

# Deep Expander Networks: Efficient Deep Networks from Graph Theory

Ameya Prabhu\*      Girish Varma\*      Anoop Namboodiri

Center for Visual Information Technology  
Kohli Center on Intelligent Systems, IIT Hyderabad, India  
[ameya.pandurang.prabhu@gmail.com](mailto:ameya.pandurang.prabhu@gmail.com), {[girish.varma](mailto:girish.varma@iiit.ac.in), [anoop](mailto:anoop@iiit.ac.in)}@iiit.ac.in  
<https://github.com/DrImpossible/Deep-Expander-Networks>

**Abstract.** Efficient CNN designs like ResNets and DenseNet were proposed to improve accuracy vs efficiency trade-offs. They essentially increased the connectivity, allowing efficient information flow across layers. Inspired by these techniques, we propose to model connections between filters of a CNN using graphs which are simultaneously sparse and well connected. Sparsity results in efficiency while well connectedness can preserve the expressive power of the CNNs. We use a well-studied class of graphs from theoretical computer science that satisfies these properties known as Expander graphs. Expander graphs are used to model connections between filters in CNNs to design networks called X-Nets. We present two guarantees on the connectivity of X-Nets: Each node influences every node in a layer in logarithmic steps, and the number of paths between two sets of nodes is proportional to the product of their sizes. We also propose efficient training and inference algorithms, making it possible to train deeper and wider X-Nets effectively.

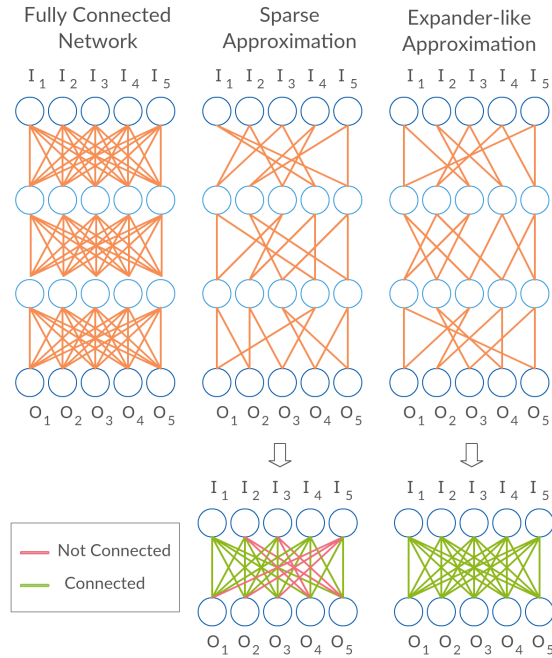
Expander based models give a 4% improvement in accuracy on MobileNet over grouped convolutions, a popular technique, which has the same sparsity but worse connectivity. X-Nets give better performance trade-offs than the original ResNet and DenseNet-BC architectures. We achieve model sizes comparable to state-of-the-art pruning techniques using our simple architecture design, without any pruning. We hope that this work motivates other approaches to utilize results from graph theory to develop efficient network architectures.

## 1 Introduction

Convolutional Neural Networks (CNNs) achieve state-of-the-art results in a variety of machine learning applications[1,2,3,4]. However, they are also computationally intensive and consume a large amount of computing power and runtime memory. After the success of VGG Networks [5], there has been significant interest in designing compact neural network architectures due to the wide range of applications valuing mobile and embedded devices based use cases.

---

\* indicates these authors contributed equally to this work.



**Fig. 1.** Popular sparse approximations are agnostic to the global information flow in a network, possibly creating disconnected components. In contrast, expander graph-based models produce sparse yet highly connected networks.

ResNet[6] and DenseNet-BC[3] directed the focus of efficient designs of convolutional layers on increasing connectivity. Additional connectivity with residual connections to previous layers provided efficient information flow through the network, enabling them to achieve an order of magnitude reduction in storage and computational requirements. We take inspiration from these approaches, to focus on designing highly connected networks. We explore making networks efficient by designing sparse networks that preserve connectivity properties. Recent architectures like MobileNet[7] improves the efficiency by an order of magnitude over a ResNet. However, in order to achieve this, they sparsify a network by removing several connections from a trained network, reducing their accuracies in the process. We ask a basic question: If we try to maximize the connectivity properties and information flow, can we achieve the same efficiency gains with minimal loss in accuracy? It is essential that the connections allow information to flow through the network easily. That is, each output node must at least have the capacity to be sensitive to features of previous layers. As we can see from Fig.1, traditional model compression techniques such as pruning can aggravate the problem, since they can prune the neuron connections of a layer, while being agnostic of global connectivity of the network. A necessary condition for having good representational power is efficient information flow through the network,

which is particularly suited to be modeled by graphs. We propose to make the connections between neurons (filters in the case of CNNs) according to specific graph constructions known as expander graphs. They have been widely studied in spectral graph theory [8] and pseudorandomness [9], and are known to be sparse but highly connected graphs. Expander graphs have a long history in theoretical computer science, also being used in practice in computer networks, constructing error correcting codes, and in cryptography (for a survey, see [10]).

