

MaskConnect: Connectivity Learning by Gradient Descent

Karim Ahmed and Lorenzo Torresani

Department of Computer Science, Dartmouth College
karim@cs.dartmouth.edu, LT@dartmouth.edu

Abstract. Although deep networks have recently emerged as the model of choice for many computer vision problems, in order to yield good results they often require time-consuming architecture search. To combat the complexity of design choices, prior work has adopted the principle of modularized design which consists in defining the network in terms of a composition of topologically identical or similar building blocks (a.k.a. modules). This reduces architecture search to the problem of determining the number of modules to compose and how to connect such modules. Again, for reasons of design complexity and training cost, previous approaches have relied on simple rules of connectivity, e.g., connecting each module to only the immediately preceding module or perhaps to all of the previous ones. Such simple connectivity rules are unlikely to yield the optimal architecture for the given problem.

In this work we remove these predefined choices and propose an algorithm to learn the connections between modules in the network. Instead of being chosen a priori by the human designer, the connectivity is learned simultaneously with the weights of the network by optimizing the loss function of the end task using a modified version of gradient descent. We demonstrate our connectivity learning method on the problem of multi-class image classification using two popular architectures: ResNet and ResNeXt. Experiments on four different datasets show that connectivity learning using our approach yields consistently higher accuracy compared to relying on traditional predefined rules of connectivity. Furthermore, in certain settings it leads to significant savings in number of parameters.

Keywords: Connectivity Learning, Image Categorization

1 Introduction

Deep neural networks have emerged as one of the most prominent models for problems that require the learning of complex functions and that involve large amounts of training data. While deep learning has recently enabled dramatic performance improvements in many application domains, the design of deep architectures is still a challenging and time-consuming endeavor. The difficulty lies in the many architecture choices that impact—often significantly—the performance of the system. In the specific domain of image categorization, which is the focus of this paper, significant research effort has been invested in the

empirical study of how depth, filter sizes, number of feature maps, and choice of nonlinearities affect performance [1,2,3,4,5,6]. Recently, several authors have proposed to simplify the architecture design by defining convolutional neural networks (CNNs) in terms of composition of topologically identical or similar building blocks or *modules*. This strategy was arguably first popularized by the VGG nets [7] which were built by stacking a series of convolutional layers having identical filter size (3×3). Other examples are ResNets [8] which are constructed by stacking residual blocks of fixed topology, ResNeXt models [9] which use multi-branch residual block modules, DenseNets [10] which use dense blocks as building blocks, or Multi-Fiber networks [11] which use parallel branches (“fibers”) connected by routers (“transistors”).

While the principle of modularized design has greatly simplified the challenge of building effective architectures for image analysis, the choice of how to combine and aggregate the computations of these building blocks still rests on the shoulders of the human designer. To avoid a combinatorial explosion of options, prior work has relied on simple, uniform rules of aggregation and composition. For example, in ResNets and DenseNets each building block is connected only to the preceding one, via identity mapping, convolution or pooling. ResNeXt models [9] use a set of simplifying assumptions: the branching factor C (also referred to as *cardinality*) is fixed to the same constant in all layers of the network, all branches of a module are fed the same input, and the outputs of parallel branches are aggregated by a simple additive operation that provides the input to the next module. While these simple rules of connectivity render network design more manageable, they are unlikely to yield the optimal connectivity for the given problem.

In this paper we remove these predefined choices and propose an algorithm that learns to combine and aggregate building blocks of a neural network by directly optimizing connectivity of modules with respect to the given task. In this new regime, the network connectivity naturally arises as a result of training rather than being hand-defined by the human designer. While in principle this involves a search over an exponential number of connectivity configurations, our method can efficiently optimize the training loss with respect to connectivity using a variant of backpropagation. This is achieved by means of *connectivity masks*, i.e., learned binary parameters that act as “switches” determining the final connectivity in our network. The masks are learned together with the convolutional weights of the network, as part of a joint optimization with respect to the given loss function for the problem.

We evaluate our method on the problem of multi-class image classification using two popular modular architectures: ResNet and ResNeXt. We demonstrate that models with our learned connectivity consistently outperform the networks based on predefined rules of connectivity for the same budget of residual blocks (and parameters). An interesting byproduct of our approach is that, in certain settings, it can automatically identify modules that are superfluous, i.e., unnecessary or detrimental for the end objective. At the end of the optimization, these unused modules can be pruned away without impacting the learned hypothesis while reducing substantially the runtime and the number of parameters to store.

By recasting the training procedure as an optimization over learning weights *and* connectivity, our method effectively searches over a larger space of solutions. This yields networks achieving higher accuracy than those constrained to use pre-defined connectivities. The average training time overhead is moderate, ranging between 13% (for ResNet models) and 39% (for ResNeXt models) compared to learning using fixed connectivity which, however, yields lower accuracy. Finally we point out that, although our experiments are carried out using ResNet and ResNeXt models, our approach is general and applicable without major modifications to other forms of network architectures and other tasks beyond image categorization. In principle our method can also be used to learn connectivity among layers of a traditional (i.e., non-modular) neural network or a CNN. However, modern networks typically include a very large number of layers (hundreds or even thousands [12]), which would make our approach very costly. Learning connectivity among modules is more manageable as each module encapsulates many layers and thus the total number of modules is typically small even for deep networks.

