TuneComp: Joint Fine-Tuning and Compression for Large Foundation Models

Chen, Xiangyu; Liu, Jing; Wang, Ye; Brand, Matthew; Wang, Pu; Koike-Akino, Toshiaki TR2025-079 June 07, 2025

Abstract

o reduce model size during post-training, compression methods, including knowledge distillation, low-rank approximation, and pruning, are often applied after fine- tuning the model. However, sequential fine-tuning and compression sacrifices performance, while creating a larger than necessary model as an intermediate step. In this work, we aim to reduce this gap, by directly constructing a smaller model while guided by the downstream task. We propose to jointly fine-tune and compress the model by gradually distilling it to a pruned low-rank structure. Experiments demonstrate that joint fine-tuning and compression significantly outperforms other sequential compression methods.

IEEE Conference on Computer Vision and Pattern Recognition (CVPR) workshop on Efficient and On-Device Generation 2025

© 2025 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Mitsubishi Electric Research Laboratories, Inc. 201 Broadway, Cambridge, Massachusetts 02139

TuneComp: Joint Fine-tuning and Compression for Large Foundation Models

Xiangyu Chen, Jing Liu, Ye Wang, Matthew Brand, Pu (Perry) Wang, Toshiaki Koike-Akino Mitsubishi Electric Research Laboratories (MERL), Cambridge, MA 02139, USA

{xiachen, jiliu, yewang, brand, pwang, koike}@merl.com

Abstract

To reduce model size during post-training, compression methods, including knowledge distillation, low-rank approximation, and pruning, are often applied after finetuning the model. However, sequential fine-tuning and compression sacrifices performance, while creating a larger than necessary model as an intermediate step. In this work, we aim to reduce this gap, by directly constructing a smaller model while guided by the downstream task. We propose to jointly fine-tune and compress the model by gradually distilling it to a pruned low-rank structure. Experiments demonstrate that joint fine-tuning and compression significantly outperforms other sequential compression methods.

1. Introduction

Large transformer-based models have demonstrated superior performance in many tasks, including natural language processing and computer vision. However, their enormous size requires substantial computational power and memory for fine-tuning and deployment to effectively adapt to downstream tasks. To tackle these, Parameter-Efficient Fine-Tuning (PEFT) was widely explored to reduce memory consumption in fine-tuning. In deployment, various post-training compression techniques have been employed in these fine-tuned models, including weight quantization, knowledge distillation, and network pruning. However, sequentially fine-tuning and compression often suffers from significant performance loss.

Joint fine-tuning and distillation has emerged as a promising alternative [18]. This approach was first proposed in computer vision tasks to jointly fine-tune and prune convolutional neural networks iteratively [35]. Recently, PC-LoRA [18] uses a decay factor to reduce the influence of pretrained teacher branch smoothly, discarding the switching between fine-tuning and compression. The inclusion of fine-tuning in the compression process leads to much higher compression ratios, while maintaining comparable accuracy. However, the potential of the progressive compression pipeline is under-explored. The structure of the



Figure 1. The proposed joint fine-tuning and compression pipeline, where compression involves low-rank approximation as well as the pruning/sparsification of the low-rank structures.



Figure 2. Comparison between our proposed TuneComp, jointly fine-tuning and compression, with other compression strategies.

student branch with a basic low-rank decomposition can be further compressed with pruning techniques. Further, the performance can be enhanced through more careful design of the transition from the teacher to the student branch.

Extending the progressive compression pipeline, our work proposes **TuneComp** as shown in Fig. 1. In this framework, each linear projection layer is split into two parallel branches. The first branch is the frozen weights of the pretrained model, while the second is the trainable

and compressed student branch based on low-rank approximation. During fine-tuning, the student branch gradually replaces the teacher branch, which reduces its influence on the output to zero, while maintaining the total power from both branches equal to 1. We also employ novel extensions of activation-aware initialization. To further compress the model, pruning is jointly applied on top of this low-rank decomposition. The entire pipeline integrates fine-tuning, knowledge distillation, low-rank decomposition, and pruning, and directly producing a compact and efficient model.

