Recall Distortion in Neural Network Pruning and the Undecayed Pruning Algorithm

Anonymous Author(s)
Affiliation
Address
email

Abstract

Pruning techniques have been successfully used in neural networks to trade accuracy for sparsity. However, the impact of network pruning is not uniform: prior work has shown that the recall for underrepresented classes in a dataset may be more negatively affected. In this work, we study such relative distortions in recall by hypothesizing an intensification effect that is inherent to the model. Namely, that pruning makes recall relatively worse for a class with recall below accuracy and, conversely, that it makes recall relatively better for a class with recall above accuracy. In addition, we propose a new pruning algorithm aimed at attenuating such effect. Through statistical analysis, we have observed that intensification is less severe with our algorithm but nevertheless more pronounced with relatively more difficult tasks, less complex models, and higher pruning ratios. More surprisingly, we conversely observe a de-intensification effect with lower pruning ratios.

1 Introduction

Back in the old days, network pruning was important to reduce the size of neural networks [22, 50, 31, 38, 23, 24]. With new advances being increasingly reliant on overwhelmingly large and costly models [3, 60, 9], network pruning is back to the game despite the gains in computing power.

Network pruning reduces the complexity of large models by setting a substantial number of parameters to zero and then using gradient descent to fine-tune the sparser model. There is typically a trade-off between making the model sparser—i.e., having fewer parameters, since those set to zero can simply be discarded—and keeping the model as accurate as possible [8], although substantial gains in sparsity may be obtained with little impact to accuracy. However, recent studies have found that the side effects of network pruning on model performance are not evenly distributed [27, 52, 28].

Those studies have considered the impact of network pruning on recall—i.e., the number of correct predictions per class—based on dataset representability. Namely, they have found that the degradation in performance is influenced by the uneven representation of classes [27, 52] and features [28] as well as class complexity [52]. Hence, they corroborate long-standing concerns that unbalanced datasets may lead to models that are less accurate in and potentially harm to minoritized groups [10, 6].

In this work, we complement those prior studies by posing recall distortion as inherent to network pruning even when the networks are trained and fine-tuned on datasets that are seemingly balanced; investigating when and how such distortions manifest; and how to reduce their effect.

More specifically, this paper presents the following contributions:

(i) We develop a statistical study of model-level distortion to investigate how recall after by network pruning due to pruning ratios, data complexity, model complexity, and pruning algorithms.

Submitted to 36th Conference on Neural Information Processing Systems (NeurIPS 2022). Do not distribute.

(ii) We observe an intensification effect, meaning that classes with recall below accuracy are more severely affected and conversely those with recall above accuracy are less severely affected by network pruning. The intensification correlates with excessive pruning ratios as well as more complex data and models, and is more pronounced with some pruning algorithms.

(iii) More surprisingly, we observe that otherwise network pruning has a corrective effect on recall differences, hence implying a de-intensification effect under moderate use.

(iv) We introduce a new gradient-based pruning method for networks trained with weight decay, Undecayed Pruning, which attenuates the intensification effect observed with other methods.

2 Related work

The use of large models has been justified and further encouraged by findings that the overparameterized regime may avoid the classic bias–variance trade-off [77, 5] and lead to better convergence during training [41, 61]. However, these models come with steep environmental footprint and hardware needs [60]. Not surprisingly, they represent an relevant application of network pruning [19, 72].

Network pruning has been motivated by parameter redundancy in models [12] and found to improve robustness against adversarial manipulation [68]. The amount of pruning that is tolerable by an architecture may depend on the task in which it is trained [43]. As observed on a recent survey [8], most approaches to network pruning are guided by either (i) removing the parameters with smallest absolute value [22, 50, 31, 21, 20, 40, 16, 14, 19, 63, 45]; or (ii) removing the parameters with smallest impact on the model [38, 23, 24, 36, 49, 13, 73, 76, 4, 39, 65, 44, 66, 70, 57, 74]. We may regard exact pruning as a special case of (ii) in which the model is not affected [55, 58, 56, 17]. Alternative approaches to network pruning include quantization [30, 47, 2], knowledge distillation [34, 42, 78, 67, 64, 62, 32, 71], and the use regularization during training [75, 53, 15].

In this work, we contribute to the study of recall distortion due to network pruning. Along those lines, the performance degradation due to network pruning has been found to be uneven and to vary based on the disproportional representation of classes [27, 52] and features [28] as well as class complexity [52]. In addition, the work led by Hooker [27, 28] studied more extensively the role of challenging samples in pruning distortion. We further comment on how these prior analysis compare to ours at the end of Section 5. Other general approaches in the literature have include analyzing the variance among all classes [7] or only including the correct predictions of the original model and using the largest recall difference as the pruning bias [33]. In the context of natural language processing, Xu et al. [72] propose a similar metric: label loyalty to reflect the resemblance of the labels predicted when compressing BERT. Besides the difference in the label predicted, some works [72, 33] borrow the idea of knowledge distillation [26] to evaluate the bias by the KL-divergence of the output probability (the output of the softmax layer) between the model before and after pruning and then modify the loss function to mitigate such bias when the model is fine-tuned after pruning.

Such disproportional impact on recall across classes as observed in [27, 52] relates to fairness in machine learning. The lack of representativeness and the biased context in which data is collected as well as the lack of transparency may lead to models making life-altering decisions that negatively affect minoritized groups, such as in criminal justice [51, 10, 54]. These concerns have motivated an extensive discussion on fairness in machine learning [11] and on the proper assessment of datasets [18], models [48], and the circumstances in which they are applied [1]. The effect of compression on fairness is proposed to study the influence of a sensitive attribute on prediction in [28, 59]. In this work, we focus on studying the algorithmic bias [46, 29] of network pruning, that is, the bias that is either (i) not present in the input data nor the original model; or (ii) deteriorated or relieved by pruning. By subscribing to this line of study, our work aims at preventing the potential negative societal impacts of network pruning.

3 Pruning algorithms

We propose Undecayed Pruning (UP) by considering the interplay between the classic representatives of pruning methods described in Section 2: Magnitude Pruning (MP) and Gradient Pruning (GP).

Magnitude Pruning MP is a simple but rather (mysteriously) effective technique of selecting among the parameters $\theta = \bar{\theta}$ the one with smallest absolute value to be pruned next from the network:

$$i = \arg\min_i \left\{|\bar{\theta}_i|\right\}$$

A parameter having a smaller value may not necessarily imply lesser importance to the model [23] (hence the mystery part). However, it is worth noticing that this technique is usually applied to networks that were trained with regularization. Since we cannot add $L0$ regularization to the loss function for directly minimizing the number of nonzero parameters as this would make the loss function not differentiable for gradient descent, we can use $L1$ or $L2$ as a proxy, which induces the values of the parameters to be as small as possible. Although no parameter ends up being zero, some parameters become so small that making them equal to zero implies an almost negligible change.

