the deep neural network to quantify the difficulty of a training sample. This is given by the variance in the prediction of a sample from multiple sub-networks of the original large network obtained by model pruning. Second, we leverage this sample-specific measure for smoothing and regularizing different families of compression techniques. The major contributions of this work can be summarized as follows:

- We perform a comprehensive analysis to evaluate the OOD generalization ability and robustness of compressed models for NLU tasks.
- We further analyze plausible reasons for the low generalizability of compressed models and demonstrate connections to shortcut learning.
- We propose a mitigation framework for regularizing model compression, termed as RMC, which smoothes the knowledge distillation training based on the estimated sample difficulties.
- We perform experiments to demonstrate that our RMC framework improves OOD generalization while not sacrificing the standard in-distribution task performance on multiple NLU tasks.

2 Related Work

Shortcut Learning and Mitigation. Recent studies indicate that PLMs tend to exploit biases and

artifacts in the dataset as shortcuts for prediction, rather than acquiring higher level semantic understanding and reasoning for NLU tasks (Niven and Kao, 2019; Du et al., 2021; McCoy et al., 2019a). There are some preliminary work to mitigate the bias of general PLMs, including product-of-experts (Clark et al., 2019; He et al., 2019; Sanh et al., 2021), re-weighting (Schuster et al., 2019; Yaghoobzadeh et al., 2019; Utama et al., 2020), adversarial training (Stacey et al., 2020), posterior regularization (Cheng et al., 2021), etc.

Robustness in Model Compression. Current practice for evaluating model compression performance focuses mainly on standard benchmark performance (Zhu et al., 2020; Wang et al., 2021). In the computer vision domain, previous work shows that compressed models perform poorly in Compression Identified Exemplars (CIE) (Hooker et al., 2019), and compression amplifies algorithmic bias towards certain demographics (Hooker et al., 2020). The most similar work to ours are two concurrent work (Xu et al., 2021a; Li et al., 2021) that investigate the performance of compressed models beyond standard benchmarks for natural language understanding tasks. However, both work mainly focus on evaluating the robustness of compressed models with respect to the scenario of adversarial attacks, i.e., TextFooler (Jin et al., 2020), and the unified adversarial framework (Li et al., 2021). In contrast, we comprehensively characterize the robustness of BERT compression in OOD test sets to probe the OOD generalizability of the compression techniques. Besides, we use insights from this robustness analysis to design a generalizable and robust model compression framework.

3 Are Compressed Models Robust?

We perform a comprehensive analysis to evaluate the robustness of compressed language models.

3.1 Compression Techniques

We consider two popular families of compression, namely, knowledge distillation and model pruning.

Knowledge Distillation: The objective here is to train a small-size model by mimicking the behavior of the larger teacher model using knowledge distillation (Hinton et al., 2015). In this work, we focus on task-agnostic distillation. In particular, we consider DistilBERT (Sanh et al., 2019) and MiniLM (Wang et al., 2020) distilled from BERT-base. For a fair comparison, we select compressed

models with similar capacities (66M parameters in this work). In order to evaluate the impact of compression techniques on model robustness, we also consider similar capacity smaller models without using knowledge distillation. These are obtained via simple truncation where we retain the first 6 layers of the large model, and via pre-training a smaller 6-layer model from scratch.

Iterative Magnitude Pruning: This is a task-specific unstructured pruning method (Sanh et al., 2020). During the fine-tuning process for each downstream task, the weights with the lowest magnitude are removed until the pruned model reaches the target sparsity. Note that we utilize the standard pruning technique, rather than the LTH-based pruning (lottery ticket hypothesis) that uses re-winding (Chen et al., 2020). We also consider different pruning ratios to obtain pruned models with different levels of sparsity.

Structured Pruning: This method family is based on the hypothesis that there is redundancy in the attention heads (Prasanna et al., 2020; Voita et al., 2019; Bian et al., 2021; Chen et al., 2021). We also consider task-specific pruning. During the fine-tuning process for each task, it prunes the whole attention heads based on their importance to the model predictions. Please refer to Sec. A in Appendix for more details. We prune around 20% attention heads in total (i.e., 28 attention heads). Further pruning increases the sparsity with significant degradation of the model’s performance on in-distribution development sets.

3.2 Evaluation Datasets

To evaluate the robustness of the compressed models introduced in the last section, we use three NLU tasks, including MNLI, FEVER, and QQP¹. Please refer to Sec. B in Appendix for more details.

- **MNLI** (Williams et al., 2018): This is a natural language inference task. In this work, we report the accuracy metric on the matched subset. We use HANS (McCoy et al., 2019b) as the adversarial test set, which contains 30,000 synthetic samples. Models that exploit shortcut features have been shown to perform poorly on the HANS test set.

¹MNLI, FEVER, and QQP are the three most widely used datasets to evaluate the shortcut learning/bias behavior and OOD generalization of PLMs in the literature (Tu et al., 2020; He et al., 2019; Clark et al., 2019; Schuster et al., 2019)

- **FEVER** (Thorne et al., 2018): This is a fact verification dataset. Recent studies indicate that there are strong shortcuts in the claims (Utama et al., 2020). To facilitate the robustness and generalization evaluation of fact verification models, two symmetric test sets (i.e., Sym v1 and Sym v2) were created, where bias exists in the symmetric pairs (Schuster et al., 2019). Both OOD test sets have 712 samples.
- **QQP:** The task is to predict whether a pair of questions is semantically equivalent. We consider the OOD test set PAWS-qqp, which contains 677 test samples generated from QQP corpus (Zhang et al., 2019; Yang et al., 2019). Besides, we also consider the PAWS-wiki OOD test set, which consists of 8,000 test samples generated from Wikipedia pages.

For all three tasks, we employ accuracy as the evaluation metric and evaluate the performance of the compressed models on both the in-distribution development set and the OOD test set.