2 Related Work

Despite their wide adoption, deep networks often require laborious model search in order to yield good results. As a result, significant research effort has been devoted to the design of algorithms for automatic model selection. However, most of this prior work falls within the genre of hyper-parameter optimization [13,14,15] rather than architecture or connectivity learning. Evolutionary search has been proposed as an interesting framework to learn both the structure as well as the connections in a neural network [16,17,18,19,20,21,22,23,24]. Architecture search has also been recently formulated as a reinforcement learning problem with impressive results [25]. Several authors have proposed learning connectivity by pruning unimportant weights from the network [26,27,28,29,30]. However, these prior methods operate in stages where initially a network with full connectivity is learned and then connections are greedily removed according to an importance criterion. Compare to all these prior approaches, our work provides the advantage of learning the connectivity by direct global optimization of the loss function of the problem at hand rather than by greedy optimization of an auxiliary proxy criterion or by costly evolutionary search. Our technical approach shares similarities with the “Shake-Shake” regularization [31]. This procedure was demonstrated on two-branch ResNeXt models and consists in randomly scaling tensors produced by parallel branches during training while at test time the network uses uniform weighting of tensors. Conversely, our algorithm *learns* an optimal binary scaling of the parallel tensors with respect to the training objective and uses the resulting network with sparse connectivity at test time. While our algorithm is limited to optimizing the connectivity structure within a predefined architecture, Adams et al. [32] proposed a nonparametric Bayesian approach that searches over an infinite network using MCMC. Our approach can be viewed as a middle ground between two extremes: using hand-defined networks versus learning/searching the full

architecture from scratch. The advantage is that our connectivity learning can be done without adding a significant training time overhead (only 13-39% depending on the architecture) compared to using fixed connectivity. The disadvantage is that the space of models considered by our approach is a lot more constrained than in the case of general architecture search. Saxena and Verbeek [33] introduced convolutional neural fabric which are learnable 3D trellises that locally connect response maps at different layers of a CNN. Similarly to our work, they enable optimization over an exponentially large family of connectivities, albeit different from those considered here. Finally, our approach is also related to conditional computation methods [34,35,36,37,38,39,40,41,42,43], which learn to drop out blocks of units. However, unlike these techniques, our algorithm learns a fixed, sparse connectivity that does *not* change with the input and thus it keeps the runtime cost and the number of used parameters constant.

3 Technical Approach

3.1 Modular architecture

We begin by defining the modular architecture that will be used by our framework. In order to present our method in its full generality, we will describe it in the context of a *general* modular architecture, which we will then instantiate in the form of the two models used in our experiments (ResNet and ResNeXt).

We assume that the general modular architecture consists of a stack of L modules. (When using ResNet the modules will be residual blocks, while for ResNeXt each module will consist of multiple parallel branches.) We denote with \mathbf{x}_j the input to the j -th module for $j = 1, \dots, L$. The input of each module is an activation tensor computed from one of the previous modules. We assume that the module implements a function $\mathcal{G}(\cdot)$ parameterized by learnable weights θ_j . The weights may for example represent the coefficients of convolutional filters. Thus, the output \mathbf{y}_j computed by the j -th module is given by $\mathbf{y}_j = \mathcal{G}(\mathbf{x}_j; \theta_j)$. In prior modular architectures, such as ResNet, ResNeXt and DenseNet, the connectivity between modules is hand-defined a priori according to a very simple rule: the input of a module is the output of the preceding module. In other words, $\mathbf{x}_j \leftarrow \mathbf{y}_{j-1}$. While this makes network design straightforward, it greatly limits the topology of architectures considered for the given task. In the next subsection we describe how to parameterize the architecture to remove these constraints and to enable connectivity learning in modular networks.

3.2 Masked architecture

We now introduce learnable *masks* defining the connectivity in the network. Specifically, we want to allow each module j to take input from one or more of the preceding modules $k = 1, \dots, j - 1$. To achieve this we define for each module a binary mask vector that controls the input pathway of that module. The binary mask vectors are learned jointly with the weights of the network. Let

$\mathbf{m}_j = [m_{j,1}, m_{j,2}, \dots, m_{j,j-1}]^\top \in \{0, 1\}^{j-1}$ be the binary mask vector defining the *active* input connections feeding the j -th module. If $m_{j,k} = 1$, then the activation volume produced by the k -th module is fed as input to the j -th module. If $m_{j,k} = 0$, then the output from the k -th module is ignored by the j -th module. The tensors from the *active* input connections are all added together (in an element-wise fashion) to form the input to the module. Thus, if we denote again with \mathbf{y}_k the output activation tensor computed by the k -th module, the input \mathbf{x}_j to the j -th module will be given by the following equation:

$$\mathbf{x}_j = \sum_{k=1}^{j-1} m_{j,k} \cdot \mathbf{y}_k \quad (1)$$

Then, the output of this module will be obtained through the usual computation, i.e., $\mathbf{y}_j = \mathcal{G}(\mathbf{x}_j; \theta_j)$. We note that under this model we no longer have predefined connectivity among modules. Instead, the mask \mathbf{m}_j now determines *selectively* for each module which outputs from the previous modules will be aggregated and form the input to the block. In this paper we constrain the aggregations of outputs from the active connections to be in the form of simple additions as this does not require new parameters. When different modules yield feature maps of different sizes, we use zero-padding shortcuts to increase the dimensions of feature tensors to the largest size (as in [8]). These shortcuts are parameter free. We leave to future work the investigation of more sophisticated, parameterized aggregation schemes.

We point out that depending on the constraints defined over \mathbf{m}_j , different interesting models can be realized. For example, by introducing the constraint that $\sum_k m_{j,k} = 1$ for each block j , then each module will receive input from only one of the preceding modules (since each $m_{j,k}$ must be either 0 or 1). At the other end of the spectrum, if we set $m_{j,k} = 1$ for all modules j, k , then all connections would be active. In our experiments we will demonstrate that the best results are typically achieved for values in between these two extremes, i.e., by connecting each module to K previous modules where K is an integer-valued hyperparameter such that $1 < K < (j - 1)$. We refer to this hyperparameter as the *fan-in* of a module. As discussed in the next section, the mask vector \mathbf{m}_j for each block is learned simultaneously with all the other weights in the network via backpropagation. Finally, we note that it may be possible for a module in the network to become unused. This happens when, as a result of the optimization, module k is such that $m_{j,k} = 0$ for all j . In this case, at the end of the optimization, we prune the module in order to reduce the number of parameters to store and to speed up inference (note that this does not affect the function computed by the network). In the next subsection we discuss our method for jointly learning the weights and the masks in the network.

