

# NodeDrop: A Method for Finding Sufficient Network Architecture Size

Louis Jensen (*equal contribution*)  
Department of Computer Science  
Boston University  
Boston, MA  
ljensen@bu.edu

Jacob Harer (*equal contribution*)  
Department of Computer Science  
Boston University  
Boston, MA  
jharer@bu.edu

Sang Chin  
Department of Computer Science  
Boston University  
Boston, MA  
spchin@bu.edu

**Abstract**—Determining an appropriate number of features for each layer in a neural network is an important and difficult task. This task is especially important in applications on systems with limited memory or processing power. Many current approaches to reduce network size either utilize iterative procedures, which can extend training time significantly, or require very careful tuning of algorithm parameters to achieve reasonable results. In this paper we propose NodeDrop, a new method for eliminating features in a network. With NodeDrop, we define a condition to identify and guarantee which nodes carry no information, and then use regularization to encourage nodes to meet this condition. We find that NodeDrop drastically reduces the number of features in a network while maintaining high performance. NodeDrop reduces the number of parameters by a factor of 114x for a VGG like network on CIFAR10 without a drop in accuracy.

**Index Terms**—Machine Learning, Deep Learning

## I. INTRODUCTION

A prime difficulty in neural network design is the appropriate tuning of network architectures. Choosing a size for each layer of a neural network is usually done by rough estimate, trial, and error. This imprecise process can often lead to network designs larger than needed to perform a particular task. Although the capacity for training large and complex networks grows with improving graphics processing unit (GPU) technology, designing too large a network can result in applications impracticable for general hardware use. Mobile devices and embedded systems limit compute, memory, and storage consumption, and as a result can only run small, minimally designed networks. A designer aiming to create such a minimal network is faced with the time-consuming task of manually tuning the number of neurons in each layer. This tuning process can result in many extended tuning experiments in order to balance the space and performance of the neural network.

The issues involved with using deep neural networks (DNNs) on constrained systems has inspired significant research. One interesting area of research is the design of systems which can automatically prune a network's parameters. Ideally these techniques can still maintain high performance while pruning as many parameters as possible, ensuring the network can fit on smaller systems. Many state-of-the-art methods for network pruning generally involve an iterative process of repeatedly pausing training, pruning parameters, and resuming

training in order for the network to reconverge. Such iterative procedures can lead to long training times. Other techniques use regularization in order to eliminate nodes. The final performance of these networks is often highly variable with the hyper-parameters of the algorithm. Thus, while these techniques do offer parameter reduction benefits, the network designer will still be faced with similar difficulties as before: a time-consuming training process and a potential hyper-parameter tuning headache.

We address the problem of parameter reduction with our novel NodeDrop technique, which prunes the network during training. The NodeDrop technique only drops nodes which carry no information and drops them fluidly during the training process.

First, we formally define the conditions necessary to guarantee a neuron carries no information. We then propose a simple variant of  $L_1$  regularization which drives nodes toward this condition. Second, we extend the NodeDrop technique to networks which use batch normalization [1]. We test our technique on modern architectures for the MNIST, CIFAR10, CIFAR100, and ImageNet datasets, and show that we are able to drop a significant number of nodes without a loss in performance. Our method requires no iterative retraining and only a modest increase in training time. We demonstrate effective results with a wide range of hyperparameters, indicating our method does not require precise hyperparameter tuning. At best case we produce a network which reduces the number of parameters by 93.27, 99.12, 87.82, and 83.47 percent for MNIST, CIFAR10, CIFAR100, and ImageNet respectively.

## II. RELATED WORKS

### A. Pruning

Network pruning comprises a set of techniques which take a pretrained network and then prune off connections using some heuristic. This is usually followed by a retraining of the network and sometimes by an iterative process of pruning and retraining the network several times. Pruning techniques first appeared in the 1990s, with the first instances using second order gradients of connections to determine which neurons should be pruned [2]–[4]. More recent approaches have taken on a wide array of methods for determining which connections should be pruned. These approaches include correlation [5]–[7],

regularization [8], [9], particle filtering on misclassification rate [10], low rank approximation [11], vector quantization [12] and tensor decomposition [13].

All network pruning techniques suffer from extended training time due to the iterative retraining of the network. This can lengthen training times significantly, and often makes tuning the various parameters in each method a lengthy chore.

### B. Regularization

A more recently developed approach to network parameter reduction is to disable parameters through regularization. A majority of these techniques have focused on the sparsification of network connections using a group sparsity approach [14]–[17]. This involves grouping the weights for every neuron and attempting to sparsify each group by penalizing its  $L_2$  norm. These techniques require all weights to be driven to zero before a node can be guaranteed to carry no information. In practice nodes are removed based on a threshold since this guarantee is difficult to meet. Because of this, regularization methods can be difficult to use as they require very precise tuning of the regularization and threshold terms.

The most similar technique to ours, Liu et al. [18], uses  $L_1$  regularization to drive the scale parameter in batch norm,  $\gamma$ , towards zero. This is similar in principle to our own experiments with batch norm. However, Liu et al. requires retraining after pruning in order to reconverge. We provide a more absolute condition to guarantee a node is off, eliminating the need for a retraining procedure and making node removal a more fluid process.