3.3 Evaluation Setup

In this work, we use the uncased BERT-base as the teacher network, and study the robustness of its compressed variants. The final model consists of the BERT-base encoder (or its compressed variants) with a classification head (a linear layer on top of the pooled output). Recent studies indicate that factors such as learning rate and training epochs could have a substantial influence on robustness (Tu et al., 2020). In particular, increasing training epochs can help improve the generalization of the OOD test set. In this work, we focus on the relative robustness of compressed models compared to the uncompressed teacher, rather than their *absolute accuracies*. For a fair comparison, we unify the experimental setup for all models. We use Adam optimizer with weight decay (Loshchilov and Hutter, 2017), where the learning rate is fixed as $2e-5$, and we train all models for 5 epochs on all datasets. We perform the experiments using PyTorch and use the pre-trained models from the Huggingface model pool (Wolf et al., 2019). We report the average results over three runs for all experiments.

3.4 Relative Robustness Metric

As we later demonstrate, with increase in compression ratio or model sparsity, the performance of the smaller models degrades for both in-domain and OOD test sets. To compare the gap between

Sparsity	#Param	MNLI			FEVER				QQP			Average \mathcal{F}_{bias}	
		DEV	HANS	\mathcal{F}_{bias}	DEV	Sym1	Sym2	\mathcal{F}_{bias}	DEV	paws _{wiki}	paws _{qqp}		\mathcal{F}_{bias}
BERT-base	109M	84.2	59.8	-	86.2	58.9	64.5	-	90.9	48.9	34.7	-	-
20%	87.2M	84.4	55.5	1.182	86.5	57.0	64.6	1.045	90.7	47.2	33.5	1.037	1.088
40%	65.4M	84.0	54.7	1.204	86.4	57.2	64.0	1.051	90.5	46.6	32.4	1.049	1.101
60%	43.6M	83.4	52.8	1.266	86.3	56.9	63.3	1.068	90.2	45.9	31.8	1.061	1.132
70%	32.7M	81.8	52.2	1.249	85.9	56.6	63.3	1.063	89.5	45.4	30.7	1.065	1.127

Table 1: Accuracy comparison (in percent) and relative bias \mathcal{F}_{bias} (the smaller the better) for models with iterative magnitude pruning with different levels of sparsity. The last column indicates the average \mathcal{F}_{bias} values over three tasks. *Pruned models have relatively higher degradation in OOD test set compared to the development set.*

Sparsity	#Param	MNLI			FEVER				QQP			Average \mathcal{F}_{bias}	
		DEV	HANS	\mathcal{F}_{bias}	DEV	Sym1	Sym2	\mathcal{F}_{bias}	DEV	paws _{wiki}	paws _{qqp}		\mathcal{F}_{bias}
BERT-base	109M	84.2	59.8	-	86.2	58.9	64.5	-	90.9	48.9	34.7	-	-
DistilBERT	66M	82.3	51.2	1.289	84.5	51.9	60.4	1.183	89.9	48.1	34.6	1.006	1.159
MiniLM	66M	83.1	51.4	1.309	84.2	53.4	60.7	1.137	89.9	46.8	31.0	1.039	1.162
Truncated-l6	66M	80.8	51.6	1.247	84.4	52.6	60.4	1.163	90.0	46.0	32.4	1.056	1.155
Pretrained-l6	66M	81.6	52.2	1.229	85.8	54.7	62.6	1.115	90.0	46.4	33.9	1.045	1.130

Table 2: Accuracy comparison (in percent) and relative bias \mathcal{F}_{bias} (the smaller the better) of compressed models with knowledge distillation. *Distilled models have relatively higher degradation in OOD test set compared to the development set.* Except BERT-base, all other models have 66M parameters.

Models	Attention heads	DEV	HANS	\mathcal{F}_{bias}
BERT-base	144	84.2	59.8	-
BERT-116heads-v1	116	84.1	55.5	1.172
BERT-116heads-v2	116	84.2	53.7	1.250
BERT-116heads-v3	116	84.0	55.3	1.176

Table 3: Accuracy comparison (in percent) and relative bias \mathcal{F}_{bias} (the smaller the better) of compressed models with structured pruning. *Pruned models have relatively higher degradation in OOD test set compared to the development set.* All compressed models have been pruned 28 attention heads.

in-distribution task performance and OOD generalizability, we define a new metric that measures this performance gap of the compressed models with respect to the uncompressed BERT-base (teacher). First, we calculate the accuracy gap between in-distribution development set and OOD test set as $\frac{F_{dev} - F_{OOD}}{F_{dev}}$ for BERT-base (denoted by $\Delta_{BERT-base}$); and its compressed variant (denoted by $\Delta_{compressed}$). Second, we compute the relative bias as the ratio between the accuracy gap of the compressed model with respect to BERT-base: $\mathcal{F}_{bias} = \frac{\Delta_{compressed}}{\Delta_{BERT-base}}$. Here $\mathcal{F}_{bias} > 1$ indicates that the compressed model is more biased than BERT-base with the degree of bias captured in a larger value of \mathcal{F}_{bias} . Since FEVER has two OOD test sets, we use the overall accuracy of sym1 and sym2 to calculate \mathcal{F}_{bias} . Similarly, the OOD accuracy for QQP is the overall accuracy on PAWS-wiki and PAWS-qqp.

3.5 Experimental Observations

We report the performance of accuracy and the relative bias measure \mathcal{F}_{bias} for iterative magnitude pruning in Table 1, knowledge distillation in Table 2 and structured pruning in Table 3. We have the following key observations.