2. Related Work

2.1. Joint fine-tuning and compression

From fine-tuning to compression, there are multiple pipelines to select [5, 8, 29]. One widely used is fine-tuning downstream tasks first and then distilling to small models for deployment. In this path, researches either focus on improving fine-tuning to closing the gap between source domain of pretrained dataset and the new domain from downstream tasks, hence making the fine-tuned model generalize better on the given the tasks. For example, a notable line of work is Parameter-Efficient Fine-Tuning (PEFT), such as LoRA [16], LoHA [42], LoKr [42], KronA [12], SuperLoRA [6], LoDA [27] and CorDA [40]. Another line aims to compress the post-training model to achieve similar performance to the original model while significantly reducing model size. Work in this line includes pruning [1, 13, 19, 23, 30], parameter efficient structures [7, 33, 37, 44], knowledge distillation [22], quantization [2, 3, 10, 26, 32, 43] and combined [17, 20, 25]. Others compress the model first and fine-tune the small model for the target tasks [4, 9, 15, 28, 34, 36, 39, 41]. However, these sequential pipelines from whichever direction, finetune then compress or compress then fine-tune, introduced an intermediate agent, either the large fine-tuned model or the small compressed model, restricting the interplay between fine-tuning and compression, and often with significantly degraded performance. Therefore, some other works tried to combine finetuning and compression. However, this line of work is still underexplored. [35] proposes to finetune and compress the model alternatively, which is timeconsuming and suboptimal. PC-LoRA [18] eliminated the alternative switching between teacher and student model with a decay factor that gradually decreasing with time, thus gradually reducing the dependence on teacher model.

2.2. Low-Rank initialization

Initialization of low-rank approximation, a parameter efficient structure $\mathbf{W} \approx \mathbf{BA}$, has been studied in many scenarios. As adapters in parameter-efficient fine-tuning, LoRA [16] proposed either \mathbf{A} or \mathbf{B} should be initialized as zero matrix, while the other one to be random Gaussian matrix.

While [14] argues that Gaussian initialization for the A matrix can lead to better performance by introducing more instability. Weight-aware initialization is also studied. PiSSA [31] instead initialized them with SVD decomposition. Nystrom initialization [24] combined both the pretrained matrix W_0 with Gaussian initialization. CorDA [40] further integrated covariance of activations, leading to activationaware initialization. However, all these methods target in the adapter-initialization in fine-tuning. For compression, ASVD [44] also claimed that task information should be integrated in the initialization to emphasize different features given different tasks and hence proposed general activationaware initialization based on the statistics of activations. However, low-rank initialization in joint fine-tuning and compression lacks investigation, given PC-LoRA [18] simply used Gaussian initialization for the low-rank matrices.

3. Method

3.1. Basic low-rank approximation via SVD

Low-rank approximation uses a pair of low-rank matrices $\mathbf{A} \in \mathbb{R}^{r \times d_{\text{in}}}$ and $\mathbf{B} \in \mathbb{R}^{d_{\text{out}} \times r}$ to approximate a matrix $\mathbf{W} \in \mathbb{R}^{d_{\text{out}} \times d_{\text{in}}}$, via $\mathbf{W} \approx \mathbf{BA}$, where typically the rank $r \ll \min\{d^{\text{in}}, d^{\text{out}}\}$. Compared to the original weight matrix W, the parameters used for the low-rank approximation reduces from d^2 to 2dr when $d^{in} = d^{out} = d$. However, this approximation becomes less accurate with smaller rank. A natural way to realize the low-rank approximation is to use the singular value decomposition (SVD), which decomposes a matrix W as $W = USV^T$, where U, V are the left and right singular vectors, respectively, and S is a diagonal matrix of non-negative singular values that are sorted in descending order, by convention. Truncating the SVD to keep only the largest r singular values and corresponding vectors, yields the optimal rank-r approximation of W (in terms of minimum Frobenius norm error), $\mathbf{W} \approx \mathbf{U}_r \mathbf{S}_r \mathbf{V}_r^T$, where \mathbf{U}_r and \mathbf{V}_r denote the first r columns of U and V, respectively, and \mathbf{S}_r is the top-left $r \times r$ block of \mathbf{S} , corresponding to the largest (most important) singular values. The low-rank approximation BA can be formed by

$$\mathbf{B} = \mathbf{U}_r \mathbf{S}_r^{1/2}, \ \mathbf{A} = \mathbf{S}_r^{1/2} \mathbf{V}_r^T.$$
(1)