Our technique falls within the regularization category. Key differences in our approach involve special regularization of the bias for each neuron and, most importantly, a condition for node removal which absolutely guarantees no effect on network output. Our condition is also more relaxed, utilizing the “dead” region in a node’s activation function, instead of requiring the node’s weights to be zero.

### C. Other approaches

Several other approaches have appeared which do not fit into the categories of the previous two subsections. Many of these approaches focus on reducing precision as opposed to reducing the number of parameters [19]–[22]. As such, these approaches are largely orthogonal to our own work, and can be used in conjunction with our work in order to compound the reduction on memory and computation. One example of this approach is quantized and binarized neural networks [19], which take this approach to new levels by using  $\{-1, 1\}$  weights and an XOR to replace multiplication.

An additional noteworthy paper is that of Molchanov et al. [23]. They achieve impressive results by sparsifying a network’s connections during training using variational dropout. Again, in theory this work should be usable in conjunction with our own.

## III. METHODS

### A. NodeDrop Condition

In this section we describe the condition for identifying useless nodes in a network. Nodes in a neural network carry information by outputting values from some distribution. A node can only be useful if that node sometimes outputs a non-zero value. A node which is guaranteed to always output a constant value is a node which can only be used as an extra bias node for future layers, thus to ensure a node is completely useless it will be important to include bias in our analysis. Moreover, if a node is guaranteed to always output the constant zero, this node is entirely useless and can be removed from the network without impact. This occurs in activation functions with a flat zero region. The popular rectified linear unit (ReLU) activation function contains such a flat zero region. This flat zero region causes the observed “Dying ReLU” effect, in which nodes become stuck in this flat region with zero gradients. We can therefore design a condition to identify when a node is useless by taking advantage of this effect.

We propose the NodeDrop condition.

- 1) Given a node with input vector  $\vec{x} \in [0, 1]^n$ , a weight vector  $\vec{w} \in \mathbb{R}^n$ , bias  $b \in \mathbb{R}$ , and an activation function  $\sigma$  such that  $\sigma(v) = 0 \forall v \leq 0$ .
- 2) We wish to find the condition under which this node is dead,  $\sigma(\vec{w} \cdot \vec{x} + b) = 0$  for all inputs  $\vec{x}$ .
- 3) Since  $\sigma(v) = 0 \forall v \leq 0$ , we simply need to find the condition under which  $\vec{w} \cdot \vec{x} + b \leq 0$ . We have constrained the inputs to be within  $[0, 1]$ ,  $\vec{x} \in [0, 1]^n$ , so we have:

$$\vec{w} \cdot \vec{x} + b \leq \|\max(\vec{w}, \vec{0})\|_1 + b \leq \|\vec{w}\|_1 + b$$

- 4) Then,  $\|\max(w_i, 0)\|_1 + b \leq 0 \Rightarrow \sigma(\vec{w} \cdot \vec{x} + b) = 0$

This leaves us with the NodeDrop condition:

$$\|\max(w_i, 0)\|_1 + b \leq 0 \tag{1}$$

This condition can be applied to a fully connected layer, or in broader contexts such as filter weights of a convolutional layer. Because nodes which satisfy this condition are guaranteed to always output zero, they can be dropped from the network without affecting its output. Note that the weaker condition  $\|\vec{w}\|_1 + b \leq 0$  can also be used, but identifies fewer nodes that can be dropped.

The constraint on the activation function  $\sigma(v) = 0 \forall v \leq 0$  can be achieved using the standard ReLU activation of  $\max(0, x)$ . However ReLU does not guarantee that the output will fall between 0 and 1, a necessary condition if we want to apply NodeDrop to the next layer in the network. In the following section we will discuss an activation function for which the NodeDrop condition can be applied to both a layer and its following layer.

### B. Activation Function

Supposing we want to apply NodeDrop to many or all layers of the network, we must use an activation function which possesses the appropriate flat zero region  $(-\infty, 0]$ , and whose outputs are always between 0 and 1. The flat zero region

guarantees the NodeDrop condition can be applied to the layer preceding activation, and the  $[0, 1]$  constraint on the output allows for the NodeDrop condition to be applied to the layer immediately following activation. These necessary constraints are reiterated below.

$$\sigma(v) = 0 \quad \forall v \leq 0 \quad (2)$$

$$\sigma(v) \in [0, 1] \quad \forall v \quad (3)$$

If the outputs of a layer are guaranteed between 0 and 1 after activation, the inputs of the next layer will satisfy the conditions assumed in proving the NodeDrop condition. Many activation functions can satisfy these conditions, but none of the most popular activation functions satisfy both together. For example, the popular ReLU function satisfies the first condition in equation 2 but not equation 3. Conversely, the popular sigmoid activation function satisfies equation 3 but not equation 2.