Iterative Magnitude Pruning: First, for slight and mid-level sparsity, the pruned models have comparable and sometimes even better performance on the in-distribution development set. Consider FEVER as an example, where the compressed model preserves the accuracy on the in-distribution set even at 60% sparsity². However, the generalization accuracy on the OOD test set has a substantial drop. This indicates that the development set fails to capture the generalizability of the pruned models. Second, as the sparsity increases, the generalization accuracy on the OOD test set substantially decreases while dropping to random guess for tasks such as MNLI. Third, at high levels of sparsity (e.g. 70%), both development and OOD test set performances are significantly affected. In general, we observe $\mathcal{F}_{bias} > 1$ for all levels of sparsity in Table 1. Note that we limit the maximum sparsity at 70% after which the training is unstable with a significant performance drop even on the development set (Liang et al., 2021). As in the previous cases, there is substantial accuracy drop on the OOD test

²Here, 60% sparsity indicates that 40% parameters are remaining after pruning.

set compared to the development set (e.g., 7.6% vs 1.9% degradation respectively for the MNLI task).

Knowledge Distillation: Similar to pruning, we observe a higher accuracy drop in the OOD test set compared to the in-distribution development set for distilled models. Consider DistilBERT performance on MNLI as an example with 1.9% accuracy drop in development set compared to 8.6% drop in the OOD test set. This can also be validated in Table 2, where all \mathcal{F}_{bias} values are larger than 1, depicting that all the distilled models are less robust than BERT-base. Another interesting observation is that distilled models, i.e., DistilBERT and MiniLM, have higher bias \mathcal{F}_{bias} compared to the pre-trained models, i.e., Pretrained-l6 and Truncated-l6, as we compare their average \mathcal{F}_{bias} values in Table 2. This indicates that the compression process plays a significant role in the low generalizability and robustness of the distilled models.

Structured Pruning: Recent studies have reported the super ticket phenomenon (Liang et al., 2021). The authors observe that, when the BERT-base model is slightly pruned, the accuracy of the pruned models improves on in-distribution development set. However, we observe that this finding does not hold for OOD test sets. From Table 3, we observe that all pruned models are less robust than BERT-base, with \mathcal{F}_{bias} much larger than 1.

4 Attribution of Low Robustness

In this section, we explore the factors that lead to low robustness of compressed models. Previous work has demonstrated that the performance of different models on the GLUE benchmark (Wang et al., 2018) tends to correlate with the performance on MNLI, making it a good representative of natural language understanding tasks in general (Phang et al., 2018; Liu et al., 2020). For this reason, we choose the MNLI task for a study.

For the MNLI task, we consider the dataset splits from (Gururangan et al., 2018). The authors partition the development set into easy/shortcut³ and hard subsets. In this experiment, we use pruned models with varying sparsity to investigate the reason for the low robustness of the compressed models. We have the following key observations.

Observation 1: The compressed models tend to overfit the easy/shortcut samples and generalize poorly on the hard ones. The performance of

³We use ‘easy’ and ‘shortcut’ interchangeably in this work.




Figure 1: Pruned model performance on hard vs easy / shortcut samples with varying sparsity, where x-axis denotes the sparsity level.

pruned models at five sparsity levels (ranging between [0.2 – 0.85]) on the easy and hard samples for the MNLI task is illustrated in Figure 1. It demonstrates that the accuracy on the hard samples is much lower compared to the accuracy on the easy ones. As the sparsity increases, we observe a larger accuracy drop on the hard samples compared to the easy ones. In particular, the accuracy gap between the two subsets is 22.7% at the sparsity of 0.85, much higher than the 16.1% accuracy gap at the sparsity of 0.4. These findings demonstrate that the compressed models overfit on the easy samples, while generalizing poorly on the hard ones. Furthermore, this phenomenon is amplified at higher levels of sparsity for the pruned models.

Observation 2: Compressed models tend to assign overconfident predictions to easy samples. One of the potential reasons is that compressed models are more prone to capture spurious correlations between shortcut features in training samples with certain class labels for their predictions (Geirhos et al., 2020; Du et al., 2021).

4.1 Variance-based Difficulty Estimation

Based on the above observations, we propose a variance-based metric to quantify the difficulty degree of each sample. For each sample in the development set, we calculate its loss at five different levels of pruning sparsity as shown in Figure 1. We further calculate the variance of the above losses for each sample and rank them based on the variance. Finally, we assign the samples with low variance to the ‘easy’ subset and rest to the ‘hard’ one. Comparing our variance-based proxy annotation with the ground truth annotation in (Gururangan et al., 2018) gives an accuracy of 82.8%. This indicates that the variance-based estimation leveraging pruning sparsity is a good indicator of sample difficulty. This motivates our design of the mitigation technique introduced in the next section.

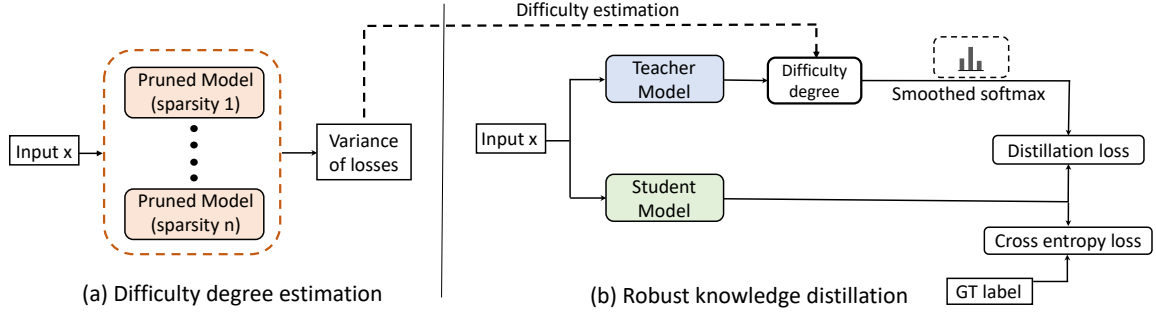


Figure 2: RMC framework for bias mitigation with two-stage training. In the first stage, we feed the training samples to pruned models at different levels of sparsity (ranging from $[0.2 - 0.85]$) as introduced in Section 4.1); compute corresponding losses and their variance to estimate the difficulty degree of each training sample. In the second stage, we use the difficulty degree to regularize the teacher network for robust model compression.

