

Deep Anchored Convolutional Neural Networks

Jiahui Huang Kshitij Dwivedi Gemma Roig
Singapore University of Technology and Design

jiahui.huang@sutd.edu.sg, kshitij.dwivedi@mymail.sutd.edu.sg, gemma.roig@sutd.edu.sg

Abstract

Convolutional Neural Networks (CNNs) have been proven to be extremely successful at solving computer vision tasks. State-of-the-art methods favor such deep network architectures for its accuracy performance, with the cost of having massive number of parameters and high weights redundancy. Previous works have studied how to prune such CNNs weights. In this paper, we go to another extreme and analyze the performance of a network stacked with a single convolution kernel across layers, as well as other weights sharing techniques. We name it Deep Anchored Convolutional Neural Network (DACNN). Sharing the same kernel weights across layers allows to reduce the model size tremendously, more precisely, the network is compressed in memory by a factor of L , where L is the desired depth of the network, disregarding the fully connected layer for prediction. The number of parameters in DACNN barely increases as the network grows deeper, which allows us to build deep DACNNs without any concern about memory costs. We also introduce a partial shared weights network (DACNN-mix) as well as an easy-plug-in module, coined regulators, to boost the performance of our architecture. We validated our idea on 3 datasets: CIFAR-10, CIFAR-100 and SVHN. Our results show that we can save massive amounts of memory with our model, while maintaining a high accuracy performance.

1. Introduction

Since the famous AlexNet [17] outperformed all its competitors on ILSVRC2012 challenge [3] in 2012, Convolutional Neural Networks (CNNs) dominated almost all other approaches in computer vision tasks in the past 6 years [30, 27, 28, 31]. He *et al.* [11] proposed Residual Networks, which allows to build extremely deep CNNs while still keep them optimizeable. Since then, the general trend is that CNNs have grown deeper and wider to achieve better performance [2, 32, 12]. As a result, current deep networks often come with vast amount of parameters and highly redundant weights [4], while the performance

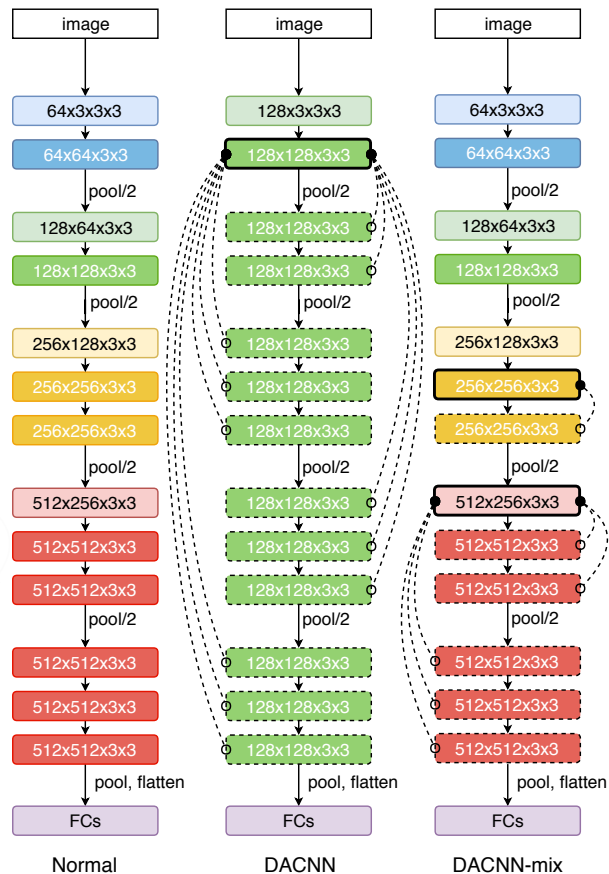


Figure 1. Example of a network with 13 convolutional layers. Kernels connected by dash lines represents they share the same set of weights, color codes denotes different kernel dimensions. **Left:** Normal Architecture (14.9 million parameters) as reference. **Middle:** Plain DACNN Architecture (0.16 million parameters). **Right:** Mixed DACNN (4.73 million parameters). For more detailed model comparison, please refer to Table 6.

gained is very limited compared to the number of parameters increased. For instance, ResNet-101 has only 1.1% gain in accuracy compared to ResNet-50 [11] on ImageNet classification task, while the number of parameters almost doubled. This unbalance between model size increment and

performance boost has become a severe problem yet to be tackled.

To compress CNNs, network weights pruning techniques have been introduced to remove some of the unnecessary filters [9, 10, 8, 13, 21, 29, 33]. The aim of such techniques is to have smaller models without compromising the accuracy performance. Hao *et al.* [21] proposed to prune weights according to their summed absolute weights. Huang *et al.* [13] introduced a pruning agent to help analyze which filters are to be removed. These pruning methods adapt a “subtraction” fashion to reduce the number of parameters, which means that once the original architecture is set, the performance of the model can barely increase since they consist on cutting off filters from it. If one needs further improvements on the model, the only way is to rebuild the original architecture and re-run the pruning algorithm again. Moreover, as far as we know, these techniques are not in-built in existing deep learning libraries and hence not widely adopted in most of the applications.