One option is a clamped ReLU activation function,  $\min(1, \max(0, x))$ . This has two flat regions,  $\sigma(v \geq 1) = 1$ ,  $\sigma(v \leq 0) = 0$ , and an intermediate region where  $\sigma(v \in [0, 1]) = v$ . This does satisfy both of the NodeDrop conditions; however, we found that having two regions with zero gradients can lead to too many nodes being “stuck” at either 0 or 1 even at network initialization. Thus, we propose the SoftClampedReLU activation function, which is a combination between ReLU and inverted SoftPlus activations:

$$\sigma(v) = \max(0, 1 - \frac{1}{\beta} \log(1 + e^{\beta(1-x)})) \quad (4)$$

Intuitively this activation is much like a ClampedReLU, but has a soft gradient in the upper region. This upper region is not perfectly flat and so values do not become stuck at  $\sigma(v) = 1$ . The  $\sigma(v \leq 0) = 0$  lower region is still perfectly flat, satisfying our flat region condition, equation 2. This activation function is shown in figure 1. In our experiments we used  $\beta = 10.0$  and observed no decline in performance with the SoftClampedReLU activation compared to ReLU.

### C. Regularization

The NodeDrop condition for identifying and eliminating useless nodes is powerful, but without encouragement, most trained networks will possess very few nodes satisfying the NodeDrop condition. Therefore, we add regularization during training to encourage features to satisfy the NodeDrop condition.

To encourage  $\|\max(\vec{w}, \vec{0})\|_1 + b = 0$  we can directly penalize its distance from zero:

$$\lambda \|\max(\vec{w}, \vec{0})\|_1 + b$$

However, this is too close to the boundary of our dead region. Alternatively to encourage  $\max(\vec{w}, \vec{0}) + b \leq 0$ , we could penalize it directly:

$$\lambda (\|\max(\vec{w}, \vec{0})\|_1 + b)$$




Fig. 1. SoftClampedRelu activation function. Shown with  $\beta = 10$ .

However, this causes the bias to tend toward negative infinity.

Instead we penalize the distance from a negative constant,  $-C$ , given:

$$\lambda \|\max(\vec{w}, \vec{0})\|_1 + b + C = \lambda \sum_i \max(w_i, 0) + b + C$$

This encourages  $\|\max(\vec{w}, \vec{0})\|_1 + b = -C$ , safely within the “dead” region, and without tending to negative infinity. As such, the choice of  $C$  is largely arbitrary; in our experiments we found a value of 1.0 worked well, though other values worked just as well.

We can also write our regularization term as a small modification to standard  $L_1$  regularization. For the case where  $\|\max(\vec{w}, \vec{0})\|_1 + b \geq 0$ , when a node is on,  $\lambda \|\max(w_i, 0)\|_1 + b + C = \lambda (\|\max(\vec{w}, \vec{0})\|_1 + \|b + C\|_1)$ . We use this modified  $L_1$  regularization given as:

$$\lambda (\|\max(\vec{w}, \vec{0})\|_1 + \|b + C\|_1) \quad (5)$$

This is normal  $L_1$  regularization with two adjustments. We use the  $L_1$  norm of  $\max(\vec{w}, \vec{0})$  instead of  $\vec{w}$  since this is a tighter bound given that  $x \geq 0$ . Instead of penalizing the bias as  $\|b\|_1$ , we penalize  $\|b + C\|_1$ . This modified  $L_1$  regularization encourages biases to take bias values near  $b = -C$ , and weight values near 0. It should be noted that penalizing the bias is important as the bias is a part of the NodeDrop conditions guarantee.

We use  $L_1$  regularization because  $L_2$  regularization does not work well in our context. For  $L_2$  regularization on both the weights and the bias, it is cheaper to use multiple weights as a bias rather than the bias itself. That is, when  $\sum_i^n w_i = b$ , then  $\|w\|_2 < \|b\|_2$ . This becomes worse when we penalize the distance of the bias from  $-C$  rather than from zero, making the normal case of an active node with bias around  $C$  quite costly. This encourages the network to use many nodes in the previous layer as an alternative to a bias, preventing us from

removing those nodes even if they carry no information beyond that of a bias.

#### D. Extension to Batch Normalization

Many state of the art networks utilize batch normalization or one of its alternatives [1], [24], [25]. In the following section we build upon the NodeDrop condition proposed earlier, designing a NodeDrop condition which will hold in a network with batch normalization. Batch normalization is given as follows:

$$\mu = \frac{1}{m} \sum_{i=1}^m x_i \quad (6)$$

$$\sigma^2 = \frac{1}{m} \sum_{i=1}^m (x_i - \mu)^2 \quad (7)$$

$$\hat{x}_i = \frac{x_i - \mu}{\sqrt{\sigma^2 + \epsilon}} \quad (8)$$

$$y_i = \gamma \hat{x}_i + \beta \quad (9)$$

where the sum is over the batch of size  $m$ , and both  $\gamma$  and  $\beta$  are learned parameters. Batch normalization is usually applied between the output of a layer and an activation function.