**Main Contributions:** i.) We propose to represent neuronal connections in deep networks using expander graphs (see Section 3). We further prove that X-Nets have strong connectivity properties (see Theorem 1). ii.) We provide memory-efficient implementations of Convolutional (X-Conv) layers using sparse matrices and propose a fast expander-specific algorithm (see Section 4). iii.) We empirically compare X-Conv layers with grouped convolutions that have the same level of sparsity but worse connectivity. X-Conv layers obtain a 4% improvement in accuracy when both the techniques are applied to the MobileNet architecture trained on Imagenet (see Section 5.1). iv.) We also demonstrate the robustness of our approach by applying the technique to some of the state of the art models like DenseNet-BC and ResNet, obtaining better performance trade-offs (see Section 5.2). v.) Additionally, our simple design achieves comparable compression rates to even the state-of-the-art trained pruning techniques. (see Section 5.3). vi.) Since we enforce the sparsity before the training phase itself, our models are inherently compact and faster to train compared to pruning techniques. We leverage this and showcase the performance of wider and deeper X-Nets (see Section 5.5).

## 2 Related Work

Our approach lies at the intersection of trained pruning techniques and efficient layer design techniques. We present a literature survey regarding both the directions in detail.

### 2.1 Efficient Layer Designs

Currently there is extensive interest in developing novel convolutional layers/blocks and effectively leveraging them to improve architectures like [11,7,12]. Such micro-architecture design is in a similar direction as our work. In contrast, approaches like [4] try to design the macro-architectures by connecting pre-existing blocks. Recent concurrent work has been on performing architecture searches effectively [13,14,15,16]. Our work is complementary to architecture search techniques as we can leverage their optimized macro-architectures.

Another line of efficient architecture design is Grouped Convolutions: which was first proposed in AlexNet[1], recently popularized by MobileNets[7] and Xception[17] architectures. This is currently a very active area of current research, with a lot of new concurrent work being proposed [18,19,20].

It is interesting to note that recent breakthroughs in designing accurate deep networks [6,3,21] were mainly by introducing additional connectivity to enable

the efficient flow of information through deep networks. This enables the training of compact, accurate deep networks. These approaches, along with Grouped Convolutions are closely related to our approach.

## 2.2 Network Compression

Several methods have been introduced to compress pre-trained networks as well as train-time compression. Models typically range from low-rank decomposition [22,23,24] to network pruning [25,26,27,28].

There is also a major body of work that quantizes the networks at train-time to achieve efficiency [29,30,31,32,33,34,35]. The problem of pruning weights in train-time have been extensively explored [36,37] primarily from weight-level [38,39,40,41] to channel-level pruning [42,37,36]. Weight-level pruning has the highest compression rate while channel-level pruning is easier to practically exploit and has compression rates almost on par with the former. Hence, channel-level pruning is currently considered superior [42]. Channel-level pruning approaches started out with no guidance for sparsity [43] and eventually added constraints [44,45,46], tending towards more structured pruning.

However, to the best of our knowledge, this is the first attempt at constraining neural network connections by graph-theoretic approaches to improve deep network architecture designs. Note that we do not prune weights during training.

## 3 Approach

Recent breakthroughs in CNN architectures like ResNet[47] and DenseNet-BC[3] are ideas based on increasing connectivity, which resulted in better performance trade-offs. These works suggest that connectivity is an important property for improving the performance of deep CNNs. In that vein, we investigate ways of preserving connectivity between neurons while significantly sparsifying the connections between them. Such networks are expected to preserve accuracy (due to connectivity) while being runtime efficient (due to the sparsity). We empirically demonstrate this in the later sections.

### 3.1 Graphs and Deep CNNs

We model the connections between neurons as graphs. This enables us to leverage well-studied concepts from Graph Theory like Expander Graphs. Now, we proceed to formally describe the connection between graphs and Deep CNNs.

**Linear Layer defined by a Graph:** Given a bipartite graph  $G$  with vertices  $U, V$ , the Linear layer defined by  $G$ , is a layer with  $|U|$  input neurons,  $|V|$  output neurons and each output neuron  $v \in V$  is only connected to the neighbors given by  $G$ . Let the graph  $G$  be sparse, having only  $M$  edges. Then this layer has only  $M$  parameters as compared to  $|V| \times |U|$ , which is the size of typical linear layers.

**Convolutional Layer defined by a Graph:** Let a Convolutional layer be defined as a bipartite graph  $G$  with vertices  $U, V$  and a window size of  $c \times c$ .

This layer takes a 3D input with  $|U|$  channels and produces a 3D output with  $|V|$  channels. The output channel corresponding to a vertex  $v \in V$  is computed only using the input channels corresponding to the neighbors of  $v$ . Let  $G$  be sparse, having only  $M$  edges. Hence the kernel of this convolutional layer has  $M \times c \times c$  parameters as compared to  $|V| \times |U| \times c \times c$ , which is the number of parameters in a vanilla CNN layer.

### 3.2 Sparse Random Graphs

We want to constrain our convolutional layers to form a sparse graph  $G$ . Without any prior knowledge of the data distribution, we take inspiration from randomized algorithms and propose choosing the neighbours of every output neuron/channel uniformly and independently at random from the set of all its input channels. It is known that a graph  $G$  obtained in this way belongs to a well-studied category of graphs called Expander Graphs, known to be sparse but well connected.

