HOW NOT TO GIVE A FLOP

COMBINING REGULARIZATION AND PRUNING FOR EFFICIENT INFERENCE

Tai Vu, Emily Wen, Roy Nehoran
Department of Computer Science
Stanford University
{ taivu, ewen22, royn }@stanford.edu

ABSTRACT

The challenge of speeding up deep learning models during the deployment phase has been a large, expensive bottleneck in the modern tech industry. In this paper, we examine the use of both regularization and pruning for reduced computational complexity and more efficient inference in Deep Neural Networks (DNNs). In particular, we apply mixup and cutout regularizations and soft filter pruning to the ResNet architecture, focusing on minimizing floating-point operations (FLOPs). Furthermore, by using regularization in conjunction with network pruning, we show that such a combination makes a substantial improvement over each of the two techniques individually.

1 Introduction

Deep Convolutional Neural Networks (CNNs) have achieved great results and become one of the leading methods for sophisticated tasks, such as computer vision and natural language processing, in recent years (Krizhevsky et al., 2012; Simonyan & Zisserman, 2014). However, they are also very computationally expensive, which has hindered their implementation in devices with limited hardware resources. This is largely due to the neural networks being over-parameterized and having significant amounts of redundant information stored in their layers (Tzelepis et al., 2019). Furthermore, these computationally expensive DNNs consume a great deal of energy, and thus also be fiscally expensive to use. In order for DNNs to be widely utilized, it is important for their computational inefficiency to be alleviated without sacrificing accuracy. This not only reduces the amount of hardware resources needed, but also the energy and cost required to deploy them.

One popular technique harnessed in previous work for boosting model performance is regularization, which has been widely studied in the deep learning domain. Several common regularization approaches are parameter norm penalty, early stopping, data augmentation and the addition of noise during training time. An example of this is mixup, a data augmentation method that adds convex combinations of pairs of training instances and their labels (Zhang et al., 2017). By training neural networks on such combinations of inputs, mixup improves the generalization errors of cutting-edge network architectures, alleviates the memorization of corrupt labels, and increases the sensitivity to adversarial examples. Likewise, the cutout approach in DeVries & Taylor (2017) randomly masks out contiguous sections of the input images to CNNs, thereby improving robustness and yielding higher accuracy level (DeVries & Taylor, 2017). Alternatively, the technique dropout drops, or temporarily removes, random units in a neural network at intermediate layers (Srivastava et al., 2014). However, previous works have emphasized the use of regularization for preventing overfitting and guaranteeing that the resulting models generalize well to unseen data. They seemed to overlook the computational complexity of such regularized models in the inference process. Particularly, few papers have tested whether regularization speeds up the calculations of CNNs or introduces computational overhead during inference time.

Another feasible measure for high-performance model inference is network pruning, which aims to reduce the number of parameters. Recent research has harnessed magnitude-based pruning by eliminating all weights of low magnitude and retraining the networks to fine tune the remaining weights (Han et al., 2015a; 2016). This approach has been shown to significantly compress the size of common CNN architectures like AlexNet and VGG-16 with minimal loss of predictive accuracy. Nevertheless, most papers have focused on limiting the memory usage of deep learning

models and embedding them on mobile devices. There is little research on using network pruning for reducing floating-point calculations, improving computational processes, and minimizing energy consumption. Furthermore, most works have not examined the combination of pruning and advanced regularization methods like mixup and cutout to attain better performance. Since mobile phones also have limited processing powers, looking into this problem could improve the ability to run complex neural networks on them in the future.

In this paper, we examine the use of both regularization and pruning for reduced computational complexity and more efficient inference in Deep Neural Networks (DNNs). In particular, we apply two regularization methods called mixup and cutout and a pruning technique named soft filter pruning to the ResNet architecture, focusing on minimizing floating-point operations (FLOPs). Furthermore, by utilizing regularization in conjunction with pruning unimportant convolutional filters below some threshold in the network, we show that the combination contributes to a marked improvement over either of the two methods individually.

Our main focus in this paper is looking into whether pruning methods and regularization techniques could work in conjunction, rather than interfere with each other, to reduce computational inefficiency and keep accuracy rates comparable. On top of that, we determine the impact of applying regularization for the purpose of improving efficiency, a context that has been relatively unexplored. Along with this, we demonstrate that combining the methods produces much higher accuracies than just pruning. Overall, through this paper, we innovate a method that allows ResNet, and potentially other similar architectures, to be used in a less computationally expensive fashion. This more efficient implementation also results in important reductions in the energy, and thus cost, to deploy DNNs.

