
Pruning Pretrained Encoders with a Multitask Objective

Patrick Xia*

Johns Hopkins University
paxia@cs.jhu.edu

Richard Shin

Microsoft Semantic Machines
richard.shin@microsoft.com

Abstract

The sizes of pretrained language models make them challenging and expensive to use when there are multiple desired downstream tasks. In this work, we adopt recent strategies for model pruning during finetuning to explore the question of whether it is possible to prune a single encoder so that it can be used for multiple tasks. We allocate a fixed parameter budget and compare pruning a single model with a multitask objective against the best ensemble of single-task models. We find that under two pruning strategies (element-wise and rank pruning), the approach with the multitask objective outperforms training models separately when averaged across all tasks, and it is competitive on each individual one. Additional analysis finds that using a multitask objective during pruning can also be an effective method for reducing model sizes for low-resource tasks.

1 Introduction

NLP models typically use pretrained text encoders like BERT (Devlin et al., 2019), which perform well when finetuned across many downstream NLP tasks. At the same time, these models are often overparameterized for the downstream task, leading to a surge of interest in reducing encoder size while retaining most of its performance on downstream tasks (Sun et al., 2020; Sanh et al., 2020a).

Meanwhile, there has been interest in adapting a *single* model to *multiple* downstream tasks through the use of a small number of additional, task-specific parameters (Houlsby et al., 2019; Shin et al., 2020; Hu et al., 2021). These techniques are useful for efficiently sharing large base models during training by freezing the underlying encoder and finetuning parameters dependent on the target task. Switching between tasks is also cheap: only a small component is changed. These methods for extending models of size N with ϵ parameters per task can perform well on t tasks at the cost of $N + t\epsilon$ instead of tN parameters, and architectural innovations have led to smaller ϵ (Karimi Mahabadi et al., 2021a,b). In practice, this allows for deployment of models for multiple tasks at the cost of a single model in terms of memory or disk space.

In this work, we explore the possibility of further reducing the number of parameters used by these multitask models to be substantially smaller than N . Specifically, we aim to do so by pruning *multitask* models using a multitask training objective. In addition to the benefits described above, we aim for the best of both worlds: a substantially pruned model that also performs well on multiple tasks. We also ask whether individual task performances can be improved by leveraging data from the other tasks, which is a strategy employed by general-purpose language modeling (Liu et al., 2019; Aghajanyan et al., 2021).

We present an empirical study specifically for pruning in the multitask scenario. In this paper:

- We extend both structured and unstructured pruning methods to the multitask setting.

*Work done while an intern at Microsoft.

- Under both methods, we find that a multitask model consistently outperforms a combination of single-task models for a given budget.
- Using a multitask objective does not necessarily lead to a loss in performance on any individual task, and in some cases, can improve upon a single-task objective.
- A multitask objective enables improved performance for tasks with smaller dataset sizes.

2 Approach

2.1 Pruning Methods

We explore 8 settings for model pruning, varying pruning method (magnitude vs. movement), what gets pruned (element-wise (unstructured) pruning vs. rank (structured) pruning), and where (global vs. local). Table 1 summarizes the differences.

Magnitude vs. movement pruning We consider two pruning heuristics. *Magnitude pruning* is a well-studied method which incrementally removes the smallest-magnitude weights during training, until only the top- k largest weights remain at completion (LeCun et al., 1990; Han et al., 2015; Zhu and Gupta, 2018). In this work, k refers to a *fraction* of the weights, and so it directly controls the sparsity level and size of the pruned model.

Movement pruning (Sanh et al., 2020b) is a first-order method which prunes by using learned importance scores that correspond to the direction of growth of the weights. Formally, for a weight \mathbf{W} , this method associates with it a set of scores \mathbf{S} . In the forward pass, a binary mask $\mathbf{M} = \text{Top}_k(\mathbf{S})$ is applied to the weight, and so for an input \mathbf{x} , $(\mathbf{W} \odot \mathbf{M})\mathbf{x}$ is used instead, where \odot is the element-wise product. In the backward pass, all weights are updated as Top_k is ignored and its gradient is approximated using a straight-through estimator (Bengio et al., 2013). Sanh et al. (2020b) showed this to be effective during finetuning when the model’s pretraining objective is different from that of a downstream task.

