

Improving Generalization of Pre-trained Language Models via Stochastic Weight Averaging

Peng Lu^{1,3}, Ivan kobyzev¹, Mehdi Rezagholizadeh¹, Ahmad Rashid^{1,3},
Ali Ghodsi², Philippe Langlais³

¹Huawei Noah's Ark Lab

² Department of Statistics and Actuarial Science, University of Waterloo

³ RALI/DIRO, Université de Montréal, Canada

{peng.lu1,ivan.kobyzev,mehdi.rezagholizadeh,ahmad.rashid}@huawei.com
ali.ghodsi@uwaterloo.ca, felipe@iro.umontreal.ca

Abstract

Knowledge Distillation (KD) is a commonly used technique for improving the generalization of compact Pre-trained Language Models (PLMs) on downstream tasks. However, such methods impose the additional burden of training a separate teacher model for every new dataset. Alternatively, one may directly work on the improvement of the optimization procedure of the compact model towards better generalization. Recent works observe that the flatness of the local minimum correlates well with better generalization. In this work, we adapt Stochastic Weight Averaging (SWA), a method encouraging convergence to a flatter minimum, to fine-tuning PLMs. We conduct extensive experiments on various NLP tasks (text classification, question answering, and generation) and different model architectures and demonstrate that our adaptation improves the generalization without extra computation cost. Moreover, we observe that this simple optimization technique is able to outperform the state-of-the-art KD methods for compact models.

1 Introduction

Large Pre-trained Language Models (PLMs) like BERT (Devlin et al., 2019) and GPT (Brown et al., 2020) have achieved state-of-the-art (SOTA) performance on a number of NLP applications. A large model pre-trained on a massive text corpus learns general-purpose features of the language and can be easily adapted to different downstream tasks by fine-tuning with task-specific objectives (Devlin et al., 2019; Liu et al., 2019; Kaplan et al., 2020). Multiple papers (Brown et al., 2020; Zeng et al., 2021; Wang et al., 2022) demonstrated that the improvement in pre-training can be transferred favorably to most downstream tasks with little or no data, which leads to considerable efforts to scale up PLMs and increase the size of the pre-training dataset. However, a recent work (Abnar et al.,

2021) empirically shows that as the upstream accuracy increases, the performance of downstream tasks gradually saturates. Furthermore, although PLMs perform well in a task-agnostic few-shot setting, it is still desirable to fine-tune them to achieve SOTA performance, especially for models with limited capacity. This encourages more efforts to improve the generalization of PLMs through additional techniques like Knowledge Distillation or better optimization.

Knowledge Distillation (Hinton et al., 2015) is an effective method to improve the generalization of compact PLMs. One can use a larger model (a teacher) to fine-tune a smaller model (a student) by making a student learn the teacher's output. This technique has helped in achieving SOTA performance on multiple Natural Language Understanding (NLU) problems for compact models (Sun et al., 2020; Jiao et al., 2020; Sanh et al., 2019). However, KD is not an efficient method as it requires training a separate teacher model for every new task. Alternatively, some works focus on boosting the performance of the optimization algorithm directly. Several directions have been explored in the literature: adding noise to the gradient descent process (Kleinberg et al., 2018; Orvieto et al., 2022), applying continuation optimization (Gulcehre et al., 2017), steering optimization towards flat optima (Foret et al., 2021; Bisla et al., 2022), averaging checkpoints (Izmailov et al., 2018; Wortsman et al., 2022). Encouraged by the success of Stochastic Weight Averaging (SWA) (Izmailov et al., 2018) for computer vision and Sharpness-Aware Minimization (SAM) for NLP (Bahri et al., 2022), we investigate whether SWA can find better, generalizable minima for fine-tuning PLMs while maintaining training efficiency. We conduct extensive experiments for PLM fine-tuning including seven NLU tasks, two question answering tasks, and one summarization task with different model architectures. Our contributions are

as follows: We adapt SWA to NLP and demonstrate that our adaptation can be applied to fine-tuning PLMs without extra computational cost on both classification and generation tasks. Furthermore, we outline several practical recipes that are different from the ones in computer vision including using a high constant learning rate in SWA (instead of a cyclical one) and averaging over steps instead of epochs.