2 RELATED WORK

Deep Neural Networks (DNNs) have achieved great outcomes and become one of the cuttingedge technologies for sophisticated tasks such as computer vision, natural language processing, and speech recognition. However, they are currently difficult to implement for mobile applications and embedded systems because they require a huge amount of storage. As a result, creating DNNs that are smaller and more efficient is very useful to make them more widely accessible for use within mobile applications.

2.1 PRUNING

Some research has illustrated that there is significant redundancy in the parameterization of many neural networks (Denil et al., 2013), which consumes memory resources and slows down computation speeds. Thus, many research papers have emphasized reducing the number of parameters as a feasible way of compressing deep learning models and reducing their complexity.

While limiting the quantity of parameters proves to be a viable approach, most recent studies have been focused on popular network architectures including fully connected networks and CNNs. There is little research on less familiar structures like RNNs, LSTM, GANs, and so on. Additionally, many research papers have tested their models on classic datasets like MNIST and ImageNet and simple CNNs like LeNet or AlexNet, with little consideration of other datasets and more advanced architectures such as ResNet, GoogleNet, UNet, DenseNet, etc.

While pruning parameters can reduce the sizes of the models, thereby minimizing memory usage and energy consumption, they do not necessarily increase computation speeds. In fact, the majority of the parameters removed are from the fully connected layers where the computation costs are low, whereas convolutional operations tend to dominate computation complexity (Li et al., 2016). Furthermore, high-performance inference on pruned models that contain sparse connections typically requires purpose-built hardware for performing sparse matrix-vector operations. The representation of sparse matrices also results in additional storage overheads, which raises the networks net memory footprint (Zhu & Gupta, 2017).

2.2 REGULARIZATION

Because convolutional neural networks tackle complex learning problems, they are large and thus susceptible to overfitting. Regularization techniques can mitigate this issue and thus boost accuracy rates. Moreover, this practice can counteract the decrease in accuracy resulting from using pruning. Some common regularization approaches are parameter norm penalty, early stopping, data augmentation, and the addition of noise during training time.

Dropout is a common regularization technique utilized for fully-connected layers. It is implemented by dropping, or temporarily removing, random units in a neural network (Srivastava et al., 2014). This is comparable to adding noise to a model's hidden units. While dropout removes parameters at intermediate layers, the technique cutout randomly occludes sections of the input images (DeVries & Taylor, 2017). By doing so, the model is able to focus more on less prominent features of the image rather than important features, i.e. those with high activations in intermediate layers of the network. This practice improves robustness and yields higher accuracy levels. Another variation of regularization, mixup, applies data augmentation (Zhang et al., 2017). It enlarges the training dataset by taking the linear interpolations of two examples and their labels from the training data. By including these constructed examples in training, the neural network becomes more robust against adversarial examples and lessens the memorization of corrupt labels. Overall, mixup improves the generalization errors of these models on popular datasets.

Additionally, Structured Sparsity Learning (SSL) combines structural regularization and locality optimization to improve accuracy on large networks (Wen et al., 2016). It employs group Lasso regularization to alter the structure of the filters, channels, filter shapes, and depth structures within DNNs to construct compressed networks. By doing so, it directly learns a more compact representation of larger DNNs and reduces computational costs.

Previous works have emphasized the use of regularization for preventing overfitting and guaranteeing that the resulting models generalize well to unseen data. Furthermore, the impact of combining pruning methods and regularization techniques has not been thoroughly explored by previous studies, but will be scrutinized in this paper.

3 METHODS

3.1 REGULARIZATION

We harnessed mixup, a data augmentation method that adds convex combinations of pairs of examples and their labels to the training data (Zhang et al., 2017). Specifically, this technique uses new training examples of the form:

$$\tilde{x} = \lambda x_i + (1 - \lambda) x_j$$
, where x_i, x_j are input vectors. $\tilde{y} = \lambda y_i + (1 - \lambda) y_j$, where y_i, y_j are one-hot label encodings.

Here, (x_i, y_i) and (x_j, y_j) are two instances randomly drawn from the training set, and $\lambda \in [0, 1]$. The purpose of mixup is to teach the convolutional neural network to map a linear combination of inputs to a linear combination of outputs. Hence, by training neural networks on such combinations of inputs, mixup encourages the linear behaviors of the models, improves the generalization errors of cutting-edge network architectures, alleviates the memorization of corrupt labels, and increases sensitivity to adversarial examples (Zhang et al., 2017).