There are extensions to both methods which use an minimum threshold rather than top- k ; however, controlling the sparsity then requires a sweep per task to discover the best threshold, which is expensive especially in the multitask setting.

Unstructured vs. structured pruning In prior work, comparisons for which weights to prune are done *element-wise* (e.g. top- k), which can lead to arbitrary, unstructured, sparsity patterns. It can be difficult to realize practical efficiency gains from doing so as arbitrary locations in weight matrices will be set to zero. Storing such sparse matrices does not lead to immediate gains and sparse matrix multiplication is not always faster, especially on GPUs (Gale et al., 2020). As such, other work considers structured pruning of entire rows or columns of the matrices, which makes it much easier to realize efficiency gains (Fan et al., 2021; Lagunas et al., 2021). We explore an alternative structured pruning approach, *rank pruning* (Yang et al., 2020).

Starting with a weight $\mathbf{W} \in \mathbb{R}^{m \times n}$, we use SVD to approximate $\mathbf{W} = \mathbf{U}\Sigma\mathbf{V}$ where $\mathbf{U} \in \mathbb{R}^{m \times k}$, $\Sigma \in \mathbb{R}^{k \times k}$ is diagonal, and $\mathbf{V} \in \mathbb{R}^{k \times n}$. Initially, $k = \min(m, n)$. By pruning values from the diagonal of Σ , we reduce its rank from k to k' . Consequently, entire rows or columns from \mathbf{U} and \mathbf{V} are pruned, resulting in three reduced matrices: \mathbf{U}' , Σ' , and \mathbf{V}' .

The resulting matrices can be stored as $\mathbf{U}'\Sigma' \in \mathbb{R}^{m \times k'}$ and $\mathbf{V}' \in \mathbb{R}^{k' \times n}$, which uses $k'(m + n)$ parameters. If $k' < \frac{k}{2}$, this is smaller than the original $m \times n$ size of \mathbf{W} .² Similarly, decomposing \mathbf{W} and pruning in this way would lead to faster inference – for an input $x \in \mathbb{R}^{m \times l}$, $\mathbf{W}^\top x$ costs $O(lmn)$ while $((\mathbf{U}'\Sigma')\mathbf{V}')^\top x$ only needs $O(lk'(m + n))$.

The elements in Σ are initially the singular values of \mathbf{W} , and so their magnitudes are interpretable as the importance of a particular dimension. Yang et al. (2020) prune directly on the magnitudes while enforcing orthogonal regularizers to constrain Σ to be close to the singular values. In contrast, we only use SVD as initialization and use no additional regularizers. Since it is not obvious that

²If $k' > \frac{k}{2}$, we can recover the unfactorized form, $\mathbf{W}' = \mathbf{U}'\Sigma'\mathbf{V}'$. Thus, rank pruning would never use more parameters or computation than element-wise pruning.

Variable	Choices	How it changes pruning objective
Selection	Magnitude	Remove model weights with smallest magnitudes
	Movement	Learned importance scores per weight
Structure	Element-wise Rank	Prune parameters independently from each other SVD then prune diagonal entries (rank)
Scope	Global	top- k computed across entire model
	Local	top- k computed per weight matrix

Table 1: Summary of the 8 settings for model pruning we explored. There are 3 variables with 2 choices each, leading to $2^3 = 8$ combinations. In all settings, the pruning methods all iteratively prune the model to keep a constant fraction of parameters.

interpreting the weights as singular values is valid after finetuning, we also try movement pruning as it is empirically effective for finetuning.

Global vs. local pruning We experiment with pruning both *locally* (each parameter is pruned to a target sparsity) and *globally* (the entire model is pruned to a target sparsity). Pruning globally allows for more aggressive pruning of less useful model components or layers. However, it could also be harmful if the weights are not globally comparable. In particular, each Transformer layer has self-attention and feedforward weights that are not normalized to the same scale, especially after SVD.