Notation: we will use S_{-r} to denote the bottom-right $r \times r$ block of S, and U_{-r} and V_{-r} to denote the corresponding last r columns of U and V, respectively.

3.2. Activation-aware low-rank approximation

Focusing on only approximating the pretrained weights is suboptimal as it does not consider the effect of the activation statistics on the overall performance of the model. Ultimately, the objective is to approximate **WX**, given a calibration set of l sample activations arranged in the matrix $\mathbf{X} \in \mathbb{R}^{d_{\text{in}} \times l}$. Following the methodology of [38, 44], we adopt an invertible transform matrix C to adapt the weights to the pattern of activations X. In this framework, SVD is performed on WC, as given below:

$$\mathbf{W} = (\mathbf{W}\mathbf{C})\mathbf{C}^{-1} = \mathrm{SVD}(\mathbf{W}\mathbf{C})\mathbf{C}^{-1} = (\mathbf{U}\mathbf{S}\mathbf{V}^T)\mathbf{C}^{-1}$$
(2)

Keeping only the top r singular values/vectors, we have the low-rank approximation $\mathbf{W} \approx \hat{\mathbf{W}} := (\mathbf{U}_r \mathbf{S}_r \mathbf{V}_r^T) \mathbf{C}^{-1}$. As derived by [38], an optimal choice, to minimize

$$\|\hat{\mathbf{W}}\mathbf{X} - \mathbf{W}\mathbf{X}\|_F^2 \tag{3}$$

is to set C equal to the Cholesky decomposition of the covariance $\mathbf{X}\mathbf{X}^T$. This choice of C serves to whiten the activations, i.e., $\mathbf{C}^{-1}\mathbf{X}(\mathbf{C}^{-1}\mathbf{X})^T = \mathbf{I}$. Similarly, setting C to be the square root (in terms of the singular values) of the covariance, i.e., $\mathbf{C} = (\mathbf{X}\mathbf{X}^T)^{1/2}$, also achieves the optimum, which we called **RootCorDA**, in contrast to CorDA [40] that sets $\mathbf{C} = (\mathbf{X}\mathbf{X}^T)$.

In RootCorDA, one can set $\mathbf{B} = \mathbf{U}_r$ and $\mathbf{A} = \mathbf{S}_r \mathbf{V}_r^T \mathbf{C}^{-1}$ with $\mathbf{C} = (\mathbf{X}\mathbf{X}^T)^{1/2}$, which achieves the global optimal of Eq. (3) given $\hat{\mathbf{W}} = \mathbf{B}\mathbf{A}$.

3.3. Prune low-rank structure during distillation

To further compress the low-rank structure, we propose to further simultaneously prune the low-rank matrices during the low-rank distillation process via hard shrinkage. Specifically, given the target pruning ratio ρ , in the forward process of each iteration, we dynamically calculate the corresponding pruning threshold for each low-rank matrix, and zero out elements with magnitudes smaller than the threshold. For example, for the low-rank matrix **A**, we have prune(**A**) = HardShrink(**A**, ρ), where the Hard-Shrink function sets the smallest ρ percent of the elements in **A** (in terms of the magnitude) to zero.

At the end of our joint fine-tuning and compression iterations, the original weight matrix \mathbf{W} is approximated/replaced by the product of two pruned low-rank matrices, i.e., prune(\mathbf{B})prune(\mathbf{A}).