In addition, we considered the cutout approach proposed in DeVries & Taylor (2017). This method randomly masks out contiguous square regions of the input images to CNNs, thereby augmenting

```
# y1, y2 should be one-hot vectors

for (x1, y1), (x2, y2) in zip(loader1, loader2):
    lam = numpy.random.beta(alpha, alpha)
    x = Variable(lam * x1 + (1. - lam) * x2)
    y = Variable(lam * y1 + (1. - lam) * y2)
    optimizer.zero_grad()
    loss(net(x), y).backward()
    optimizer.step()

(a) One epoch of mixup training in PyTorch.

ERM

mixup

(b) Effect of mixup (\alpha = 1) on a toy problem. Green: Class 0. Orange: Class 1. Blue shading indicates p(y = 1|x).
```

Figure 1: Illustration of mixup, from Zhang et al. (2017).

the training data with partially occluded versions of existing samples. In particular, cutout applies a fixed-size zero-mask to a random portion of each input during each epoch. In this way, the technique encourages the neural network to take more of the full image context into account instead of relying on specific visual features. Thus, cutout improves robustness and yields higher accuracy levels (DeVries & Taylor, 2017).



Figure 2: Cutout applied to images from the CIFAR 10 dataset, from DeVries & Taylor (2017).

3.2 PRUNING

We made use of soft filter pruning, a technique that simultaneously prunes unimportant filters and trains the model (He et al., 2018). By simultaneously pruning and training the model, this method allows almost all of the layers to be pruned at once rather than pursuing the greedy layer by layer pruning methodology. More specifically, this algorithm can be broken down into four main parts.

First, we select filters from each layer that will be pruned. The number of filters chosen is equal to the product of the number of total filters and a pruning hyperparameter that is used for every layer. In particular, we set the pruning rate hyperparameter to be 0.9. Within a layer, a filter is selected with probability proportional to its l_p norm. After selection, these filters are set to 0, i.e. are pruned. Then, the network is trained for one epoch, where the pruned filters may be updated to be non-zero due to backpropagation. These previous steps are repeated until convergence. At convergence, the zero filters are removed from the network. This generates the compact model.

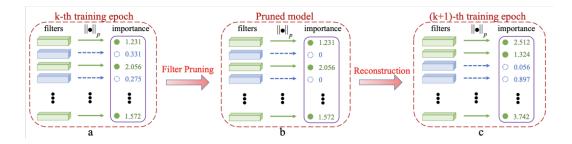


Figure 3: Steps of soft filter pruning, from He et al. (2018).

4 IMPLEMENTATION

We leveraged PyTorch and PyTorch Lightning for implementing and training the models (Facebook, 2017a; Falcon et al., 2019). Additionally, we harnessed Weights & Biases for saving and visualizing the results (Wandb, 2019).

Link to our project:

https://github.com/taivu1998/ResNet-Regularization-Pruning

5 EXPERIMENTS

Pruning methods and regularization techniques have been used independently to improve the computational efficiency of ResNet models. However, there has not been much work done in determining the computational complexity of the models when both techniques are implemented together. We show that on the ResNet architecture and the CIFAR 10 dataset, pruning methods and regularization techniques implemented together work in conjunction, rather than interfere with each other, to reduce computational inefficiency more than when either is implemented independently.

5.1 EXPERIMENT DESIGN

To investigate the impact of using pruning and regularization together, we considered four different types of networks:

- 1. Control network: This is the standard ResNet architecture, with no additions. We will be testing this model in order to have a point of comparison with the regularized (2) and pruning (3) networks. Using this, we will be able to isolate the effect that these methods have on the control network individually to gain better insight on our results from the combination.
- 2. Network with regularization: This network has either mixup or cutout on top of the control network, with all other parameters fixed. Using this network, we will be able to directly compare to the combined (4) network, in order to determine whether achieves an equivalent (±5%) accuracy with a lower FLOP count.
- 3. Network with pruning: This is the control network with soft filter pruning. We will be able to use this network to also compare to the combined (4) network, and whether we can achieve better results.
- 4. Network with both pruning and regularization: This network has both regularization (mixup or cutout) and soft filter pruning combined. The goal of this network is to reduce the number of FLOPs from the regularization (2) network, without decreasing accuracy by more than 5%.

