

Learning Time/Memory-Efficient Deep Architectures with Budgeted Super Networks

Tom Véniat* and Ludovic Denoyer*,†

*Sorbonne Université, LIP6, F-75005, Paris, France

†Criteo Research

{tom.veniat, ludovic.denoyer}@lip6.fr

Abstract

We propose to focus on the problem of discovering neural network architectures efficient both in terms of prediction quality and cost. For instance, our approach is able to solve the following tasks: learn a neural network able to predict well in less than 100 milliseconds or learn an efficient model that fits in a 50 Mb memory. Our contribution is a novel family of models called *Budgeted Super Networks*. They are learned using gradient descent techniques applied on a budgeted learning objective function which integrates a maximum authorized cost where this cost can be of different nature. We present a set of experiments on computer vision problems and analyze the ability of our technique to deal with three different costs: the computation cost, the memory consumption cost, and also a distributed computation cost. We particularly show that our model can discover neural network architectures that have a better accuracy than the ResNet and CNF architectures on CIFAR-10 and CIFAR-100, at a lower cost.

1. Introduction

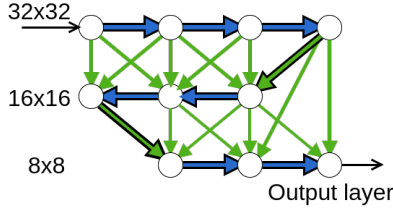
In the Deep Learning community, finding the best Neural Network architecture for a given task is a key problem that is mainly addressed *by hand* or using validation techniques. For instance, in computer vision, this has lead to particularly well-known models like GoogleNet [28] or ResNet [11]. More recently, there is a surge of interest in developing techniques able to automatically discover efficient neural network architectures. Different algorithms have been proposed including evolutionary methods [27, 20, 23] or reinforcement learning-based approaches [32]. But in all cases, this selection is usually based solely on a final predictive performance of the model such as the accuracy.

When facing real-world problems, this predictive perfor-

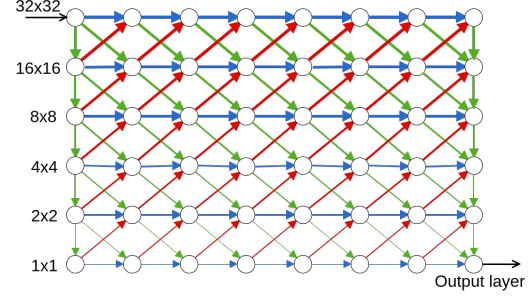
mance is not the only measure that matters. Indeed, learning a very good predictive model with the help of a cluster of GPUs might lead to a neural network that can be incompatible with low-resource mobile devices. Another example concerns distributed models in which one part of the computation is made *in the cloud* and the other part is made *on the device*. In such situations, an efficient architecture would have to predict accurately while minimizing the amount of exchanged messages between the cloud and the device. One important research direction is thus to propose models that can learn to take into account the inference cost in addition to the quality of the prediction.

We formulate this issue as a problem of automatically learning a neural network architecture under budgeted constraints. To tackle this problem, we propose a *budgeted learning approach* that integrates a maximum cost directly in the learning objective function. The main originality of our approach with respect to state-of-the-art is the fact that it can be used with any type of costs, existing methods being usually specific to particular constraints like inference speed or memory consumption – see Section 5 for a complete review of state-of-the-art. In our case, we investigate the ability of our method to deal with three different costs: (i) the *computation cost* reflecting the inference speed of the resulting model, (ii) the *memory consumption cost* that measures the final size of the model, and the (iii) *distributed computation cost* that measures the inference speed when computations are distributed over multiple machines or processors.

Our model called *Budgeted Super Network* (BSN) is based on the following principles: (i) the user provides a (big) *Super Network* (see Section 2) that defines a very large set of possible final network architectures as well as a maximum authorized cost. (ii) Since finding the best architecture that satisfies the cost constraint is an intractable combinatorial problem (Section 3.1), we relax this optimization



(a) **ResNet Fabric:** The ResNet Fabric is a super network that includes the ResNet model as a particular sub-graph. Each row corresponds to a particular size and number of feature maps. Each edge represents a simple *building block* (as described in [11]) i.e two stacked convolution layers + a shortcut connection. We use projection shortcuts (with 1x1 convolutions) for all connections going across different feature map sizes (green edges). Note that here, the subgraph corresponding to the bold edges is a ResNet-20. By increasing the width of the ResNet Fabric, we can include different variants of ResNets (from ResNet-20 with width 3 up to Resnet-110 with a width of 18).