Gradient Pruning GP approximates the impact of modifying the parameters of a neural network by estimating the first-order variation of the loss function $\mathcal{L}$ on the training set between the trained parameters $\theta = \bar{\theta}$ to a new set of parameters θ through the Taylor series around $\theta = \bar{\theta}$:

$$\mathcal{L}(\theta) = \mathcal{L}(\bar{\theta}) + (\theta - \bar{\theta})\nabla\mathcal{L}(\bar{\theta}) + O(\|\theta - \bar{\theta}\|^2)$$

By assuming that $O(\|\theta - \bar{\theta}\|^2) \approx 0$ for small changes, we estimate the change to the loss function by pruning a single parameter i—i.e., with $\tilde{\theta}$ such that $\tilde{\theta}_i = 0$ and $\tilde{\theta}_j = \bar{\theta}_j$ for $j \neq i$—as follows:

$$\mathcal{L}(\tilde{\theta}) - \mathcal{L}(\bar{\theta}) \approx (\tilde{\theta} - \bar{\theta})\nabla\mathcal{L}(\bar{\theta}) = -\bar{\theta}_i\nabla_i\mathcal{L}(\bar{\theta})$$

Given the approximate nature of estimating the impact on the training set, which may not necessarily reflect on the test set, small absolute perturbations are often preferred to negative but larger variations. Hence, the choice of the parameter i to be pruned next is effectively framed as follows:

$$i = \arg\min_i \left\{|-\bar{\theta}_i\nabla_i\mathcal{L}(\bar{\theta})|\right\}$$

Undecayed Pruning We understand that simultaneously using regularization and GP is conflicting because sparsity is induced in two different ways. However, whereas GP has a more direct proxy to model impact, using some amount of regularization tends to be beneficial during training. Hence, we approach this by isolating the effect of regularization from the variation of the loss function. Due to its greater popularity, we consider $L2$ regularization–i.e,, weight decay–in what follows. Let us denote by $\mathcal{T}$ the loss function of a model after deducting weight decay with hyperparameter ε:

$$\mathcal{T}(\theta) = \mathcal{L}(\theta) - \frac{\varepsilon}{2}\|\theta\|^2$$

With $\nabla_i\mathcal{T}(\theta) = \nabla_i\mathcal{L}(\theta) - \varepsilon\theta_i$, we estimate the change to the alternative loss function $\mathcal{T}$ by pruning a single parameter i—i.e., with $\tilde{\theta}$ such that $\tilde{\theta}_i = 0$ and $\tilde{\theta}_j = \bar{\theta}_j$ for $j \neq i$—as follows:

$$\mathcal{T}(\tilde{\theta}) - \mathcal{T}(\bar{\theta}) \approx (\tilde{\theta} - \bar{\theta})\nabla\mathcal{T}(\bar{\theta}) = -\bar{\theta}_i\nabla_i\mathcal{T}(\bar{\theta}) = -\bar{\theta}_i\nabla_i\mathcal{T}(\bar{\theta}) = -\bar{\theta}_i\nabla_i\mathcal{L}(\bar{\theta}) + \varepsilon\bar{\theta}_i^2$$

In other words, discounting weight decay is equivalent to a specific balance of the criteria for MP and GP: the first term is equivalent to GP; the second term is curiously equivalent to MP; and the latter is prioritized in proportion to the weight decay hyperparameter ε. Moreover, if the neural network training converges to a local optimum, in which case $\nabla\mathcal{L}(\bar{\theta}) = 0$, then UP is equivalent to MP. Similar to the case of GP, we choose the parameter i to prune based on the absolute impact on $\mathcal{T}$:

$$i = \arg\min_i \left\{|-\bar{\theta}_i\nabla\mathcal{L}(\bar{\theta}) + \varepsilon\bar{\theta}_i^2|\right\}$$

4 Model properties

Let $A(m)$ be the accuracy of an unpruned model m as measured on the test data. In other words, $A(m)$ is the number of correct predictions divided by the number of samples. For a *pruning ratio* t, meaning that t is the number of parameters before pruning divided by the number of parameters after pruning, let $A_t(m)$ denote the accuracy of model m on the test data after pruning. For simplicity, we may assume that $t = 1$ if t is absent from the notation and omit m if always referring to the same model, and thus we may assume that $A = A(m) = A_1(m)$. The same applies to other metrics.

Similarly, let $R^c(m)$ be the *recall* for class c of an unpruned model m on the test data. In other words, R^c is the number of correct predictions for class c divided by the number of samples for class c. For a pruning ratio t, let $R_t^c(m)$ denote the recall for class c of model m on the test data after pruning.

We are particularly interested in the how these metrics differ for each class. Namely, let

$$B_t^c(m) = R_t^c(m) - A_t(m)$$

denote the *recall balance*[1]. When $B_t^c(m) > 0$, we say that model m at pruning ratio t *overperforms* for class c; and when $B_t^c(m) < 0$ we may say that m *underperforms* for class c. Finally, let

$$\bar{B}_t^c(m) = \frac{B_t^c(m)}{A_t(m)} = \frac{R_t^c(m) - A_t(m)}{A_t(m)}$$

denote the *normalized recall balance*. The further away this value is from 1, the more pronounced is the difference in performance between class c and the other classes in model m at the pruning ratio t.

Proposition 4.1. *If a dataset with set of classes $\mathbb{C}$ is* balanced*, meaning that each class $c \in \mathbb{C}$ has the same number of samples, then for any model m and pruning ratio t it holds that*

$$\sum_{c \in \mathbb{C}} B_t^c(m) = 0.$$

Proof. If the dataset is balanced, then it follows that $A_t = \frac{\sum_{c \in \mathbb{C}} R_t^c}{|\mathbb{C}|}$ and hence $\sum_{c \in \mathbb{C}} R_t^c = |\mathbb{C}| A_t$. Therefore, $0 = \sum_{c \in \mathbb{C}} (R_t^c) - |\mathbb{C}| A_t = \sum_{c \in \mathbb{C}} (R_t^c - A_t) = \sum_{c \in \mathbb{C}} B_t^c$. □

Corollary 4.1.1. *If a dataset is balanced, then for any model m and pruning ratio t it holds that*

$$\sum_{c \in \mathbb{C}} \bar{B}_t^c(m) = 0.$$

Proof. Immediate from dividing $\sum_{c \in \mathbb{C}} B_t^c(m)$ by $A_t(m)$ due to Proposition 4.1. □

Example Consider a model in which $A = 80\%$ with two classes X and Y such that $R^X = 90\%$ and $R^Y = 70\%$, and after pruning we have $A_t = 60\%$, $R_t^X = 70\%$, and $R_t^Y = 50\%$. Note that the recall balance remains the same for each class before and after pruning: $B^X = B_t^X = 10\%$ and $B^Y = B_c^Y = -10\%$. However, the predictions of the pruned network for class Y are as good as a guess. In turn, the normalized recall balance is the same in absolute value when the two classes are compared, but that value increases due to pruning: $\bar{B}^X = -\bar{B}^Y = \frac{1}{8}$ and $\bar{B}_t^X = -\bar{B}_t^Y = \frac{1}{6}$.