Our work focuses on addressing network memory compression problem in an “addition” fashion. First, we propose a novel architecture that stacks a single convolution kernel over all layers (Figure 1, middle), and we extend it to partially share weights between pooling layers (Figure 1, right). We call it *deep-anchored-convolutional neural network (DACNN)*. We also introduce an “easy-plug-in” way to add few extra parameters into the DACNN base model to boost the performance. Because the number of extra parameters introduced is determined by the model designer, this method provides an easy control of the trade-off between model size and model performance. In addition, the idea is easily realizable in code and it is applicable to most of the existing architectures.

We provide a detailed analysis on different DACNN architectures, as well as discuss how to efficiently add extra parameters to DACNN to achieve better performance with high memory compression rates. We demonstrate the efficiency and efficacy of our proposed method on CIFAR-10, CIFAR-100 [16] and SVHN [6] datasets.

2. Related Works

In this section we revisit network pruning techniques, as well as specific network models, namely ShaResNet [1], SqueezeNet [14] and residual adapters [25], which are related to our work.

Network Pruning. Over the past years, network pruning has become a popular topic to compress the model size of neural networks. Han *et al.* [10] developed a method that replaces weights below a threshold with zeros. It forms a sparse matrix with less parameters, and then trains it for several iterations to achieve promising compression versus accuracy results. They further introduced quantization and

huffman encoding into their Deep Compression [9] pruning method. Huang *et al.* [13] proposed a data-driven pruning method by introducing a pruning agent to remove unnecessary CNN filters. They use reinforcement learning to train the agent to prune the network while retains the network with a desired performance. Many of these pruning methods require backbone framework modification of the model, which reduce their applicability. Some of these methods even require dedicated hardware support [8]. As a result, network pruning methods are not adopted on most of the existing DNN architectures. In addition, due to the nature of pruning, performance of the network can barely increase, and reconstruction of the base model and re-running of the pruning algorithm are required if one wants to improve the network performance.

ShaResNet. Boulch [1] proposed sharing weights among residual blocks to reduce the number of parameters without losing much performance. More concretely, a basic residual block [11] is composed of 2 convolution operations with filter size 3×3 , ShaResNet uses shared weights to replace all the second convolution kernel within blocks that operate in the same spatial resolution (between 2 pooling layers). Thus, nearly half of the parameters from convolution can be cut off. A similar technique is applied to bottleneck blocks in deeper ResNets. Despite achieving promising results, this method is not flexible, as only one convolution is shared across blocks, and it’s difficult to cut more parameters or increase performance from this architecture.

SqueezeNet. Iandola *et al.* [14] introduced SqueezeNet, which consist on replacing some of the 3×3 convolution filters with 1×1 filters as well as reducing the number of input channels to 3×3 filters. They also pushed the downsampling of activation maps towards the end of the architecture to improve accuracy. With this approach, they were able to achieve same performance compared to AlexNet [17] while using $50\times$ less parameters. Yet, their method is not applicable to very deep networks such as ResNets [11], because it requires carefully designed structure for every layer, and downsampling in later layers in deep networks greatly increases computational cost.

Residual Adapters. Residual adapters were introduced by Rebuffi *et al.* [25] as a technique for multi-task learning. They plug task-specific residual adapter modules (banks of 1×1 convolution kernels) into residual blocks of the network. For different task domains, only these adapters varies while the rest parameters (90%) remains the same. Since these 1×1 convolution kernels are relatively small in size and they helps to regulate convolutional layer expressions, we introduce them as extra parameters to DACNNs to in-

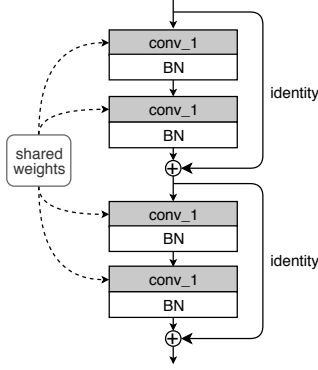


Figure 2. DACNN base block with residual learning: parameters for convolution are shared in all layers, while parameters for Batch Normalization are free.

crease performance efficiently. In our work we call these 1×1 convolution kernels *regulators*.

3. Anchored Weights Convolution

Here, we introduce to notation that we use in the rest of the paper, as well as the components of our proposed anchored weights convolution architecture.

Notation. Let us consider X_0 as an input image to the convolutional network with L number of layers. $F_l(\cdot)$ denotes the transformation of each layer l , which can be a combination of convolution [19] (weights represented by W_l^{conv}), batch normalization [15] (weights represented by W_l^{bn}), non-linear activation (ReLU [23] in our case) or pooling [20]. We refer the output of layer l as X_l .

Weights Sharing. In most of the CNN architectures, transformation function $F_l(\cdot)$ have different parameters for each layer: W_l^{conv} and W_l^{bn} . A transformation of each layer can be represented as:

$$X_l = F_l(X_{l-1}, \{W_l^{conv}, W_l^{bn}\}) \quad (1)$$

In our DACNN architecture (figure 1), we first set a constant C as the number of activation map channels for all layers, then we use one layer of transformation to expand a 3-channel-image to a desired activation map with C channels:

$$X_1 = F_1(X_0, \{W_1^{conv}, W_1^{bn}\}) \quad (2)$$

From the second layer on wards, we internalize a set of global convolution weights W_G^{conv} of shape $(C \times C \times filtersize)$, and those are applied to every transformation function throughout the entire network:

$$X_l = F_l(X_{l-1}, \{W_G^{conv}, W_l^{bn}\}) \quad (for\ l \geq 2) \quad (3)$$

In this way, only weights for the first convolution W_1^{conv} , weights for global convolution W_G^{conv} and weights for each layer’s batch normalization W_l^{bn} need to be initialized and trained, greatly reducing the total number of parameters.