(b) **Convolutional Neural Fabric [25]:** Each row corresponds to a particular resolution of feature maps. The number of features map is constant across the whole network. Each edge represents a convolution layer. The color of an edge represents the difference between input and output maps resolutions. Blue edges keep the same resolution, green edges decrease the resolution (stride > 1) and red edges increase the resolution (upsampling). Feature maps are aggregated (by addition) at each node before being sent to following layers.

Figure 1: This figure illustrates the two Super Networks on top of which cost-constrained architectures will be discovered. The ResNet Fabric is a generalization of ResNets, while CNF has been proposed in [25]. In the two cases, our objective is to discover architectures that are efficient both in prediction quality and cost, by sampling edges over these super networks.

problem and propose a stochastic model (called *Stochastic Super Networks* – Section 3.2) that can be optimized using policy gradient-inspired methods (Section 3.3). We show that the optimal solution of this stochastic problem corresponds to the optimal constrained network architecture (Proposition 1) validating our approach. At last, we evaluate this model on various computer vision tasks. We particularly show that, by taking inspiration from the *Residual Networks* (ResNet) and *Convolutional Neural Fabrics* (CNF), our model is able to discover new neural network architectures that outperform these baselines at a lower computation/memory/distributed cost (Section 4) on CIFAR-10 and CIFAR-100. The related work is presented in Section 5.

2. Super Networks

We consider the classical supervised learning problem defined by an input space \mathcal{X} and an output space \mathcal{Y} . In the following, input and output spaces correspond to multi-dimensional real-valued spaces (vectors, matrices or larger tensors). The training set is denoted as $\mathcal{D} = \{(x^1, y^1), \dots, (x^\ell, y^\ell)\}$ where $x^i \in \mathcal{X}$, $y^i \in \mathcal{Y}$ and ℓ is the number of supervised examples. At last, we consider a model $f : \mathcal{X} \rightarrow \mathcal{Y}$ that predicts an output given a particular input.

We first describe a family of models called *Super Net-*

works (S-network)¹ since our contribution presented in Section 3 will be a (stochastic) extension of this model. Note that the principle of Super Networks is not new and similar ideas have been already proposed in the literature under different names, e.g Deep Sequential Neural Networks [5], Neural Fabrics [25], or even PathNet [8] – and for different tasks.

A Super Network is composed of a set of layers connected together in a direct acyclic graph structure. Each edge is a (small) neural network, and the Super Network corresponds to a particular combination of these neural networks and defines a computation graph. Examples of S-networks are given in Figure 1. More formally, let us denote l_1, \dots, l_N a set of layers, N being the number of layers, such that each layer l_i is associated with a particular representation space \mathcal{X}_i which is a multi-dimensional real-valued space. l_1 will be the *input layer* while l_N will be the *output layer*. We also consider a set of (differentiable) functions $f_{i,j}$ associated to each possible pair of layers such that $f_{i,j} : \mathcal{X}_i \rightarrow \mathcal{X}_j$. Each function $f_{i,j}$ will be referred as a *module* in the following: it takes data from \mathcal{X}_i as inputs and transforms these data to \mathcal{X}_j . Note that each $f_{i,j}$ will make complex disk/memory/network operations and their computation will have consequences on the inference speed

¹The name *Super Network* comes from [8] which presents an architecture close to ours for a completely different objective.

of the S-network. Each $f_{i,j}$ module is associated with parameters in θ , θ being implicit in the notation for sake of clarity.

On top of this structure, a particular architecture $E = \{e_{i,j}\}_{(i,j) \in [1;N]^2}$ is a binary adjacency matrix over the N layers such that E defines a directed acyclic graph (DAG) with a single source node l_1 and a single sink node l_N . Different matrices E will thus correspond to different super network architectures. A S-network will be denoted (E, θ) in the following, θ being the parameters of the different modules, and E being the architecture of the super network.

Predicting with S-networks: The computation of the output $f(x, E, \theta)$ given an input x and a super network (E, θ) is made through a classic forward algorithm, the main idea being that the output of modules $f_{i,j}$ and $f_{k,j}$ leading to the same layer l_j will be added in order to compute the value of l_j . Let us denote $l_i(x, E, \theta)$ the value of layer l_i for input x , the computation is recursively defined as:

$$\text{Input: } l_1(x, E, \theta) = x$$

$$\text{Layer Computation: } l_i(x, E, \theta) = \sum_k e_{k,i} f_{k,i}(l_k(x, E, \theta)) \quad (1)$$

Note that learning the parameters θ can be made using classical back-propagation and gradient-descent techniques.

3. Learning Cost-constrained architectures

Our main idea is the following: we now consider that the structure E of the S-network (E, θ) describes not a single neural network architecture but a set of possible architectures. Indeed, each sub-graph of E (subset of edges) corresponds to a S-network. A sub-graph of edges will be denoted $H \odot E$ where H corresponds to a binary matrix in responsible for selecting the edges in E , \odot denoting the Hadamard product between H and E . Our objective will thus be to identify the best matrix H such that the corresponding S-network $(H \odot E, \theta)$ will be a network efficient both in terms of predictive quality and in terms of computation/memory/. cost.