To achieve a similar NodeDrop condition for batch normalization as in equation 1, we would like to determine when  $y_i \leq 0$ . We similarly require an activation with a flat zero region, but no longer require an input  $\vec{x}$  between 0 and 1. Therefore, for our batch normalization NodeDrop (NodeDrop-BN) technique we are able to use the popular ReLU activation function. Our NodeDrop-BN condition is given in the following lemma.

**Lemma III.1.**  $|\gamma|\sqrt{m} + \beta \leq 0 \implies y_i \leq 0$ .

*Proof.*

$$\begin{aligned} \text{given } \hat{x}_i &= \frac{x_i - \mu}{\sqrt{\sigma^2 + \epsilon}} \\ \implies \hat{x}_i^2 &= \frac{(x_i - \mu)^2}{\sigma^2 + \epsilon} \\ \implies \sum_{i=0}^m \hat{x}_i^2 &= \frac{\sum_{i=0}^m (x_i - \mu)^2}{\sigma^2 + \epsilon} \\ &= \frac{m\sigma^2}{\sigma^2 + \epsilon} \\ \implies \sum_{i=0}^m \hat{x}_i^2 &\leq m \\ \implies -\sqrt{m} &\leq \hat{x}_i \leq \sqrt{m} \end{aligned}$$

Together  $|\hat{x}_i| \leq \sqrt{m}$  and  $y_i = \gamma \hat{x}_i + \beta$  imply  $y_i \leq |\gamma|\sqrt{m} + \beta$ . Therefore

$$|\gamma|\sqrt{m} + \beta \leq 0 \implies y_i \leq 0$$

This gives us the NodeDrop-BN condition:

$$|\gamma|\sqrt{m} + \beta \leq 0 \quad (10)$$

Traditionally, batch normalization stores a running mean,  $\mu$ , and variance,  $\sigma^2$ , during training. These stored values are

then used during testing. Our condition guarantees a node is always off during training, but does not guarantee a node will always be off during testing. We make the assumption that a node which is always off during training should also be off during testing. Thus, we can safely remove these nodes without impact. We experimentally validate this assumption in section IV.

The condition in equation 10 implies that so long as we use an activation function where  $\sigma(v) = 0 \forall v \leq 0$  (for example ReLU), we can determine if a node is off using only the batch normalization parameters,  $\gamma$  and  $\beta$ , and the training batch size,  $m$ . Following the same methodology for regularization as in equation 5, we define our NodeDrop-BN  $L_1$  regularization term as:

$$|\gamma|\sqrt{m} + |\beta + C| \quad (11)$$

$L_2$  regularization is generally applied to the layers before batch normalization. Unlike the vanilla NodeDrop regularization,  $L_2$  regularization does not interfere with NodeDrop-BN technique, because the  $L_2$  regularization applied to layers before a batch normalization has no effect on the output of the batch normalization layer. NodeDrop-BN builds upon the vanilla NodeDrop and is more practically suited for most applications than NodeDrop due to the power of batch normalization. Please test the code implementing both NodeDrop and NodeDrop-BN at the link provided: <https://github.com/ljensen3/NodeDrop>.

## IV. EXPERIMENTS

Having established a theoretical basis for the NodeDrop condition and regularization technique, we will now establish NodeDrop's practical viability as a method for shrinking networks. The NodeDrop technique requires two hyperparameters:  $C$  and  $\lambda$ . The  $C$  value is unimportant, and can be set to almost any positive value without impacting results or parameter reduction. However, the  $\lambda$  parameter is crucial in determining the balance between learning the objective and dropping nodes. Therefore, we closely examine the effect that choosing different  $\lambda$  values has on both network performance and parameter reduction. We test many  $\lambda$  values on the MNIST and CIFAR10 datasets. We also test a few  $\lambda$  values on the CIFAR100 dataset.

The network initialization size should affect the number of nodes dropped. We show that if a network starts near optimal size, NodeDrop will maintain accuracy and only drop what few nodes it can. Furthermore, we show that if a network is grossly oversized at initialization, NodeDrop will drop many nodes and converge towards the same size as a smaller network initialization. This result is desirable, as it demonstrates NodeDrop is largely unaffected by poor layer size choices. NodeDrop uses  $\lambda$  to determine the balance between performance and number of nodes utilized. Therefore, a network architect using NodeDrop can afford to initialize a large network, and remain confident that NodeDrop will eliminate needless nodes. Using the MNIST dataset, we demonstrate this ability by showing that networks will converge to the same size from multiple initialization sizes, for a fixed  $\lambda$ .