**Expander Graph:** A bipartite expander with degree  $D$  and spectral gap  $\gamma$ , is a bipartite graph  $G = (U, V, E)$  ( $E$  is the set of edges,  $E \subseteq U \times V$ ) in which:

**1.) Sparsity:** Every vertex in  $V$  has only  $D$  neighbors in  $U$ . We will be using constructions with  $D \ll |U|$ . Hence the number of edges is only  $D \times |V|$  as compared to  $|U| \times |V|$  in a dense graph.

**2.) Spectral Gap:** The eigenvalue with the second largest absolute value  $\lambda$  of the adjacency matrix is bounded away from  $D$  (the largest eigenvalue). Formally  $1 - \lambda/D \geq \gamma$ .

**Random expanders:** A random bipartite expander of degree  $D$  on the two vertex sets  $U, V$ , is a graph in which for every vertex  $v \in V$ , the  $D$  neighbors are chosen independently and uniformly from  $U$ . It is a well-known result in graph theory that such graphs have a large spectral gap ([9]). Similar to random expanders, there exist several explicit expander constructions. More details about explicit expanders can be found in the supplementary section. We now proceed to give constructions of deep networks that have connections defined by an expander graph.


**Expander Linear Layer (X-Linear):** The Expander Linear (X-Linear) layer is a layer defined by a random bipartite expander  $G$  with degree  $D$ . The expander graphs that we use have values of  $D \ll |U|$ , while having an expansion factor of  $K \approx D$ , which ensures that the layer still has good expressive power.

**Expander Convolutional Layer (X-Conv):** The Expander Convolutional (X-Conv) layer is a convolutional layer defined by a random bipartite expander graph  $G$  with degree  $D$ , where  $D \ll |U|$ .

**Deep Expander Networks (X-Nets):** Given expander graphs

$$G_1 = (V_0, V_1, E_1), G_2 = (V_1, V_2, E_2), \dots, G_t = (V_{t-1}, V_t, E_t)$$

, we define the Deep Expander Convolutional Network (Convolutional X-Net or simply X-Net) as a  $t$  layer deep network in which the convolutional layers are replaced by X-Conv layers and linear layers are replaced by X-Linear layers defined by the corresponding graphs.



**Fig. 2.** The proposed fast convolution algorithm for X-Conv layer. We represent all the non-zero filters in the weight matrix of the X-Conv layer as a compressed dense matrix of  $D$  channels. The algorithm starts by selecting  $D$  channels from input (with replacement) using a mask created while initializing the model. The output is computed by convolving these selected channels with the compressed weight matrices.

### 3.3 Measures of Connectivity

In this subsection, we describe some connectivity properties of Expander graphs (see [9], for the proofs). These will be used to prove the properties of sensitivity and mixing of random walks in X-Nets .

**Expansion:** For every subset  $S \subseteq V$  of size  $\leq \alpha|V|$  ( $\alpha \in (0, 1)$  depends on the construction), let  $N(S)$  be the set of neighbors. Then  $|N(S)| \geq K|S|$  for  $K \approx D$ . That is, the neighbors of the vertices in  $S$  are almost distinct. It is known that random expanders have expansion factor  $K \approx D$  (see Theorem 4.4 in [9]).

**Small Diameter:** The diameter of a graph is the length of the longest path among all shortest paths. If  $G(U, V, E)$  is a  $D$ -regular expander with expansion factor  $K > 1$  and diameter  $d$ , then  $d \leq O(\log n)$ . This bound on the diameter implies that for any pair of vertices, there is a path of length  $O(\log n)$  in the graph.

**Mixing of Random Walks:** Random walks in the graph quickly converge to the uniform distribution over nodes of the graph. If we start from any vertex and keep moving to a random neighbor, in  $O(\log n)$  steps the distribution will be close to uniform over the set of vertices.

### 3.4 Sensitivity of X-Nets

X-Nets have multiple layers, each of which have connections derived from an expander graph. We can guarantee that the output nodes in such a network are sensitive to all the input nodes.

**Theorem 1 (Sensitivity of X-Nets).** *Let  $n$  be the number of input as well as output nodes in the network and  $G_1, G_2, \dots, G_t$  be  $D$  regular bipartite expander graphs with  $n$  nodes on both sides. Then every output neuron is sensitive to every input in a Deep X-Net defined by  $G_i$ 's with depth  $t = O(\log n)$ .*

*Proof.* For every pair of input and output  $(u, v)$ , we show that there is a path in the X-Net. The proof is essentially related to the fact that expander graphs have diameter  $O(\log n)$ . A detailed proof can be found in the supplementary material.

Next, we show a much stronger connectivity property known as mixing for the X-Nets. The theorem essentially says that the number of edges between subsets of input and output nodes is proportional to the product of their sizes. This result implies that the connectivity properties are uniform and rich across all nodes as well as subsets of nodes of the same size. Simply put, all nodes tend to have equally rich representational power.