Our choice of normalizing through dividing by $A_t(m)$ rather than by $R_t^c(m)$ aims at avoiding outliers from small recall values as well as undefined results if a high pruning ratio leads to zero recall. Moreover, an extension of Proposition 4.1 such as Corollary 4.1.1 would not be possible otherwise.

We may obtain similar properties by replacing recall $R^c(m)$ with *precision* $P^c(m)$. In other words, P^c is the number of correct predictions for class c divided by the total number of predictions for class c. By extension we may also consider the effect of pruning on the *F-score*, in which case we would consider replacing R^c with the harmonic mean of recall and precision, i.e., $2\frac{R^c P^c}{R^c + P^c}$. However, we have preliminary found the effect of pruning more expressive on recall than on precision or F-score.

5 Measuring intensification

For a given pruning ratio t and class c of a model m, we consider the following *intensification ratio*:

$$I_t^c(m) := \frac{\bar{B}_t^c(m)}{\bar{B}^c(m)} \equiv \frac{\text{Normalized recall balance } \textbf{after} \text{ pruning}}{\text{Normalized recall balance } \textbf{before} \text{ pruning}}$$

[1]For a web search on April 22, 2022, there were no hits for either "recall minus accuracy" or "subtract accuracy from recall" and very few references for "difference between recall and accuracy".

This metric can be used to evaluate if pruning widens the performance gap between classes, which happens when $I_t^c > 1$. For a class c in which the original model overperforms ($\bar{B}^c > 0$), a ratio greater than 1 after pruning implies that $|\bar{B}_t^c| > |\bar{B}^c|$ since $\bar{B}^c > 0 \wedge I_t^c > 1 \rightarrow \bar{B}_t^c > \bar{B}^c > 0$. Similarly, for a class c in which the original model underperforms ($\bar{B}^c > 0$), a ratio greater than 1 after pruning *also* implies that $|\bar{B}_t^c| > |\bar{B}^c|$ since $\bar{B}^c < 0 \wedge I_t^c > 1 \rightarrow \bar{B}_t^c < \bar{B}^c < 0$. In other words, we can use the same metric to determine if pruning comparatively improves and worsens normalized recall balance for classes in which the model respectively overperforms and underperforms.

Example (cont.) For both classes X and Y, we obtain the same intensification ratio $I_t^X = I_t^Y = \frac{4}{3}$. Hence, pruning the model leads to a greater disparity in relative performance for those classes.

Since $I_t^c = \dfrac{B_t^c}{B^c}\dfrac{A}{A_t}$ and we may assume $A_t < A$ for sufficiently large t, one could argue that $\dfrac{A}{A_t} > 1$ alone leads to $I_t^c > 1$. However, our experiments show quite the opposite in the case of random pruning, which can be regarded as a less careful selection of weights to prune. Moreover, we consider a ratio above 1 from another perspective. Namely, $I_t^c > 1$ implies $\dfrac{B_t^c}{B^c} > \dfrac{A_t}{A}$ indicating the relative drop in recall balance (if indeed it drops at all) is less severe than the relative drop in accuracy.

We summarize intensification across classes for a model m at pruning ratio t by estimating the slope $\alpha_t(m)$ of a simple linear regression between normalized recall balance before and after pruning:

$$\bar{B}_t^c(m) = \alpha_t(m)\bar{B}^c(m)$$

The slope $\alpha_t(m)$ is fit using all classes in $\mathcal{C}$. We omit the intercept term from this regression as it would always be zero in the balanced datasets that we study due to Corollary 4.1.1. Consequently, the ordinary least squares estimate of the slope becomes a weighted mean of the intensification ratios I_t^c:

$$\alpha_t = \frac{\sum\limits_{c\in\mathbb{C}} \bar{B}^c \bar{B}_t^c}{\sum\limits_{c\in\mathbb{C}} (\bar{B}^c)^2} = \frac{\sum\limits_{c\in\mathbb{C}} (\bar{B}^c)^2 I_t^c}{\sum\limits_{c\in\mathbb{C}} (\bar{B}^c)^2}$$

In this weighted mean, less weight is given to classes in which normalized recall balance before pruning is closer to 0. We believe this is an appropriate model-level summary, since it de-emphasizes classes for which recall and accuracy are nearly the same before pruning. Hence, we conclude that pruning induces an overall intensification on m if $\alpha_t(m) > 1$ and de-intensification if $\alpha_t(m) < 1$.

In addition, within each model there is a statistical dependency between the set of per-class ratios $\{I_t^c \mid c \in \mathbb{C}\}$ obtained after pruning. However, the model-level summaries $\alpha_t(m)$ are independent across random initializations associated with training each model m (conditional on the dataset), which allows us to use simpler statistical inference methods as described in Section 6.2. In our statistical analysis, we evaluate $\mathbb{E}[\alpha_t]$ as the hypothetical expectation across infinitely many m.

Comparatively, the statistical analysis in Hooker et al. [27] would be equivalent to evaluating $\mathbb{E}[I_t^c]$, as it is based on $\dfrac{R_t^c}{A_t} = I_t^c - 1$. A multiple linear regression with $\dfrac{R_t^c}{A_t}$ as dependent variable and $\dfrac{R^c}{A}$ as one of the independent variables is fit by Paganini [52] over a large number of models and classes.

6 Experimental setting

We seek to understand when $\mathbb{E}[\alpha_t] > 1$, or alternatively $\mathbb{E}[\alpha_t] > 1$, by evaluating the impact of network pruning algorithms on models trained on seemingly balanced datasets and through varying but well-known architectures. Hence, we consider the dependence of the intensification ratio at the model level by denoting it as $\alpha_{t,P}^{D,M}$ for a dataset $D \in \mathbb{D}$, an architecture $M \in \mathbb{M}$, a pruning ratio $t \in \mathbb{T}$, and a pruning algorithm $P \in \mathbb{P}$. In what follows, we may omit the indices if they are constant.

6.1 Computational details

We use some combinations of (i) models based on ResNet-$\{20, 32, 44, 56, 110\}$ [25] and an adaptation of LeNet5 [37] (ii) trained in MNIST [37], Fashion-MNIST [69], CIFAR-10, and CIFAR-100 [35], and then (iii) pruned using MP, GP, and Random Pruning (RP) from ShrinkBench [8] and our

implementation of UP in ShrinkBench, all of which tested with (iv) pruning ratios 2, 4, 10, 20, and 50. For each combination of model and dataset evaluated, we have trained 30 models, which are the same used for pruning at each pruning ratio. The datasets were chosen due to their popularity and equal representation across classes. The architectures and pruning ratios were chosen based on preliminary experiments aiming for good accuracy and also to prune to up to a ratio with evidence of distortion. MNIST and Fashion-MNIST are only trained in LeNet5. When the dataset or model does not vary, we use CIFAR-10 and ResNet-56 due to their intermediary complexity. When the pruning algorithm does not vary, we use MP due to its popularity and seemingly better performance than GP [8].