Batch Normalization. Batch normalization [15] was first introduced as a technique to improve the performance and stability of deep neural networks. As we will show in the sequel, it is a crucial component in our architecture for achieving a good accuracy performance, as it allows scaling the activation map A_l of each layer:

$$X_l = W_l^{bn} A_l. \quad (4)$$

Since the parameters of batch normalization for each layer are different, we can obtain different transformation functions across layers, and thus distinguish our work from simply stacking convolution kernels. In section 4, we further discuss that scaling with batch normalization is a crucial operation in our architecture for performance.

Residual Learning. ResNet [11] is an architecture composed of residual blocks. Each block’s output is an element-wise addition of input and activation:

$$X_l = X_{l-1} + F_l(X_{l-1}, \{W_l^{conv}, W_l^{bn}\}). \quad (5)$$

In our case, we adapt the idea of residual learning by using residual blocks with shared convolution weights (Figure 2):

$$X_l = X_{l-1} + F_l(X_{l-1}, \{W_G^{conv}, W_l^{bn}\}) \quad (6)$$

In this way, we are able to increase the depth of the network keeping it optimizable. In our case, for all DACNN architectures that are deeper than 17 layers, residual learning is applied.

Mixed Architecture. Since convolution kernels may behave differently when receptive field changes, we also adapted our approach of sharing weights on layers that only operates at same spatial resolution. In other words, instead of sharing one convolution weights throughout the entire network, we separate the network into sections by pooling [20] layers, and weights will only be shared within each section (Figure 3, left). In this way, number of channels can be expanded as the network goes deeper. As a trade-off, more parameters will be needed:

- One transition layer for each channel expansion (one more 1×1 convolutional layer needed if residual learning is adopted).
- One convolutional layer for each section (as shared weights).

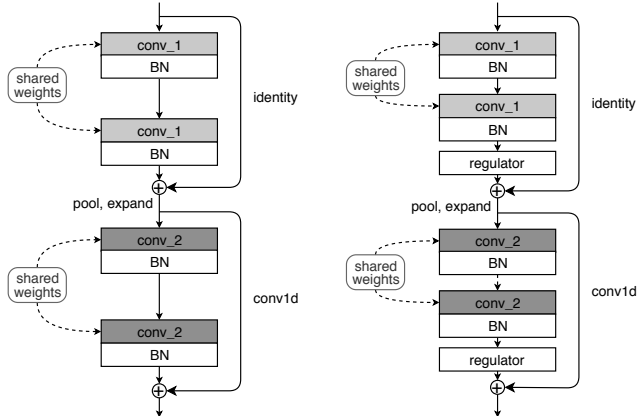


Figure 3. **Left:** Base block for mixed DACNN, parameters are only shared between blocks operates in the same spatial resolution (between pooling layers). **Right:** Base block for mixed DACNN with regulators, one regulator is appended to each block with free parameters.

- Batch normalization [15] weights for all layers.

For example, if the number of channels expands 4 times in an architecture: $(3 \rightarrow 64 \rightarrow 128 \rightarrow 256 \rightarrow 512)$, 8 sets of convolution weights need to be initialized: 4 for expansion and 4 for section weights sharing.

Regulators. We also provide an easy-plug-in way to improve the performance of DACNNs, which we call regulators. A single regulator is constructed with a 1×1 convolution kernel [19], a batch normalization [15] layer and a ReLU [23] activation layer, as illustrated in Figure 3, right. All parameters in a regulator are not shared and it can be plugged in anywhere of the network as long as dimension matches, it helps to regulate the output of each convolution layer with shared weights. In deep architectures that adapt residual learning, we argue that 1 regulator for each residual block is enough to achieve a desirable performance. We will also provide detailed experiment results on how many and where to add these regulators in the later section.

Implementation details. For plain DACNNs, we use a $3 \times 128 \times 3 \times 3$ convolution kernel to expand a 3-channel image to a 128-channel activation map, and then followed by a $128 \times 128 \times 3 \times 3$ kernel stacked L times, where L is the desired depth of the network, in addition, all convolution kernels above are followed by a batch normalization [15] layer (free parameter) and a ReLU [23] activation layer. For DACNNs that are deeper than 17 layers, residual learning is adopted to keep them optimizable. For mixed DACNNs, $(3 \rightarrow 64 \rightarrow 128 \rightarrow 256 \rightarrow 512)$ is adopted as channel expansion pattern for all architectures.

	w.o BN	w BN
DACNN14 (VGG based)	97.48	42.29
DACNN18 (ResNet based)	97.17	38.63

Table 1. TOP-1 error (%) on CIFAR-100 dataset. We compare DACNN network without batch normalization (BN), and DACNN with free BN weights for every layer. VGG and ResNet are selected as base template for 14 and 18 layer DACNN, respectively. Both DACNN architectures have 128 as the fixed number of channels. The results shows that batch norm is crucial to our architecture.

4. Experiments

In this section, we provide detailed analysis and results of DACNN on different perspectives. Experiments are conducted on Cifar 10, Cifar 100 [16] and SVHN [6] datasets.