TABLE I  
MNIST NETWORK ARCHITECTURES: NUMBER OF FEATURES BY LAYER

NETWORK NAME	LAYER 1 CONV2D $3 \times 3$	LAYER 2 CONV2D $3 \times 3$ MAXPOOL $2 \times 2$	LAYER 3 CONV2D $3 \times 3$	LAYER 4 CONV2D $3 \times 3$ MAXPOOL $2 \times 2$	LAYER 5 DENSE	LAYER 6 OUTPUT
DENSE160	16	16	32	32	64	10
DENSE240	24	24	48	48	96	10
DENSE320	32	32	64	64	128	10
DENSE480	48	48	96	96	192	10
DENSE640	64	64	128	128	256	10




Fig. 2. In the right and center figures, the  $\lambda$  parameter values plotted on the y-axis are on a logarithmic scale. We note that the performance and parameter reduction both maintain desirable levels for a large range of  $\lambda$  values (over several orders of magnitude). This indicates the ease of tuning the NodeDrop technique. In the leftmost figure, networks of different starting size converge to nearly the same size for a given  $\lambda$ . The dashed diagonal line represents networks without pruning. Note that increased initialization size has a slight effect on final size, as indicated by the slight upward slopes. This effect is greater for larger  $\lambda$ . Please note MNIST accuracy’s in the right most figure, demonstrating maintained performance.

Many pruning methods require an increase in training time to be effective. The NodeDrop technique does not delay performance or accuracy convergence, but in order to allow the number of network nodes to converge, one must train for a longer time. We examine the training time required for this convergence with experiments on the CIFAR10 datasets.

Most importantly we test to ensure NodeDrop maintains performance and effectively drops nodes. We find that NodeDrop regularization does not affect a network’s performance for a large swath of  $\lambda$  values, only reducing testing accuracy if extreme  $\lambda$  values are chosen.

Furthermore, we demonstrate that NodeDrop is able to drop more than 100x parameters from popular networks such as VGG16, while continuing to maintain classification accuracy on the CIFAR10 dataset. We test NodeDrop network performance and parameter reduction on MNIST, CIFAR10, and CIFAR100.

#### A. MNIST Experiments

The MNIST dataset [26] provides an opportunity to perform a large number of experiments because of the dataset’s rapid accuracy convergence. Thus, we used this dataset to sweep across  $\lambda$  values for five differently sized, but otherwise similar, network architectures, as shown in table I. We demonstrate

NodeDrop’s ability to rapidly converge to similarly sized networks from different starting sizes.

For all MNIST experiments we used a simple network design: four convolution layers and a single fully connected layer. We used  $3 \times 3$  filters in all convolution layers, and performed max-pooling after every second convolution layer. We varied the width of the layers in order to test the effects of changing network initialization size. We did not investigate the effects of changing network depth, but suspect that prudent selection of network depth remains important. The network architectures are described in table I. The following consistent hyperparameters were used across all MNIST runs: learning rate =  $1.0 \times 10^{-3}$ , batch size = 1024, optimizer = *Adam*, loss function = *cross entropy*, epochs = 480. Figure 2 shows the relationships found between accuracy,  $\lambda$  values, and initialization sizes.

1) *Choosing Lambda*: Choosing an appropriate value for the NodeDrop’s  $\lambda$  parameter remains an important task. In order to prove that the NodeDrop technique remains robust for many selections of  $\lambda$ , we tested five different network initialization sizes to observe differences in convergence across  $\lambda$  values. The network architectures and hyperparameters are discussed in section IV-A. We tested ten different  $\lambda$  values between  $\lambda = 1.0 \times 10^{-8}$  and  $\lambda = 1.0 \times 10^{-3}$ .

Our results indicate that easy tuning is a benefit of the




Fig. 3. Results on CIFAR10 for VGG with and without Batch Normalization over a spread of  $\lambda$  choices. Top Left: Classification error for VGG without Batch Normalization. Top Right: Final parameters after training using NodeDrop. Bottom Left: Classification error for VGG with Batch Normalization. Bottom Right: Final parameters after training using NodeDrop-BN. For both NodeDrop and NodeDrop-BN, a range of  $\lambda$  values are acceptable. Baseline accuracy and network size is indicated by the dashed lines.

TABLE II

RESULTS FOR CIFAR10 DATASET SHOWN. SEE [9], [18] TO COMPARE SOME SIMILAR EXPERIMENTS FOR OTHER NETWORK REDUCTION TECHNIQUES.

NETWORK	$\lambda$	TEST ERROR	PARAMETERS	PRUNED %	FACTOR	NODES	PRUNED %
VGG 16 w/o BN	BASELINE	13.01	15.04M	0.0	1.0	4736	0.0
	$1.0 \times 10^{-6}$	14.14	0.45M	97.00	33.28	1115	76.46
	$1.0 \times 10^{-5}$	13.27	0.31M	97.96	48.98	859	81.9
	$3.2 \times 10^{-5}$	<b>13.76</b>	<b>0.13M</b>	<b>99.12</b>	<b>114.00</b>	<b>612</b>	<b>87.08</b>
	$1.0 \times 10^{-4}$	90.00	0.0M	100.0	-	0	100.0
VGG 16	BASELINE	6.50	15.04M	0.0	1.0	4736	0.0
	$1.0 \times 10^{-6}$	6.88	8.88M	40.7	1.69	3624	23.48
	$1.0 \times 10^{-5}$	7.36	1.39M	90.75	10.81	1164	75.42
	$3.2 \times 10^{-5}$	<b>7.41</b>	<b>0.61M</b>	<b>95.96</b>	<b>24.76</b>	<b>751</b>	<b>84.14</b>
	$1.0 \times 10^{-4}$	20.16	0.10M	99.35	152.84	308	93.50
DENSENET40 w/o BN	BASELINE	14.94	1.04M	0.0	1.0	456	0.0
	$1.0 \times 10^{-6}$	15.21	0.66M	35.69	1.55	363	20.39
	$1.0 \times 10^{-5}$	14.74	0.41M	60.47	2.54	291	36.18
	$1.0 \times 10^{-4}$	<b>14.99</b>	<b>0.08M</b>	<b>91.96</b>	<b>12.43</b>	<b>154</b>	<b>66.22</b>
DENSENET40	BASELINE	6.80	1.05M	0.0	1.0	456	0.0
	$1.0 \times 10^{-6}$	7.13	0.99M	4.19	1.04	447	1.97
	$1.0 \times 10^{-5}$	6.75	0.98M	5.67	1.06	443	2.85
	$1.0 \times 10^{-4}$	<b>7.79</b>	<b>0.55M</b>	<b>47.12</b>	<b>1.89</b>	<b>333</b>	<b>26.73</b>