3.4. Joint fine-tuning, distillation and compression

3.4.1. Distill from teacher branch to student branch

Similar to PC-LoRA [18], a progressive compression pipeline is employed. Specifically, the output of each linear layer can be denoted as $\mathbf{Y} = \alpha \mathbf{Y}_{\text{teacher}} + \alpha' \mathbf{Y}_{\text{student}}$, where α decreases from one to zero according to

$$\alpha_t = \begin{cases} 1 - \sin\left(\frac{\pi t}{2T}\right) & \text{if } t \le T\\ 0 & \text{otherwise} \end{cases}$$
(4)

where t is the current iteration number, and T is the number of decaying iterations, which is set as 80% of total number of iterations. In PC-LoRA [18], $\alpha' = 1$ and $\mathbf{Y}_{\text{teacher}} = \mathbf{W}_0 \mathbf{X} + \mathbf{b}_0$, $\mathbf{Y}_{\text{student}} = \mathbf{W} \mathbf{X} + \mathbf{b}$. Instead, to maintain the total power from both branches to be equal to one, we propose to use $\alpha'_t = \sqrt{1 - \alpha_t^2}$.

Our overall pipeline becomes

$$\mathbf{Y} = \alpha_t \underbrace{(\mathbf{W}_0 \mathbf{X} + \mathbf{b}_0)}_{\text{teacher}} + \sqrt{1 - \alpha_t^2} \underbrace{(\mathbf{W} \mathbf{X} + \mathbf{b})}_{\text{student}}, \quad (5)$$

where decay factor α_t is set as in Eq. (4).

3.4.2. Layer-wise regularization

To generalize better on downstream tasks, a regularization term is added to the overall loss:

$$L_{\text{total}} = L_{\text{task}}\{y, \hat{y}\} + \gamma L_{\text{feat}}\{F_t, F_s\}$$
(6)

where $L_{\text{feat}}\{F_t, F_s\} = \frac{1}{m} \sum_{i=1}^m \text{MSE}(F_{t_i}, F_{s_i})$. Here y, \hat{y} are the ground truth and outputs of the network, F_{s_i} and F_{t_i} are outputs of *i*-th linear layer from student and teacher model respectively.

Instead of setting γ as a constant 0.2 in [18], we propose to decrease γ from 1 to 0 during the iterations, similar to Eq. (4). Our goal is to provide more layer-wise guidance at the beginning of the distillation iterations, and gradually relax this guidance.

4. Experiments

To validate our joint fine-tuning and compression method, we evaluated it on the Vision Transformer (ViT) [11] model for image classification.

4.1. Settings

We evaluated our TuneComp pipeline on the transfer learning task from ImageNet1K [14] to CIFAR100 [21] using the ViT-Base model as in [6]. We use training set of CI-FAR100 as the fine-tuning dataset, and test on the testing set of CIFAR100. Similar to [6], OneCycleLR scheduler is used and learning rate ranges from 7×10^{-5} to 3×10^{-3} , and the best classification accuracy is recorded. The rank of the low-rank matrices of the student branch is tested from {32, 64, 128, 256}. Note that smaller rank corresponds to higher compression rate.

4.2. Comparison of compression strategies

We compare the proposed TuneComp with some widely used pipelines:

- Fine-tune: baseline, no compression at all.
- *Fine-tune* → *distill*: first full (dense) fine-tune each linear layer, and then compress the model by distillation.
- $Distill \rightarrow fine-tune$: distill to a small model first, and then fine-tune the small model.
- *TuneComp* (joint fine-tune and compress): progressively distill, prune, and fine-tune simultaneously, where the pruning ratios are tested between 0 to 90%.