5 Mitigation Framework

In this section, we propose a general bias mitigation framework (see Figure 2), termed as RMC (Robust Model Compression), to improve the robustness of compressed models on downstream tasks. Our RMC framework follows the philosophy of task-specific knowledge distillation (Sanh et al., 2020; Jiao et al., 2020), but with explicit regularization of the teacher network leveraging sample uncertainty. This prevents the compressed model from overfitting in the easy samples that contain shortcut features and helps improve its robustness. This regularized training is implemented in two stages.

5.1 Quantifying Sample Difficulty

In the first stage, our objective is to quantify the difficulty degree of each training sample.

Variance Computation: Following the observations obtained in Section 4.1, we first use iterative magnitude pruning to obtain a series of pruned models from BERT-base with different levels of sparsity and then we use the losses of the pruned models at different levels of sparsity to compute their variance v_i for each training sample x_i : $v_i = \frac{\sum_{t=1}^n (l_{i,t} - \bar{l}_i)^2}{n}$. We choose five sparsity levels, i.e., $n = 5$, that are diverse enough to reflect the difficulty degree of each training sample. Here, samples with high variance correspond to hard ones.

Difficulty Degree Estimation: Based on the variance v_i for each training sample x_i , we can estimate its difficulty degree as:

$$d_i = \alpha + \frac{1 - \alpha}{V_{\max} - V_{\min}} \cdot (v_i - V_{\min}), \quad (1)$$

where V_{\min} and V_{\max} denote the minimum and maximum values of the variances, respectively. Equation 1 is used to normalize the variance of

the training samples in the range of $[\alpha, 1]$, where $d_i = 1$ denotes the most difficult training sample, according to our criteria of loss variance. Samples with d_i closer to α are treated as shortcut/biased samples. Prior work (Niven and Kao, 2019) show that the bias behavior of the downstream training set can be attributed to data collection and annotation biases. Since the bias level is different for each dataset, we assign a different α in Equation 1 to each training set to reflect its bias level.

5.2 Robust Knowledge Distillation

In the second stage, we fine-tune BERT-base on the downstream tasks to obtain the softmax probability for each training sample. We then use the difficulty degree of the training samples (discussed in the previous section) to smooth the teacher predictions. The instance-level smoothed softmax probability is used to guide the training of compressed models through regularized knowledge distillation.

Smoothing Teacher Predictions: We smooth the softmax probability \hat{y}_i^T from the teacher network, according to the difficulty degree d_i of each training sample x_i . The smoothed probability is given as:

$$s_{i,j} = \frac{(\hat{y}_i^T)_j^{d_i}}{\sum_{k=1}^K (\hat{y}_i^T)_k^{d_i}}, \quad (2)$$

where K denotes the total number of class labels. We perform instance-level smoothing for each training sample x_i . If the difficulty degree of a training sample $d_i = 1$, then the softmax probability s_i for the corresponding sample from the teacher is unchanged. In contrast, at the other extreme as $d_i \rightarrow \alpha$, we increase the regularization to encourage the compressed model to assign less over-confident predictions to the sample. The difficulty degree range is $[\alpha, 1]$ rather than $[0, 1]$ to avoid over-smoothing of the teacher predictions.

Sparsity	#Param	MNLI			FEVER				QQP			Average \mathcal{F}_{bias}	
		DEV	HANS	\mathcal{F}_{bias}	DEV	Sym1	Sym2	\mathcal{F}_{bias}	DEV	paws _{wiki}	paws _{qqp}		\mathcal{F}_{bias}
BERT-base	110M	84.2	59.8	-	86.2	58.9	64.5	-	90.9	48.9	34.7	-	-
40% – Vanilla	65.4M	84.0	54.7	1.204	86.4	57.2	64.0	1.051	90.5	46.6	32.4	1.049	1.101
– Distil	65.4M	84.1	56.2	1.145	86.3	58.4	64.5	1.013	90.5	47.3	33.2	1.032	1.063
– Smooth	65.4M	84.2	56.5	1.135	86.2	60.7	65.8	0.937	90.7	47.2	33.8	1.036	1.036
– Focal	65.4M	84.0	56.7	1.122	86.4	59.4	65.2	0.981	90.7	46.2	32.1	1.060	1.054
– JTT	65.4M	83.8	56.3	1.132	86.2	58.1	64.9	1.008	90.4	47.3	33.7	1.030	1.057
– RMC	65.4M	84.2	58.6	1.049	86.1	61.9	66.4	0.897	90.4	47.6	34.3	1.023	0.990

Table 4: Generalization accuracy comparison (in percent) and the corresponding \mathcal{F}_{bias} values for iterative magnitude pruning at 40% sparsity with different *mitigation methods*. The last column indicates average \mathcal{F}_{bias} over three tasks.

Smoothness-Induced Model Compression: We employ the smoothed softmax probability s_i from BERT-base to supervise the training of the compressed models, where the overall loss function is:

$$\mathcal{L}(x) = (1 - \lambda) * \mathcal{L}_1(y_i, \hat{y}_i^S) + \lambda * \mathcal{L}_2(s_i, \hat{y}_i^S), \quad (3)$$

where y_i is the ground truth and \hat{y}_i^S is the probability of the compressed model. \mathcal{L}_1 denotes the cross-entropy loss, and \mathcal{L}_2 represents the knowledge distillation loss with KL divergence. Hyperparameter λ manages the trade-off between learning from hard label y_i and softened softmax probability s_i . Among the different families of compression techniques introduced in Section 3.1, we directly fine-tune the distilled models using Equation 3. For iterative magnitude pruning, we use Equation 3 to guide the pruning during the fine-tuning process.