**Theorem 2 (Mixing in X-Nets).** *Let  $n$  be the number of input as well as output nodes in the network and  $G$  be  $D$  regular bipartite expander graph with  $n$  nodes on both sides. Let  $S, T$  be subsets of input and output nodes in the X-Net layer defined by  $G$ . The number of edges between  $S$  and  $T$  is  $\approx D|S||T|/n$*

*Proof.* A detailed proof is provided in the supplementary material.

## 4 Efficient Algorithms

In this section, we present efficient algorithms of X-Conv layers. Our algorithms achieve speedups and save memory in the training as well as the inference phase. This enables one to experiment with significantly wider and deeper networks given memory and runtime constraints. We exploit the structured sparsity of expander graphs to design fast algorithms. We propose two methods of training X-Nets, both requiring substantially less memory and computational cost than their vanilla counterparts:

- 1) Using Sparse Representations
- 2) Expander-Specific Fast Algorithms.

### 4.1 Using Sparse Representation

The adjacency matrices of expander graphs are highly sparse for  $D \ll n$ . Hence, we can initialize a sparse matrix with non-zero entries corresponding to the edges of the expander graphs. Unlike most pruning techniques, the sparse connections are determined before training phase, and stay fixed. Dense-Sparse convolutions are easy to implement, and are supported by most deep learning libraries. CNN libraries like Cuda-convnet[48] support such random sparse convolution algorithms.


**Algorithm 1:** Fast Algorithm for Convolutions in X-Conv Layer

- 
- 1: For every vertex  $v \in \{1, \dots, n\}$ , let  $N(v, i)$  denote the  $i$ th neighbor of  $v$  ( $i \in \{1, \dots, D\}$ ).
  - 2: Let  $K_v$  be the  $c \times c \times D \times 1$  sized kernel associated with the  $v$ th output channel.
  - 3: Let  $O_v[x, y]$  be the output value of the  $v$ th channel at the position  $x, y$ .
  - 4: **for**  $v=1$  to  $n$  **do**
  - 5:      $O_v[x, y] = K_v * \text{Mask}_{N(v,1), \dots, N(v,D)}(I)[x, y]$ .
- 

**4.2 X-Net based Fast Dense Convolution**

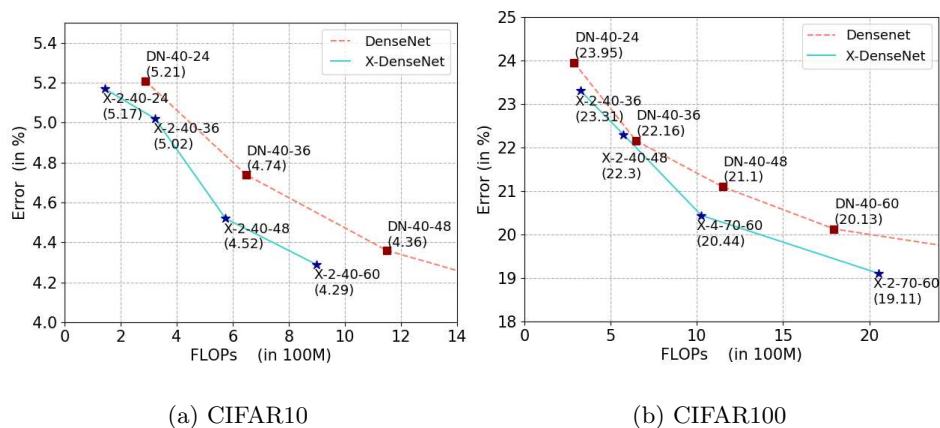
Next, we present fast algorithms that exploit the sparsity of expander graphs.

**X-Conv:** In an X-Conv layer, every output channel is only sensitive to *out*rom input channels. We propose to use a mask to select  $D$  channels of the input, and then convolve with a  $c \times c \times D \times 1$  kernel, obtaining a single channel per filter in the output. The mask is obtained by choosing  $D$  samples uniformly (without replacement) from the set  $\{1, \dots, N\}$ , where  $N$  is the number of input channels. The mask value is 1 for each of the selected  $D$  channels and 0 for others (see Algorithm 4.1). This is illustrated in Figure 2. There has been recent work about fast CUDA implementations called Block-Sparse GPU Kernels[49], which can implement this algorithm efficiently.

**5 Experiments and Results**

**Fig. 3.** Comparison between Grouped convolutions and X-Conv using MobileNet architecture trained on ImageNet. X- $d$  or G- $d$  represents the  $1 \times 1$  conv layers are compressed by  $d$  times using X-Conv or Groups. We observe X-MobileNets beat Group-MobileNet by 4% in accuracy on increasing sparsity.





**Fig. 4.** We show the error as a function of #FLOPs during test-time (below) for DenseNet-BC with X-DenseNet-BCs on CIFAR10 and CIFAR100 datasets. We observe X-DenseNet-BCs achieve better performance tradeoffs over DenseNet-BC models. For each datapoint, we mention the X-C-D-G notation (see Section 5.2) along with the accuracy.

In this section, we benchmark and empirically demonstrate the effectiveness of X-Nets on a variety of CNN architectures. Our code is available at: <https://github.com/DrImpossible/Deep-Expander-Networks>.