INITIALIZATION	r = 32	r = 64
$\mathbf{B} = 0, \mathbf{A} = \mathbf{G}$ Aussian	65.92	69.07
$\mathbf{B} = \mathbf{G}$ AUSSIAN, $\mathbf{A} = \mathbf{G}$ AUSSIAN	70.13	73.14
$\mathbf{B} = 0, \mathbf{A} = \mathbf{W}_0 \times \text{Gaussian}$	72.73	74.18
$\mathbf{B} = \mathbf{U}_{-r}, \mathbf{A} = \mathbf{V}_{-r}$	70.91	72.37
$\mathbf{B} = \mathbf{U}_{-r} \mathbf{S}_{-r}, \mathbf{A} = \mathbf{V}_{-r}$	70.35	68.40
$\mathbf{B} = \mathbf{U}_{-r}, \mathbf{A} = \mathbf{S}_{-r} \mathbf{V}_{-r}$	73.01	72.57
$\mathbf{B} = \mathbf{U}_{-r} \mathbf{S}_{-r}^{1/2}, \mathbf{A} = \mathbf{S}_{-r}^{1/2} \mathbf{V}_{-r}$	70.93	72.42
$\mathbf{B} = \mathbf{U}_r, \mathbf{A} = \mathbf{V}_r$	75.03	79.61
$\mathbf{B} = \mathbf{U}_r \mathbf{S}_r, \mathbf{A} = \mathbf{V}_r$	78.01	83.73
$\mathbf{B} = \mathbf{U}_r, \mathbf{A} = \mathbf{S}_r \mathbf{V}_r$	77.96	83.82
$\mathbf{B} = \mathbf{U}_r \mathbf{S}_r^{1/2}, \mathbf{A} = \mathbf{S}_r^{1/2} \mathbf{V}_r$	78.74	83.89
CorDA	76.70	83.34
ROOTCORDA (PROPOSED)	78.99	84.62

Table 1. Comparison of the classification accuracies by different low-rank matrices initialization methods for TuneComp.

• Distill only: compression only without fine-tuning.

The accuracy of the *Distill only* method is very poor (below 40%) and hence omitted from the figure. As shown by the Pareto fronts plotted in Fig. 2, *TuneComp* significantly outperforms baseline methods. *Distill* \rightarrow *fine-tune* performs second best. Further, while *FT* \rightarrow *distill* can even outperform the state-of-the-art joint fine-tuning and compression method PC-LoRA at low compression rate, it is still dominated by our proposed method.

4.3. Effect of low-rank matrix initialization

We evaluated different initialization methods for the lowrank matrices **B** and **A** listed below.

(a) For Gaussian initialization, we evaluated:

- **B** = **0**, **A** = Gaussian [14];
- $\mathbf{B} = \text{Gaussian}, \mathbf{A} = \text{Gaussian};$

The results, found in Tab. 1, show that for the joint finetuning and compression case, initializing both low-rank matrices as Gaussian results in better performance, especially when the rank is low.

(b) For weight-aware initialization, we compared with Nystrom initialization [24] where $\mathbf{B} = \mathbf{0}$, $\mathbf{A} = \mathbf{W} \times$ Gaussian. We also compared with many variants of SVD based initializations listed in Tab. 1.

We can see from Tab. 1 that, for SVD based low-rank decomposition, selecting top r principle vectors works better than least r singular vectors, and is much better than simple Gaussian initialization and Nystrom initialization. Symmetrically assigning singular values into left and right singular vectors is often a good choice.

(c) For the activation-aware initialization, we evaluated:

- CorDA initialization from [40];
- Proposed RootCorDA initialization.

The last few rows in Tab. 1 shows that CorDA in [40] cannot consistently outperform the weight-aware SVD initial-



Figure 3. Performance of TuneComp under different pruning ratio $\rho \in \{0\%, 20\%, 40\%, 60\%, 80\%, 90\%, 95\%\}.$

Table 2. Classification accuracy of TuneComp compressed model using different regularization weighting strategies.

REGULARIZATION	r = 32	r = 64	r = 128	r = 256
CONSTANT [18]	78.39	84.27	88.12	90.80
Dynamic	80.41	86.11	89.49	91.69

ization (Symmetric singular values with top-r singular vectors). However, our RootCorDA outperforms CorDA.