3.3 MaskConnect: learning to connect

We refer to our learning algorithm as MaskConnect. It performs joint optimization of a given learning objective ℓ with respect to both the weights of the network

(θ) as well as the masks (\mathbf{m}). Since in this paper we apply our method to the problem of image categorization, we use the traditional multi-class cross-entropy objective for the loss ℓ . However, our approach can be applied without change to other loss functions and other tasks benefitting from connectivity learning.

In MaskConnect the weights have real values, as in traditional networks, while the masks have binary values. This renders the optimization challenging. To learn these binary parameters, we adopt a modified version of backpropagation, inspired by the algorithm proposed by Courbariaux et al. [44] to train neural networks with binary weights. During training we store and update a real-valued version $\tilde{\mathbf{m}}_j \in [0, 1]^{j-1}$ of the masks, with entries clipped to lie between 0 and 1.

In general, the training via backpropagation consists of three steps: 1) forward propagation, 2) backward propagation, and 3) parameters update. At each iteration, we stochastically binarize the real-valued masks into binary-valued vectors $\mathbf{m}_j \in \{0, 1\}^{j-1}$ which are then used for the forward propagation and backward propagation (steps 1 and 2). Instead, during the parameters update (step 3), the method updates the real-valued masks $\tilde{\mathbf{m}}_j$. The weights θ of the convolutional and fully connected layers are optimized using standard backpropagation. We discuss below the details of our mask training procedure, under the constraint that at any time there can be only K active entries in the binary mask \mathbf{m}_j , where K is a predefined integer hyperparameter with $1 \leq K \leq j - 1$. In other words, we impose the following constraints:

$$m_{j,k} \in \{0, 1\} \quad \forall j, k, \quad \text{and} \quad \sum_{k=1}^{j-1} m_{j,k} = K \quad \forall j.$$

These constraints imply that each module receives input from exactly K previous modules.

Forward Propagation. During the forward propagation, our algorithm first normalizes the real-valued entries in the mask of each block j to sum up to 1, such that $\sum_{k=1}^{j-1} \tilde{m}_{j,k} = 1$. This is done so that $\text{Mult}(\tilde{m}_{j,1}, \tilde{m}_{j,2}, \dots, \tilde{m}_{j,j-1})$ defines a proper multinomial distribution over the $j - 1$ possible input connections into module j . Then, the binary mask \mathbf{m}_j is stochastically generated by drawing K distinct samples $a_1, a_2, \dots, a_K \in \{1, \dots, (j - 1)\}$ from the multinomial distribution over the connections. Finally, the entries corresponding to the K samples are activated in the binary mask vector, i.e., $m_{j,a_k} \leftarrow 1$, for $k = 1, \dots, K$. The input activation volume to the module j is then computed according to Eq. 1 from the sampled binary masks. We note that the sampling from the Multinomial distribution ensures that the connections with largest $\tilde{m}_{j,k}$ values will be more likely to be chosen, while at the same time the stochasticity of this process allows different connectivities to be explored, particularly during early stages of the learning when the real-valued masks still have fairly uniform distributions.

Backward Propagation. In the backward propagation step, the gradient $\partial\ell/\partial y_k$ with respect to each output is obtained via back-propagation from $\partial\ell/\partial x_j$ and the binary masks $m_{j,k}$.

Mask Update. In the parameter update step our algorithm computes the gradient with respect to the binary masks for each module. Then, using these computed gradients and the given learning rate, it updates the real-valued masks via gradient descent. At this time we clip the updated real-valued masks to constrain them to remain within the valid interval $[0, 1]$ (as in [44]).

Pseudocode for our training procedure is given in the supplementary material. After joint training over θ and \mathbf{m} , we have found beneficial to (1) freeze the binary masks to the top- K values for each mask (i.e., by setting as active connections in \mathbf{m}_j those corresponding to the K largest values in $\tilde{\mathbf{m}}_j$) and then (2) fine-tune the weights θ of the network with respect to these fixed binary masks.

In the next subsections we discuss how we instantiated our general approach for the two architectures considered in our experiments: ResNet and ResNeXt.

3.4 MaskConnect applied to ResNet

The application of our algorithm to ResNets is quite straightforward. ResNets are modular networks obtained by stacking residual blocks. A residual block implements a residual function $\mathcal{F}(\cdot)$ with reference to the layer input. Figure 1(a)(left) illustrates an example of these modular components where the 3 layers in the block implement the residual function $\mathcal{F}(\mathbf{x}; \theta)$. A shortcut connections adds the residual block output $\mathcal{F}(\mathbf{x})$ to its input \mathbf{x} . Thus the complete function $\mathcal{G}(\cdot)$ implemented by a residual block computes $\mathcal{G}(\mathbf{x}; \theta) = \mathcal{F}(\mathbf{x}; \theta) + \mathbf{x}$. The ResNets originally introduced in [8] use a hand-defined connectivity that passes the output of a block to the immediately subsequent block, i.e., $\mathbf{x}_{j+1} \leftarrow \mathcal{F}(\mathbf{x}_j; \theta_j) + \mathbf{x}_j$. Here we propose to use MaskConnect to learn the input connections for each individual residual block in the network. This changes the input provided to block $j + 1$ in the network to be $\mathbf{x}_{j+1} \leftarrow \sum_{k=1}^j m_{j+1,k} [\mathcal{F}(\mathbf{x}_k; \theta_k) + \mathbf{x}_k]$, where binary parameters $m_{j+1,k}$ are learned automatically by our approach simultaneously with the weights θ subject to the constraint that $\sum_{k=1}^j m_{j+1,k} = K$. This implies that under our model each residual block now receives input from exactly K out of the preceding blocks. The output tensors from the K selected blocks are aggregated using element-wise addition and passed as input to the module. Our experiments present results for varying values of fan-in hyperparameter K , which controls the density of connectivity.