6 Mitigation Performance Evaluation

In this section, we conduct experiments to evaluate the robustness of our RMC mitigation framework.

6.1 Experimental Setup

For all experiments, we follow the same setting as in Section 3.3, and the same evaluation datasets as in Section 3.2. We use the OOD test set exclusively for evaluation. We compute the variance of samples (outlined in Section 4.1) in the in-distribution development set to split it into a shortcut and hard subset. The relative robustness between the hard and easy subset is used to tune the hyperparameter α in Equation 1, where we set α as 0.5, 0.3, 0.2 for MNLI, FEVER, and QQP, respectively. The weight λ in Equation 3 is fixed as 0.9 for all experiments.

6.2 Baseline Methods

We consider the following five baselines. Please refer to Sec. C in Appendix for more details.

- **Vanilla:** This only fine-tunes the base encoder without any regularization.
- **Distil** (Task-Specific Knowledge Distillation) (Sanh et al., 2020): This first fine-tunes BERT-base on the downstream NLU tasks. The softmax probability from the fine-tuned BERT-base is used as the supervision signal for distillation.
- **Smooth** (Global Smoothing) (Müller et al., 2019): This performs global smoothing for *all* training samples with task-specific knowledge distillation, where we use the same level of regularization as in RMC ($d_i = 0.9$ in Equation 2). In contrast, RMC uses instance-level smoothing.
- **Focal** (Focal Loss) (Lin et al., 2017): Compared to cross-entropy loss, focal loss has an additional regularizer to reduce the weight for easy samples and assign a higher weight to hard samples bearing less-confident predictions.
- **JTT** (Just Train Twice) (Liu et al., 2021): This is a re-weighting method, which first trains the BERT-base model using standard cross-entropy loss for several epochs, and then trains the compressed model while up-weighting the training examples that are misclassified by the first model, i.e., hard samples.

6.3 Mitigation Performance Analysis

We compare our RMC framework with the above baselines and have the following key observations.

Iterative Magnitude Pruning: Table 4 shows the mitigation results of accuracy and relative bias \mathcal{F}_{bias} . All mitigation methods are performed with pruned models at 40% sparsity. We observe that task-specific knowledge distillation only slightly improves accuracy on the OOD test set compared to Vanilla tuning, since the teacher model itself is not robust for downstream tasks (Niven and Kao, 2019). Global smoothing further improves generalization accuracy compared to prior methods. Our

Sparsity	#Param	MNLI			FEVER				QQP			Average \mathcal{F}_{bias}	
		DEV	HANS	\mathcal{F}_{bias}	DEV	Sym1	Sym2	\mathcal{F}_{bias}	DEV	paws _{wiki}	paws _{qqp}		\mathcal{F}_{bias}
BERT-base	110M	84.2	59.8	-	86.2	58.9	64.5	-	90.9	48.9	34.7	-	-
MiniLM – Vanilla	66M	83.1	51.4	1.309	84.2	53.4	60.7	1.137	89.9	46.8	31.0	1.039	1.162
– Distil	66M	83.1	53.7	1.221	83.8	56.5	61.0	1.052	89.6	46.7	31.8	1.037	1.103
– Smooth	66M	82.7	53.8	1.206	83.7	56.9	62.1	1.017	89.4	46.8	32.2	1.032	1.085
– Focal	66M	83.2	55.6	1.145	83.8	54.7	61.4	1.081	90.3	46.8	33.2	1.041	1.089
– JTT	66M	82.8	55.7	1.129	83.5	53.8	61.7	1.085	90.1	47.0	32.9	1.034	1.083
– RMC	66M	83.7	57.8	1.068	85.3	58.0	63.3	1.017	90.5	47.0	33.4	1.038	1.041

Table 5: Generalization accuracy and the \mathcal{F}_{bias} values comparison of different training strategies with and without mitigation on in-distribution development set and OOD test set using MiniLM as the compressed encoder.

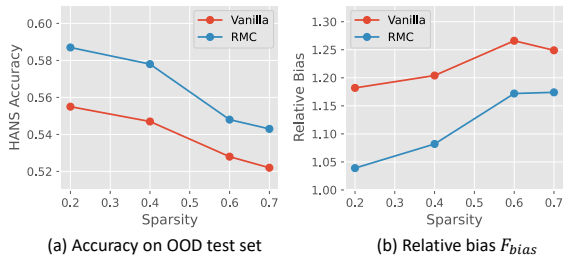


Figure 3: RMC mitigation performance for iterative magnitude pruning at different levels of pruning sparsity for MNLI task.

RMC framework obtains the best accuracy on OOD test set across all the tasks on aggregate. RMC further reduces the average relative bias \mathcal{F}_{bias} by 10% over Vanilla tuning, as shown in Table 4, indicating the benefits of uncertainty-based sample-wise smoothing in terms of improving model robustness. For the MNLI task, we also illustrate the mitigation performance of our RMC framework for different levels of sparsity in Figure 3. We observe that RMC consistently improves accuracy on OOD HANS while reducing the relative bias \mathcal{F}_{bias} for all levels of sparsity over the Vanilla method.

Knowledge Distillation: Table 5 shows the mitigation results of accuracy and relative bias \mathcal{F}_{bias} . We observe that RMC significantly improves over MiniLM for OOD generalization leveraging smoothed predictions from BERT-base teacher. With instance-level smoothing in RMC, the generalization accuracy for the compressed model on the OOD test set is significantly closer to BERT-base teacher compared to the other methods. We also decrease the relative bias \mathcal{F}_{bias} in Table 5 by 10.4% over Vanilla tuning. On the QQP task, RMC simultaneously improves the performance of compressed model on both the in-distribution development set and the two OOD test sets.