The LeNet5 models are trained with SGD optimizer for 30 epochs, with batch size of 128 and learning rate of 0.01, and then fine-tuned for another 15 epochs after pruning. The ResNet models are trained with SGD optimizer for 60 epochs, with batch size of 128, a decreasing learning rate schedule, and weight decay of 0.0005. These values were selected based on preliminary testing. For all models, the weights for the epoch with the greatest top-1 accuracy are saved during training and fine-tuning for testing. All experiments are run on GPUs of Nividia GeForce 3060ti and 3090.

6.2 Statistical methods

We are interested in comparing the distributions of $\alpha_{t,P}^{D,M}$ across different scenarios. Each $\alpha_{t,P}^{D,M}$ is calculated using the test set recall metrics from just one run of training and a pair of models, corresponding to before and after pruning at ratio t. Thus we frame our statistical analyses as inferences about $\mathbb{E}[\alpha_{t,P}^{D,M}]$, where the expectation is taken over neural network models trained with different random seeds. We have tested varying M, D, or P along with t. Hence, we evaluate the impact of the pruning ratio along every other dimension. For model complexity through α_t^M, we use CIFAR-10 trained on ResNet-$\{20, 32, 44, 56, 110\}$ and pruned with MP. For dataset complexity through α_t^D, we use all the datasets with their default model and pruned with MP. For pruning algorithm through $\alpha_{t,P}$, we use CIFAR-10 trained on ResNet-56 and pruned with all pruning methods. Due to the encouraging results with UP, we have repeated the model complexity with UP for comparison.

Figures Each scatterplot matrix shows the normalized recall balances before vs after pruning, at several t (rows) and several of either M, D, or P (columns). Each point corresponds to one class c for one model m. One regression line and numeric summaries are overlaid in each subplot: $\hat{\alpha}$ is an average slope using all m together, r^2 is the corresponding coefficient of determination, and $\bar{A}$ is the average accuracy across all m. Boxplots summarize how the corresponding model-level slopes $\alpha_{t,P}^{D,M}(m)$ vary across m, and how their distribution changes with t and either M, D, or P.

Confidence intervals At each boxplot, we calculate a t-based 99% confidence interval (CI) for its $\mathbb{E}[\alpha_{t,P}^{D,M}]$. We chose 99% confidence to achieve 95% family-wide confidence (within each family of 5 pruning ratios at a given M, D, or P) after a Bonferroni correction for simultaneously reporting 5 dependent CIs. The labels for pruning ratios in the boxplots include $<$, $>$, or ? to denote whether the corresponding CI is below 1 (de-intensification), above 1 (intensification), or overlaps 1. The plots with each CI can be found in the Appendix. Figures 7 (Appendix), 8 (Appendix), and 2 indicate that within each M, D, or P, low ratios tend to have CIs entirely below 1, some moderate ratios have CIs that overlap with 1, and high ratios tend to have CIs entirely above 1. The exception is P=RP, where most CIs overlap with 1.

Hypothesis tests by t, M, D, or P We carry out t-tests for each set of hypotheses in Section 7. Tests comparing pairs of M or of D are independent-samples tests, but tests comparing pairs of t or of P are paired-samples tests: for each m there is a natural pairing between two αs using the same uncompressed m. Each p-value reported in Tables 1 through 9 (Appendix) has been multiplied by the number of rows in its column of the table, as a Bonferroni multiple-testing correction for simultaneously evaluating all the rows in that column. When a reported p-value is below the usual significance level 0.05, we have evidence that intensification differs between the pairs in that setting.

7 Analysis

7.1 The influence of the pruning ratio

We carry out one-sided paired-samples t-tests between the following pairs of hypotheses:

$$H_0^{M,t,i} : \mathbb{E}[\alpha_{t_i}^M] \geq \mathbb{E}[\alpha_{t_{i+1}}^M] \qquad H_0^{D,t,i} : \mathbb{E}[\alpha_{t_i}^D] \geq \mathbb{E}[\alpha_{t_{i+1}}^D] \qquad H_0^{P,t,i} : \mathbb{E}[\alpha_{t_i}^P] \geq \mathbb{E}[\alpha_{t_{i+1}}^P]$$

$$H_a^{M,t,i} : \mathbb{E}[\alpha_{t_i}^M] < \mathbb{E}[\alpha_{t_{i+1}}^M] \qquad H_a^{D,t,i} : \mathbb{E}[\alpha_{t_i}^D] < \mathbb{E}[\alpha_{t_{i+1}}^D] \qquad H_a^{P,t,i} : \mathbb{E}[\alpha_{t_i}^P] < \mathbb{E}[\alpha_{t_{i+1}}^P]$$

In all of those, $t_i < t_{i+1}$ are consecutive pruning ratios in $\mathcal{T}$.

Figures 7 (Appendix) and 3 suggest that at each M, intensification tends to remain similar or become stronger at higher pruning ratios. Tables 1 and 2 (Appendix) confirm strong evidence of this trend between most consecutive pairs of ratios, for both MP and UP, at most model sizes, except for ratios 2 vs 4 on some larger architectures.

Figure 8 (Appendix) and Figure 5 (Appendix) suggest that at each D, intensification tends to remain similar or become stronger at higher pruning ratios. Table 3 (Appendix) confirms that we have strong evidence of increasing intensification for MNIST when going from pruning ratio 20 to 50; for Fashion, when going from 10 to 20 and from 20 to 50; and for both CIFAR-10 and CIFAR-100, for every consecutive pair of ratios.

Figure 2 and Figure 6 (Appendix) suggest that at each P except RP, intensification tends to remain similar or become stronger at higher pruning ratios. Table 4 (Appendix) confirms strong evidence of this trend between most consecutive pairs of ratios at most P. However, for RP, the average slope actually appears to decrease instead during several ratio increases.

7.2 The influence of model complexity

For each ratio $t \in \mathcal{T}$, we carry out a one-sided independent-samples t-test of

$$H_0^{M,i,t} : \mathbb{E}[\alpha_t^{M_i}] \leq \mathbb{E}[\alpha_{t_i}^{M_{i+1}}]$$

$$H_a^{M,i,t} : \mathbb{E}[\alpha_t^{M_i}] > \mathbb{E}[\alpha_{t_i}^{M_{i+1}}]$$

for M_i and M_{i+1} as consecutive models in ResNet-$\{20, 32, 44, 56, 110\}$ on CIFAR-10, using MP and UP. The direction of these tests was chosen because, a priori, we expected smaller models (on the same dataset) to be more susceptible to intensification.