### 5.1 Comparison with Grouped Convolution

First, we compare our Expander Convolutions (X-Conv) against Grouped Convolutions (G-Conv). We choose G-Conv as it is a popular approach, on which a lot of concurrent works [18] have developed their ideas. G-Conv networks have the same sparsity as X-Conv networks but lack only the connectivity property. This will test whether increasing connectivity increases accuracy, i.e. does a graph without good connectivity properties provides worse accuracy? We choose MobileNet as the base model for this experiment, since it is the state-of-the-art in efficient CNN architectures. We compare X-Conv against grouped convolutions using MobileNet-0.5 on the ImageNet classification task. We replace the  $1 \times 1$  convolutional layers in MobileNet-0.5 with X-Conv layers forming X-MobileNet-0.5. Similarly, we replace them with G-Conv layers to form Group-MobileNet-0.5. Note that we perform this only in layers with most number of parameters (after the 8th layer as given in Table 1 of [7]). We present our results in Figure 3. The reference original MobileNet-0.5 has an error of 36.6% with a cost of 150M FLOPs. Additional implementation details are given in the supplementary material.


We can observe that X-MobileNets beat Group-MobileNets by over 4% in terms of accuracy when we increase sparsity. This also demonstrates that X-

Conv can be used to further improve the efficiency of even the most efficient architectures like MobileNet.

## 5.2 Comparison with Efficient CNN Architectures

In this section, we test whether Expander Graphs can improve the performance trade-offs even in state-of-the-art architectures such as DenseNet-BCs [3] and ResNets[47] on the ImageNet [50] dataset. We additionally train DenseNet-BCs on CIFAR-10 and CIFAR-100 [51] datasets to demonstrate the robustness of our approach across datasets.

Our X-ResNet-C-D is a  $D$  layered ResNet that has every layer except the first and last replaced by an X-Conv layer that compresses connections between it and the previous layer by a factor of  $C$ . We compare across various models like ResNets-34,50,101. Similarly, our X-DenseNet-BC-C-D-G architecture has depth  $D$ , and growth rate  $G$ . We use DenseNet-BC-121-32,169-32,161-48,201-32 as base models. These networks have every layer except the first and last replaced by an X-Conv layer that compresses connections between it and the previous layer by a factor of  $C$ . More details are provided in the supplementary material.



**Fig. 5.** We show the error as a function of #FLOPs to compare between ResNet and X-ResNet on the ImageNet dataset. We observe X-ResNets achieve better performance tradeoffs over original ResNet models.

Model	Accuracy	#FLOPs
ResNet		(in 100M)
X-ResNet-2-34	69.23%	35
X-ResNet-2-50	<b>72.85%</b>	<b>40</b>
ResNet-34	71.66%	70
X-ResNet-2-101	<b>74.87%</b>	<b>80</b>
ResNet-50	74.46%	80
ResNet-101	75.87%	160
DenseNet-BC		
X-DenseNet-BC-2-121	70.5%	28
X-DenseNet-BC-2-169	71.7%	33
X-DenseNet-BC-2-201	72.5%	43
X-DenseNet-BC-2-161	<b>74.3%</b>	<b>55</b>
DenseNet-BC-121	73.3%	55
DenseNet-BC-169	74.8%	65
DenseNet-BC-201	75.6%	85
DenseNet-BC-161	76.3%	110

**Table 1.** Results obtained by ResNet and DenseNet-BC models on ImageNet dataset, ordered by #FLOPs. or each datapoint, we use the X-C-D-G notation (see Section 5.2) along with the accuracy.

## 5.3 Comparison with Pruning Techniques

We plot the performance tradeoff of X-ResNets against ResNets in Figure 5 . We achieve significantly better performance tradeoffs compared to the original model. More specifically, we can reduce the #FLOPs in ResNets by half while incurring only 1-1.5% decrease in accuracy. Also, we can compare models with similar #FLOPs or accuracy with the help of Table 1. We observe that X-ResNet-2-50 has 43% fewer FLOPs than ResNet-34, but achieves a 1% improvement in

Method	Accuracy	#Params	Training
Li et al. [37]	93.4%	5.4M	✗
Liu et al. [42]	93.8%	2.3M	✗
X-VGG16-1	93.4%	1.65M (9x)	✓
X-VGG16-2	93.0%	1.15M (13x)	✓
VGG16-Orig	94.0%	15.0M	-

**Table 2.** Comparison with other methods on CIFAR-10 dataset using VGG16 as the base model. We significantly outperform popular compression techniques, achieving similar accuracies with upto 13x compression rate.

accuracy against it. Similarly, X-DenseNet-BC-2-161 has similar #FLOPs as DenseNet-BC-121, but achieves a 1% improvement in accuracy.

To further prove the robustness of our approach on DenseNet-BC, we test the same on CIFAR10 and CIFAR100, and plot the tradeoff curve in Figure 4. We observe that we can achieve upto 33% compression keeping accuracy constant on CIFAR-10 and CIFAR-100 datasets.

We compare our approach with methods which prune the weights during or after training. Our method can be thought of as constraining the weight matrices with a well studied sparse connectivity pattern even before the training starts. This results in fast training for the compact X-Conv models, while the trained pruning techniques face the following challenges:

- 1) Slow initial training due to full dense model.
- 2) Several additional phases of pruning and retraining.