3.5 MaskConnect applied to multi-branch ResNeXt

The adaptation of MaskConnect to ResNeXt architectures is slightly more complex, as ResNeXt is based on a multi-branch topology. ResNeXt was motivated by the observation that it is beneficial to arrange residual blocks not only along the depth dimension but also to implement parallel multiple threads of computation feeding from the same input layer. The outputs of the parallel residual blocks are then summed up together with the original input and passed on to the next module. The resulting multi-branch module is illustrated in Figure 1(b)(left). More formally, let $\mathcal{F}(\mathbf{x}; \theta_j^{(i)})$ be the transformation implemented by the j -th

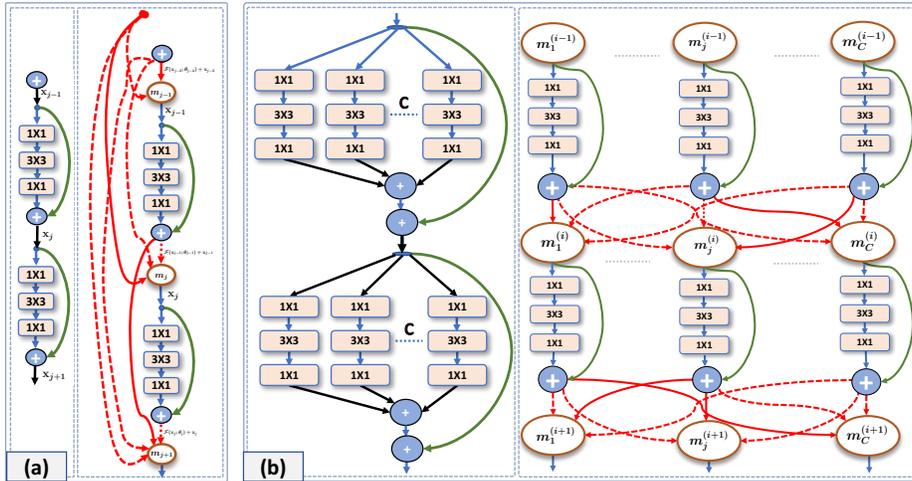


Fig. 1: Application of MaskConnect to two forms of modular network: **(a)** ResNet [45] and **(b)** multi-branch ResNeXt [9]. In traditional ResNet (a)(left) the connections between blocks are fixed (black links) so that each block receives input from only the preceding block. Our approach (a)(right) learns the optimal input connections (solid red links) for each individual block from a collection of potential connections (solid and dotted red links). Similarly, in traditional ResNeXt (b)(left) each module consists of C parallel residual blocks which are all aggregated and fed to the next module (black links). MaskConnect (b)(right) replaces the fixed aggregation points of ResNeXt with learnable masks \mathbf{m} defining the active input connections (solid red links) for each individual residual block.

residual block in module i -th of the network, where $j = 1, \dots, C$ and $i = 1, \dots, L$, with L denoting the total number of modules stacked on top of each other to form the complete network. The hyperparameter C is called the cardinality of the module and defines the number of parallel branches within each module. The hyperparameter L controls the total depth of the network. Then, in traditional ResNeXt, the output of the i -th module is computed as:

$$\mathbf{y}_i = \mathbf{x}_i + \sum_{j=1}^C \mathcal{F}(\mathbf{x}_i; \theta_j^{(i)}) \quad (2)$$

In [9] it was experimentally shown that increasing the cardinality C is a more effective way of improving accuracy compared to increasing depth or the number of filters. In other words, given a fixed budget of parameters, ResNeXt nets were shown to consistently outperform single-branch ResNets.

However, in an attempt to ease network design, a couple of restrictive limitations were embedded in the architecture of ResNeXt modules: (1) the C parallel

feature extractors in each module operate on the same input; (2) the number of active branches is constant at all depth levels of the network.

MaskConnect allows us to remove these restrictions without adding any significant burden on the process of manual network design, with the exception of a single additional integer hyperparameter (K) for the entire network. As in ResNeXt, our proposed architecture consists of a stack of L multi-branch modules, each containing C parallel feature extractors. However, differently from ResNeXt, each branch in a module can take a different input. The input pathway of each branch is controlled by a binary mask vector. Let $\mathbf{m}_j^{(i)} = [m_{j,1}^{(i)}, m_{j,2}^{(i)}, \dots, m_{j,C}^{(i)}]^\top \in \{0, 1\}^C$ be the binary mask vector defining the *active* input connections feeding the j -th residual block in module i . We note that under this model we no longer have fixed aggregation nodes summing up *all* outputs computed from a module. Instead, the mask $\mathbf{m}_j^{(i)}$ now determines *selectively* for each block which branches from the previous module will be aggregated to form the input to the next block. Under this new scheme, the parallel branches in a module receive different inputs and as such are likely to yield more diverse features.