Figures 1, 3, 7 (Appendix) and 4 (Appendix) suggest that at most ratios, smaller model sizes tend to have more intensification than larger model sizes. Tables 5 and 6 (Appendix) confirm strong evidence of this trend between most consecutive pairs of model sizes (except 56 vs 110 with MP) at high t, but rarely at low t.

7.3 The influence of dataset complexity

For each ratio $t \in \mathcal{T}$, we carry out a one-sided independent-samples t-test of

$$H_0^{D,i,j,t} : \mathbb{E}[\alpha_t^{D_i}] \geq \mathbb{E}[\alpha_t^{D_j}]$$

$$H_a^{D,i,j,t} : \mathbb{E}[\alpha_t^{D_i}] < \mathbb{E}[\alpha_t^{D_j}]$$

for the following (D_i, D_j) pairs: (MNIST, CIFAR-10); (Fashion, CIFAR-10); and (CIFAR-10, CIFAR-100), using MP. These pairs were chosen because, a priori, we expected "more complex" datasets to experience more intensification at a given pruning ratio. We judged that MNIST and Fashion have comparable complexity (10 black-and-white classes), but both are less complex than CIFAR-10 (10 color-image classes), which is less complex than CIFAR-100 (100 color-image classes).

Figure 8 (Appendix) and Figure 5 (Appendix) suggest that at most ratios, MNIST and Fashion tend to have shallower slopes (less intensification) than CIFAR-10, which has shallower slopes than CIFAR-100. Table 7 (Appendix) confirms strong evidence of this trend between MNIST and CIFAR-10 at the smaller ratios; between Fashion and CIFAR-10 only at moderate ratios; and between CIFAR-10 and CIFAR-100 only at the higher ratios.

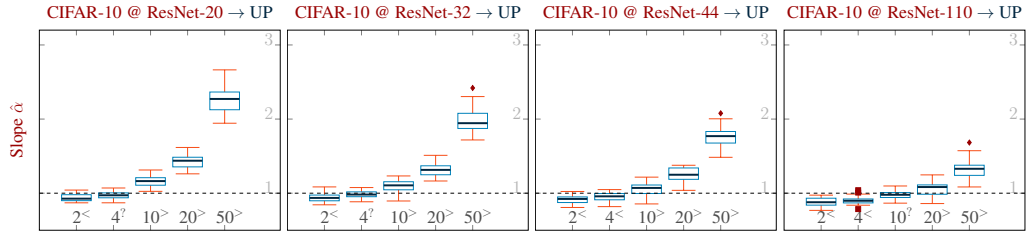

Figure 1: Boxplots of $\alpha_{t,P}^{D,M}(m)$ across m, at each t within each M for P=UP. Superscripts $<$, $>$, or ? denote where 99% CIs were below 1, above 1, or overlapped 1. See Section 6.2.

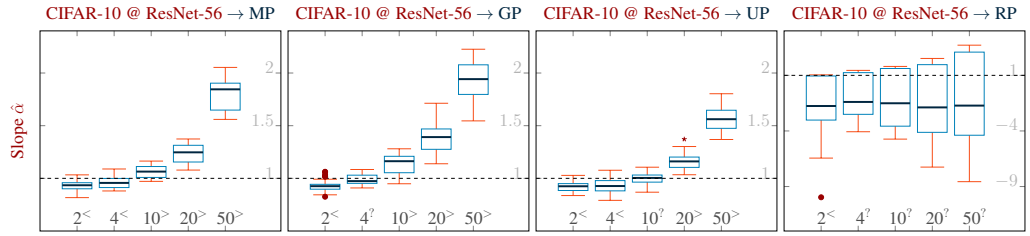

Figure 2: Boxplots of $\alpha_{t,P}^{D,M}(m)$ across m, at each t within each P. Superscripts $<$, $>$, or ? denote where 99% CIs were below 1, above 1, or overlapped 1. See Section 6.2.

7.4 The influence of the pruning algorithm

For each ratio $t \in \mathcal{T}$, we carry out a two-sided paired-samples t-test of

$$H_0: \quad \mathbb{E}[\alpha_t^{P_i}] = E[\alpha_t^{P_j}]$$

$$H_a: \quad \mathbb{E}[\alpha_t^{P_i}] \neq E[\alpha_t^{P_j}]$$

for pruning algorithms $P_i, P_j \in \{\text{MP}, \text{GP}, \text{UP}, \text{RP}\}$ on CIFAR-10 with ResNet-56. Without clear a priori expectations for which algorithms would experience more intensification than others, we compare all pairs and use two-sided tests. We also compared just MP and UP at all models in ResNet-$\{20, 32, 44, 56, 110\}$.

Figure 2 and Figure 6 (Appendix) suggest a trend in which, at most ratios, UP tends to have shallower but still positive slopes (less intensification) than MP or GP, while RP always has negative slopes. Table 8 (Appendix) confirms we have evidence that the average slopes for MP, GP, and UP differ from each other at higher pruning ratios, but not necessarily at the lowest ratios. We also have strong evidence that all three methods differ from RP, at each ratio. Also, comparing Figure 3 to Figure 4 (Appendix) suggests that UP tends to have less intensification than MP at each t and M. Table 9 (Appendix) confirms this for most ratios and models.

8 Conclusion

In this work, we have found evidence that network pruning may cause recall distortion even in models trained in seemingly balanced datasets. Our results show a statistically significant effect of the pruning ratio on intensifying the normalized recall balance (the difference between recall and accuracy divided by accuracy) for underperforming classes (recall below accuracy) as well as overperforming classes (recall above accuracy). This effect also increases with dataset complexity and decreases with model complexity. The effect is absent under Random Pruning (RP), hence providing further evidence that accuracy reduction alone does not lead to intensification as considered in our work and similarly framed in prior work [27, 52, 28]. Curiously, we have observed a de-intensification effect whenever the conditions for network pruning are not as extreme. Thus, the distortions produced by network pruning may have a corrective effect if the conditions are not too extreme.