Hence they achieve the compactness and runtime efficiency only in test time. Nevertheless we show similar sparsity can be achieved by our approach without explicitly pruning. We benchmark on VGG16 and AlexNet architectures since most previous results in the pruning literature have been reported on these architectures. In Table 2, we compare two X-VGG-16 models against existing pruning techniques. We achieve comparable accuracies to the previous state-of-the-art model with 50% fewer parameters and #FLOPs. Similarly, in Table 3 we compare X-AlexNet with trained pruning techniques on the Imagenet dataset. Despite having poor connectivity due to parameters being concentrated only in the last three fully connected layers, we achieve similar accuracy to AlexNet model using only 7.6M-9.7M parameters out of 61M, comparable to the state-of-the-art pruning techniques which have upto 3.4M-5.9M parameters. Additionally, it is possible to improve compression by applying pruning methods on our compact architectures, but pruning X-Nets is out of the scope of our current work.

Method	Accuracy	#Params	Training Speedup?
Network Pruning			
Collins et al.[52]	55.1%	15.2M	$\times$
Zhou et al. [45]	54.4%	14.1M	$\times$
Han et al. [31]	57.2%	6.7M	$\times$
Han et al. [31]	57.2%	6.7M	$\times$
Srinivas et al. [44]	56.9%	5.9M	$\times$
Guo et al. [41]	56.9%	3.4M	$\times$
X-AlexNet-1	55.2%	7.6M	$\checkmark$
X-AlexNet-2	56.2%	9.7M	$\checkmark$
AlexNet-Orig	57.2%	61M	-

**Table 3.** Comparison with other methods on ImageNet-2012 using AlexNet as the base model. We are able to achieve comparable accuracies using only 9.7M parameters.

#### 5.4 Stability of Models

We give empirical evidence as well as a theoretical argument regarding the stability of our method. For the vanilla DNN training, the weights are randomly initialized, and randomized techniques like dropouts, augmentation are used. Hence there is some randomness present and is well accepted in DNN literature prior to our method. We repeat experiments on different datasets (Imagenet and CIFAR10) and architectures (VGG, DenseNet and MobileNet0.5) to empirically show that the accuracy of expander based models has variance similar to vanilla DNN training over multiple runs.

We repeated the experiments with independent sampling of random expanders on the VGG and DenseNet baselines on the CIFAR10 dataset. The results can be seen in Table 4. It is noted that the accuracy values changes only by less than 0.3% across runs and the standard deviation of expander method is also comparable to the vanilla DNN training.

We also repeated experiments of our main result, which is the comparison with grouped convolutions on ImageNet dataset. We rerun the experiment with MobileNet0.5 feature extractor twice with Groups and the expander method. As can be seen from Table 5, the accuracy variations are comparable between the two models, and it is less than 1%.

A theoretical argument also concludes that choosing random graphs doesn't degrade stability. It is a well known result (See Theorem 4.4 in [9]) in random graph theory, that graphs chosen randomly are well connected with overwhelmingly high probability (with only inverse exponentially small error, due to the Chernoff's Tail bounds) and satisfies the Expander properties. Hence the chance that for a specific run, the accuracy gets affected due to the selection of a particularly badly connected graph is insignificant.

Model	Accuracy%	Max%	Min%
VGG	93.96±0.12	94.17	93.67
X-VGG-1	93.31±0.18	93.66	93.06
X-VGG-2	92.91±0.19	93.26	92.69
XDNetBC-40-24	94.41±0.19	94.63	94.18
XDNetBC-40-36	94.98±0.14	95.21	94.84
XDNetBC-40-48	95.49±0.15	95.65	95.28
XDNetBC-40-60	95.75±0.07	95.81	95.68

**Table 4.** The accuracies (mean  $\pm$  stddev) of various models over 10 training runs on CIFAR-10 dataset.

MobileNet Variant	Mean Accuracy	Range (Max-Min)
Base	63.39%	0.11%
G2	57.45 %	0.06 %
X2	58.22 %	0.14 %
G4	49.41 %	0.55 %
X4	54.00 %	0.53 %
G8	45.13 %	0.03 %
X8	49.23 %	0.60 %
G16	39.03 %	0.64 %
X16	44.63 %	0.18 %

**Table 5.** The mean accuracy and range of variation over 2 runs of MobileNet0.5 variants on ImageNet dataset.

## 5.5 Training Wider and Deeper networks

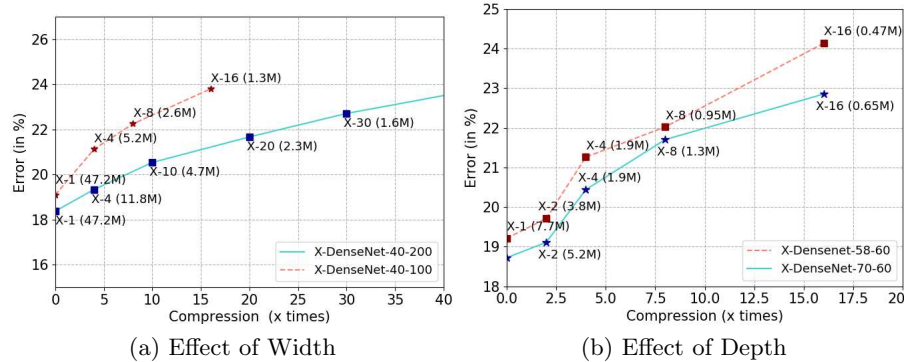
Since X-Nets involve constraining the weight matrices to sparse connectivity patterns before training, the fast algorithms can make it possible to utilize memory and runtime efficiently in training phase. This makes it possible to train significantly deeper and wider networks. Note the contrast with pruning techniques, where it is necessary to train the full, bulky model, inherently limiting the range of models that can be compressed.