For our control networks, we implemented ResNet 20, ResNet 32, ResNet 44, ResNet 56, and ResNet 110 based on Facebook's models (Facebook, 2017b). We considered two regularization techniques, mixup and cutout, and trained ten different types of networks with regularization: either mixup or cutout applied on each of the five control models (DeVries & Taylor, 2017; Zhang et al.,

2017). For pruning, we harnessed soft filter pruning (He et al., 2018). Each of the five control models then had pruning added. Lastly, we constructed our fourth class of networks by combining soft filter pruning and each of mixup and cutout to each of the control models.

5.2 Dataset

We used the CIFAR 10 dataset, which has 60,000 color images (Krizhevsky, 2009). Each of them is an RGB image of size $32 \times 32 \times 3$. These images are equally split into ten different classes: airplane, automobile, bird, cat, deer, dog, frog, horse, ship, and truck. From these examples, 10,000 of them are allocated as test images. The goal of these networks is to correctly allocate as many of these pictures into the correct class of the ten. The accuracy score is determined by dividing the number of correct classifications over the total number of images in the test set.

5.3 EVALUATION

When implementing pruning and regularization techniques to our model, we say that our model does not affect accuracy if the accuracy is at most 5% lower than that of the control network with the same size. There are several other variables that we controlled for, including the number of epochs each model was run for, the learning rate, and the dataset we were testing on. Specifically, we trained our networks on the CIFAR 10 data for 200 epochs with a learning rate of 0.1. By having these control variables, we are able to minimize the differences between our models and ensure that our results are the effect of changes in our independent variable.

Our dependent variable is the computational efficiency of our ResNet-based models on the CIFAR 10 dataset. We determined the computational efficiency of each by calculating the number of floating-point operations (FLOPs). FLOPs is the number of floating point operations carried out by a model, and is a proxy for measuring its computational efficiency. It solely depends on the architecture of the model, with no regards to how the model is trained. We calculated FLOPs manually based on the channel and width of each layer of the network using the method from He et al. (2018). FLOPs allow us to precisely determine the computational cost of a model in a consistent way. The goal of our research is to reduce the cost of running deep neural networks. By measuring FLOPs and working on reducing them, we think that the results of our research can be used to limit the cost of running DNNs in real industry use cases. Throughout our paper we refer to the FLOP count of a model as its cost.

Our independent variable is which technique is implemented on the standard ResNet model (Facebook, 2017b). More specifically, we had ResNet models with the following methods implemented on top of them: mixup, cutout, soft filter pruning, mixup and soft filter pruning combined, and cutout and soft filter pruning combined. For each of these cases we also altered the number of layers in the ResNet architecture (20, 32, 44, 56, 110) to get a broader view of how these techniques affect the different network sizes. We show consistent results across all network sizes to help make our results more robust and more widely applicable to more use cases. This helps minimize the risks related to the effects of certain network sizes on the metrics that we are measuring.

The goal of our evaluation is to illustrate that we can reduce the computational complexity of artificial neural networks during the inference phase without any significant loss of accuracy. In particular, for each type of the ResNet structure, we show that a combination of regularization and pruning can decrease its number of FLOPs while keeping its accuracy level comparable. To accomplish this, we calculated both the FLOP count and the accuracy score of every combination of regularization and pruning.

We expect that for each control model (i.e. ResNet 20, 32, 44, 56, 110), adding regularization will improve accuracy and leave the FLOP count constant, while adding pruning will decrease both FLOPs and accuracy. The goal of the research is to show that the fourth class of networks reduces the number of FLOPs without sacrificing accuracy. We tested four different sizes of each of the five network types (20, 32, 44, 56, 110), in order to get a curve of five different points of FLOPs and accuracy. Using this, we can show how the curves move on the graph with network changes. Our goal is for the (4) curve to have both lower FLOP counts than (2) and higher accuracies than (1) and (3) across all network sizes.

Table 1: Results: Accuracy and MegaFLOPs of all methods and network sizes.