4.1. DACNN Analysis

We conducted thorough experiments on CIFAR-100 dataset [16] to evaluate our DACNN. We analyze the importance of batch normalization, the role of the depth of the architecture with respect to the number of parameters, as well as all the proposed additional components described in the previous section.

Dataset and training settings. We use CIFAR-100 dataset [16] to perform the analysis. The dataset contains 60,000 32×32 color images in 100 different classes (600 images per class). It is split into training set and test set with the ratio of 5:1. All models in this section are trained 90 epochs with random horizontal flip as data augmentation [17], for preprocessing, we normalize the data using the channel means and standard deviations as in [17]. The networks are updated with ADAGRAD [5] optimizer with learning rate set to 0.1 and decreases to 0.01 from 45th epoch onwards.

Importance of Batch Normalization. Here we analyze the impact of batch normalization (BN) [15] in our DACNN architectures. We trained 2 models on CIFAR-100 dataset: a 14-layer plain DACNN (VGG [28] based) and a 18-layer plain DACNN with residual learning (ResNet [11] based), they are trained both with and without BN for each layer. As we show in the results in Table 1, the network performs poorly without batch normalization in both tested models, giving an error of almost random guessing. Our hypothesis of this behavior is that BN helps scale the output feature map after every shared filter, thus introduces some divergence in to the network rather than simply stacking weights. Therefore, for the rest of the experiments, all DACNN configurations are equipped with free batch norm parameters as a default setting.

# layers	TOP-1 err. (%)	# param (M)
3	56.91 (54.41)	0.164 (0.45)
5	40.11 (33.74)	0.164 (0.74)
7	42.22 (32.66)	0.164 (1.03)
9	42.40 (32.67)	0.164 (1.33)
11	42.37 (30.36)	0.165 (1.63)
14	42.29 (30.44)	0.165 (2.06)
18	38.63 (28.31)	0.166 (2.66)
34	38.88 (27.22)	0.168 (5.02)

Table 2. TOP-1 test error (%) on CIFAR-100: Layers vs # parameters. Data inside parentheses are for architectures without sharing weights for comparison. All models here have 128 number of channels through out the network, for architectures deeper than 17 layers, residual learning is applied. For DACNNs, number of parameters barely increase as the network goes deeper.

Impact of Depth in DACNN. Here, we provide comparison between DACNNs of different depths, we also compare them to networks without sharing the kernel weights with the same structures. As in the batch normalization experiment, all networks have 128 as fixed number of channels, pooling layers are inserted in between convolutions. In addition, for architectures deeper than 17 layers, residual learning is adopted to the entire network.

Results are shown in Table 2. The error rate of plain DACNN drops as the network goes deeper up to the 5-th layer. The next drop is when we introduce residual learning into DACNNs (see the 18-layer network). Compare to networks with free weights, model sizes of DACNNs do not increase much as the network grows. However, we observe that simply stacking plain DACNN kernels and increasing the depth barely benefit the performance of our architecture, we explore further improvements which are discussed in the sequel.

Mixed DACNN. Next we evaluate mixed architectures for DACNNs. We choose VGG, ResNet18 and ResNet34 [11, 28] as our base architectures. As we introduced earlier, mixed DACNN architectures require extra parameters whenever channel dimension expands, since all architectures above share the same channel expansion pattern, the number of parameters required are almost the same (for ResNet based architectures, one more 1×1 convolution is needed for each shortcut expansion).

As shown in results of Table 3, with greatly reduced number of parameters, performance of mixed DACNNs are comparable to VGG and ResNets with same number of layers. Mixed DACNN18 has only 0.41% performance drop compared to ResNet18, but model number of parameters is reduced by 55%. Although increasing the depth doesn't benefit performance of mixed DACNNs, it doesn't increase the number of parameters neither. In later section, we ar-

Template	TOP-1 err. (%)	param (M)
VGG	30.20	14.90
DACNN-14(MIX)	30.54	4.73
ResNet-18	27.31	11.13
DACNN-18(MIX)	27.72	4.90
ResNet-34	26.10	21.24
DACNN-34(MIX)	27.70	4.91

Table 3. Classification results on CIFAR-100 dataset. We choose VGG, ResNet18, and ResNet34 as our base models. Mixed DACNNs obtained desirable performances with much less parameters compared to original VGG and ResNets.

SECTION	TOP-1 err. (%)	extra param
NONE	38.63	-
section 1	38.28	32K
section 2	38.09	32K
section 3	37.46	32K
section 4	38.07	32K
ALL	35.39	128K

Table 4. Allocations of regulators. Plain DACNN-18 was selected as the testing model, the network is cut into 4 sections by 3 pooling layers, and regulators are appended to residual blocks in each section separately. ALL denotes all sections, **bold** denotes overall best result and **blue** denotes best result on single section.

gue that deeper DACNNs have higher capacity for improvements.

Regulators. Here we examine the effectiveness of regulators (1×1 convolution kernels) in our DACNN architecture. First we test regulators on a plain 18 layer DACNN with residual learning, we separate the network into 4 sections by pooling layers, and experiment to add regulators to different sections separately (note that we only append one regulator into each residual block). Results are shown on Table 4. We observe that the performance of the network increases by a considerable margin as we append regulators to different sections. Since each regulator is a 1×1 convolution kernel, only few extra parameters are added to the network. According to the results, appending regulators to the 3rd section gives the best efficiency, we dropped 1.17% on the error rate compared to plain DACNN18 using only 32K parameters, and by appending regulators to all sections, we are able decrease the error rate by 3.24%.