We have also introduced a gradient-based pruning algorithm that discounts the effect of weight decay, which we interpret as more aligned with a magnitude-based approach to pruning. Whereas the new

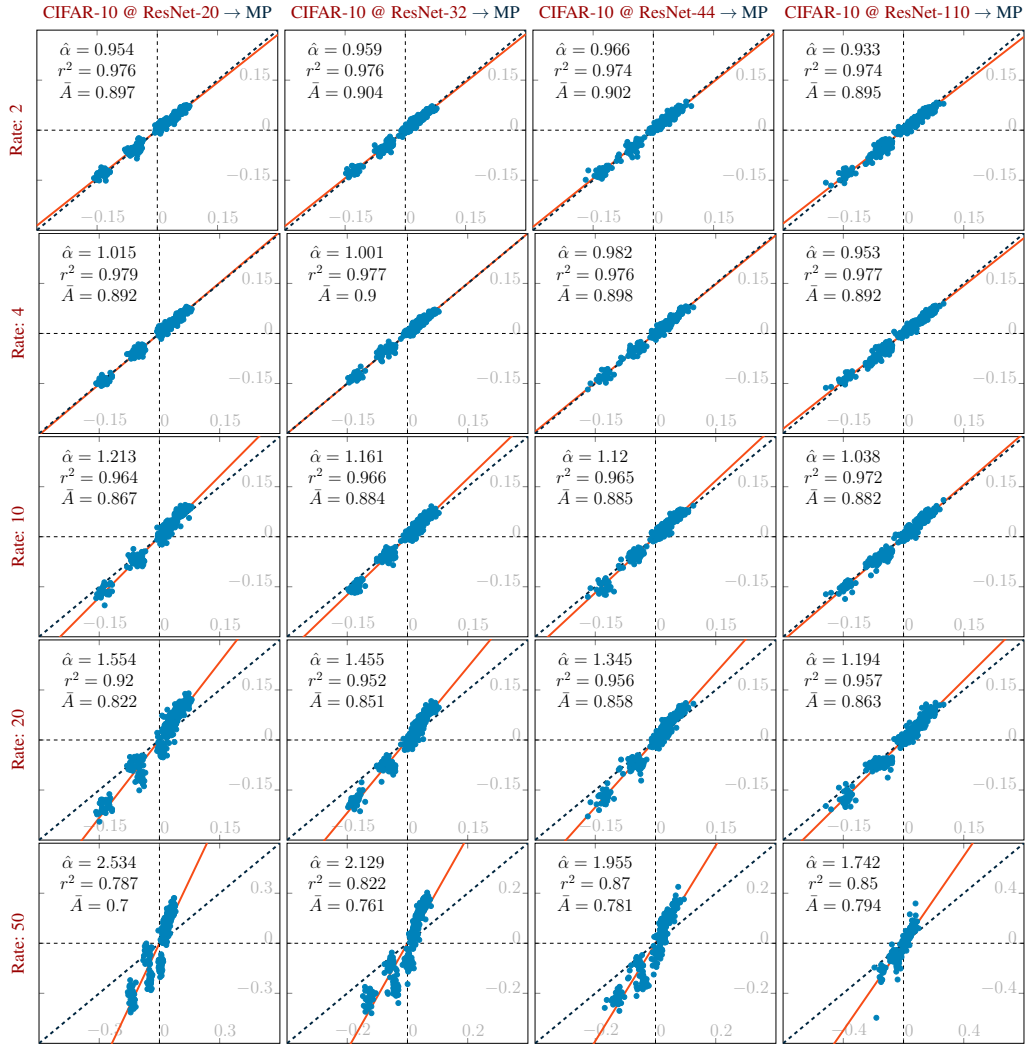

Figure 3: Scatterplot matrix of $\bar{B}^c(m)$ (x-axis) vs $\bar{B}^c_t(m)$ (y-axis), at several values of t (rows) and M (columns) for $P =$ MP. Each scatterplot point corresponds to one c for one m. See Section 6.2.

algorithm is not designed with a focus on avoiding intensification, it succeeds in doing so while also attenuating further intensification. We conjecture that its benefit comes from better aligning of impact-based pruning in presence of weight decay.

One possible limitation of our study is that it does not include a vast list of pruning algorithms. However, it includes the classic methods conventionally used and sheds light on how they can be combined through our new algorithm.

References

[1] Fairness and abstraction in sociotechnical systems, author=Andrew D. Selbst and Danah Boyd and Sorelle A. Friedler and Suresh Venkatasubramanian and Janet Vertesi, booktitle=FAT*, year=2019.

[2] Raziel Alvarez, Rohit Prabhavalkar, and Anton Bakhtin. On the efficient representation and execution of deep acoustic models. *arXiv preprint arXiv:1607.04683*, 2016.

[3] D. Amodei, D. Hernandez, G. Sastry, J. Clark, G. Brockman, and I. Sutskever. AI and compute. https://openai.com/blog/ai-and-compute/, 2018. Accessed: 2022-05-16.

[4] Cenk Baykal, Lucas Liebenwein, Igor Gilitschenski, Dan Feldman, and Daniela Rus. Data-dependent coresets for compressing neural networks with applications to generalization bounds. In *ICLR*, 2019.

[5] Mikhail Belkin, Daniel Hsu, Siyuan Ma, and Soumik Mandal. Reconciling modern machine learning practice and the bias-variance trade-off. *PNAS*, 116(32):15849–15854, 2019.

[6] Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. On the dangers of stochastic parrots: Can language models be too big? In *FAccT*, 2021.

[7] Cody Blakeney, Nathaniel Huish, Yan Yan, and Ziliang Zong. Simon says: Evaluating and mitigating bias in pruned neural networks with knowledge distillation. *arXiv preprint arXiv:2106.07849*, 2021.

[8] D. Blalock, J. Ortiz, J. Frankle, and J. Guttag. What is the state of neural network pruning? In *MLSys*, 2020.

[9] Rishi Bommasani, Drew A. Hudson, Ehsan Adeli, Russ Altman, Simran Arora, Sydney von Arx, Michael S. Bernstein, Jeannette Bohg, Antoine Bosselut, Emma Brunskill, Erik Brynjolfsson, Shyamal Buch, Dallas Card, Rodrigo Castellon, Niladri Chatterji, Annie Chen, Kathleen Creel, Jared Quincy Davis, Dora Demszky, Chris Donahue, Moussa Doumbouya, Esin Durmus, Stefano Ermon, John Etchemendy, Kawin Ethayarajh, Li Fei-Fei, Chelsea Finn, Trevor Gale, Lauren Gillespie, Karan Goel, Noah Goodman, Shelby Grossman, Neel Guha, Tatsunori Hashimoto, Peter Henderson, John Hewitt, Daniel E. Ho, Jenny Hong, Kyle Hsu, Jing Huang, Thomas Icard, Saahil Jain, Dan Jurafsky, Pratyusha Kalluri, Siddharth Karamcheti, Geoff Keeling, Fereshte Khani, Omar Khattab, Pang Wei Koh, Mark Krass, Ranjay Krishna, Rohith Kuditipudi, Ananya Kumar, Faisal Ladhak, Mina Lee, Tony Lee, Jure Leskovec, Isabelle Levent, Xiang Lisa Li, Xuechen Li, Tengyu Ma, Ali Malik, Christopher D. Manning, Suvir Mirchandani, Eric Mitchell, Zanele Munyikwa, Suraj Nair, Avanika Narayan, Deepak Narayanan, Ben Newman, Allen Nie, Juan Carlos Niebles, Hamed Nilforoshan, Julian Nyarko, Giray Ogut, Laurel Orr, Isabel Papadimitriou, Joon Sung Park, Chris Piech, Eva Portelance, Christopher Potts, Aditi Raghunathan, Rob Reich, Hongyu Ren, Frieda Rong, Yusuf Roohani, Camilo Ruiz, Jack Ryan, Christopher Ré, Dorsa Sadigh, Shiori Sagawa, Keshav Santhanam, Andy Shih, Krishnan Srinivasan, Alex Tamkin, Rohan Taori, Armin W. Thomas, Florian Tramèr, Rose E. Wang, William Wang, Bohan Wu, Jiajun Wu, Yuhuai Wu, Sang Michael Xie, Michihiro Yasunaga, Jiaxuan You, Matei Zaharia, Michael Zhang, Tianyi Zhang, Xikun Zhang, Yuhui Zhang, Lucia Zheng, Kaitlyn Zhou, and Percy Liang. On the opportunities and risks of foundation models. *arXiv*, 2108.07258, 2021.