Method	Model	MegaFLOPs	Accuracy (%)
Control	ResNet 20 40.55		91.63
Control	ResNet 32 68.86		92.11
Control	ResNet 44	97.17	92.54
Control	ResNet 56	125.49	92.49
Control	ResNet 110	252.89	92.58
Mixup	ResNet 20	40.55	92.67
Mixup	ResNet 32	68.86	93.39
Mixup	ResNet 44	97.17	94.16
Mixup	ResNet 56	125.49	94.15
Mixup	ResNet 110	252.89	94.71
Cutout	ResNet 20	40.55	93.00
Cutout	ResNet 32	68.86	93.15
Cutout	ResNet 44	97.17	93.10
Cutout	ResNet 56	125.49	93.01
Cutout	ResNet 110	252.89	92.76
Pruning	ResNet 20	34.37	91.09
Pruning	ResNet 32	58.58	91.40
Pruning	ResNet 44	82.78	92.23
Pruning	ResNet 56	106.99	91.42
Pruning	ResNet 110	215.92	91.88
Mixup & Pruning	ResNet 20	34.37	91.94
Mixup & Pruning	ResNet 32	58.58	93.06
Mixup & Pruning	ResNet 44	82.78	93.12
Mixup & Pruning	ResNet 56	106.99	93.80
Mixup & Pruning	ResNet 110	215.92	93.04
Cutout & Pruning	ResNet 20	34.37	92.87
Cutout & Pruning	ResNet 32	58.58	93.28
Cutout & Pruning	ResNet 44	82.78	94.12
Cutout & Pruning	ResNet 56	106.99	94.52
Cutout & Pruning	ResNet 110	215.92	94.57

6 RESULTS

6.1 Control

We ran all five sizes of the control network, which are recorded in Table 1 and visualized in Figure 4. We can see that these models have decent performances on the CIFAR 10 dataset, with the predictive accuracies ranging from 91.63% to 92.58%. These statistics are similar to those in He et al. (2016).

When it comes to computational cost, the number of FLOPs increases as the size of the network rises, with 40.6 megaFLOPs for ResNet 20 and 252.9 megaFLOPs for ResNet 110. This result is completely reasonable because larger networks have more weights and require more computations when processing a certain input. We have verified that the FLOP counts we obtained on our control models (i.e. with no techniques implemented) are consistent with the calculations from He et al. (2018) in addition to those from another paper by Li et al. (2016).

Meanwhile, it is evident that the accuracy level tends to go up with increasing network sizes, but the increase is no more than 1%. This means that it is not worthwhile to utilize a network larger than ResNet 32, because beyond that, the cost rises dramatically without a considerable accuracy increase. For the rest of the analysis, we have continued running on all network sizes, in order to observe whether our results apply in other cases.

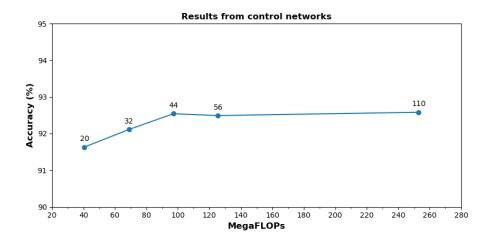


Figure 4: Results from control networks. Five different labeled network sizes of ResNet with corresponding MegaFLOPs and accuracy when run on CIFAR 10.

Method	Model	Increase in Accuracy (%)
Mixup	ResNet 20	1.04
Mixup	ResNet 32	1.28
Mixup	ResNet 44	1.62
Mixup	ResNet 56	1.66
Mixup	ResNet 110	2.13
Cutout	ResNet 20	1.37
Cutout	ResNet 32	1.04
Cutout	ResNet 44	0.56
Cutout	ResNet 56	0.52
Cutout	ResNet 110	0.18

Table 2: Increase in accuracy of regularized vs. control networks.

6.2 REGULARIZATION

When implementing mixup and cutout on all five of our control models (Table 1 and Figure 5), there are consistent increases in accuracy level as shown in Table 2. In fact, all cutout models undergo accuracy increases of 0.18-1.37%. Likewise, mixup seems to have a greater effect, as it raises the accuracy scores of all five networks by 1.04-2.13%. These figures illustrate that both mixup and cutout are highly effective at improving the generalization error and enhancing the accuracy level of the ResNet architecture. While the goal of this paper is to reduce FLOPs while keeping accuracy consistent, this accuracy increase could help overcome decreases that could result from combining this with pruning.

Additionally, for mixup, the improvement tends to become more significant with increasing network sizes. Specifically, ResNet 110 experiences a 2.13% accuracy increase by adding mixup. Therefore, the mixup models show an upward trend in accuracy levels as we go from ResNet 20 to ResNet 110. However, with regards to cutout, larger networks demonstrate less remarkable effects, with the accuracy increases of 1.37% and 0.18% for ResNet 20 and ResNet 110 respectively. At the same time, all five cutout models have comparable performances with no more than 0.4% differences in accuracy

Furthermore, the addition of regularization does not affect the FLOP counts of the models. This fact is to be expected, as regularization just changes the values of, not the number of parameters in a neural network.