Models	DEV	HANS	Hard (H)	Easy (E)	Gap (E-H)
MiniLM–Vanilla	83.1	51.4	73.2	90.9	17.7
MiniLM–RMC	83.7	57.8	74.9	90.6	15.7
40%–Vanilla	84.0	54.7	74.9	91.0	16.1
40%–RMC	84.2	58.6	75.9	90.3	14.4

Table 6: Our RMC framework improves accuracy of the compressed models on the hard samples and reduces overfitting on the shortcut/easy samples, leading to reduced performance gap between the two subsets.

6.4 Further Analysis on Robust Mitigation

In this section, we further investigate the reasons for the improved generalization performance with RMC with an analysis on the MNLI task. Table 6 shows the accuracy performance of RMC for model pruning and distillation on the shortcut/easy and hard samples. We observe RMC to improve the model performance on the under-represented hard samples, where it reduces the generalization gap between the hard and shortcut/easy subset by 10.6% at 0.4 level of sparsity and by 11.3% for knowledge distillation. This analysis demonstrates that RMC reduces the overfitting of the compressed models on the easy samples and encourages them to learn more from the hard ones, thus improving the generalization on the OOD test sets.

7 Conclusions

In this work, we conduct a comprehensive study of the robustness challenges in compressing large PLMs when fine-tuning in downstream NLU datasets. Furthermore, we propose a general mitigation framework with instance-level smoothing for robust model compression. Experimental analysis demonstrates our framework to improve the generalization and OOD robustness of compressed models for different compression techniques, while not sacrificing the in-distribution performance.

Limitations

First, we study the shortcut learning/bias problem and OOD generalization of model compression techniques, exclusively focusing on the two most widely used families of compression techniques, including knowledge distillation and pruning. Our empirical analysis indicates that these two families of compression techniques suffer from the low generalization issue. However, other types of compression technique, such as matrix decomposition and quantization, are not discussed in this work. Studying the whole compression techniques is a challenging topic and will be investigated in our future research. Second, our RMC framework needs to calculate the variance of losses for each training sample, thus requiring additional training time. Training efficiency can be further improved by implementing parallel training or more efficient ways of calculating sample difficulty, which will also be studied in our future research.

References

- Yuchen Bian, Jiaji Huang, Xingyu Cai, Jiahong Yuan, and Kenneth Church. 2021. On attention redundancy: A comprehensive study. In *North American Chapter of the Association for Computational Linguistics: Human Language Technologies (NAACL)*.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in Neural Information Processing Systems (NeurIPS)*.
- Tianlong Chen, Yu Cheng, Zhe Gan, Lu Yuan, Lei Zhang, and Zhangyang Wang. 2021. [Chasing sparsity in vision transformers: an end-to-end exploration](#).
- Tianlong Chen, Jonathan Frankle, Shiyu Chang, Sijia Liu, Yang Zhang, Zhangyang Wang, and Michael Carbin. 2020. The lottery ticket hypothesis for pre-trained bert networks. *34th Conference on Neural Information Processing Systems (NeurIPS)*.
- Hao Cheng, Xiaodong Liu, Lis Pereira, Yaoliang Yu, and Jianfeng Gao. 2021. Posterior differential regularization with f-divergence for improving model robustness. *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Christopher Clark, Mark Yatskar, and Luke Zettlemoyer. 2019. Don't take the easy way out: Ensemble based methods for avoiding known dataset biases. *Empirical Methods in Natural Language Processing (EMNLP)*.
- Alexander D'Amour, Katherine Heller, Dan Moldovan, Ben Adlam, Babak Alipanahi, Alex Beutel, Christina Chen, Jonathan Deaton, Jacob Eisenstein, Matthew D Hoffman, et al. 2020. Underspecification presents challenges for credibility in modern machine learning. *arXiv preprint arXiv:2011.03395*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Bert: Pre-training of deep bidirectional transformers for language understanding. *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Mengnan Du, Fengxiang He, Na Zou, Dacheng Tao, and Xia Hu. 2022. Shortcut learning of large language models in natural language understanding: A survey. *arXiv preprint arXiv:2208.11857*.
- Mengnan Du, Varun Manjunatha, Rajiv Jain, Ruchi Deshpande, Franck Dernoncourt, Jiuxiang Gu, Tong Sun, and Xia Hu. 2021. Towards interpreting and mitigating shortcut learning behavior of nlu models. *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Robert Geirhos, Jörn-Henrik Jacobsen, Claudio Michaelis, Richard Zemel, Wieland Brendel, Matthias Bethge, and Felix A Wichmann. 2020. Shortcut learning in deep neural networks. *Nature Machine Intelligence*, 2(11):665–673.
- Suchin Gururangan, Swabha Swayamdipta, Omer Levy, Roy Schwartz, Samuel R Bowman, and Noah A Smith. 2018. Annotation artifacts in natural language inference data. *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- He He, Sheng Zha, and Haohan Wang. 2019. Unlearn dataset bias in natural language inference by fitting the residual. *2019 EMNLP workshop*.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. 2015. Distilling the knowledge in a neural network. *NeurIPS Deep Learning Workshop*.
- Sara Hooker, Aaron Courville, Gregory Clark, Yann Dauphin, and Andrea Frome. 2019. What do compressed deep neural networks forget? *arXiv preprint arXiv:1911.05248*.
- Sara Hooker, Nyalleng Moorosi, Gregory Clark, Samy Bengio, and Emily Denton. 2020. Characterising bias in compressed models. *arXiv preprint arXiv:2010.03058*.
- Shaoyi Huang, Dongkuan Xu, Ian EH Yen, Sung-en Chang, Bingbing Li, Shiyang Chen, Mimi Xie, Hang Liu, and Caiwen Ding. 2021. Sparse progressive distillation: Resolving overfitting under pretrain-and-finetune paradigm. *arXiv preprint arXiv:2110.08190*.
- 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. *Findings of EMNLP*.