2 Related Work

KD for PLMs Knowledge Distillation (KD) is a commonly used technique for improving the generalization of a compact PLM (Sun et al., 2020; Jiao et al., 2020; Sanh et al., 2019). KD can be thought of as an instance-specific label regularization for the training process (Yuan et al., 2020; Zhang and Sabuncu, 2020). However, KD is not an efficient technique because it adds the burden of fine-tuning an extra teacher model for every new task. Moreover, a recent work (Kobyzev et al., 2022) shows that without additional techniques and resources like data augmentation, re-weighting or intermediate layer distillation (Wu et al., 2021; Rashid et al., 2021; Lu et al., 2021; Sun et al., 2019) pure label regularization cannot outperform base fine-tuning for PLMs.

Geometry of the Loss Landscape Loss landscape analysis is a popular method for investigating the model convergence problem (Li et al., 2018; Hao et al., 2019; Fort et al., 2019). It has been widely studied that neural networks tend to generalize better when they converge to flatter local minima (Hochreiter and Schmidhuber, 1994; Keskar et al., 2017; Dosovitskiy et al., 2021). Some works inspect the curvature and flatness of the loss surface via Hessian spectrum analysis (Ghorbani et al., 2019; Sagun et al., 2018) and show the positive correlation between local sharpness and the generalization gap. Such observations have inspired research on improving generalization from a geometric perspective (Foret et al., 2021; Kwon et al., 2021; Bisla et al., 2022). Another recently proposed method, Stochastic Weight Averaging (SWA) (Izmailov et al., 2018), also improves generalization and is easy to implement without significant computational overhead. The original paper explains the better generalization by empirically showing the flatness of the minimum discovered by SWA. However, all the analysis was done on only computer vision tasks and models. To the best of our

Algorithm 1: Stochastic Weight Averaging

```

input : Scheduler  $\mathcal{L}_r(i)$ , interval length  $K$ , number
        of iterations  $N$ , weights  $W$ 
output :  $W_{\text{SWA}}$ 
1  $W \leftarrow \hat{W}$ ;
2  $W_{\text{SWA}} \leftarrow W$ ;
3 for  $i \leftarrow 1$  to  $N$  do
4    $\eta = \mathcal{L}_r(i)$ ; // step size
5    $W \leftarrow W - \eta \nabla \mathcal{L}_i(W)$ ;
6   if  $\text{Mod}(i, K) = 0$  then
7      $n_{\text{model}} \leftarrow \frac{i}{K}$ ; // # of models
8      $W_{\text{SWA}} \leftarrow \frac{1}{n_{\text{model}}+1} (W_{\text{SWA}} \times n_{\text{model}} + W)$ ;
9   end
10 end
11 return  $W_{\text{SWA}}$ 

```

knowledge, our paper is the first work adapting SWA for fine-tuning PLMs.

Flatness for PLMs A recent work (Bahri et al., 2022) applies Sharpness-Aware Minimization (SAM) (Foret et al., 2021) to PLM fine-tuning on several NLP tasks. It shows that the flatness-aware optimizer can substantially improve the generalization of language models, but it takes $1.25\times$ computational cost.

3 Revisiting SWA

We start by briefly reminding the Stochastic Weight Averaging (Izmailov et al., 2018), a method which implements efficient weight ensembling during model optimization. First, a neural network is trained with SGD till it converges to a region close to a local minimum of the training loss landscape. Then, in the second stage, SWA averages the checkpoints along the SGD trajectory using a designated learning rate schedule for better exploration of the parameter space. The original paper considers Cyclical and High Constant learning rate schedules (CHC). During the update, SWA collects model parameters for every K step and averages them for the model ensemble. Cha et al. (2021) empirically confirms that SWA can avoid convergence to a sharp local minimum. The training process is shown in Algorithm 1.