2.2 Multitask extension

Extending these methods to the multitask setting is straightforward. We train separate, unpruned classification heads for each task while sharing a common set of pruned encoder weights (and learned importance scores, in the case of movement pruning) across all tasks.

Alternatively, one could selectively prune or modify weights (and learn importance scores) based on the specific task (Liang et al., 2021; Guo et al., 2021). In preliminary experiments, we explore this idea by considering three methods of creating the task-specific mask: 1) *shared*, in which $\mathbf{M}_{t_i} = \mathbf{M}_{t_j}$ for all tasks t_i, t_j , and so there would only be one set of scores, \mathbf{S} ; 2) *separate*, in which for a task t , a separate scores \mathbf{S}_t is learned and updated, which can then produce a corresponding \mathbf{M}_t ; 3) *hybrid*, in which both a shared \mathbf{S} and a task-specific \mathbf{S}_t are learned and the corresponding mask $\mathbf{M}_t = \max(\text{Top}_k(\mathbf{S}), \text{Top}_k(\mathbf{S}_t))$.

However, we found both *separate* and *hybrid* to be expensive or slow during training when there are too many tasks, as a separate \mathbf{S}_t needs to be stored, possibly outside of limited GPU memory. At the same time, we found that they did not lead to performance improvement. Especially under *separate*, there was minimal sharing of parameters between tasks. Furthermore, this idea cannot be easily extended to magnitude pruning. Therefore, for the remainder of the study, we focus on the simplest *shared* setting.

In multitask pruning, each task is sampled uniformly at random and optimized for that task’s objective. While there may be some benefits to smoothing the sampling based on the dataset size (e.g. see Section 3.5.2 of Raffel et al. (2020)), we focus on the simplest setting for this proof-of-concept study.

2.3 Model and Datasets

Like prior work (Jiao et al., 2020; Sun et al., 2020; Sanh et al., 2020b), we prune the BERT-base model (Devlin et al., 2019) and evaluate on a variety of downstream NLP benchmarks. We report pruning fraction relative only to the 12-layer Transformer, and we consider the embedding and output classification layers a fixed cost. This is also the treatment in Sanh et al. (2020b). We use their code with the same hyperparameters for all main experiments, with the exception of rank pruning for which we increase the learning rate of Σ to 5×10^{-3} . Like that work, we also primarily benchmark on MNLI (Williams et al., 2018), SQuADv1.1 (Rajpurkar et al., 2016), and QQP (Iyer et al., 2017).

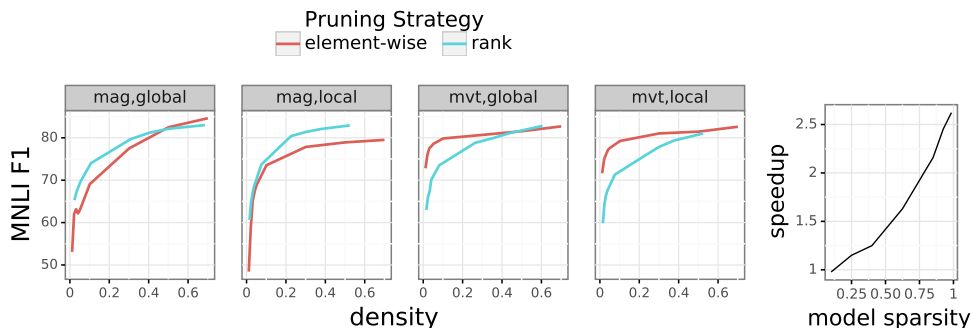


Figure 1: Left: The performance on MNL (dev) across [magnitude, movement] and [global, local] pruning strategies. Within each plot, we show the performance of pruned element-wise and rank-pruned models. Right: Comparison of the runtime of a model using rank pruning relative to an entirely dense model, showing that the structured sparsity ensured by rank pruning can lead to practical benefits.