NodeDrop technique. We found that  $\lambda$  selections across orders of magnitude yielded desirable results, as shown in figure 2. For  $\lambda > 10^{-4}$  we noticed a drop in MNIST accuracy, and for  $\lambda < 1.0 \times 10^{-7}$  we judged there to be a significant sacrifice in parameter reduction. Choosing appropriate  $\lambda$  will always be dependent on both application and loss function. Because of these MNIST experiments, we expect that the NodeDrop technique is robust for a large range of  $\lambda$  selections. For a network designer using the popular cross-entropy loss objective function, as we did, we would suggest  $\lambda = 1.0 \times 10^{-5}$ .

2) *Network Sizes*: In the previous section (IV-A1) we experimentally observed that tuning the  $\lambda$  parameter of the NodeDrop technique should not cause a network designer grief. In this section, we will experimentally observe that choosing initialization layer sizes should also prove easy. We

use the same experiments from the previous section (IV-A1), but instead plot the effect of initializing with differently sized networks. This plot, shown in figure 2, demonstrated that the NodeDrop technique will converge to a similar “equilibrium” from many differently sized initialization networks. The size of the final network is instead mostly dependent on  $\lambda$ . A network designer should err towards too large a network in order to ensure desirable performance.

### B. CIFAR10 and CIFAR100 Experiments

1) *Dataset*: The CIFAR dataset [27] consists of 32x32 colored natural images. Both CIFAR10 and CIFAR100 are designed for classification, containing 10 and 100 classes respectively. There are 50,000 training images and 10,000 testing images for both. We adopt a standard data augmentation

TABLE III

RESULTS FOR CIFAR100 DATASET SHOWN. SEE [9], [18] TO COMPARE SOME SIMILAR EXPERIMENTS FOR OTHER NETWORK REDUCTION TECHNIQUES.

NETWORK	$\lambda$	TEST ERROR	PARAMETERS	PRUNED %	FACTOR	NODES	PRUNED %
VGG 16	BASELINE	27.65	15.04M	0.0	1.0	4736	0.0
	$1.0 \times 10^{-6}$	27.69	9.78M	34.99	1.54	3914	17.35
	$1.0 \times 10^{-5}$	<b>28.04</b>	<b>1.83M</b>	<b>87.82</b>	<b>8.21</b>	<b>1623</b>	<b>65.73</b>
	$1.0 \times 10^{-4}$	38.49	0.46M	96.93	32.58	729	84.6
DENSENET40	BASELINE	26.5	1.05M	0.0	1.0	456	0.0
	$1.0 \times 10^{-6}$	26.92	1.05M	2.27	1.02	451	1.09
	$1.0 \times 10^{-5}$	<b>27.01</b>	<b>1.03M</b>	<b>4.74</b>	<b>1.05</b>	<b>445</b>	<b>2.41</b>
	$1.0 \times 10^{-4}$	29.38	0.744M	31.12	1.45	376	17.54

TABLE IV  
IMAGENET CLASSIFICATION RESULTS

NETWORK	$\lambda$	TEST ERROR	PARAMETERS	PRUNED %	FACTOR	NODES	PRUNED %
VGG 19	BASELINE	33.79	143.65M	0.0	1.0	14696	0.0
	$1.0 \times 10^{-5}$	34.85	23.75M	83.47	6.05	6670	54.61




Fig. 4. Accuracy stabilizes after less than 100 epochs in this CIFAR10 run, indicating the NodeDrop technique does not delay performance convergence. Training for another 400 epochs helps maximize parameter reduction. scheme where the training images are shifted and mirrored horizontally [18], [28].

2) *Architectures and Training*: We implement our technique on two standard models, VGG [29] and DenseNet [30]. Our VGG network is a slight variant of the standard VGG16 model. We follow the standard modification of VGG for CIFAR [18], [23], by removing the 3 final fully connected layers of size 4096 and instead using only a single fully connected layer of size 512. We train the network using SGD with momentum of 0.9. The network is trained for 200 epochs with an initial learning rate of 0.1 which is decayed by 0.1 at epochs 80 and 130. We tested both with and without batch normalization, and discovered that batch normalization is necessary for the large VGG16 initialization when applied to the more difficult CIFAR100 dataset. Therefore results without batch normalization are excluded for CIFAR100.