4.4. The effect of pruning

We studied the influence of integrating pruning into the compression. Specifically, during joint fine-tuning and distillation of the low-rank structure, pruning on the low-rank structure is simultaneously applied to each matrix **B** and **A**, and the pruning ratio can be predefined. Fig. 3 demonstrates the performance of TuneComp under different pruning ratios. The blue dashed line corresponds to no pruning (ratio = 0%), while other lines correspond to different pruning ratios. We can see that pruning 20% (yellow dashed line) or 40% (green dashed line) consistently outperforms the no pruning case, in terms of the accuracy-efficiency tradeoff. However, pruning 60% or more may not have a better accuracy-efficiency tradeoff.

4.5. Effect of regularization decay

Finally, we compare our proposed regularization decay strategy in Sec. 3.4.2 with the constant regularization $\gamma = 0.2$ [18]. The results are shown in Tab. 2. It is clear that our proposed decay strategy achieves better performance.

5. Conclusion

We proposed a method that simultaneously performs finetuning, knowledge distillation, low-rank approximation, and pruning. Experiments show that it significantly outperforms baselines, especially when compression rate is high.

References

- Guangji Bai, Yijiang Li, Chen Ling, Kibaek Kim, and Liang Zhao. Sparsellm: Towards global pruning of pre-trained language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 2
- [2] Ron Banner, Yury Nahshan, and Daniel Soudry. Post training 4-bit quantization of convolutional networks for rapiddeployment. Advances in Neural Information Processing Systems, 32, 2019. 2
- [3] Yaohui Cai, Zhewei Yao, Zhen Dong, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. Zeroq: A novel zero shot quantization framework. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 13169–13178, 2020. 2
- [4] Beidi Chen, Tri Dao, Eric Winsor, Zhao Song, Atri Rudra, and Christopher Ré. Scatterbrain: Unifying sparse and lowrank attention. *Advances in Neural Information Processing Systems*, 34:17413–17426, 2021. 2
- [5] Feiyang Chen, Ziqian Luo, Lisang Zhou, Xueting Pan, and Ying Jiang. Comprehensive survey of model compression and speed up for vision transformers. *Journal of Information*, *Technology and Policy*, pages 1–12, 2024. 2
- [6] Xiangyu Chen, Jing Liu, Ye Wang, Pu Wang, Matthew Brand, Guanghui Wang, and Toshiaki Koike-Akino. Super-LoRA: Parameter-efficient unified adaptation for large vision models. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR) Workshops, pages 8050–8055, 2024. 2, 3
- [7] Xiangyu Chen, Ye Wang, Matthew Brand, Pu Perry Wang, Jing Liu, and Toshiaki Koike-Akino. Slaying the hydra: Parameter-efficient hyper networks with low-displacement rank adaptation. In *Adaptive Foundation Models: Evolving AI for Personalized and Efficient Learning*, 2024. 2
- [8] Hongrong Cheng, Miao Zhang, and Javen Qinfeng Shi. A survey on deep neural network pruning: Taxonomy, comparison, analysis, and recommendations. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2024. 2
- [9] Krzysztof Marcin Choromanski, Valerii Likhosherstov, David Dohan, Xingyou Song, Andreea Gane, Tamas Sarlos, Peter Hawkins, Jared Quincy Davis, Afroz Mohiuddin, Lukasz Kaiser, et al. Rethinking attention with performers. In *International Conference on Learning Representations*, 2021. 2
- [10] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. Gpt3. int8 (): 8-bit matrix multiplication for transformers at scale. In Advances in Neural Information Processing Systems, 2022. 2
- [11] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020. 3
- [12] Ali Edalati, Marzieh Tahaei, Ivan Kobyzev, Vahid Partovi Nia, James J Clark, and Mehdi Rezagholizadeh. KronA: Parameter efficient tuning with Kronecker adapter. In

NeurIPS'23 Workshop on on Efficient Natural Language and Speech Processing, 2023. 2