Wide-DenseNets<sup>1</sup> offered a better accuracy-memory-time trade-off. We increase the width and depth of these networks to train significantly wider and deeper networks. The aim is to study whether leveraging the effectiveness of X-Nets in this fashion can lead to better accuracies. We widen and deepen the DenseNet-BC-40-60 architecture, increasing the growth rate from 60 to 100 and 200 respectively and compare the effect of increasing width on these new models. Similarly, we increase the depth from 40 to 58 and 70 to obtain deeper networks. We benchmark these approaches using CIFAR-100 dataset and present the results in Figure 6.

We have two interesting observations. First, the deeper X-DenseNet-BC-70-60 significantly outperforms X-DenseNet-BC-58-60 and wider X-DenseNet-40-200 outperforms X-DenseNet-BC-40-100 with fewer parameters for a wide range of  $C$  values (Expander degree).

The second interesting observation is the decreasing slope of the curves. This indicates that expander graph modeling seems to be effective on wider and deeper X-Nets i.e X-DenseNet-BC models suffer lesser penalty with increasing depth and width compression. This enables X-Nets to work at high compression rates of 30x, compressing DenseNet-BC-40-200 model from 19.9B FLOPs to 0.6B FLOPs with only 4.3% drop in accuracy. We hope this preliminary investigation holds significant value in alleviating the constraint of GPU memory and resources.

<sup>1</sup> <https://github.com/liuzhuang13/DenseNet#wide-densenet-for-better-timeaccuracy-and-memoryaccuracy-tradeoff>



**Fig. 6.** We show the performance tradeoff obtained on training significantly wider and deeper networks on CIFAR-100 dataset. Every datapoint is X-C specified along with the number of parameters, C being the compression factor. We show that training wider or deeper networks along with more compression using X-Nets achieve better accuracies with upto two-thirds of the total parameter and FLOPs on CIFAR-100 dataset.

## 6 Conclusion

We proposed a new network layer architecture for deep networks using expander graphs that give strong theoretical guarantees on connectivity. The resulting architecture (X-Net) is shown to be highly efficient in terms of both computational requirements and model size. In addition to being compact and computationally efficient, the connectivity properties of the network allow us to achieve significant improvements over the state-of-the-art architectures in performance on a parameter or run-time budget. In short, we show that the use of principled approaches that sparsify a model while maintaining global information flows can help in developing efficient deep networks.

To the best of our knowledge, this is the first attempt at using theoretical results from graph theory in modeling connectivity to improve deep network architectures. We believe that the field of deep networks can gain significantly from other similar explorations.

## References

1. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: NIPS. (2012) 1097–1105
2. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: ECCV, Springer (2016) 630–645
3. Huang, G., Liu, Z., Weinberger, K.Q., van der Maaten, L.: Densely connected convolutional networks. CVPR (2017)
4. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: CVPR. (June 2015) 1–9

5. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. *CoRR* **abs/1409.1556** (2014)
6. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: *CVPR*. (June 2016)
7. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017)
8. Spielman, D.A.: Spectral graph theory and its applications. In: *FOCS 2007*. (Oct 2007) 29–38
9. Vadhan, S.P.: Pseudorandomness. *Foundations and Trends in Theoretical Computer Science* **7**(13) (2012) 1–336
10. Hoory, S., Linial, N., Wigderson, A.: Expander graphs and their applications. *BULL. AMER. MATH. SOC.* **43**(4) (2006) 439–561
11. Iandola, F.N., Han, S., Moskewicz, M.W., Ashraf, K., Dally, W.J., Keutzer, K.: Squeezenet: Alexnet-level accuracy with 50x fewer parameters and <0.5mb model size. *ICLR* (2017)
12. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. *arXiv preprint arXiv:1709.01507* (2017)
13. Liu, H., Simonyan, K., Vinyals, O., Fernando, C., Kavukcuoglu, K.: Hierarchical representations for efficient architecture search. *arXiv preprint arXiv:1711.00436* (2017)
14. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. *arXiv preprint arXiv:1707.07012* (2017)
15. Zhong, Z., Yan, J., Liu, C.L.: Practical network blocks design with q-learning. *arXiv preprint arXiv:1708.05552* (2017)
16. Liu, C., Zoph, B., Shlens, J., Hua, W., Li, L.J., Fei-Fei, L., Yuille, A., Huang, J., Murphy, K.: Progressive neural architecture search. *arXiv preprint arXiv:1712.00559* (2017)
17. Chollet, F.: Xception: Deep learning with depthwise separable convolutions. *CVPR* (2017)
18. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. *arXiv preprint arXiv:1707.01083* (2017)
19. Sandler, M., Howard, A., Zhu, M., Zhmoginov, A., Chen, L.C.: Inverted residuals and linear bottlenecks: Mobile networks for classification, detection and segmentation. *arXiv preprint arXiv:1801.04381* (2018)
20. Huang, G., Liu, S., van der Maaten, L., Weinberger, K.Q.: Condensenet: An efficient densenet using learned group convolutions. *arXiv preprint arXiv:1711.09224* (2017)
21. Xie, S., Girshick, R., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: *CVPR, IEEE* (2017) 5987–5995
22. Sainath, T.N., Kingsbury, B., Sindhvani, V., Arisoy, E., Ramabhadran, B.: Low-rank matrix factorization for deep neural network training with high-dimensional output targets. In: *ICASSP, IEEE* (2013) 6655–6659
23. Novikov, A., Podoprikin, D., Osokin, A., Vetrov, D.P.: Tensorizing neural networks. In: *NIPS*. (2015) 442–450
24. Masana, M., van de Weijer, J., Herranz, L., Bagdanov, A.D., Malvarez, J.: Domain-adaptive deep network compression. *Network* **16** (2017) 30
25. Blundell, C., Cornebise, J., Kavukcuoglu, K., Wierstra, D.: Weight uncertainty in neural networks. *ICML* (2015)
26. Liu, B., Wang, M., Foroosh, H., Tappen, M., Pensky, M.: Sparse convolutional neural networks. In: *CVPR*. (2015) 806–814