CHC learning rate schedule In every cycle, the learning rate anneals from η_{\max} to η_{\min} . The learning rate for each training iteration i is given by:

$$\begin{aligned}
 t_i &= \text{mod}(i-1, K) + 1, \\
 \eta_i &= \left(1 - \frac{t_i}{K}\right)\eta_{\max} + \frac{t_i}{K}\eta_{\min},
 \end{aligned}$$

where K is the cycle length and $\eta_{\max} > \eta_{\min}$ are the maximum and minimum learning rate, respectively.

Table 1: Performance on GLUE test sets. For reference we give several knowledge distillation techniques for comparison.

Method	CoLA	MRPC	RTE	SST-2	QNLI	QQP	MNLI	Avg.
DistilRoBERTa	53.3	87.6	72.2	93.3	90.4	88.9	83.5	81.3
+KD	54.3	86	74.1	93.1	91.1	89.5	83.6	81.7
+TAKD	53.2	86.7	74.2	93.2	91	89.4	83.8	81.6
+AKD	54	88	73.7	93.6	90.8	89.7	83.8	81.9
+SAM	47.5	88.8	73.6	94.1	89.6	89.1	83.8	80.9
+SWA	53.7	89.5	73.6	93.7	91.4	89.3	84.3	82.2

When $\eta_{\max} = \eta_{\min}$, the learning rate becomes constant.

SWA for PLMs We accommodate the SWA method for PLM fine-tuning. We observe that several modifications are required compared to SWA for computer vision. First, we noticed that a high constant learning rate performs better than a cyclical learning rate. Then, since the fine-tuning needs much less epochs than training from scratch, we average over steps instead of epochs as in the original SWA for computer vision. Besides that, for most of the NLU tasks we average a larger number of components and start averaging (stage 2 of SWA) earlier: after 50% of the training instead of 75% in the original method. The latter could be because the PLM is closer to the local minimum than a randomly initialized model and hence allows more exploration.

4 Experiments

In this section, we introduce the setup of our experiments, i.e., the benchmarks and corresponding evaluation methodology and discuss experimental results. All experiments were performed on single NVIDIA Tesla (V100) GPU. We also give an example of flatness analysis in the Appendix A.

4.1 Text Classification

Data. We conduct experiments on 7 classification tasks from the GLUE benchmark (Wang et al., 2019). These datasets can be broadly divided into 3 families of problems. Single sentence tasks which include linguistic acceptability (CoLA) and sentiment analysis (SST-2), pairwise classification which include Similarity (MRPC) and paraphrasing (QQP) and inference tasks which include Natural Language Inference (MNLI and RTE) and Question Answering NLI (QNLI). Following prior works (Devlin et al., 2019; Liu et al., 2019), we

report Matthews correlation for CoLA, F1 for QQP and MRPC and accuracy for the rest.

Experimental Setup. We choose the pre-trained DistilRoBERTa (6-layer) (Sanh et al., 2019) as the base model for our experiments. We use a RoBERTa-Large model (24-layer) (Liu et al., 2019) fine-tuned on each of the tasks as the corresponding teacher model in the KD baseline. All models are trained with the AdamW optimizer with the default settings (Loshchilov and Hutter, 2019) for 30 epochs¹. To explore the best hyper-parameters, we used a grid search over the learning rate $\in \{1e-5, 2e-5, 3e-5\}$, batch size $\in \{8, 16, 32, 64\}$, KD interpolation (α) $\in \{0.1, 0.2, 0.5\}$, KD temperature $\in \{1, 5, 10\}$ and the hyper-parameter $\rho \in \{0.05, 0.15\}$ for SAM.

We compare SWA with SAM (Bahri et al., 2022) and three KD techniques: vanilla KD, Annealing KD (AKD) (Jafari et al., 2021), and TAKD (Mirzadeh et al., 2020).