We also analyze the effect of using regulators to models with different depth. In this experiment, regulators are appended to all sections to give better performance. Results are shown on Table 5. As expected, deeper networks give better performance since they have more residual blocks to fit in regulators. On a 34-layer DACNN, we are able to obtain about 1.5% drop on the error rate compare to a 18-

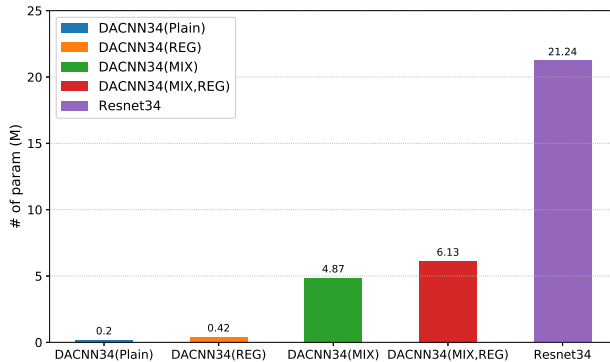
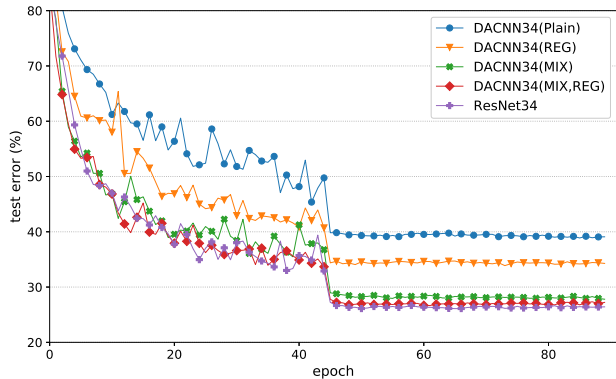


Figure 4. **Left:** TOP-1 test error (%) on CIFAR-100 dataset. Solid curves denote DACNN and its variations, dotted curve denotes ResNet-34 as comparison. **Right:** Number of parameters of these architectures. Mixed DACNN-34 and mixed DACNN-34 with regulators achieved similar performance as ResNet34 while using much less parameters.

model	TOP-1 err. (%)	extra param (M)
DACNN-18 (REG)	35.39	0.32
DACNN-24 (REG)	34.86	0.52
DACNNt-34 (REG)	33.88	0.68

Table 5. The effect of regulators on plain DACNNs. Extra parameters are computed with respect to a plain DACNN-18. All networks are trained and tested on CIFAR-100 dataset. REG denotes regulators, **bold** denotes our best result.

layer DACNN with more regulators appended (0.34M more parameters).

Lastly, we combine everything above together. We apply both mixed structure and regulators to DACNNs, the results can be found at the bottom of Table 6. In comparison with ResNet-18, Mixed DACNN-34 with regulators obtains better accuracy while using only half number of the parameters; Comparing to ResNet-34, mixed DACNN-34 with regulators are 0.56% lower in accuracy, but the model size is 4× smaller.

Model Efficiency. Here, we evaluate model efficiency by considering model size and performance. Architectures with higher performance and less parameters will be considered as high efficiency models.

We provide results of different architectures on CIFAR-100, as well as their number of parameters on Table 6. On Figure 4, we plotted the testing curves and model sizes of 34-layer network architectures with different configurations. From the results, we observe that plain DACNNs are extremely small in model size but their accuracy are not competitive. Yet, as we introduce mixed structure and regulators into the model, the boost in accuracy is tremendous, while the number of parameters of resulting models are still smaller than VGG and ResNets [28, 11] by a large margin.

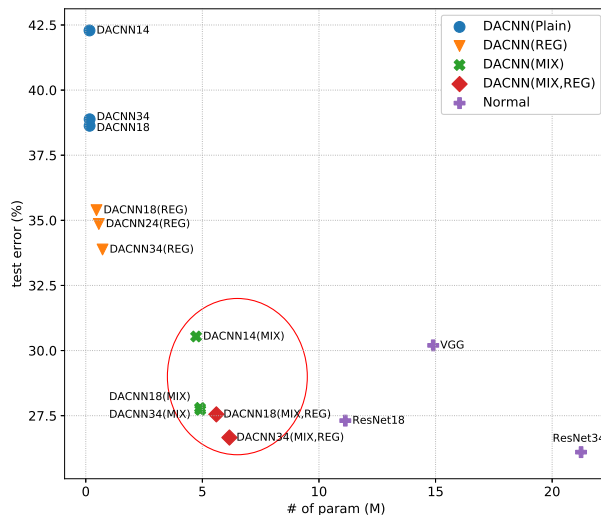


Figure 5. Model Efficiency. This is a plot of TOP-1 error rate on CIFAR-100 vs. number of parameters. Architectures with high model efficiency are at the bottom left of the graph. Mixed DACNNs and mixed DACNNs with regulators (circled with red) give highest model efficiency among all other models in our experiment.