[10] Joy Buolamwini and Timnit Gebru. Gender shades: Intersectional accuracy disparities in commercial gender classification. In *FAT**, 2018.

[11] Sam Corbett-Davies and Sharad Goel. The measure and mismeasure of fairness: A critical review of fair machine learning. *arXiv*, 1808.00023, 2018.

[12] M. Denil, B. Shakibi, L. Dinh, M. Ranzato, and N. Freitas. Predicting parameters in deep learning. In *NeurIPS*, 2013.

[13] X. Dong, S. Chen, and S. Pan. Learning to prune deep neural networks via layer-wise optimal brain surgeon. In *NeurIPS*, 2017.

[14] Bryn Elesedy, Varun Kanade, and Yee Whye Teh. Lottery tickets in linear models: An analysis of iterative magnitude pruning, 2020.

[15] Angela Fan, Edouard Grave, and Armand Joulin. Reducing transformer depth on demand with structured dropout. *arXiv preprint arXiv:1909.11556*, 2019.

[16] J. Frankle and M. Carbin. The lottery ticket hypothesis: Finding sparse, trainable neural networks. In *ICLR*, 2019.

[17] I. Ganev and R. Walters. Model compression via symmetries of the parameter space. 2022.

[18] Timnit Gebru, Jamie Morgenstern, Briana Vecchione, Jennifer Wortman Vaughan, Hanna Wallach, Hal Daumé III, and Kate Crawford. Datasheets for datasets. *Communications of the ACM*, 64(12):86–92, 2021.

[19] M. Gordon, K. Duh, and N. Andrews. Compressing BERT: Studying the effects of weight pruning on transfer learning. In *Rep4NLP Workshop*, 2020.

[20] S. Han, H. Mao, and W. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and Huffman coding. In *ICLR*, 2016.

[21] S. Han, J. Pool, J. Tran, and W. Dally. Learning both weights and connections for efficient neural network. In *NeurIPS*, 2015.

[22] S. Hanson and L. Pratt. Comparing biases for minimal network construction with back-propagation. In *NeurIPS*, 1988.

[23] B. Hassibi and D. Stork. Second order derivatives for network pruning: Optimal Brain Surgeon. In *NeurIPS*, 1992.

[24] B. Hassibi, D. Stork, and G. Wolff. Optimal brain surgeon and general network pruning. In *IEEE International Conference on Neural Networks*, 1993.

[25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016.

[26] Geoffrey Hinton, Oriol Vinyals, Jeff Dean, et al. Distilling the knowledge in a neural network. *arXiv preprint arXiv:1503.02531*, 2(7), 2015.

[27] S. Hooker, A. Courville, G. Clark, Y. Dauphin, and A. Frome. What do compressed deep neural networks forget? *arXiv*, 1911.05248, 2019.

[28] S. Hooker, N. Moorosi, G. Clark, S. Bengio, and E. Denton. Characterising bias in compressed models. *arXiv*, 2010.03058, 2020.

[29] Sara Hooker. Moving beyond "algorithmic bias is a data problem". *Patterns*, 2(4):100241, 2021.

[30] Benoit Jacob, Skirmantas Kligys, Bo Chen, Menglong Zhu, Matthew Tang, Andrew Howard, Hartwig Adam, and Dmitry Kalenichenko. Quantization and training of neural networks for efficient integer-arithmetic-only inference. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2704–2713, 2018.

[31] S. Janowsky. Pruning versus clipping in neural networks. *Physical Review A*, 1989.

[32] Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*, 2019.

[33] Vinu Joseph, Shoaib Ahmed Siddiqui, Aditya Bhaskara, Ganesh Gopalakrishnan, Saurav Muralidharan, Michael Garland, Sheraz Ahmed, and Andreas Dengel. Going beyond classification accuracy metrics in model compression. *arXiv preprint arXiv:2012.01604*, 2020.

[34] Harini Kannan, Alexey Kurakin, and Ian Goodfellow. Adversarial logit pairing. *arXiv preprint arXiv:1803.06373*, 2018.

[35] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, University of Toronto, 2009.

[36] V. Lebedev and V. Lempitsky. Fast ConvNets using group-wise brain damage. In *CVPR*, 2016.

[37] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 1998.

[38] Y. LeCun, J. Denker, and S. Solla. Optimal brain damage. In *NeurIPS*, 1989.

[39] N. Lee, T. Ajanthan, and P. Torr. SNIP: Single-shot network pruning based on connection sensitivity. In *ICLR*, 2019.

[40] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. Graf. Pruning filters for efficient convnets. In *ICLR*, 2017.

[41] Hao Li, Zheng Xu, Gavin Taylor, Christoph Studer, and Tom Goldstein. Visualizing the loss landscape of neural nets. In *NeurIPS*, 2018.

[42] Yawei Li, Shuhang Gu, Christoph Mayer, Luc Van Gool, and Radu Timofte. Group sparsity: The hinge between filter pruning and decomposition for network compression. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 8018–8027, 2020.

[43] L. Liebenwein, C. Baykal, B. Carter, D. Gifford, and D. Rus. Lost in pruning: The effects of pruning neural networks beyond test accuracy. In *MLSys*, 2021.

[44] Lucas Liebenwein, Cenk Baykal, Harry Lang, Dan Feldman, and Daniela Rus. Provable filter pruning for efficient neural networks. In *ICLR*, 2020.

[45] Shiwei Liu, Tianlong Chen, Xiaohan Chen, Zahra Atashgahi, Lu Yin, Huanyu Kou, Li Shen, Mykola Pechenizkiy, Zhangyang Wang, and Decebal Constantin Mocanu. Sparse training via boosting pruning plasticity with neuroregeneration. In *NeurIPS*, 2021.