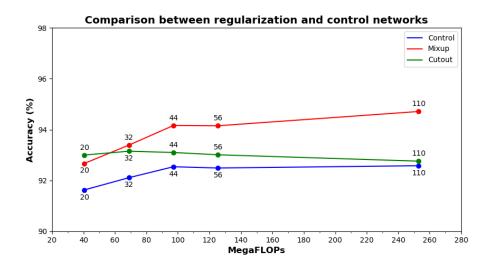


Figure 5: Results from regularized networks compared with control networks. Five different labeled network sizes of ResNet with corresponding MegaFLOPs and accuracy when run on CIFAR 10.

Model	Decrease in Accuracy (%)	Decrease in FLOPs (%)
ResNet 20	0.16	15.2
ResNet 32	0.71	14.9
ResNet 44	0.31	14.8
ResNet 56	1.07	14.7

14.6

0.70

Table 3: Pruning effect on accuracy and FLOPs.

6.3 PRUNING

ResNet 110

As shown in Figure 6 and Table 3, soft filter pruning reduces the accuracy scores by 1.07% in the worst case across all five network sizes, which is much less than the threshold that we set of a maximum 5% decrease. This is a minimal accuracy decrease that may be recovered by combining this technique with regularization. More interestingly, we can see that the effect of pruning becomes more significant as we go from ResNet 20 to ResNet 44, before the accuracy decreases. One possible explanation is that for ResNet 44, filter pruning sharply reduces the numbers of parameters and prevents overfitting. Larger networks like ResNet 56 and ResNet 110 still have a vast quantity of weights after being pruned, so they are still prone to overfitting and generalize less competently to unseen data.

With regards to FLOP counts, all 5 model sizes experience roughly 15% declines (Table 3). These are sharp decreases in the computational costs that do not significantly affect the accuracies. Thus, soft filter pruning shows promise for a method that could help achieve the goal of creating DNNs that perform more efficient inference.

6.4 COMBINED

The results of adding both soft filter pruning and regularization to the baseline models are illustrated in Table 4, Figure 7, and Figure 8. We expected the predictive accuracies of the combined models to increase back to the range of just the regularization networks.

As for the models with mixup and pruning implemented, their accuracy scores are consistently higher than those of the pruning networks and the control networks, with 0.31 - 2.38% differences.

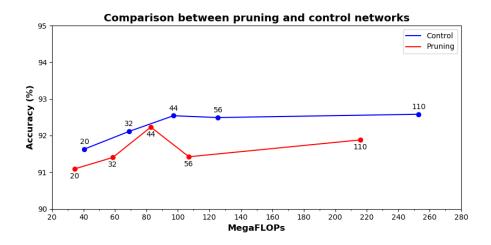


Figure 6: Results from pruning networks compared with control networks. Five different labeled network sizes of ResNet with corresponding MegaFLOPs and accuracy when run on CIFAR 10.

		•	
Method	Model	Increase in Accuracy (%)	Decrease in FLOPs (%)
Mixup & Pruning	ResNet 20	0.31	15.2
Mixup & Pruning	ResNet 32	0.95	14.9
Mixup & Pruning	ResNet 44	0.58	14.8
Mixup & Pruning	ResNet 56	1.31	14.7
Mixup & Pruning	ResNet 110	0.46	14.6
Cutout & Pruning	ResNet 20	1.24	15.2
Cutout & Pruning	ResNet 32	1.17	14.9
Cutout & Pruning	ResNet 44	1.58	14.8
Cutout & Pruning	ResNet 56	2.03	14.7
Cutout & Pruning	ResNet 110	1.99	14.6

Table 4: Increase in accuracy of combined vs. control networks.

These accuracy levels are also slightly lower than the figures for the mixup models, with decreases of no more than 1.67% (less than the threshold of 5%).

In terms of the models with cutout and implemented, there results are even more notable. Indeed, there accuracy scores also exceed those of the pruning and control models by about 1.17-3.10%. More remarkably, their accuracy levels tend to surpass the figures for the cutout models, with rises of more than 1.81% for ResNet 32, 44, 44, 56, and 110. In addition, the effect of combining cutout and pruning becomes increasingly significant for larger networks, which is demonstrated by an upward trend in the accuracy levels of these combined models.