For instance, mixed DACNN-34 with regulators has 15.1 million parameters fewer than ResNet-34. Figure 5 illustrates a plot of parameter efficiency, in which architectures with high model efficiency are expected to be plotted on the bottom left of the graph. As shown on the figure, DACNNs with mixed structure and regulators are much more efficient than plain DACNNs and normal architectures (VGG, ResNets [28, 11]).

method	TOP-1 err. (%)	# param (M)
VGG (1-fc)	31.20	14.90
ResNet18	27.31	11.13
ResNet34	26.10	21.24
<hr/>		
DACNN-14 (plain)	42.20	0.16
DACNN-18 (plain)	38.63	0.16
DACNN-34 (plain)	38.88	0.17
<hr/>		
DACNN-18 (REG)	35.39	0.46
DACNN-24 (REG)	34.86	0.56
DACNN-34 (REG)	33.88	0.72
<hr/>		
DACNN-14 (MIX)	30.54	4.73
DACNN-18 (MIX)	27.72	4.90
DACNN-34 (MIX)	27.80	4.91
<hr/>		
DACNN-18 (MIX, REG)	27.55	5.60
DACNN-34 (MIX, REG)	26.66	6.10

Table 6. Classification results on **CIFAR-100** dataset. All DACNN variations are shown on this table, together with VGG and ResNets for comparison. Accuracy of DACNN-34 is only a small margin lower than ResNet34, while cutting off 70% of the parameters.

4.2. Classification Results on CIFAR-10 and SVHN

To validate our method on other datasets, we also trained DACNNs on CIFAR-10 and SVHN [16, 6]. CIFAR-10 has similar configuration as CIFAR-100 but with only 10 classes. SVHN is a house number recognition dataset obtained from Google Street View images, there are 73,257 images in its training set, and 26,032 images in the test set.

Mixed DACNN-18 with regulators and mixed DACNN-34 with regulators are selected in this experiment, we also trained ResNet-18 and ResNet-34 for comparison. On CIFAR-10, models are trained 90 epochs without data augmentation, learning rate was set to 0.1 and decreases to 0.01 at 45th epoch. On SVHN dataset, models are trained 60 epochs, following common practice, [7, 12, 26, 22] no data augmentation is applied. Learning rate starts from 0.1, decreases as a factor of 10 for every 20 epochs, we use ADA-GRAD [5] as our optimizer in both cases.

The results are shown on Table 7, and Figure 6 is a plot of testing curves of both networks on both datasets. With this experiment, we validate the competitiveness and effective of our architecture as results are comparable to those with the architecture with free weights, and using much less number of parameters.

We also compare our method with 2 other pruning techniques: Agent Pruning by Huang *et al.* [13], and sparse matrix proposed by Han *et al.* [10], results are shown on Table 8. DACNNs are able to achieve high compression ratios with low accuracy drops compared to the other 2 methods. In addition, DACNNs are easier for implementation and deployment, while the rest two require data-driven fine tuning or additional software/hardware support.

model	CIFAR-10	SVHN
ResNet-18	6.47%	4.38%
ResNet-34	6.18%	4.08%
DACNN-18 (MIX,REG)	7.26%	4.37%
DACNN-34 (MIX,REG)	7.23%	4.06%

Table 7. TOP-1 errors on **CIFAR-10** and **SVHN** dataset. Results shows that DACNNs even outperformed ResNets on SVHN dataset.

method	Accuracy drop (%)	Prune Ratio (%)
Agent Pruning	0.3	27.1
Agent Pruning	1.0	37.0
Agent Pruning	1.7	67.9
SM Pruning	1.3	27.1
SM Pruning	6.9	37.0
SM Pruning	6.5	67.9
DACNN(M,R)	0.79	49.6
DACNN(plain)	10.2	98.5

Table 8. Classification results on **CIFAR-10** dataset. Agent Pruning [13], SM Pruning denotes sparse matrix Pruning [10] and (M,R) denotes Mixed Structure and Regulators respectively. ResNet 18 was selected as base architecture.

4.3. Filter Visualization

Here, we visualize the filters of DACNN. We apply similar technique as network fooling [24] for filter visualization: First input an image of random noise, then we optimize the input image with respect to each convolution layer using backpropagation [18].

We train a 5 layer plain DACNN on CIFAR-10 (32×32) and plot the results of some filters on Figure 7. The reason why we choose plain DACNN here, is that we want to demonstrate the output features can still be very diversified across layers even without implementations of regulators and mixed-architecture. This also illustrates the importance of batch normalization to scale the filter responses.

5. Conclusion

We introduced a new convolutional neural network architecture, which we refer it as *Deep Anchored Convolutional Neural Network (DACNN)*. It shares weights for convolution kernels across layers while keep parameters for Batch Normalization free. Due to high weights reuse, number of parameter of DACNN barely increases as the network goes deeper. Thus, it is a novel way for model size compression.

Since we observe that simply increasing the depth of DACNNs contributes little to the performance, we also propose two ways to improve the performance of DACNNs: Mixed Structure and Regulators.