- [13] Elias Frantar and Dan Alistarh. Sparsegpt: massive language models can be accurately pruned in one-shot. In *Proceedings* of the 40th International Conference on Machine Learning, pages 10323–10337, 2023. 2
- [14] Soufiane Hayou, Nikhil Ghosh, and Bin Yu. The impact of initialization on lora finetuning dynamics. arXiv preprint arXiv:2406.08447, 2024. 2, 3, 4
- [15] Yen-Chang Hsu, Ting Hua, Sungen Chang, Qian Lou, Yilin Shen, and Hongxia Jin. Language model compression with weighted low-rank factorization. In *International Conference on Learning Representations*, 2022. 2
- [16] Edward J Hu, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, Weizhu Chen, et al. Lora: Lowrank adaptation of large language models. In *International Conference on Learning Representations*, 2022. 2
- [17] Yating Huang, Yunzhe Hao, Jiaming Xu, and Bo Xu. Compressing speaker extraction model with ultra-low precision quantization and knowledge distillation. *Neural Networks*, 154:13–21, 2022. 2
- [18] Injoon Hwang, Haewon Park, Youngwan Lee, Jooyoung Yang, and SunJae Maeng. PC-loRA: Low-rank adaptation for progressive model compression with knowledge distillation. arXiv preprint arXiv:2406.09117, 2024. 1, 2, 3, 4
- [19] Berivan Isik, Hermann Kumbong, Wanyi Ning, Xiaozhe Yao, Sanmi Koyejo, and Ce Zhang. Gpt-zip: Deep compression of finetuned large language models. In Workshop on Efficient Systems for Foundation Models@ ICML2023, 2023.
- [20] Mengyu Ji, Gaoliang Peng, Sijue Li, Feng Cheng, Zhao Chen, Zhixiong Li, and Haiping Du. A neural network compression method based on knowledge-distillation and parameter quantization for the bearing fault diagnosis. *Applied Soft Computing*, 127:109331, 2022. 2
- [21] Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009. 3
- [22] Minae Kwon, Sang Michael Xie, Kalesha Bullard, and Dorsa Sadigh. Reward design with language models. In *The Eleventh International Conference on Learning Representations*, 2023. 2
- [23] Woosuk Kwon, Sehoon Kim, Michael W Mahoney, Joseph Hassoun, Kurt Keutzer, and Amir Gholami. A fast posttraining pruning framework for transformers. *Advances in Neural Information Processing Systems*, 35:24101–24116, 2022. 2
- [24] Bingcong Li, Liang Zhang, Aryan Mokhtari, and Niao He. On the crucial role of initialization for matrix factorization. arXiv preprint arXiv:2410.18965, 2024. 2, 4
- [25] Benedetta Liberatori, Ciro Antonio Mami, Giovanni Santacatterina, Marco Zullich, and Felice Andrea Pellegrino. Yolo-based face mask detection on low-end devices using pruning and quantization. In 2022 45th Jubilee International Convention on Information, Communication and Electronic Technology (MIPRO), pages 900–905. IEEE, 2022. 2
- [26] Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang,

Chuang Gan, and Song Han. Awq: Activation-aware weight quantization for on-device llm compression and acceleration. *Proceedings of Machine Learning and Systems*, 6:87–100, 2024. 2