As before, different constraints over $\mathbf{m}_j^{(i)}$ will give rise to different forms of architecture. By introducing the constraint that $\sum_k m_{j,k}^{(i)} = 1$ for all blocks j , then each residual block will receive input from only one branch (since each $m_{j,k}^{(i)}$ must be either 0 or 1). If instead we set $m_{j,k}^{(i)} = 1$ for all blocks j, k in each module i , then all connections would be active and we would obtain again the fixed ResNeXt architecture. In our experiments we present results obtained by varying the fan-in hyperparameter K such that $1 < K < C$. We also note that it may be possible for a residual block in the network to become unused, as a result of the optimization over the mask values. Thus, at any point in the network the total number of active parallel threads can be any number smaller than or equal to C . This implies that a variable branching factor is learned adaptively for the different depths in the network.

4 Experiments

We tested our approach on the task of image categorization using two different examples of modularized architecture: ResNet [8] and ResNeXt [9]. We used the following datasets for our evaluation: CIFAR-10 [46], CIFAR-100 [46], Mini-ImageNet [47], as well as the full ImageNet [48]. In this paper we include the results achieved on CIFAR-100 and ImageNet [48], while the results for CIFAR-10 [46] and Mini-ImageNet [47] (showing consistent improvements up to nearly 4% over fixed connectivity) can be found in the supplementary material.

4.1 CIFAR-100

CIFAR-100 contains images of size 32x32. It consists of 50,000 training images and 10,000 test images. Each image is labeled as belonging to one of 100 possible classes.

Table 1: CIFAR-100 accuracies achieved by models trained using the connectivity of **ResNet** [45] (Fixed-Prev), a fixed random connectivity (Fixed-Random), and the connectivity learned by our approach (Learned)

Model	Connectivity	Accuracy (%)
ResNet-38	Fixed-Prev, K=1 [45]	68.54
	Fixed-Random, K=10	62.67
	Learned, K=10	70.40
ResNet-74	Fixed-Prev, K=1 [45]	70.64
	Fixed-Random, K=15	66.93
	Learned, K=15	72.81
ResNet-110	Fixed-Prev, K=1 [45]	71.21
	Fixed-Random, K=20	67.22
	Learned, K=20	73.15

CIFAR-100 results based on the *ResNet* architecture.

Effect of fan-in (K). The fan-in hyperparameter (K) defines the number of *active* input connections feeding each residual block. We study the effect of the fan-in on the performance of models built and trained using our proposed approach. We use residual blocks consisting of two 3×3 convolutional layers. We use a model obtained by stacking $L = 18$ residual blocks with total depth of $D = 2 + 2L = 38$ layers. We trained and tested this architecture using different fan-in values: $K = 1, \dots, 17$. All these models have the same learning capacity as varying K does not affect the number of parameters. The results are shown in Figure 2.

We notice that the best accuracy is achieved using $K = 10$. Using a very low or very high fan-in yields lower accuracy. However, the algorithm does not appear to be overly sensitive to the fan-in hyperparameter, as a wide range of values for K (from $K = 7$ to $K = 13$) produce accuracy close to the best.

Varying the model. We trained several ResNet models differing in depth, using both MaskConnect as well as the traditional predefined connectivity. For these experiments we use a stack of L residual blocks with two 3×3 convolutional layers for each block. We choose $L \in \{18, 36, 54\}$ to build networks with depths $D = 2 + 2L$ equal to 38, 74, and 110 layers, respectively. We show the classification

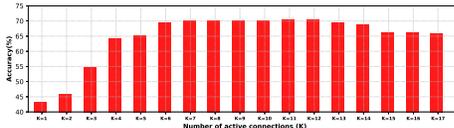


Fig. 2: Varying the fan-in (K), i.e., the number of learned active connections to each residual block. The plot reports accuracy achieved by MaskConnect on CIFAR-100 using a **ResNet-38** architecture ($L = 18$ blocks). All models have the same number of parameters (0.57M). The best accuracy is achieved at $K = 10$.

accuracy achieved by different models in Table 1. We report the results achieved using MaskConnect with fan-in $K = 10$, $K = 15$, $K = 20$ for models of depth $D = 38$, $D = 74$, $D = 110$, respectively. Fixed-Prev denotes the performance of ResNet, where each block is connected to only the previous block ($K = 1$). We also include the accuracy achieved by choosing a random connectivity (Fixed-Random) using the same fan-in values K as our approach and training the parameters while keeping the random connectivity fixed. This baseline is useful to show that our model achieves higher accuracy over traditional ResNet not because of the higher number of connections (i.e., $K > 1$), but rather because it learns the connectivity. Indeed, the results in Table 1 show that learning the connectivity using MaskConnect yields consistently higher accuracy than using multiple random connections or a single connection to the previous block.

CIFAR-100 results based on multi-branch *ResNeXt*.

Effect of fan-in (K). Even for ResNeXt, we start by studying the effect of the fan-in hyperparameter (K). For this experiment we use a model obtained by stacking $L = 6$ multi-branch residual modules, each having cardinality $C = 8$ (number of branches in each module). We use residual blocks consisting of 3 convolutional layers with a bottleneck implementing dimensionality reduction on the number of feature channels, as shown in Figure 1(b). The bottleneck for this experiment was set to $w = 4$. Since each residual block consists of 3 layers, the total depth of the network in terms of learnable layers is $D = 2 + 3L = 20$.

We trained and tested this architecture using different fan-in values: $K = 1, \dots, 8$. Again, varying K does not alter the number of parameters. The results are shown in Figure 3. We can see that the best accuracy is achieved by connecting each residual block to $K = 4$ branches out of the total $C = 8$ in each module. Note that when setting $K = C$, there is no need to learn the masks. In this case each mask is simply replaced by an element-wise addition of the outputs from all the branches. This renders the model equivalent to ResNeXt [9], which has fixed connectivity. Based on the results of Figure 3, in all our experiments below we use $K = 4$ (since it gives the best accuracy) but also $K = 1$ since it gives high sparsity which, as we will see shortly, implies savings in number of parameters.