Results Table 1 shows the GLUE leaderboard test results. We observe that SWA outperforms fine-tuning on all datasets, which indicates its general applicability and, on average, SWA performance is comparable to SOTA KD methods.

Table 2: Performance on SQuAD (v1 and v2) datasets.

Model	SQuAD_v1		SQuAD_v2	
	EM	F1	EM	F1
DistilRoBERTa	80.5	87.8	71.8	75.2
+SWA	81.1	88.2	72.4	75.7
RoBERTa-base	85.9	92.0	79.8	83.2
+SWA	86.1	92.2	80.1	83.4

¹Our experiments were conducted with huggingface toolkit (<https://github.com/huggingface/transformers>).

4.2 Question Answering

We conduct experiments on two Question Answering benchmarks, SQuAD v1 (Rajpurkar et al., 2016) and SQuAD v2 (Rajpurkar et al., 2018) and report the exact match and F1-score on development sets. We choose the pre-trained DistilRoBERTa (6-layer) (Sanh et al., 2019) and RoBERTa (Liu et al., 2019) as the backbone models for our QA experiments.

Table 2 shows the validation results on the two SQuAD datasets. We applied SWA to RoBERTa model with different sizes (6 and 12 layers). We observed that SWA can consistently outperform regular fine-tuning method on both versions of SQuAD and the results show that SWA provides larger improvement for a small model.

4.3 Summarization

Table 3: Performance on XSum datasets.

Model	ROUGE		
	1	2	L
T5-small	31.33	9.26	24.19
+ SWA	32.01	9.79	24.75

We conduct experiments on XSum (Narayan et al., 2018) dataset with T5-small (Raffel et al., 2020) model. The models are evaluated based on ROUGE scores (Lin, 2004). Specifically, we report ROUGE-1, ROUGE-2, and ROUGE-L, which compute the uniform, bi-gram, and longest common sequence overlap with the reference summaries.

Table 3 shows our results. We observe that SWA improves the overall scores for T5 model over fine-tuning.

5 Analysis

5.1 Constant LR vs Cyclical LR

We compare the effect of constant and cyclical learning rates on the SST-2 and MRPC datasets from GLUE on Figure 1. We can observe that a high constant learning rate performs much better than the cyclical one. We conducted an extensive grid search for $(\eta_{\max}, \eta_{\min}) \in \{(2e-5, 1e-6), (1e-5, 1e-6), (5e-6, 1e-6), (3e-6, 1e-6)\}$ and $\eta_c \in \{2e-6, 3e-6\}$ and report the averaged accuracy and corresponding standard derivations of each method. The patterns show the superior performance of high constant learning rate.




Figure 1: Accuracy of DistilRoBERTa models for SST-2 and MRPC with cyclical and high constant learning rates.

5.2 Time Analysis of SWA.

Table 4: Time analysis for DistilRoBERTa.

Method	relative time
Finetune	1×
KD	2.85×
AKD	2.51×
SAM	2.11×
SWA	1.08×

Table 4 shows the relative time analysis for DistilRoBERTa model with different methods. It is worth mentioning that, unlike KD methods, SWA doesn't need a trained teacher model, and, unlike SAM, it doesn't require double forward and backward pass computations.

6 Conclusion

In this work we have adapted the SWA algorithm to the fine-tuning of PLMs. SWA is from a family of algorithms that aim to improve generalization for deep neural networks by finding a flatter, as opposed to a sharp, minimum. We present the recipe that can adapt SWA to fine-tuning PLMs, i.e. a high constant learning rate, averaging over training steps and choosing more components for discrimination and less for generation.

On compact PLMs, our adaptation of SWA performs on par with SOTA KD methods on the GLUE benchmark, and improves on fine-tuning on question answering and text summarization. Moreover, unlike KD or SAM (another flatness seeking optimization algorithm), SWA introduces no additional computation cost. As future work, we want to explore apply to the pre-training of LMs.