27. He, Y., Zhang, X., Sun, J.: Channel pruning for accelerating very deep neural networks. *CVPR* (2017)
28. Molchanov, P., Tyree, S., Karras, T., Aila, T., Kautz, J.: Pruning convolutional neural networks for resource efficient inference. *ICLR* (2017)
29. Rastegari, M., Ordonez, V., Redmon, J., Farhadi, A.: Xnor-net: Imagenet classification using binary convolutional neural networks. In: *ECCV*. (2016)
30. Courbariaux, M., Bengio, Y.: Binarynet: Training deep neural networks with weights and activations constrained to +1 or -1. *ICML* (2016)
31. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *ICLR* (2016)
32. Wu, J., Leng, C., Wang, Y., Hu, Q., Cheng, J.: Quantized convolutional neural networks for mobile devices. In: *CVPR*. (2016) 4820–4828
33. Bagherinezhad, H., Rastegari, M., Farhadi, A.: Lcnn: Lookup-based convolutional neural network. *CVPR* (2017)
34. Zhu, C., Han, S., Mao, H., Dally, W.J.: Trained ternary quantization. *ICLR* (2017)
35. Zhou, S., Wu, Y., Ni, Z., Zhou, X., Wen, H., Zou, Y.: Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *ICLR* (2016)
36. Wen, W., Wu, C., Wang, Y., Chen, Y., Li, H.: Learning structured sparsity in deep neural networks. In: *NIPS*. (2016) 2074–2082
37. Li, H., Kadav, A., Durdanovic, I., Samet, H., Graf, H.P.: Pruning filters for efficient convnets. *ICLR* (2017)
38. Lebedev, V., Lempitsky, V.: Fast convnets using group-wise brain damage. In: *CVPR*. (2016) 2554–2564
39. Scardapane, S., Comminiello, D., Hussain, A., Uncini, A.: Group sparse regularization for deep neural networks. *Neurocomputing* **241** (2017) 81–89
40. Srinivas, S., Babu, R.V.: Data-free parameter pruning for deep neural networks. *BMVC* (2015)
41. Guo, Y., Yao, A., Chen, Y.: Dynamic network surgery for efficient dnns. In: *NIPS*. (2016) 1379–1387
42. Liu, Z., Li, J., Shen, Z., Huang, G., Yan, S., Zhang, C.: Learning efficient convolutional networks through network slimming. *ICCV* (2017)
43. Chen, W., Wilson, J., Tyree, S., Weinberger, K., Chen, Y.: Compressing neural networks with the hashing trick. In: *ICML*. (2015) 2285–2294
44. Srinivas, S., Subramanya, A., Babu, R.V.: Training sparse neural networks. In: *(CVPRW)*, *IEEE* (2017) 455–462
45. Zhou, H., Alvarez, J.M., Porikli, F.: Less is more: Towards compact cnns. In: *ECCV*, Springer (2016) 662–677
46. Yoon, J., Hwang, S.J.: Combined group and exclusive sparsity for deep neural networks. In: *ICML*. (2017) 3958–3966
47. He, K., Sun, J.: Convolutional neural networks at constrained time cost. In: *CVPR*, *IEEE* (2015) 5353–5360
48. Krizhevsky, A.: Cuda-convnet: High-performance c++/cuda implementation of convolutional neural networks. (2012)
49. Gray, S., Radford, A., Kingma, D.P.: Gpu kernels for block-sparse weights. *arXiv preprint arXiv:1711.09224* (2017)
50. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *CVPR 2009*, *IEEE* (2009) 248–255
51. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. (2009)
52. Collins, M.D., Kohli, P.: Memory bounded deep convolutional networks. *arXiv preprint arXiv:1412.1442* (2014)