With these two methods, DACNN is an efficient model compression approach adopting an “addition” fashion: First

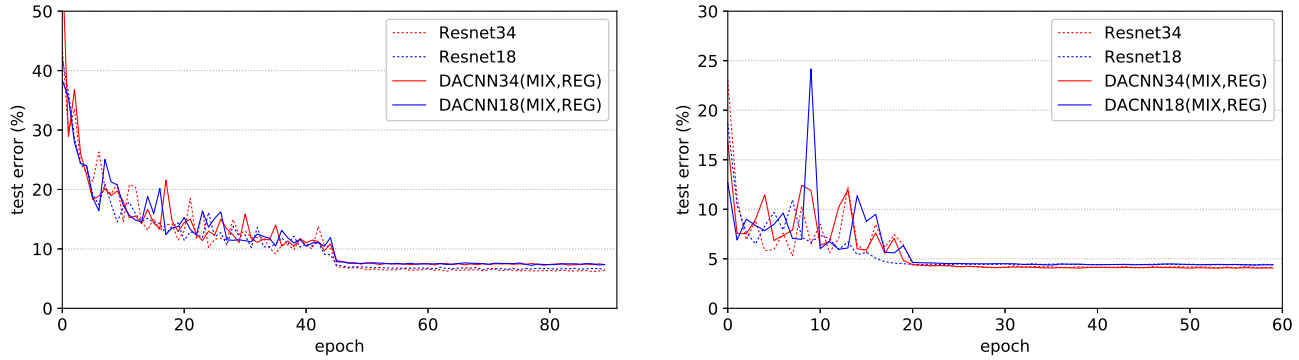


Figure 6. Training & Testing on **CIFAR-10** and **SVHN** dataset. Solid curves denotes results for DACNNs, and dashed curves denotes results for ResNets. We observe that ResNets are slightly better on CIFAR-10, while DACNNs are better on SVHN.

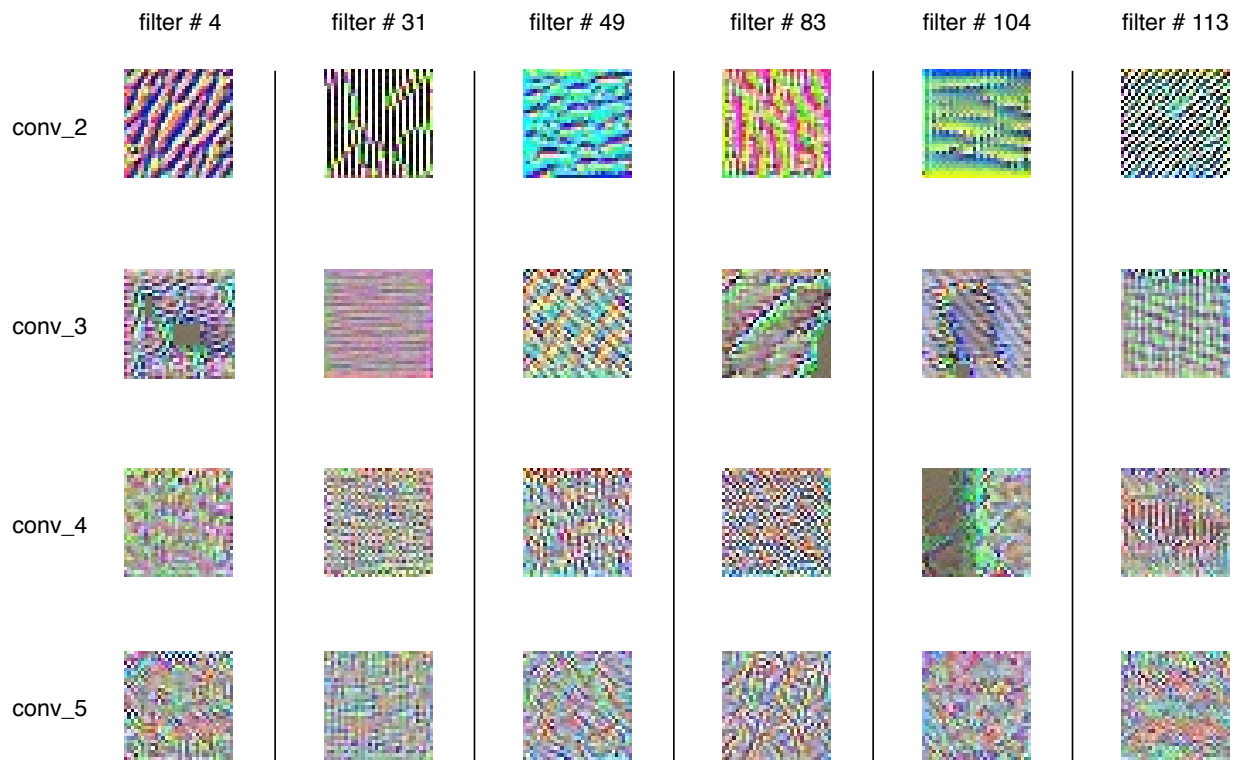


Figure 7. Filter Visualization. Each of the image above is optimized for the corresponding convolutional layer using backpropagation. We trained a 5 layer plain DACNN on CIFAR-10 dataset and selected 5 filters for demonstration.

initialize a plain DACNN to a desired depth (deeper networks have higher capacity for further improvements). Then, selectively apply mixed structure and append regulators to achieve a desirable performance. As a result, DACNNs are able to obtain similar performance with much less parameters compare to some popular architectures like VGG and ResNet [28, 11].

Acknowledgements. This work was funded by the SUTD-MIT IDC grant (IDG31800103). K.D. was also funded by SUTD Presidents Graduate Fellowship.