Acknowledgments

We thank Mindspore,² which is a new deep learning computing framework, for partial support of this work.

Limitations

In the current work, we adapt a weigh assembling method, SWA, to PLMs fine-tuning and conduct extensive experiments to show that it can improve the generalization ability of PLMs. However, due to the limited computation resources, we don't explore the SWA for PLMs with a larger model size.

References

- Samira Abnar, Mostafa Dehghani, Behnam Neyshabur, and Hanie Sedghi. 2021. [Exploring the limits of large scale pre-training](#). *CoRR*, abs/2110.02095.
- Dara Bahri, Hossein Mobahi, and Yi Tay. 2022. [Sharpness-aware minimization improves language model generalization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), ACL 2022, Dublin, Ireland, May 22-27, 2022*, pages 7360–7371. Association for Computational Linguistics.
- Devansh Bisla, Jing Wang, and Anna Choromanska. 2022. [Low-pass filtering SGD for recovering flat optima in the deep learning optimization landscape](#). In *International Conference on Artificial Intelligence and Statistics, AISTATS 2022, 28-30 March 2022, Virtual Event*, volume 151 of *Proceedings of Machine Learning Research*, pages 8299–8339. PMLR.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.
- Junbum Cha, Sanghyuk Chun, Kyungjae Lee, Han-Cheol Cho, Seunghyun Park, Yunsung Lee, and Sungrae Park. 2021. [SWAD: Domain generalization by seeking flat minima](#). In *Advances in Neural Information Processing Systems*.
- 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, NAACL-HLT 2019, Minneapolis, MN, USA, June 2-7, 2019, Volume 1 (Long and Short Papers)*, pages 4171–4186. Association for Computational Linguistics.
- Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. 2021. [An image is worth 16x16 words: Transformers for image recognition at scale](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Pierre Foret, Ariel Kleiner, Hossein Mobahi, and Behnam Neyshabur. 2021. [Sharpness-aware minimization for efficiently improving generalization](#). In *9th International Conference on Learning Representations, ICLR 2021, Virtual Event, Austria, May 3-7, 2021*. OpenReview.net.
- Stanislav Fort, Huiyi Hu, and Balaji Lakshminarayanan. 2019. [Deep ensembles: A loss landscape perspective](#). *CoRR*, abs/1912.02757.
- Behrooz Ghorbani, Shankar Krishnan, and Ying Xiao. 2019. [An investigation into neural net optimization via hessian eigenvalue density](#). In *Proceedings of the 36th International Conference on Machine Learning, ICML 2019, 9-15 June 2019, Long Beach, California, USA*, volume 97 of *Proceedings of Machine Learning Research*, pages 2232–2241. PMLR.
- Caglar Gulcehre, Marcin Moczulski, Francesco Visin, and Yoshua Bengio. 2017. [Mollifying networks](#). In *5th International Conference on Learning Representations*.
- Yaru Hao, Li Dong, Furu Wei, and Ke Xu. 2019. [Visualizing and understanding the effectiveness of BERT](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing, EMNLP-IJCNLP 2019, Hong Kong, China, November 3-7, 2019*, pages 4141–4150. Association for Computational Linguistics.
- Geoffrey E. Hinton, Oriol Vinyals, and Jeffrey Dean. 2015. [Distilling the knowledge in a neural network](#). *CoRR*, abs/1503.02531.
- Sepp Hochreiter and Jürgen Schmidhuber. 1994. [Simplifying neural nets by discovering flat minima](#). In *Advances in Neural Information Processing Systems*, volume 7. MIT Press.
- Pavel Izmailov, Dmitrii Podoprikin, Timur Garipov, Dmitry P. Vetrov, and Andrew Gordon Wilson. 2018.