Varying the models. In Table 2 we show the classification accuracy achieved with ResNeXt models of different depth and cardinality (the details of each model are listed in the Supplementary Material). For each architecture we also include the accuracy achieved with full (as opposed to learned) connectivity, which corresponds to ResNeXt. These results show that learning the connectivity produces consistently higher accuracy than using fixed connectivity, with accuracy gains of up to 2.2% compared to the state-of-the-art ResNeXt model. Furthermore, we can notice that the accuracy of models based on random connectivity (Fixed-Random) is considerably lower compared to our approach, despite having the same connectivity density ($K = 4$). This shows that the improvements of our approach over ResNeXt are not due to sparser connectivity but they are rather

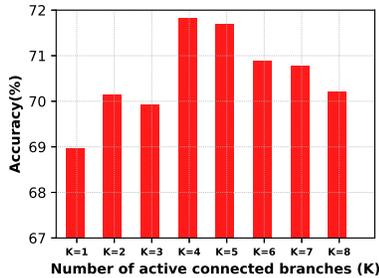


Fig. 3: Varying the fan-in (K) of our model, i.e., the number of active input branches to each residual block. The plot reports accuracy achieved on CIFAR-100 using a network stack of $L=6$ **ResNeXt** modules having cardinality $C=8$ and bottleneck width $w=4$. All models have the same number of parameters (0.28M).

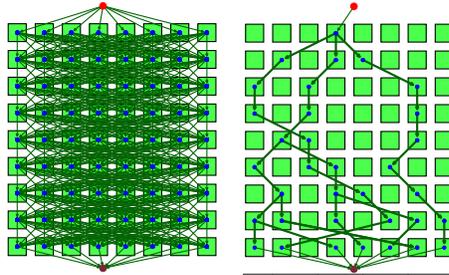


Fig. 4: A visualization of the fixed connectivity of **ResNeXt** (left) vs the connectivity learned by our method (right) using $K=1$. Each green square is a residual block, each row of $C=8$ square is a multi-branch module. Arrows indicate pathways connecting residual blocks of adjacent modules. It can be noticed that MaskConnect learns sparse connections. The squares without in/out edges are those pruned at the end of learning. This gives rise to a branching factor that varies along the depth of the net.

due to *learned* connectivity. We note that these improvements in accuracy come at little computational training cost: the average training time overhead for learning masks and weights is about 39% using our unoptimized implementation compared to learning only the weights given a fixed connectivity.

Parameter savings. Our proposed approach provides the benefit of automatically identifying residual blocks that are unnecessary. At the end of the training, the unused residual blocks can be pruned away. This yields savings in the number of parameters to store and in test-time computation. In Table 2, columns *Train* and *Test* under *Params* show the original number of parameters (used during training) and the number of parameters after pruning (used at test-time). Note that for the biggest architecture, our approach using $K=1$ yields a parameter saving of 40% compared to ResNeXt with full connectivity (20.5M vs 34.4M), while achieving the same accuracy. Thus, in summary, using fan-in $K=4$ gives models that have the same number of parameters as ResNeXt but they yield higher accuracy; using fan-in $K=1$ gives a significant saving in number of parameters and accuracy on par with ResNeXt.

Visualization of the learned connectivity. Figure 4 provides an illustration of the connectivity learned by MaskConnect for $K=1$ versus the fixed connectivity of ResNeXt for model $\{D=29, w=8, C=8\}$. While ResNeXt feeds the same

Table 2: CIFAR-100 accuracies achieved by two **ResNeXt** architectures trained using predefined full connectivity (Fixed-Full) [9], random connectivity (Fixed-Random, $K=4$), and the connectivity learned by our algorithm (Learned, $K=1$, $K=4$). Each model was trained 4 times, using different random initializations. We report the best test performance as well as the mean test performance computed from the 4 runs. We list the number of parameters used during training (Params-Train) and the number of parameters obtained after pruning the unused blocks (Params-Test). Our learned connectivity using $K=4$ produces accuracy gains of up to 2.2% compared to the strong ResNeXt model, while using $K=1$ yields results equivalent to ResNeXt but it induces a significant reduction in number of parameters at test time (e.g., a saving of 40% for model {29,64,8})

Architecture	Connectivity	Params		Accuracy (%)
		Train	Test	best (mean±std)
{Depth (D), Bottleneck width (w), Cardinality (C)}				
{29,8,8}	Fixed-Full, $K=8$ [9]	0.86M	0.86M	73.52 (73.37±0.13)
	Learned , $K=1$	0.86M	0.65M	73.91 (73.76±0.14)
	Learned , $K=4$	0.86M	0.81M	75.89 (75.77±0.12)
	Fixed-Random, $K=4$	0.86M	0.85M	72.85 (72.66±0.24)
{29,64,8}	Fixed-Full, $K=8$ [9]	34.4M	34.4M	82.23 (82.12±0.12)
	Learned , $K=1$	34.4M	20.5M	82.31 (82.15±0.15)
	Learned , $K=4$	34.4M	32.1M	84.05 (83.94±0.11)
	Fixed-Random, $K=4$	34.4M	34.3M	81.96 (81.73±0.20)

input to all blocks of a module, our algorithm learns different input pathways for each block and yields a branching factor that varies along depth.

4.2 ImageNet

Finally, we evaluate our approach on the large-scale ImageNet 2012 dataset [48], which includes images of 1000 classes. We train our approach on the training set (1.28M images) and evaluate it on the validation set (50K images).