When it comes to the computational costs, all ten networks have the same FLOP counts as the corresponding pruning models do. This is to be expected because regularization does not affect the costs. In other words, they show approximately 15% reductions in the numbers of FLOPs with trivial impacts on accuracy.

Overall, based on the aforementioned statistics, we learn that the combination of regularization and pruning techniques results in compressed models that achieve lower computational costs and higher accuracy scores. Therefore, our proposal that the two techniques can effectively combine without conflicting is validated by the data.

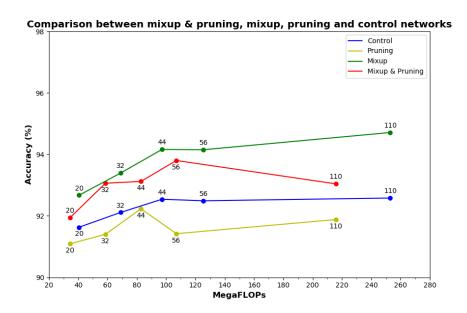


Figure 7: Results from combined methods (mixup and pruning) compared to mixup, pruning and control individually. Five different labeled network sizes of ResNet with corresponding MegaFLOPs and accuracy when run on CIFAR 10.

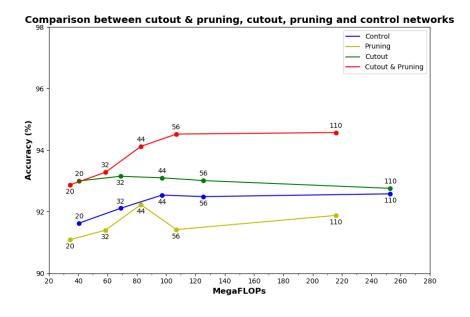


Figure 8: Results from combined methods (cutout and pruning) compared to cutout, pruning and control individually. Five different labeled network sizes of ResNet with corresponding MegaFLOPs and accuracy when run on CIFAR 10.

7 DISCUSSION

For each given network size, the accuracy of the combined model is greater than the figures for the control network and the network with pruning. However, the number of FLOPs of the combined model is the same as that of the model with pruning and much less than that of the control network. Hence, pruning is a complement to regularization, and combining them produces a superior performance with a good tradeoff between accuracy and computational cost.

It is interesting to note that the FLOP count decreases by roughly 15% when pruning is implemented. This large decrease is particularly surprising given that the pruning rate was set to 0.9, i.e. 10% of the parameters were pruned from each layer every time. Hence, we can see that soft filter pruning is particularly effective for reducing the computational complexity of ResNet during inference time.

In the future it would be interesting to conduct additional experiments with varied pruning rates such as 0.8 and 0.7. Perhaps a smaller compression rate will result in a model with fewer learnable parameters, which is less suspectible to overfitting. Hence, such a model may even have a higher accuracy level with a lower FLOP count.

In addition, we will apply a combination of regularization and pruning to larger models sizes like ResNet 18 and ResNet 34 and larger datasets like ImageNet. With more training data and larger networks that contain much more parameters, the effects of mixup, cutout, and soft filter pruning may become more significant. It is also interesting to run experiments with other popular network architectures like AlexNet, VGGNet, GoogleNet, UNet, and DenseNet and observe whether our conclusion applies to these models.

Furthermore, we are currently evaluating our models almost solely on their numbers of FLOPs, with the criteria that their accuracies must be above the specified threshold. We would also like to construct a metric that is better able to take into account percent changes in both cost and accuracy. This would allow us to develop more nuanced insights into the results of the combined models compared to the models with just regularization or pruning.

8 Conclusion

In this paper we proposed that combining pruning and regularization methods would perform better than the baseline and either technique individually. In particular, we define the performance by the reduction in the cost, or the FLOP count, of the model from the control network of the same size, such that the change in accuracy is minimal (< 5% decrease).

In comparison with the control network, models with regularization implemented increase the accuracies, while those with soft filter pruning decrease the costs. Combining both the regularization and pruning techniques results in models with higher accuracies than those with just pruning. Meanwhile, compared to models with only regularization, these combined models have both significantly lower costs (about 15%) and slightly lower (or even higher) accuracies (up to 2%). Consequently, these results substantiate our hypothesis.