²<https://www.mindspore.cn/>

- Averaging weights leads to wider optima and better generalization. In *Proceedings of the Thirty-Fourth Conference on Uncertainty in Artificial Intelligence, UAI 2018, Monterey, California, USA, August 6-10, 2018*, pages 876–885. AUAI Press.
- Aref Jafari, Mehdi Rezagholizadeh, Pranav Sharma, and Ali Ghodsi. 2021. [Annealing knowledge distillation](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume, EACL 2021, Online, April 19 - 23, 2021*, pages 2493–2504. Association for Computational Linguistics.
- 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, Online Event, 16-20 November 2020*, volume EMNLP 2020 of *Findings of ACL*, pages 4163–4174. Association for Computational Linguistics.
- Jared Kaplan, Sam McCandlish, Tom Henighan, Tom B. Brown, Benjamin Chess, Rewon Child, Scott Gray, Alec Radford, Jeffrey Wu, and Dario Amodei. 2020. [Scaling laws for neural language models](#). *CoRR*, abs/2001.08361.
- Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. 2017. [On large-batch training for deep learning: Generalization gap and sharp minima](#). In *5th International Conference on Learning Representations, ICLR 2017, Toulon, France, April 24-26, 2017, Conference Track Proceedings*. OpenReview.net.
- Robert D. Kleinberg, Yuanzhi Li, and Yang Yuan. 2018. An alternative view: When does sgd escape local minima? In *ICML*.
- Ivan Kobyzev, Aref Jafari, Mehdi Rezagholizadeh, Tianda Li, Alan Do-Omri, Peng Lu, Ali Ghodsi, and Pascal Poupart. 2022. Towards understanding label regularization for fine-tuning pre-trained language models. *arXiv preprint arXiv:2205.12428*.
- Jungmin Kwon, Jeongseop Kim, Hyunseo Park, and In Kwon Choi. 2021. [ASAM: adaptive sharpness-aware minimization for scale-invariant learning of deep neural networks](#). In *Proceedings of the 38th International Conference on Machine Learning, ICML 2021, 18-24 July 2021, Virtual Event*, volume 139 of *Proceedings of Machine Learning Research*, pages 5905–5914. PMLR.
- Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. 2018. [Visualizing the loss landscape of neural nets](#). In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems 2018, NeurIPS 2018, December 3-8, 2018, Montréal, Canada*, pages 6391–6401.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. 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](#). *CoRR*, abs/1907.11692.
- Ilya Loshchilov and Frank Hutter. 2019. [Decoupled weight decay regularization](#). In *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Peng Lu, Abbas Ghaddar, Ahmad Rashid, Mehdi Rezagholizadeh, Ali Ghodsi, and Philippe Langlais. 2021. [RW-KD: Sample-wise loss terms re-weighting for knowledge distillation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3145–3152, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Seyed Iman Mirzadeh, Mehrdad Farajtabar, Ang Li, Nir Levine, Akihiro Matsukawa, and Hassan Ghasemzadeh. 2020. Improved knowledge distillation via teacher assistant. In *AAAI*.
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing, Brussels, Belgium, October 31 - November 4, 2018*, pages 1797–1807. Association for Computational Linguistics.
- Antonio Orvieto, Hans Kersting, Frank Proske, Francis Bach, and Aurelien Lucchi. 2022. [Anticorrelated noise injection for improved generalization](#). *Computing Research Repository*, arXiv:2202.02831. Version 1.
- 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.
- Pranav Rajpurkar, Robin Jia, and Percy Liang. 2018. [Know what you don’t know: Unanswerable questions for squad](#). In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics, ACL 2018, Melbourne, Australia, July 15-20, 2018, Volume 2: Short Papers*, pages 784–789. Association for Computational Linguistics.
- 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, EMNLP 2016, Austin, Texas, USA, November 1-4, 2016*, pages 2383–2392. The Association for Computational Linguistics.