For DenseNet we implement the standard DenseNet-40 given in the original paper with  $L = 40$  and  $k = 12$ . We train the model as per the original paper with SGD and momentum 0.9.

The network is trained for 300 epochs with an initial learning rate of 0.1 and is decayed by 0.1 at epochs 150 and 225. As with VGG we found that the CIFAR100 dataset required batch normalization, but we were again able to train a variant on CIFAR10 without batch normalization.

3) *Lambda Parameter Tests*: As with the MNIST experiments, we tested a range of  $\lambda$ 's on CIFAR10 in order to determine the choices which suit the network and dataset well. Furthermore, here we test NodeDrop-BN, which was not tested in the MNIST experiments. Results for VGG on CIFAR10 with varying choices of  $\lambda$  are shown in figure 3.

For the case without batch normalization our network maintains performance and prunes a large number of nodes over many choices of  $\lambda$ . As with the MNIST case, this suggests that choosing  $\lambda$  is relatively easy. All choices of  $\lambda \leq 3.2 \times 10^{-5}$  achieved high performance with significant pruning. For  $\lambda \geq 1.0 \times 10^{-4}$  the regularization parameter proved too high, causing an entire layer to turn off, which in turn caused the network to turn off all other layers.

For NodeDrop-BN, we find that  $\lambda \leq 3.2 \times 10^{-5}$  is appropriate for maintaining performance. However, NodeDrop-BN requires more precise tuning than NodeDrop, as only  $\lambda \geq 3.2 \times 10^{-6}$  achieved desirable parameter reduction. Based on the above results we continue to recommend an initial lambda setting of  $\lambda = 1 \times 10^{-5}$  for the cross-entropy loss objective function.

4) *Network Convergence Time*: Sometimes it is important to avoid needlessly extending training time. In this section we analyze NodeDrop's effect on training time. Using  $\lambda = 10^{-5}$ , we train a network for 2000 epochs in order to observe network parameter and performance convergence over time. This experiment used the VGG16 network without batch normalization on the CIFAR10 dataset. Our results, shown in figure 4, indicate that while accuracy convergence is not

delayed by the NodeDrop technique, one will need to wait longer to maximize NodeDrop’s parameter reduction.

5) *Parameter Reduction*: Results for CIFAR10 and CIFAR100 are given in tables II and III respectively. We highlight the rows which provide the highest parameter reduction while maintaining high accuracy.

For the VGG network we are able to drop a significant number of parameters without degradation to the accuracy of the network. For NodeDrop-BN, we can prune 95 percent of the parameters for CIFAR10 and 88 percent for CIFAR100. For vanilla NodeDrop, we can prune 99 percent of the parameters on CIFAR10. This suggests that VGG is a significantly oversized network for application to the CIFAR datasets.

It is more difficult to prune nodes from the DenseNet architecture than for VGG. We are only able to prune approximately 5 percent of the parameters from DenseNet on CIFAR100. We believe this suggests that the DenseNet architecture is already well sized for CIFAR100. DenseNet starts at around 1 million parameters, which is close to the number of remaining parameters after our best case pruning of the VGG network.

### C. ImageNet Experiments

The ImageNet dataset consists of variously sized images designed for 1000 classes. For the ImageNet dataset, we used only the standard VGG19 network with batch normalization. We trained the networks for 50 epochs with the Adam optimizer to demonstrate NodeDrop’s effectiveness with equal training times. Due to the long training time of ImageNet we tested only the  $\lambda = 10^{-5}$  value and compared with a similarly trained baseline model. For the NodeDrop ImageNet experiments we started with 5 initialization epochs in which  $\lambda = 0$  before using  $\lambda = 10^{-5}$  for the remaining 45 epochs. We used these initialization epochs in order to allow for some network convergence before dropping nodes. The results for ImageNet, shown in table IV, indicate a pruning of 83.47 percent of the networks initial nodes, but there does seem to be a small one percent drop in accuracy.

## V. CONCLUSION AND OUTLOOK

In this paper, we proposed the novel NodeDrop technique for reducing parameters in neural networks. The NodeDrop technique consists of a condition for identifying nodes which are guaranteed to carry no information, and a regularization term to encourage this condition to be met. We also propose a modified version of NodeDrop, NodeDrop-BN, for use in networks with batch normalization. Experiments on the MNIST and CIFAR10 datasets show that NodeDrop does not significantly increase training time, and it facilitates network design with the easily tuneable hyperparameter  $\lambda$ . With experiments on MNIST, CIFAR10, CIFAR100, and ImageNet datasets, using VGG16, VGG19, and DenseNet architectures, we demonstrate that NodeDrop compares favorably with other parameter reduction techniques. NodeDrop reduces the number of parameters in a network by up to a factor of 114x. We hope that NodeDrop and NodeDrop-BN will prove useful in neural network design, and will help to make the implementation of neural networks on constrained systems more practical.