3 Experiments and Results

First, we benchmark each of the 8 pruning methods on a single dataset (Section 3.1). The goal is to identify task and practical tradeoffs between structured and unstructured pruning and the best settings for each. In Section 3.2, we compare a model pruned with a multitask objective against an ensemble of single-task models to explore whether a multitask objective can be effective. Finally, we ask whether a multitask objective can provide benefits as an *auxiliary* objective even in the single-task setting (Section 3.3).

3.1 Comparing Pruning Strategies

We establish the relative task performance difference between element-wise and rank pruning. We prune BERT-base to different sparsity levels while finetuning on MNL for each of the parameter combinations described in Section 2.1. We perform this in the single-task setting to compare the possible benefits that can be derived from rank pruning and to find the best configuration for both methods. Figure 1a presents both task performance (averaged F1) and inference speed.³

Task performance These results in element-wise pruning once again corroborate the superiority of global movement pruning found by Sanh et al. (2020b). In contrast, we find that for rank pruning, local magnitude pruning is generally the best strategy. The poor performance of global pruning can be explained by the fact that it is possible to prune some parameters marginally (i.e. not beyond the $\frac{k}{2}$ threshold needed to see improvements) and so some of the budget allocated towards pruning is unused. Furthermore, the magnitudes of the singular values of the feedforward and self-attention layers are not necessarily comparable, as we find that attention weights are pruned more aggressively.

Comparing between the best settings for unstructured (element-wise) and structured (rank) pruning, we find that unstructured pruning retains performance better when the parameter count is low.

Runtime comparison To compare the runtime of element-wise and rank pruning, we make a simplifying assumption that unstructured pruning does not shrink the size of any individual matrix, and so the runtime at any sparsity level would be equal to that of an unpruned model. Under this assumption, we find in Figure 1b that rank pruning offers substantial speedups and is almost immediately (at around 10% sparsity) faster than the dense model.

However, sparsifying a matrix can lead to specialized hardware and algorithmic optimizations as demonstrated by sparse multiplication libraries (Gale et al., 2020). Lagunas et al. (2021) optimize element-wise unstructured pruning in a simple manner by removing entirely pruned rows, columns or attention heads. They show that even at high sparsities (more than 90%), this strategy achieves at

³The runtime comparison was performed on a single NVIDIA RTX 6000.