- Di Jin, Zhijing Jin, Joey Tianyi Zhou, and Peter Szolovits. 2020. Is bert really robust? natural language attack on text classification and entailment. *AAAI Conference on Artificial Intelligence (AAAI)*.
- Tianda Li, Ahmad Rashid, Aref Jafari, Pranav Sharma, Ali Ghodsi, and Mehdi Rezagholizadeh. 2021. How to select one among all? an empirical study towards the robustness of knowledge distillation in natural language understanding. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 750–762.
- Chen Liang, Simiao Zuo, Minshuo Chen, Haoming Jiang, Xiaodong Liu, Pengcheng He, Tuo Zhao, and Weizhu Chen. 2021. Super tickets in pre-trained language models: From model compression to improving generalization. *59th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. 2017. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision (ICCV)*.
- Evan Z Liu, Behzad Haghgoo, Annie S Chen, Aditi Raghunathan, Pang Wei Koh, Shiori Sagawa, Percy Liang, and Chelsea Finn. 2021. Just train twice: Improving group robustness without training group information. In *International Conference on Machine Learning (ICML)*.
- Xiaodong Liu, Yu Wang, Jianshu Ji, Hao Cheng, Xueyun Zhu, Emmanuel Awa, Pengcheng He, Weizhu Chen, Hoifung Poon, Guihong Cao, et al. 2020. The microsoft toolkit of multi-task deep neural networks for natural language understanding. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 118–126.
- 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*.
- Ilya Loshchilov and Frank Hutter. 2017. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*.
- R Thomas McCoy, Junghyun Min, and Tal Linzen. 2019a. Berts of a feather do not generalize together: Large variability in generalization across models with similar test set performance. *arXiv preprint arXiv:1911.02969*.
- R Thomas McCoy, Ellie Pavlick, and Tal Linzen. 2019b. Right for the wrong reasons: Diagnosing syntactic heuristics in natural language inference. *57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Paul Michel, Omer Levy, and Graham Neubig. 2019. Are sixteen heads really better than one? *33rd Conference on Neural Information Processing Systems (NeurIPS)*.
- Pramod Kaushik Mudrakarta, Ankur Taly, Mukund Sundararajan, and Kedar Dhamdhere. 2018. Did the model understand the question? *56th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Rafael Müller, Simon Kornblith, and Geoffrey E Hinton. 2019. When does label smoothing help? In *Advances in Neural Information Processing Systems (NeurIPS)*.
- Timothy Niven and Hung-Yu Kao. 2019. Probing neural network comprehension of natural language arguments. *57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Jason Phang, Thibault FÉvry, and Samuel R Bowman. 2018. Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks. *arXiv e-prints*, pages arXiv–1811.
- Sai Prasanna, Anna Rogers, and Anna Rumshisky. 2020. When bert plays the lottery, all tickets are winning. *Empirical Methods in Natural Language Processing (EMNLP)*.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *NeurIPS Workshop*.
- Victor Sanh, Thomas Wolf, Yonatan Belinkov, and Alexander M Rush. 2021. Learning from others’ mistakes: Avoiding dataset biases without modeling them. *International Conference on Learning Representations (ICLR)*.
- Victor Sanh, Thomas Wolf, and Alexander M Rush. 2020. Movement pruning: Adaptive sparsity by fine-tuning. *34th Conference on Neural Information Processing Systems (NeurIPS)*.
- Tal Schuster, Darsh J Shah, Yun Jie Serene Yeo, Daniel Filizzola, Enrico Santus, and Regina Barzilay. 2019. Towards debiasing fact verification models. *Empirical Methods in Natural Language Processing (EMNLP)*.
- Joe Stacey, Pasquale Minervini, Haim Dubossarsky, Sebastian Riedel, and Tim Rocktäschel. 2020. Avoiding the hypothesis-only bias in natural language inference via ensemble adversarial training. *Empirical Methods in Natural Language Processing (EMNLP)*.
- Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. 2019. Patient knowledge distillation for bert model compression. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. 2020. Mobilebert: a compact task-agnostic bert for resource-limited devices. *58th Annual Meeting of the Association for Computational Linguistics (ACL)*.

- James Thorne, Andreas Vlachos, Christos Christodoulopoulos, and Arpit Mittal. 2018. Fever: a large-scale dataset for fact extraction and verification. *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Lifu Tu, Garima Lalwani, Spandana Gella, and He He. 2020. An empirical study on robustness to spurious correlations using pre-trained language models. *Transactions of the Association for Computational Linguistics (TACL)*.
- Prasetya Ajie Utama, Nafise Sadat Moosavi, and Iryna Gurevych. 2020. Towards debiasing nlu models from unknown biases. *Empirical Methods in Natural Language Processing (EMNLP)*.
- Elena Voita, David Talbot, Fedor Moiseev, Rico Senrich, and Ivan Titov. 2019. Analyzing multi-head self-attention: Specialized heads do the heavy lifting, the rest can be pruned. *57th Annual Meeting of the Association for Computational Linguistics (ACL)*.
- Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel Bowman. 2018. [GLUE: A multi-task benchmark and analysis platform for natural language understanding](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 353–355, Brussels, Belgium. Association for Computational Linguistics.
- Boxin Wang, Shuohang Wang, Yu Cheng, Zhe Gan, Ruoxi Jia, Bo Li, and Jingjing Liu. 2021. Infobert: Improving robustness of language models from an information theoretic perspective. In *International Conference on Learning Representations*.
- 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. *34th Conference on Neural Information Processing Systems (NeurIPS)*.
- Adina Williams, Nikita Nangia, and Samuel R Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi Louf, Morgan Funtowicz, et al. 2019. Huggingface’s transformers: State-of-the-art natural language processing. *arXiv preprint arXiv:1910.03771*.
- Canwen Xu, Wangchunshu Zhou, Tao Ge, Ke Xu, Julian McAuley, and Furu Wei. 2021a. Beyond preserved accuracy: Evaluating loyalty and robustness of bert compression. *Empirical Methods in Natural Language Processing (EMNLP)*.
- Dongkuan Xu, Ian EH Yen, Jinxi Zhao, and Zhibin Xiao. 2021b. Rethinking network pruning—under the pre-train and fine-tune paradigm. *North American Chapter of the Association for Computational Linguistics (NAACL)*.
- Yadollah Yaghoobzadeh, Remi Tachet, Timothy J Hazen, and Alessandro Sordoni. 2019. Robust natural language inference models with example forgetting. *arXiv e-prints*, pages arXiv–1911.
- Yinfei Yang, Yuan Zhang, Chris Tar, and Jason Baldridge. 2019. PAWS-X: A Cross-lingual Adversarial Dataset for Paraphrase Identification. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Yuan Zhang, Jason Baldridge, and Luheng He. 2019. Paws: Paraphrase adversaries from word scrambling. *arXiv preprint arXiv:1904.01130*.
- Chen Zhu, Yu Cheng, Zhe Gan, Siqi Sun, Tom Goldstein, and Jingjing Liu. 2020. [FreeLB: Enhanced adversarial training for natural language understanding](#). In *International Conference on Learning Representations*.