ImageNet results based on the *ResNet* architecture. For this experiment we use a stack of $L=16$ residual blocks with 3 convolutional layers with a bottleneck architecture. Thus, the total number of layers is $D=2 + 3L=50$. Compared to the traditional ResNet using fixed connectivity, the same network trained using MaskConnect with fan-in $K=10$ yields a top-1 accuracy gain of 1.94% (78.09% vs 76.15%).

ImageNet results based on multi-branch *ResNeXt*. In Table 3, we report the top accuracies for three different ResNeXt architectures. For these experiments we set $K=C/2$. We can observe that for all three architectures, our learned connectivity yields an improvement in accuracy over fixed full connectivity [9].

Table 3: ImageNet accuracies (single crop) achieved by different architectures using the predefined connectivity of **ResNeXt** (Fixed-Full) versus the connectivity learned by our algorithm (Learned)

Architecture	Connectivity	Accuracy	
		<i>Top-1</i>	<i>Top-5</i>
{Depth (D), Bottleneck width (w), Cardinality (C)}			
{50,4,32}	Fixed-Full, K=32 [9]	77.8	93.3
	Learned , K=16	79.1	94.1
{101,4,32}	Fixed-Full, K=32 [9]	78.8	94.1
	Learned , K=16	79.5	94.5
{101,4,64}	Fixed-Full, K=64 [9]	79.6	94.7
	Learned , K=32	79.8	94.8

5 Conclusions

In this paper we introduced an algorithm to learn the connectivity of deep modular networks. The problem is formulated as a single joint optimization over the weights and connections between modules in the model. We tested our approach on challenging image categorization benchmarks where it led to significant accuracy improvements over the state-of-the-art ResNet and ResNeXt models using fixed connectivity. An added benefit of our approach is that it can automatically identify superfluous blocks, which can be pruned after training without impact on accuracy for more efficient testing and for reducing the number of parameters to store.

While our experiments were carried out on two particular architectures (ResNet and ResNeXt) and a specific form of building block (residual block), we expect the benefits of our approach to extend to other modules and network structures. For example, it could be applied to learn the connectivity of skip-connections in DenseNets [10], which are currently based on predefined connectivity rules. In this paper, our masks perform non-parametric additive aggregation of the branch outputs. It would be interesting to experiment with learnable (parametric) aggregations of the outputs from the individual branches. Our approach is limited to learning connectivity within a given, fixed architecture. Future work will explore the use of learnable masks for full architecture discovery.

Acknowledgements. This work was funded in part by NSF award CNS-120552. We gratefully acknowledge NVIDIA and Facebook for the donation of GPUs used for portions of this work.

References

1. Glorot, X., Bordes, A., Bengio, Y.: Deep sparse rectifier neural networks. In: Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2011, Fort Lauderdale, USA, April 11-13, 2011. (2011) 315–323
2. Krizhevsky, A., Sutskever, I., Hinton, G.E.: Imagenet classification with deep convolutional neural networks. In: Advances in Neural Information Processing Systems 25, Lake Tahoe, Nevada, United States. (2012) 1106–1114
3. Sermanet, P., Eigen, D., Zhang, X., Mathieu, M., Fergus, R., LeCun, Y.: Overfeat: Integrated recognition, localization and detection using convolutional networks. In: International Conference on Learning Representations (ICLR). (2013)
4. Maas, A.L., Hannun, A.Y., Ng, A.Y.: Rectifier nonlinearities improve neural network acoustic models. Proc. ICML **30** (2013) 1
5. Zeiler, M.D., Fergus, R.: Visualizing and understanding convolutional networks. In: Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I. (2014) 818–833
6. Szegedy, C., Liu, W., Jia, Y., Sermanet, P., Reed, S.E., Anguelov, D., Erhan, D., Vanhoucke, V., Rabinovich, A.: Going deeper with convolutions. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2015, Boston, MA, USA, June 7-12, 2015. (2015) 1–9
7. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. In: International Conference on Learning Representations (ICLR). (2015)
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Computer Vision and Pattern Recognition (CVPR), 2016 IEEE Conference on. (2016)
9. Xie, S., Girshick, R.B., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. (2017)
10. Huang, G., Liu, Z., Weinberger, K.Q.: Densely connected convolutional networks. In: IEEE Conference on Computer Vision and Pattern Recognition, CVPR. (2017)
11. Chen, Y., Kalantidis, Y., Li, J., Yan, S., Feng, J.: Multi-fiber networks for video recognition. In: European Conference on Computer Vision (ECCV). (2018)
12. He, K., Zhang, X., Ren, S., Sun, J.: Identity mappings in deep residual networks. In: Computer Vision - ECCV 2016 - 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part IV. (2016) 630–645
13. Bergstra, J., Bengio, Y.: Random search for hyper-parameter optimization. Journal of Machine Learning Research **13** (2012) 281–305
14. Snoek, J., Larochelle, H., Adams, R.P.: Practical bayesian optimization of machine learning algorithms. In: Advances in Neural Information Processing Systems 25, Lake Tahoe, Nevada, United States. (2012) 2960–2968
15. Snoek, J., Rippel, O., Swersky, K., Kiros, R., Satish, N., Sundaram, N., Patwary, M.M.A., Prabhat, Adams, R.P.: Scalable bayesian optimization using deep neural networks. In: Proceedings of the 32nd International Conference on Machine Learning, ICML 2015, Lille, France, 6-11 July 2015. (2015) 2171–2180
16. Pham, H., Guan, M.Y., Zoph, B., Le, Q.V., Dean, J.: Efficient neural architecture search via parameter sharing. arXiv preprint arXiv:1802.03268 (2018)
17. Such, F.P., Madhavan, V., Conti, E., Lehman, J., Stanley, K.O., Clune, J.: Deep neuroevolution: genetic algorithms are a competitive alternative for training deep neural networks for reinforcement learning. arXiv preprint arXiv:1712.06567 (2017)