- [27] Jing Liu, Toshiaki Koike-Akino, Pu Wang, Matthew Brand, Ye Wang, and Kieran Parsons. Loda: Low-dimensional adaptation of large language models. In *NeurIPS'23 Work*shop on on Efficient Natural Language and Speech Processing, 2023. 2
- [28] Zhuang Liu, Mingjie Sun, Tinghui Zhou, Gao Huang, and Trevor Darrell. Rethinking the value of network pruning. In *International Conference on Learning Representations*, 2019. 2
- [29] Alexandre Lopes, Fernando Pereira dos Santos, Diulhio de Oliveira, Mauricio Schiezaro, and Helio Pedrini. Computer vision model compression techniques for embedded systems: A survey. *Computers & Graphics*, 123:104015, 2024. 2
- [30] Xin Men, Mingyu Xu, Qingyu Zhang, Bingning Wang, Hongyu Lin, Yaojie Lu, Xianpei Han, and Weipeng Chen. Shortgpt: Layers in large language models are more redundant than you expect. arXiv preprint arXiv:2403.03853, 2024. 2
- [31] Fanxu Meng, Zhaohui Wang, and Muhan Zhang. PiSSA: Principal singular values and singular vectors adaptation of large language models. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 2
- [32] Markus Nagel, Mart van Baalen, Tijmen Blankevoort, and Max Welling. Data-free quantization through weight equalization and bias correction. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 1325– 1334, 2019. 2
- [33] Arijit Sehanobish, Kumar Avinava Dubey, Krzysztof Marcin Choromanski, Somnath Basu Roy Chowdhury, Deepali Jain, Vikas Sindhwani, and Snigdha Chaturvedi. Structured unrestricted-rank matrices for parameter efficient finetuning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 2
- [34] Jingtong Su, Yihang Chen, Tianle Cai, Tianhao Wu, Ruiqi Gao, Liwei Wang, and Jason D Lee. Sanity-checking pruning methods: Random tickets can win the jackpot. Advances in neural information processing systems, 33:20390–20401, 2020. 2
- [35] Frederick Tung, Srikanth Muralidharan, and Greg Mori. Fine-pruning: Joint fine-tuning and compression of a convolutional network with bayesian optimization. arXiv preprint arXiv:1707.09102, 2017. 1, 2
- [36] Haotao Wang, Shupeng Gui, Haichuan Yang, Ji Liu, and Zhangyang Wang. Gan slimming: All-in-one gan compression by a unified optimization framework. In *European Conference on Computer Vision*, pages 54–73. Springer, 2020. 2
- [37] Hongyi Wang, Saurabh Agarwal, and Dimitris Papailiopoulos. Pufferfish: Communication-efficient models at no extra cost. *Proceedings of Machine Learning and Systems*, 3:365– 386, 2021. 2
- [38] Xin Wang, Yu Zheng, Zhongwei Wan, and Mi Zhang. SVD-LLM: Truncation-aware singular value decomposition

for large language model compression. *arXiv preprint* arXiv:2403.07378, 2024. 2, 3

- [39] Yunyang Xiong, Zhanpeng Zeng, Rudrasis Chakraborty, Mingxing Tan, Glenn Fung, Yin Li, and Vikas Singh. Nyströmformer: A nyström-based algorithm for approximating self-attention. In *Proceedings of the AAAI conference on artificial intelligence*, pages 14138–14148, 2021. 2
- [40] Yibo Yang, Xiaojie Li, Zhongzhu Zhou, Shuaiwen Leon Song, Jianlong Wu, Liqiang Nie, and Bernard Ghanem. CorDA: Context-oriented decomposition adaptation of large language models for task-aware parameter-efficient finetuning. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*, 2024. 2, 3, 4
- [41] Lewei Yao, Renjie Pi, Hang Xu, Wei Zhang, Zhenguo Li, and Tong Zhang. Joint-detnas: Upgrade your detector with nas, pruning and dynamic distillation. In *Proceedings of* the IEEE/CVF conference on computer vision and pattern recognition, pages 10175–10184, 2021. 2
- [42] Shih-Ying Yeh, Yu-Guan Hsieh, Zhidong Gao, Bernard BW Yang, Giyeong Oh, and Yanmin Gong. Navigating text-toimage customization: From LyCORIS fine-tuning to model evaluation. In *The Twelfth International Conference on Learning Representations*, 2024. 2
- [43] Zhihang Yuan, Lin Niu, Jiawei Liu, Wenyu Liu, Xinggang Wang, Yuzhang Shang, Guangyu Sun, Qiang Wu, Jiaxiang Wu, and Bingzhe Wu. Rptq: Reorder-based post-training quantization for large language models. arXiv preprint arXiv:2304.01089, 2023. 2
- [44] Zhihang Yuan, Yuzhang Shang, Yue Song, Qiang Wu, Yan Yan, and Guangyu Sun. ASVD: Activation-aware singular value decomposition for compressing large language models. arXiv preprint arXiv:2312.05821, 2023. 2