Further research could look into different regularization and pruning techniques and other ways of combining pruning and regularization that can boost the performance of artificial neural networks during inference time.

REFERENCES

- Wenlin Chen, James Wilson, Stephen Tyree, Killian Weinberger, and Yixin Chen. Compressing neural networks with the hashing trick. 2015.
- Maxwell D Collins and Pushmeet Kohli. Memory bounded deep convolutional networks. *arXiv* preprint arXiv:1412.1442, 2014.
- Misha Denil, Babak Shakibi, Laurent Dinh, Marc'Aurelio Ranzato, and Nando De Freitas. Predicting parameters in deep learning. In *Advances in neural information processing systems*, pp. 2148–2156, 2013.
- Emily L. Denton, Wojciech Zaremba, Joan Bruna, Yann LeCun, and Rob Fergus. Exploiting linear structure within convolutional networks for efficient evaluation. 2014.
- Terrance DeVries and Graham W. Taylor. Improved regularization of convolutional neural networks with cutout. 2017.
- Facebook. Pytorch, 2017a. URL https://pytorch.org.
- Facebook. Fb resnet torch, 2017b. URL https://github.com/facebook/fb.resnet.torch.
- William Falcon, Jirka Borovec, Ethan Harris, , Matthew Painter, Justus Schock, Nick Eggert, Jeff Ling, Jeremy Jordan, and Tullie Murrell. Pytorch lightning, 2019. URL https://github.com/PytorchLightning/pytorch-lightning.
- Yunchao Gong, Liu Liu, Ming Yang, and Lubomir Bourdev. Compressing deep convolutional networks using vector quantization. 2014.
- Song Han, Huizi Mao, and William J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. 2015a.
- Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*, pp. 1135–1143, 2015b.
- Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A. Horowitz, and William J. Dally. Eie: efficient inference engine on compressed deep neural network. 2016.
- Babak Hassibi and David G Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in neural information processing systems*, pp. 164–171, 1993.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866*, 2018.
- Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network. 2015.
- Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and; 0.5 mb model size. arXiv preprint arXiv:1602.07360, 2016.
- Animesh Koratana, Daniel Kang, Peter Bailis, and Matei Zaharia. Lit: Block-wise intermediate representation training for model compression. 2018.
- Alex Krizhevsky. Learning multiple layers of features from tiny images, 2009. URL https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, pp. 1097–1105, 2012.

- Yann LeCun, John S Denker, and Sara A Solla. Optimal brain damage. In *Advances in neural information processing systems*, pp. 598–605, 1990.
- Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. arXiv preprint arXiv:1608.08710, 2016.
- Lyken17. Thop: Pytorch-opcounter, 2019. URL https://github.com/Lyken17/pytorch-OpCounter.
- Steven J. Nowlan and Geoffrey E. Hinton. Simplifying neural networks by soft weight-sharing. 1992.
- Stuart F Oberman and Michael J Flynn. Design issues in division and other floating-point operations. *IEEE Transactions on computers*, 46(2):154–161, 1997.
- Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov. Dropout: A simple way to prevent neural networks from overfitting., 2014.
- Vivienne Sze, Yu-Hsin Chen, Tien-Ju Yang, and Joel S Emer. Efficient processing of deep neural networks: A tutorial and survey. *Proceedings of the IEEE*, 105(12):2295–2329, 2017.
- Georgios Tzelepis, Ahraz Asif, Saimir Baci, Selcuk Cavdar, and Eren Erdal Aksoy. Deep neural network compression for image classification and object detection. arXiv preprint arXiv:1910.02747, 2019.
- Wandb. Weights & biases, 2019. URL https://www.wandb.com.
- Ji Wang, Weidong Bao, Lichao Sun, Xiaomin Zhu, Bokai Cao, and S. Yu Philip. Private model compression via knowledge distillation. 2019.
- Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. 2016.
- Samuel Williams, Andrew Waterman, and David Patterson. Roofline: An insightful visual performance model for floating-point programs and multicore architectures. Technical report, Lawrence Berkeley National Lab.(LBNL), Berkeley, CA (United States), 2009.
- Zichao Yang, Marcin Moczulski, Misha Denil, Nando de Freitas, Alex Smola, Le Song, and Ziyu Wang. Deep fried convnets. 2015.
- Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. 2017.
- Chenzhuo Zhu, Song Han, Huizi Mao, and William J. Dally. Trained ternary quantization. 2016.
- Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. *arXiv* preprint arXiv:1710.01878, 2017.