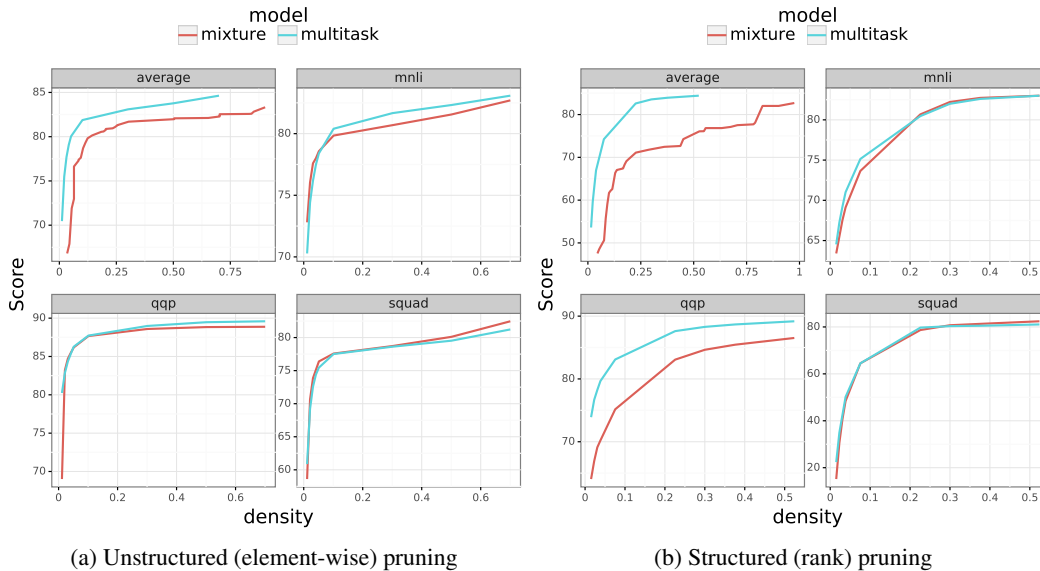


Figure 2: Given a fixed parameter budget (expressed as a fraction of a single model size), we compare a single model pruned with a multitask objective (blue) and the best combination of 3 individual task models for a given size. The red line (“mixture”) is the Pareto frontier of these combinations.

most around a $1.5\times$ speedup. Meanwhile, rank pruning achieves a comparable speedup at just 50% and almost a $2.5\times$ speedup at 90%. While rank pruning has lower F1 under a fixed parameter budget, it is a competitive option given a fixed *latency* budget.

Global rank pruning One limitation of global rank pruning is that the initialization is not normalized across weights. Specifically, we observe that the with *magnitude* as our selection heuristic, attention weights are pruned more aggressively because most of their singular values are smaller than those of the feedforward parameters. This could result in important weights in the attention layers being pruned before less important ones in the feedforward layers, and it could explain why global rank pruning performs poorly at high sparsities (as shown in Figure 1a). The opposite behavior is observed for rank *movement* pruning, where pruning globally is slightly preferred over pruning locally. This suggests that first-order information might be more comparable globally.

Future work can investigate finding a balance between using the interpretability of the magnitudes of the singular values and first-order information from the finetuning process. Alternatively, one could explore a hybrid pruning strategy: with *magnitude* pruning, separate thresholds (or k) for the attention and the feedforward parameters.

3.2 Multitask pruning

With the above baseline established, we conduct experiments in a multitask setting with MNLI, QQP, and SQuAD. Our goal is to answer: given a fixed parameter budget, is it better to train (and prune) three specialized models or one multitask model for all three tasks?

Concretely, we prune models to different sparsity levels using four separate objectives: MNLI, QQP, SQuAD, and multitask. We compare the best mixture of single-task models to a single model pruned with a multitask objective.⁴ The best mixture is determined by the Pareto frontier of all possible ensembles of models. Figure 2 shows that the multitask model outperforms the mixture on a macro-averaged 3-task metric. In addition, it matches or exceeds the performance on individual tasks.

⁴We train for 8 epochs with 2 initial and 2 final warmup epochs. This is close the average of the hyperparameters previously found for the three separate tasks.

Model	Rank	Prune %	MNLI 393K	SQuAD 87.6K	QQP 364K	QNLI 105K	SST-2 67K	MRPC 3.7K	STSB 7K	RTE 2.5K	CoLA 8.5K	Avg.
RP-3	38	7.6%	73.8	64.3	83.1	-	-	-	-	-	-	-
RP-1	38	7.6%	73.6	64.4	75.2	80.0	87.4	75.7	14.4	52.7	0	58.2
BNG	150	29.2%	-	-	-	-	91.3 ± 0.4	87.8 ± 0.6	-	-	38.7 ± 1.6	-
RP-9	38	7.6%	72.1	54.2	81.3	84.7	87.6	83.6	86.3	67.1	23.6	71.2
RP-9	76	15.0%	76.8	70.5	84.5	88.2	89.8	87.5	88.0	70.4	34.6	76.7

Table 2: Individual task performance (dev.) of a model pruned using rank pruning with the multitask objective. The size of the training set for each task is also listed. RP-9 is trained on all nine tasks, RP-3 is the three-task model from Section 3.2, and RP-1 represents 9 separate single-task baseline models. BNG (Ben Noach and Goldberg, 2020) is three separate single-task low-rank models tuned using knowledge distillation.

We find that multitask pruning outperforms the mixture with both element-wise pruning and rank pruning, suggesting that the ability to leverage a multitask objective during pruning may extend to other, novel methods for pruning. This also provides evidence that in encoders, there exist smaller shared subnetworks that can be leveraged across multiple tasks.