[46] Ninareh Mehrabi, Fred Morstatter, Nripsuta Saxena, Kristina Lerman, and Aram Galstyan. A survey on bias and fairness in machine learning. *ACM Computing Surveys (CSUR)*, 54(6):1–35, 2021.

[47] Paulius Micikevicius, Sharan Narang, Jonah Alben, Gregory Diamos, Erich Elsen, David Garcia, Boris Ginsburg, Michael Houston, Oleksii Kuchaiev, Ganesh Venkatesh, et al. Mixed precision training. *arXiv preprint arXiv:1710.03740*, 2017.

[48] Margaret Mitchell, Simone Wu, Andrew Zaldivar, Parker Barnes, Lucy Vasserman, Ben Hutchinson, Elena Spitzer, Inioluwa Deborah Raji, and Timnit Gebru. Model cards for model reporting. In *FAT**, 2019.

[49] P. Molchanov, S. Tyree, T. Karras, T. Aila, and J. Kautz. Pruning convolutional neural networks for resource efficient inference. In *ICLR*, 2017.

[50] M. Mozer and P. Smolensky. Using relevance to reduce network size automatically. *Connection Science*, 1989.

[51] Cathy O'Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown, 2016.

[52] M. Paganini. Prune responsibly. *arXiv*, 2009.09936, 2020.

[53] Carles Riera, Camilo Rey, Thiago Serra, Eloi Puertas, and Oriol Pujol. Training thinner and deeper neural networks: Jumpstart regularization. In *CPAIOR*, 2022.

[54] Cynthia Rudin, Caroline Wang, and Beau Coker. The age of secrecy and unfairness in recidivism prediction. *Harvard Data Science Review*, 2(1), 2020.

[55] T. Serra, A. Kumar, and S. Ramalingam. Lossless compression of deep neural networks. In *CPAIOR*, 2020.

[56] T. Serra, X. Yu, A. Kumar, and S. Ramalingam. Scaling up exact neural network compression by ReLU stability. In *NeurIPS*, 2021.

[57] Sidak Pal Singh and Dan Alistarh. WoodFisher: Efficient second-order approximation for neural network compression. In *NeurIPS*, 2020.

[58] G. Sourek and F. Zelezny. Lossless compression of structured convolutional models via lifting. In *ICLR*, 2021.

[59] Samuil Stoychev and Hatice Gunes. The effect of model compression on fairness in facial expression recognition. *arXiv preprint arXiv:2201.01709*, 2022.

[60] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in NLP. In *ACL*, 2019.

[61] Ruoyu Sun, Dawei Li, Shiyu Liang, Tian Ding, and Rayadurgam Srikant. The global landscape of neural networks: An overview. *IEEE Signal Processing Magazine*, 37(5):95–108, 2020.

[62] Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*, 2019.

[63] H. Tanaka, D. Kunin, D. Yamins, and S. Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. In *NeurIPS*, 2020.

[64] Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Well-read students learn better: The impact of student initialization on knowledge distillation. *arXiv preprint arXiv:1908.08962*, 13, 2019.

[65] C. Wang, R. Grosse, S. Fidler, and G. Zhang. EigenDamage: Structured pruning in the Kronecker-factored eigenbasis. In *ICML*, 2019.

[66] C. Wang, G. Zhang, and R. Grosse. Picking winning tickets before training by preserving gradient flow. In *ICLR*, 2020.

[67] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers. *Advances in Neural Information Processing Systems*, 33:5776–5788, 2020.

[68] Dongxian Wu and Yisen Wang. Adversarial neuron pruning purifies backdoored deep models. In *NeurIPS*, 2021.

[69] Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-MNIST: a novel image dataset for benchmarking machine learning algorithms. *arXiv*, 1708.07747, 2017.

[70] X. Xing, L. Sha, P. Hong, Z. Shang, and J. Liu. Probabilistic connection importance inference and lossless compression of deep neural networks. In *ICLR*, 2020.

[71] Canwen Xu, Wangchunshu Zhou, Tao Ge, Furu Wei, and Ming Zhou. Bert-of-theseus: Compressing bert by progressive module replacing. *arXiv preprint arXiv:2002.02925*, 2020.

[72] Canwen Xu, Wangchunshu Zhou, Tao Ge, Ke Xu, Julian McAuley, and Furu Wei. Beyond preserved accuracy: Evaluating loyalty and robustness of bert compression. *arXiv preprint arXiv:2109.03228*, 2021.

[73] R. Yu, A. Li, C. Chen, J. Lai, V. Morariu, X. Han, M. Gao, C. Lin, and L. Davis. NISP: Pruning networks using neuron importance score propagation. In *CVPR*, 2018.

[74] Xin Yu, Thiago Serra, Srikumar Ramalingam, and Shandian Zhe. The combinatorial brain surgeon: Pruning weights that cancel one another in neural networks. In *ICML*, 2022.

[75] Xin Yu, Zhiding Yu, and Srikumar Ramalingam. Learning strict identity mappings in deep residual networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.

[76] W. Zeng and R. Urtasun. MLPrune: Multi-layer pruning for automated neural network compression. 2018.

[77] Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. Understanding deep learning requires rethinking generalization. In *ICLR*, 2017.

[78] Wangchunshu Zhou, Canwen Xu, and Julian McAuley. Meta learning for knowledge distillation. *arXiv preprint arXiv:2106.04570*, 2021.

Checklist

The checklist follows the references. Please read the checklist guidelines carefully for information on how to answer these questions. For each question, change the default **[TODO]** to [Yes] , [No] , or [N/A] . You are strongly encouraged to include a **justification to your answer**, either by referencing the appropriate section of your paper or providing a brief inline description. For example:

- Did you include the license to the code and datasets? [Yes] See Instruction
- Did you include the license to the code and datasets? [No] The code and the data are proprietary.
- Did you include the license to the code and datasets? [N/A]

Please do not modify the questions and only use the provided macros for your answers. Note that the Checklist section does not count towards the page limit. In your paper, please delete this instructions block and only keep the Checklist section heading above along with the questions/answers below.

1. For all authors...
   (a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
   (b) Did you describe the limitations of your work? [Yes]
   (c) Did you discuss any potential negative societal impacts of your work? [Yes]
   (d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
   (a) Did you state the full set of assumptions of all theoretical results? [Yes]
   (b) Did you include complete proofs of all theoretical results? [Yes]
3. If you ran experiments...
   (a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes] See the supplemental material.
   (b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes] See Section 6.1.
   (c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes] See Section 6.1.
   (d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [Yes] See Section6.1
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
   (a) If your work uses existing assets, did you cite the creators? [Yes]
   (b) Did you mention the license of the assets? [N/A]
   (c) Did you include any new assets either in the supplemental material or as a URL? [Yes] Code
   (d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
   (e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
   (a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
   (b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
   (c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]