## REFERENCES

- [1] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *International Conference on Machine Learning (ICML)*, 2015.
- [2] B. Hassibi and D. G. Stork, "Second order derivatives for network pruning: Optimal brain surgeon," *Neural Information Processing Systems (NIPS)*, 1993.
- [3] Y. LeCun, J. S. Denker, S. Solla, R. E. Howard, and L. D. Jackel, "Optimal brain damage," *Neural Information Processing Systems (NIPS)*, 1990.
- [4] R. Reed, "Pruning algorithms-a survey," *IEEE transactions on Neural Networks*, 1993.
- [5] Y. Sun, X. Wang, and X. Tang, "Sparsifying neural network connections for face recognition," *Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [6] S. Han, H. Mao, and W. J. Dally, "Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding," *International Conference on Learning Representations (ICLR)*, 2016.
- [7] S. Srinivas and R. V. Babu, "Data-free parameter pruning for deep neural networks," *British Machine Vision Conference (BMVC)*, 2015.
- [8] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," *Neural Information Processing Systems (NIPS)*, 2015.
- [9] H. Li, A. Kadav, I. Durdanovic, H. Samet, and H. P. Graf, "Pruning filters for efficient convnets," *International Conference on Learning Representations (ICLR)*, 2017.
- [10] S. Anwar, K. Hwang, and W. Sung, "Structured pruning of deep convolutional neural networks," *ACM Journal on Emerging Technologies in Computing Systems (JETC)*, 2017.
- [11] E. L. Denton, W. Zaremba, J. Bruna, Y. LeCun, and R. Fergus, "Exploiting linear structure within convolutional networks for efficient evaluation," *Neural Information Processing Systems (NIPS)*, 2014.
- [12] Y. Gong, L. Liu, M. Yang, and L. Bourdev, "Compressing deep convolutional networks using vector quantization," *arXiv preprint arXiv:1412.6115*, 2014.
- [13] Y.-D. Kim, E. Park, S. Yoo, T. Choi, L. Yang, and D. Shin, "Compression of deep convolutional neural networks for fast and low power mobile applications," *International Conference on Learning Representations (ICLR)*, 2015.
- [14] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," *Neural Information Processing Systems (NIPS)*, 2016.
- [15] H. Zhou, J. M. Alvarez, , and F. Porikli, "Less is more: Towards compact cnns," *European Conference on Computer Vision (ECCV)*, 2016.
- [16] J. M. Alvarez and M. Salzmann, "Learning the number of neurons in deep networks," *Neural Information Processing Systems (NIPS)*, 2016.
- [17] V. Lebedev and V. Lempitsky, "Fast convnets using group-wise brain damage," *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [18] Z. Liu, J. Li, Z. Shen, G. Huang, S. Yan, and C. Zhang, "Learning efficient convolutional networks through network slimming," *International Conference on Computer Vision (ICCV)*, 2017.
- [19] M. C. I. Hubara, , D. Soudry, R. El-Yaniv, and Y. Bengio, "Binarized neural networks," *Neural Information Processing Systems (NIPS)*, 2016.
- [20] V. Vanhoucke, A. Senior, and M. Z. Mao, "Improving the speed of neural networks on cpus," *Workshop on Deep Learning and Unsupervised Feature Learning (NIPS)*, 2011.
- [21] S. Gupta, A. Agrawal, K. Gopalakrishnan, and P. Narayanan, "Deep learning with limited numerical precision," *International Conference on Machine Learning (ICML)*, 2015.
- [22] M. Rastegari, V. Ordonez, J. Redmon, and A. Farhadi, "Xnor-net: Imagenet classification using binary convolutional neural networks," *European Conference on Computer Vision (ECCV)*, 2016.
- [23] D. Molchanov, A. Ashukha, and D. P. Vetrov, "Variational dropout sparsifies deep neural networks," *International Conference on Machine Learning (ICML)*, 2017.
- [24] J. L. Ba, J. R. Kiros, and G. E. Hinton, "Layer normalization," *arXiv preprint arXiv:1607.06450*, 2016.
- [25] T. Salimans and D. P. Kingma, "Weight normalization: A simple reparameterization to accelerate training of deep neural networks," *Neural Information Processing (NIPS)*, 2016.
- [26] Y. LeCun and C. Cortes, "MNIST handwritten digit database," 1998. [Online]. Available: <http://yann.lecun.com/exdb/mnist/>
- [27] A. Krizhevsky, V. Nair, and G. Hinton, "Cifar-10 (canadian institute for advanced research)," 2009. [Online]. Available: <http://www.cs.toronto.edu/~kriz/cifar.html>
- [28] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [29] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *International Conference on Learning Representations (ICLR)*, 2014.
- [30] L. v. d. M. K. Q. W. Gao Huang, Zhuang Liu, "Densely connected convolutional networks," *Computer Vision and Pattern Recognition (CVPR)*, 2017.