3.3 Multitask training as an auxiliary pruning objective

In Section 3.2, we find that in some cases, a multitask pruning objective is helpful even when only one task (e.g. QQP) is of interest. This finding, along with prior work demonstrating the effectiveness of multitask training (Wang et al., 2019; Liu et al., 2019; Aghajanyan et al., 2021), motivates our next question: if we only care about a single task, should we still use a multitask objective?

To test this hypothesis, we expand the three tasks to 9 and prune a model with a 9-task objective, sampling each task uniformly at random. In addition to the three aforementioned tasks, we include CoLA (Warstadt et al., 2019), SST-2 (Socher et al., 2013), MRPC (Dolan and Brockett, 2005), STSB (Cer et al., 2017), and QNLI (Rajpurkar et al., 2016), and RTE (Dagan et al., 2005). For comparison, we also prune a model for each task separately. These additional models are pruned for 10 epochs, keeping other hyperparameters the same as the multitask model.

Table 2 shows the performance on each of the tasks when using local magnitude rank pruning. We find that the multitask models (RP-9) perform well on the smaller datasets of RTE, CoLA, and STSB, outperforming the single-task baseline (RP-1) and coming close larger models from prior work which relied on a hyperparameter search for each task (Ben Noach and Goldberg, 2020). In contrast, we perform no hyperparameter tuning (beyond finding the best combinations in Section 3.1). These results collectively suggest that multitask-based pruning offers a way to effectively prune models for low-resourced tasks without extensive hyperparameter search.

Complementary to our approach of using multitask training is a set of methods in which a small number of additional task-specific parameters are directly embedded in the model, such as prefix embeddings (Keskar et al., 2019; Li and Liang, 2021), adapters (Houlsby et al., 2019), diff parameters (Guo et al., 2021), and LoRa layers (Hu et al., 2021), among others. These parameters and specialized toward a specific task and can cheaply provide additional benefits. Future work can explore incorporating these task-specific modules into (or after) the pruning process.

4 Conclusion

We explore multitask pruning, which aims to perform well on multiple tasks with a given parameter budget smaller than the size of single-task model. Using methods for both unstructured (element-wise) pruning and structured (rank) pruning strategies, we find that pruning with a multitask objective outperforms combining multiple models pruned separately on individual task objectives. Additionally, we find that pruning with a multitask objective (without additional hyperparameter tuning) can help further prune tasks with low-resource datasets. We hope that these lessons learned guide future work in compressing multitask models.

Acknowledgements

We would like to thank Jason Eisner, Hao Fang, Edward Hu, Anthony Platanios, Yu Su, Sam Thomson, and Benjamin Van Durme for helpful discussions and feedback for this project.

References

- Armen Aghajanyan, Ancht Gupta, Akshat Shrivastava, Xilun Chen, Luke Zettlemoyer, and Sonal Gupta. 2021. Muppet: Massive multi-task representations with pre-finetuning.
- Matan Ben Noach and Yoav Goldberg. 2020. Compressing pre-trained language models by matrix decomposition. In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 884–889, Suzhou, China. Association for Computational Linguistics.
- Yoshua Bengio, Nicholas Léonard, and Aaron C. Courville. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *ArXiv*, abs/1308.3432.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation. In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Ido Dagan, Oren Glickman, and Bernardo Magnini. 2005. The pascal recognising textual entailment challenge. In *Machine Learning Challenges Workshop*, pages 177–190. Springer.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. BERT: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- William B Dolan and Chris Brockett. 2005. Automatically constructing a corpus of sentential paraphrases. In *Proceedings of the Third International Workshop on Paraphrasing (IWP2005)*.
- Chun Fan, Jiwei Li, Xiang Ao, Fei Wu, Yuxian Meng, and Xiaofei Sun. 2021. Layer-wise model pruning based on mutual information.
- Trevor Gale, Matei Zaharia, Cliff Young, and Erich Elsen. 2020. Sparse gpu kernels for deep learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '20*. IEEE Press.
- Demi Guo, Alexander Rush, and Yoon Kim. 2021. Parameter-efficient transfer learning with diff pruning. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4884–4896, Online. Association for Computational Linguistics.
- Song Han, Jeff Pool, John Tran, and William Dally. 2015. Learning both weights and connections for efficient neural network. In *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for NLP. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 2790–2799. PMLR.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. Lora: Low-rank adaptation of large language models. *ArXiv*, abs/2106.09685.
- Shankar Iyer, Nikhil Dandekar, and Kornel Csernai. 2017. First quora dataset release: Question pairs.