References

- [1] A. Bouch. Sharesnet: reducing residual network parameter number by sharing weights. *CoRR*, abs/1702.08782, 2017. 2

- [2] Y. Chen, J. Li, H. Xiao, X. Jin, S. Yan, and J. Feng. Dual path networks. In *NIPS*, 2017. [1](#)
- [3] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei. Imagenet: A large-scale hierarchical image database. *2009 IEEE Conference on Computer Vision and Pattern Recognition*, pages 248–255, 2009. [1](#)
- [4] M. Denil, B. Shakibi, L. Dinh, M. Ranzato, and N. de Freitas. Predicting parameters in deep learning. In *NIPS*, 2013. [1](#)
- [5] J. C. Duchi, E. Hazan, and Y. Singer. Adaptive subgradient methods for online learning and stochastic optimization. *Journal of Machine Learning Research*, 12:2121–2159, 2011. [4](#), [7](#)
- [6] I. J. Goodfellow, Y. Bulatov, J. Ibarz, S. Arnaud, and V. D. Shtet. Multi-digit number recognition from street view imagery using deep convolutional neural networks. *CoRR*, abs/1312.6082, 2013. [2](#), [4](#), [7](#)
- [7] I. J. Goodfellow, D. Warde-Farley, M. Mirza, A. C. Courville, and Y. Bengio. Maxout networks. In *ICML*, 2013. [7](#)
- [8] S. Han, X. Liu, H. Mao, J. Pu, A. Pedram, M. Horowitz, and W. J. Dally. Eie: Efficient inference engine on compressed deep neural network. *2016 ACM/IEEE 43rd Annual International Symposium on Computer Architecture (ISCA)*, pages 243–254, 2016. [2](#)
- [9] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural network with pruning, trained quantization and huffman coding. *CoRR*, abs/1510.00149, 2015. [2](#)
- [10] S. Han, J. Pool, J. Tran, and W. J. Dally. Learning both weights and connections for efficient neural networks. In *NIPS*, 2015. [2](#), [7](#)
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [8](#)
- [12] G. Huang, Y. Sun, Z. Liu, D. Sedra, and K. Q. Weinberger. Deep networks with stochastic depth. In *ECCV*, 2016. [1](#), [7](#)
- [13] Q. Huang, S. K. Zhou, S. You, and U. Neumann. Learning to prune filters in convolutional neural networks. *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*, pages 709–718, 2018. [2](#), [7](#)
- [14] F. N. Iandola, M. W. Moskewicz, K. Ashraf, S. Han, W. J. Dally, and K. Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and 1mb model size. *CoRR*, abs/1602.07360, 2016. [2](#)
- [15] S. Ioffe and C. Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *ICML*, 2015. [3](#), [4](#)
- [16] A. Krizhevsky. Learning multiple layers of features from tiny images. 2009. [2](#), [4](#), [7](#)
- [17] A. Krizhevsky, I. Sutskever, and G. E. Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012. [1](#), [2](#), [4](#)
- [18] Y. LeCun, B. E. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. E. Hubbard, and L. D. Jackel. Backpropagation applied to handwritten zip code recognition. *Neural Computation*, 1:541–551, 1989. [7](#)
- [19] Y. LeCun, L. Bottou, and P. Haffner. Gradient-based learning applied to document recognition. 1998. [3](#), [4](#)
- [20] Y. LeCun, L. Bottou, and P. Haffner. Gradient-based learning applied to document recognition. 1998. [3](#)
- [21] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf. Pruning filters for efficient convnets. *CoRR*, abs/1608.08710, 2016. [2](#)
- [22] M. Lin, Q. Chen, and S. Yan. Network in network. *CoRR*, abs/1312.4400, 2013. [7](#)
- [23] V. Nair and G. E. Hinton. Rectified linear units improve restricted boltzmann machines. In *ICML*, 2010. [3](#), [4](#)
- [24] A. M. Nguyen, J. Yosinski, and J. Clune. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 427–436, 2015. [7](#)
- [25] S.-A. Rebuffi, H. Bilen, and A. Vedaldi. Learning multiple visual domains with residual adapters. In *NIPS*, 2017. [2](#)
- [26] P. Sermanet, S. Chintala, and Y. LeCun. Convolutional neural networks applied to house numbers digit classification. *Proceedings of the 21st International Conference on Pattern Recognition (ICPR2012)*, pages 3288–3291, 2012. [7](#)
- [27] P. Sermanet, D. Eigen, X. Zhang, M. Mathieu, R. Fergus, and Y. LeCun. Overfeat: Integrated recognition, localization and detection using convolutional networks. *CoRR*, abs/1312.6229, 2013. [1](#)
- [28] K. Simonyan and A. Zisserman. Very deep convolutional networks for large-scale image recognition. *CoRR*, abs/1409.1556, 2014. [1](#), [4](#), [5](#), [6](#), [8](#)
- [29] S. Srinivas and R. V. Babu. Data-free parameter pruning for deep neural networks. In *BMVC*, 2015. [2](#)
- [30] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. In *2015 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1–9, June 2015. [1](#)
- [31] C. Szegedy, V. Vanhoucke, S. Ioffe, J. Shlens, and Z. Wojna. Rethinking the inception architecture for computer vision. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2818–2826, 2016. [1](#)
- [32] S. Zagoruyko and N. Komodakis. Wide residual networks. *CoRR*, abs/1605.07146, 2016. [1](#)
- [33] H. Zhou, J. M. Alvarez, and F. M. Porikli. Less is more: Towards compact cnns. In *ECCV*, 2016. [2](#)