- Ahmad Rashid, Vasileios Lioutas, and Mehdi Rezagholizadeh. 2021. Mate-kd: Masked adversarial text, a companion to knowledge distillation. 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 1062–1071.
- Levent Sagun, Utku Evci, V. Ugur Güney, Yann N. Dauphin, and Léon Bottou. 2018. Empirical analysis of the hessian of over-parametrized neural networks. In *6th International Conference on Learning Representations, ICLR 2018, Vancouver, BC, Canada, April 30 - May 3, 2018, Workshop Track Proceedings*. OpenReview.net.
- 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*.
- 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 and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 4323–4332, Hong Kong, China. 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, ACL 2020, Online, July 5-10, 2020*, pages 2158–2170. Association for Computational Linguistics.
- 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 *7th International Conference on Learning Representations, ICLR 2019, New Orleans, LA, USA, May 6-9, 2019*. OpenReview.net.
- Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, and Furu Wei. 2022. Deepnet: Scaling transformers to 1, 000 layers. *CoRR*, abs/2203.00555.
- Mitchell Wortsman, Gabriel Ilharco, Samir Yitzhak Gadre, Rebecca Roelofs, Raphael Gontijo-Lopes, Ari S. Morcos, Hongseok Namkoong, Ali Farhadi, Yair Carmon, Simon Kornblith, and Ludwig Schmidt. 2022. Model soups: averaging weights of multiple fine-tuned models improves accuracy without increasing inference time. *Computing Research Repository*, arXiv:2203.05482. Version 2.
- Yimeng Wu, Mehdi Rezagholizadeh, Abbas Ghaddar, Md. Akmal Haidar, and Ali Ghodsi. 2021. Universal-kd: Attention-based output-grounded intermediate layer knowledge distillation. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing, EMNLP 2021, Virtual Event / Punta Cana, Dominican Republic, 7-11 November, 2021*, pages 7649–7661. Association for Computational Linguistics.
- Li Yuan, Francis E. H. Tay, Guilin Li, Tao Wang, and Jiashi Feng. 2020. Revisiting knowledge distillation via label smoothing regularization. In *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition, CVPR 2020, Seattle, WA, USA, June 13-19, 2020*, pages 3902–3910. Computer Vision Foundation / IEEE.
- Wei Zeng, Xiaozhe Ren, Teng Su, Hui Wang, Yi Liao, Zhiwei Wang, Xin Jiang, ZhenZhang Yang, Kaisheng Wang, Xiaoda Zhang, Chen Li, Ziyang Gong, Yifan Yao, Xinjing Huang, Jun Wang, Jianfeng Yu, Qi Guo, Yue Yu, Yan Zhang, Jin Wang, Hengtao Tao, Dasen Yan, Zexuan Yi, Fang Peng, Fangqing Jiang, Han Zhang, Lingfeng Deng, Yehong Zhang, Zhe Lin, Chao Zhang, Shaojie Zhang, Mingyue Guo, Shanzhi Gu, Gaojun Fan, Yaowei Wang, Xuefeng Jin, Qun Liu, and Yonghong Tian. 2021. Pangu- α : Large-scale autoregressive pretrained chinese language models with auto-parallel computation. *CoRR*, abs/2104.12369.
- Zhilu Zhang and Mert R. Sabuncu. 2020. Self-distillation as instance-specific label smoothing. In *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*.

A Flatness analysis

To compare the flatness of fine-tuning and SWA, we compute the Hessian matrix H for two minima found by fine-tuning and SWA on MRPC training dataset. We consider two widely used measures of flatness: the largest eigenvalue of the Hessian matrix $\lambda_{\max}(H)$ and the trace of Hessian. The results are shown in Table 5. It is clear that the minimum found by SWA has smaller $\lambda_{\max}(H)$ and $\text{trace}(H)$, thus being flatter.

Table 5: Flatness of fine-tune and SWA on MRPC. To compute the Hessian, we exclude the embedding layers of DistilRoBERTa including word and position embedding, etc.

Method	$\lambda_{\max}(H)$	$\text{trace}(H)$
Fine-tune	60.22	76.4
SWA	1.19	4.72