- 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 *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4163–4174, Online. Association for Computational Linguistics.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, Mostafa Dehghani, and James Henderson. 2021a. Parameter-efficient multi-task fine-tuning for transformers via shared hypernetworks. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 565–576, Online. Association for Computational Linguistics.
- Rabeeh Karimi Mahabadi, Sebastian Ruder, and James Henderson. 2021b. Compacter: Efficient low-rank hypercomplex adapter layers.
- Nitish Shirish Keskar, Bryan McCann, Lav Varshney, Caiming Xiong, and Richard Socher. 2019. CTRL - A Conditional Transformer Language Model for Controllable Generation. *arXiv preprint arXiv:1909.05858*.
- François Lagunas, Ella Charlaix, Victor Sanh, and Alexander M. Rush. 2021. Block pruning for faster transformers.
- Yann LeCun, John Denker, and Sara Solla. 1990. Optimal brain damage. In *Advances in Neural Information Processing Systems*, volume 2. Morgan-Kaufmann.
- Xiang Lisa Li and Percy Liang. 2021. Prefix-tuning: Optimizing continuous prompts for generation. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- 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. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6524–6538, Online. Association for Computational Linguistics.
- Xiaodong Liu, Pengcheng He, Weizhu Chen, and Jianfeng Gao. 2019. Multi-task deep neural networks for natural language understanding. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4487–4496, Florence, Italy. Association for Computational Linguistics.
- 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. *Journal of Machine Learning Research*, 21(140):1–67.
- Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. 2016. SQuAD: 100,000+ questions for machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 2383–2392, Austin, Texas. Association for Computational Linguistics.
- Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2020a. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter.
- Victor Sanh, Thomas Wolf, and Alexander M. Rush. 2020b. Movement pruning: Adaptive sparsity by fine-tuning. In *NeurIPS*.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online. Association for Computational Linguistics.

- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D. Manning, Andrew Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pages 1631–1642, Seattle, Washington, USA. 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 *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2158–2170, Online. Association for Computational Linguistics.
- Alex Wang, Jan Huala, Patrick Xia, Raghavendra Pappagari, R. Thomas McCoy, Roma Patel, Najoung Kim, Ian Tenney, Yinghui Huang, Katherin Yu, Shuning Jin, Berlin Chen, Benjamin Van Durme, Edouard Grave, Ellie Pavlick, and Samuel R. Bowman. 2019. Can you tell me how to get past sesame street? sentence-level pretraining beyond language modeling. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4465–4476, Florence, Italy. Association for Computational Linguistics.
- Alex Warstadt, Amanpreet Singh, and Samuel R. Bowman. 2019. Neural network acceptability judgments. *Transactions of the Association for Computational Linguistics*, 7:625–641.
- Adina Williams, Nikita Nangia, and Samuel Bowman. 2018. A broad-coverage challenge corpus for sentence understanding through inference. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1112–1122, New Orleans, Louisiana. Association for Computational Linguistics.
- Huanrui Yang, Minxue Tang, Wei Wen, Feng Yan, Daniel Hu, Ang Li, Hai Li, and Yiran Chen. 2020. Learning low-rank deep neural networks via singular vector orthogonality regularization and singular value sparsification. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops*.
- Michael H. Zhu and Suyog Gupta. 2018. To prune, or not to prune: Exploring the efficacy of pruning for model compression.