The next sections are organized as follows: (i) First, we formalize this problem as a combinatorial problem where one wants to discover the best matrix H in the set of all possible binary matrices of size $N \times N$. Since this optimization problem is intractable, we propose a new family of models called *Stochastic Super Networks* where the edges of E are sampled following a parametrized distribution Γ before each prediction. We then show that the resulting budgeted learning problem is continuous and that its solution corresponds to the optimal solution of the initial budgeted problem (Proposition 1). We then propose a practical learning

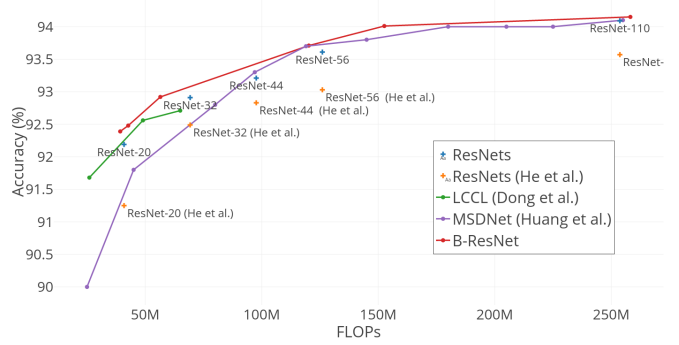


Figure 2: Accuracy/Time trade-off using B-ResNet on CIFAR-10.

algorithm to learn Γ and θ simultaneously by gradient descent.

3.1. Budgeted Architectures Learning

Let us consider H a binary matrix of size $N \times N$. Let us denote $C(H \odot E) \in \mathbb{R}^+$ the computation cost² associated to the computation of the S-network $(H \odot E, \theta)$ which can be of different nature (see Section 4). Let us also define \mathbf{C} the maximum computation cost the user would allow. For instance, when solving the problem of *learning a model with a computation time lower than 200 ms* then \mathbf{C} is equal to 200ms. We aim at solving the following soft constrained budgeted learning problem:

$$H^*, \theta^* = \arg \min_{H, \theta} \frac{1}{\ell} \sum_i \Delta(f(x^i, H \odot E, \theta), y^i) + \lambda \max(0, C(H \odot E) - \mathbf{C}) \quad (2)$$

where λ corresponds to the importance of the cost penalty. Note that the computational cost is specific to the particular infrastructure on which the model is ran. For instance, if C is the cost in milliseconds, the value of $\mathbf{C}(H \odot E)$ will not be the same depending on the device on which the model is used. While previous works make assumptions on the structure of the cost function [18], our approach does not rely on any constraint concerning $C(H \odot E)$ except the fact that this cost can be measured during training. Finding a solution to this learning problem is not trivial since it involves the computation of all possible architectures which is prohibitive ($\mathcal{O}(2^N)$ in the worst case). We explain in the next section how this problem can be solved using *Stochastic Super Networks*.

²Note that we consider that the cost only depends on the network architecture. The model could easily be extended to costs that depend on the input x to process, or to stochastic costs – see appendix

3.2. Stochastic Super Networks

Now, given a particular architecture E , we consider the following stochastic model – called **Stochastic Super Network** (SS-network) – that computes a prediction in two steps:

1. First a sub-graph H is sampled based on a distribution with parameters Γ . This operation will be denoted $H \sim \Gamma$
2. The final prediction is made using this sub-graph i.e by computing $f(x, H \odot E, \theta)$.

A SS-network is thus defined by a triplet (E, Γ, θ) , Γ and θ being both learned on a training set.

We can rewrite the budgeted learning objective of Equation 2 as:

$$\Gamma^*, \theta^* = \arg \min_{\Gamma, \theta} \frac{1}{\ell} \sum_i \mathbb{E}_{H \sim \Gamma} [\Delta(f(x^i, H \odot E, \theta), y^i) + \lambda \max(0, C(H \odot E) - \mathbf{C})] \quad (3)$$

Proposition 1 (proof in Appendix) *When the solution of Equation 3 is reached, then the models sampled following (Γ^*) and using parameters θ^* are optimal solution of the problem of Equation 2.*

Said otherwise, solving the stochastic problem will provide a model that has a good predictive performance under the given cost constraint.

Edge Sampling: In order to avoid inconsistent architectures where the input layer and the output one are not connected, we sample the edges following the given procedure: For each layer l_i visited following the DAG order of E (from the first layer to the last one) and for all $k < i$: If l_k is connected to the input layer l_1 based on the previously sampled edges, then $h_{k,i}$ is sampled following a Bernoulli distribution with probability³ $\gamma_{k,i}$. In the other cases then $h_{k,i} = 0$.