A More Details of Pruning Methods

In this section, we introduce more details about the compression techniques studied.

knowledge Distillation: For a fair comparison, we do not compare with TinyBERT (Jiao et al., 2020) and MobileBERT (Sun et al., 2020), since TinyBERT is fine-tuned with data augmentation on NLU tasks, and MobileBERT is distilled from BERT-large rather than BERT-base.

Magnitude Pruning: It is based on the overparameterization assumption of pre-trained language models (Xu et al., 2021b; Huang et al., 2021). For iterative magnitude pruning, we freeze all the embedding modules and only prune the parameters in the encoder (i.e., 12 layers of Transformer blocks). After pruning, the pruned weight values are set to 0 to reduce the amount of information to store. Unlike the LTH version, we consider standard magnitude pruning without using rewinding.

Structured Pruning: To calculate the importance, we follow (Michel et al., 2019; Prasanna et al., 2020) and calculate the expected sensitivity of the attention heads to the mask variable $\xi^{(h,l)}$: $I_h^{(h,l)} = E_{x \sim X} \left| \frac{\partial \mathcal{L}(x)}{\partial \xi^{(h,l)}} \right|$, where $I_h^{(h,l)}$ denotes the contribution score of the attention head h in layer l , $\mathcal{L}(x)$ represents the loss value for the sample x , and $\xi^{(h,l)}$ is the mask of the attention head h in layer l . After obtaining the contribution scores, the attention heads with lowest score $I_h^{(h,l)}$ are pruned.

B More on Evaluation Datasets

In this section, we introduce more details about the three benchmark datasets.

MNLI: This task aims to predict whether the relationship between the premise and the hypothesis is contradiction, entailment, or neutral. It is divided into a training set and development set with 392,702 and 9,815 samples, respectively.

FEVER: The task is to predict whether the claims support, refute, or not-have-enough-information about the evidence. Recent studies indicate that there are strong shortcuts in claims (Utama et al., 2020). It is divided into a training set and a development set with 242,911 and 16,664 samples, respectively.

QQP: It is divided into a training set and a development set with 363,846 and 40,430 samples, respectively.

C More on Comparing Baselines

In this section, we introduce more details on comparing baselines.

Distil and Smooth: For both baseline methods, we use a loss function similar to that of Equation 3. We fix the weight λ to 0.9 for all experiments, to encourage the compressed model to learn more from the probability output of the teacher network. A major difference between the two baselines is that Smooth has an additional smoothing process involved during the fine-tuning process.

Focal Loss: The original focal loss function is: $FL(p_i) = -(1 - p_i)^\gamma \log(p_i)$. Our implementation is as follows:

$$FL(p_i) = -\frac{(1 - p_i)^\gamma}{\frac{1}{N} \sum_{k=1}^N (1 - p_k)^\gamma} \log(p_i).$$

The hyperparameter γ controls the weight difference between hard and easy samples, and is fixed at 2.0 for all tasks. We use the denominator to normalize the weights within a batch, where N is the batch size. This is used to guarantee that the average weight for a batch of training samples is 1.0. As such, the weight for the easy samples would be down-weighted to lower than 1.0, and the weight for hard samples would be up-weighted to values larger than 1.0.

JTT: This is also a reweighting baseline that encourages the model to learn more from hard samples. The hyperparameter λ_{up} in (Liu et al., 2021) is set to 2.0. We also normalize the weights so that the average weight for each training sample is 1.0.

D Running Environment

For a fair evaluation of the robustness of compressed models, we run all experiments using a server with 4 NVIDIA GeForce 3090 GPUs. All experiments are implemented with the Pytorch version of the Hugging Face Transformer library.

E The Capacity Issue

One natural speculation about the low robustness of compressed models is due to their low capacity (i.e., smaller size). To disentangle the two important factors that influence model performance, i.e., low capacity and compression, we compare distilled models with Uncased-l6, which is trained only using pretraining. The results are given in Table 2. The results indicate that Uncased-l6 has better generalization ability over the MNLI and

FEVER two tasks. Take structured pruning as an example; although the three pruned models in Table 3 have the same model size, their generalization accuracy is different. These results indicate that the low robustness of compressed models is not entirely due to their low capacity, and compression plays a significant role.

F MNL I Easy and Hard Subsets

The authors train a hypothesis-only model and use it to generate predictions for the whole development set (Gururangan et al., 2018). Samples that are given correct predictions by the hypothesis-only model are regarded as easy samples, and vice versa. The easy subset contains 5488 samples, and the hard subset contains 4302 samples.