18. Salimans, T., Ho, J., Chen, X., Sidor, S., Sutskever, I.: Evolution strategies as a scalable alternative to reinforcement learning. arXiv preprint arXiv:1703.03864 (2017)
19. Liu, H., Simonyan, K., Vinyals, O., Fernando, C., Kavukcuoglu, K.: Hierarchical representations for efficient architecture search. arXiv preprint arXiv:1711.00436 (2017)
20. Xie, L., Yuille, A.L.: Genetic cnn. In: ICCV. (2017) 1388–1397
21. Wierstra, D., Gomez, F.J., Schmidhuber, J.: Modeling systems with internal state using evoluno. In: Genetic and Evolutionary Computation Conference, GECCO 2005, Proceedings, Washington DC, USA, June 25-29, 2005. (2005) 1795–1802
22. Floreano, D., Dürr, P., Mattiussi, C.: Neuroevolution: from architectures to learning. *Evolutionary Intelligence* **1**(1) (2008) 47–62
23. Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y.L., Le, Q.V., Kurakin, A.: Large-scale evolution of image classifiers. CoRR **abs/1703.01041** (2017)
24. Fernando, C., Banarse, D., Blundell, C., Zwols, Y., Ha, D., Rusu, A.A., Pritzel, A., Wierstra, D.: Pathnet: Evolution channels gradient descent in super neural networks. CoRR **abs/1701.08734** (2017)
25. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. In: International Conference on Learning Representations (ICLR). (2017)
26. LeCun, Y., Denker, J.S., Solla, S.A.: Optimal brain damage. In: Advances in Neural Information Processing Systems 2, [NIPS Conference, Denver, Colorado, USA, November 27-30, 1989]. (1989) 598–605
27. Han, S., Mao, H., Dally, W.J.: Deep compression: Compressing deep neural network with pruning, trained quantization and Huffman coding. In: International Conference on Learning Representations (ICLR). (2015)
28. Han, S., Pool, J., Tran, J., Dally, W.J.: Learning both weights and connections for efficient neural network. In: Advances in Neural Information Processing Systems 28, Montreal, Quebec, Canada. (2015) 1135–1143
29. Guo, Y., Yao, A., Chen, Y.: Dynamic network surgery for efficient dnns. In: Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain. (2016) 1379–1387
30. Han, S., Pool, J., Narang, S., Mao, H., Tang, S., Elsen, E., Catanzaro, B., Tran, J., Dally, W.J.: DSD: regularizing deep neural networks with dense-sparse-dense training flow. In: International Conference on Learning Representations (ICLR). (2016)
31. Gastaldi, X.: Shake-shake regularization. CoRR **abs/1705.07485** (2017)
32. Adams, R.P., Wallach, H.M., Ghahramani, Z.: Learning the structure of deep sparse graphical models. In: Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics, AISTATS 2010, Chia Laguna Resort, Sardinia, Italy, May 13-15, 2010. (2010) 1–8
33. Saxena, S., Verbeek, J.: Convolutional neural fabrics. In: Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain. (2016) 4053–4061
34. Wu, Z., Nagarajan, T., Kumar, A., Rennie, S., Davis, L.S., Grauman, K., Feris, R.: Blockdrop: Dynamic inference paths in residual networks. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. (2018) 8817–8826
35. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation. arXiv preprint arXiv:1308.3432 (2013)

36. Bengio, E., Bacon, P.L., Pineau, J., Precup, D.: Conditional computation in neural networks for faster models. arXiv preprint arXiv:1511.06297 (2015)
37. Bengio, Y.: Deep learning of representations: Looking forward. In: International Conference on Statistical Language and Speech Processing, Springer (2013) 1–37
38. Shazeer, N., Mirhoseini, A., Maziarz, K., Davis, A., Le, Q., Hinton, G., Dean, J.: Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. arXiv preprint arXiv:1701.06538 (2017)
39. Davis, A., Arel, I.: Low-rank approximations for conditional feedforward computation in deep neural networks. arXiv preprint arXiv:1312.4461 (2013)
40. Eigen, D., Ranzato, M., Sutskever, I.: Learning factored representations in a deep mixture of experts. arXiv preprint arXiv:1312.4314 (2013)
41. Denoyer, L., Gallinari, P.: Deep sequential neural network. arXiv preprint arXiv:1410.0510 (2014)
42. Cho, K., Bengio, Y.: Exponentially increasing the capacity-to-computation ratio for conditional computation in deep learning. arXiv preprint arXiv:1406.7362 (2014)
43. Almahairi, A., Ballas, N., Cooijmans, T., Zheng, Y., Larochelle, H., Courville, A.: Dynamic capacity networks. In: International Conference on Machine Learning. (2016) 2549–2558
44. Courbariaux, M., Bengio, Y., David, J.: Binaryconnect: Training deep neural networks with binary weights during propagations. In: Advances in Neural Information Processing Systems 28, Montreal, Quebec, Canada. (2015) 3123–3131
45. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. CoRR **abs/1512.03385** (2015)
46. Krizhevsky, A.: Learning multiple layers of features from tiny images (2009) Technical Report <https://www.cs.toronto.edu/~kriz/learning-features-2009-TR.pdf>.
47. Vinyals, O., Blundell, C., Lillicrap, T., Kavukcuoglu, K., Wierstra, D.: Matching networks for one shot learning. In: Advances in Neural Information Processing Systems 29, Barcelona, Spain. (2016) 3630–3638
48. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Li, F.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR 2009), 20-25 June 2009, Miami, Florida, USA. (2009) 248–255