3.3. Learning Algorithm

We consider a generic case where the cost-function $C(H \odot E)$ is unknown and can be observed at the end of the computation of the model over an input x . Note that this case also includes stochastic costs where the cost is a random variable, caused by some network latency during distributed computation for example. We now describe the case where C is deterministic, its extension to stochastic values being simple (see appendix).

³Note that $\gamma_{k,i}$ is obtained by applying a *logistic* function over a continuous parameter value, but this is made implicit in our notations.

Let us denote $\mathcal{D}(x, y, \theta, E, H)$ the quality of the model $(H \odot E, \theta)$ on a given training pair (x, y) :

$$\mathcal{D}(x, y, \theta, E, H) = \Delta(f(x, H \odot E, \theta), y) + \lambda \max(0, C(H \odot E) - \mathbf{C}) \quad (4)$$

We propose to use a policy gradient inspired algorithm as in [5, 1] for learning. In that case, let us denote $\mathcal{L}(x, y, E, \Gamma, \theta)$ the expectation of \mathcal{D} over the possible sampled matrices H :

$$\mathcal{L}(x, y, E, \Gamma, \theta) = \mathbb{E}_{H \sim \Gamma} \mathcal{D}(x, y, \theta, E, H) \quad (5)$$

The gradient of \mathcal{L} can be written as⁴:

$$\begin{aligned} \nabla_{\theta, \Gamma} \mathcal{L}(x, y, E, \Gamma, \theta) &= \sum_H P(H|\Gamma) [(\nabla_{\theta, \Gamma} \log P(H|\Gamma)) \mathcal{D}(x, y, \theta, E, H)] \\ &\quad + \sum_H P(H|\Gamma) [\nabla_{\theta, \Gamma} \Delta(f(x, H \odot E, \theta), y)] \end{aligned} \quad (6)$$

The first term corresponds to the gradient over the log-probability of the sampled structure H while the second term is the gradient of the prediction loss given the sampled structure $H \odot E$.

Learning can be made using back-propagation and stochastic-gradient descent algorithms as it is made in Deep Reinforcement Learning models. Note that in practice, in order to reduce the variance of the estimator, the update is made following:

$$\begin{aligned} \nabla_{\theta, \Gamma} \mathcal{L}(x, y, E, \Gamma, \theta) &\approx (\nabla_{\theta, \Gamma} \log P(H|\Gamma)) (\mathcal{D}(x, y, \theta, E, H) - \tilde{\mathcal{D}}) \\ &\quad + \nabla_{\theta, \Gamma} \Delta(f(x, H \odot E, \theta), y) \end{aligned} \quad (7)$$

where H is sampled following Γ , and where $\tilde{\mathcal{D}}$ is the average value of $\mathcal{D}(x, y, \theta, E, H)$ computed on a batch of learning examples.

4. Experiments

4.1. Implementation

We study two particular super network architectures: (i) **ResNet Fabric (RNF)**: the first one is called ResNet Fabric (see Figure 1a) and is used for **image classification**. This super network is inspired by the ResNet[11] architecture on which extra modules (i.e. edges) have been added. The underlying idea is that a particular sub-graph of the ResNet Fabric corresponds exactly to a ResNet model. We thus aims at testing the ability of our approach to discover ResNet-inspired efficient architectures, or at least to converge to a ResNet model that is known to be efficient. (ii)

⁴details provided in appendix

Convolutional Neural Fabric (CNF): The CNF architecture has been proposed in [25]. It is a generic architecture that can be used for both **image classification** and **image segmentation**. The layers of the CNF super network are organized in a $W \times H$ matrix (see Figure 1b). We always use $W = 8$ when running our budgeted algorithm. Different values of W (as in [25]) are used as baselines.

Image classification has been tested on CIFAR-10 and CIFAR-100 [16] while the image segmentation has been performed on the Part Label dataset [15].

For these two architectures denoted B-CNF and B-ResNet, we consider three different costs functions: the first one is the (i) *computation cost* computed as the number of operations⁵ made by the resulting network as used in [7] or [14]. Note that this cost is highly correlated with the execution time⁶. The second one is the *memory consumption cost*, measured as the number of parameters of the resulting models. At last, the third cost (iii) is the *distributed computation cost* which is detailed in Section 4.3 and corresponds to the ability of a particular model to be efficiently computed over distributed architectures.

4.2. Experimental Protocol and Baselines

We have learned our model with various values of C . For the image classification problem, since we directly compare to ResNet, we have chosen values of C that corresponds to the cost of the ResNet-20/32/44/56/110 architectures. This allows us to compare the performance of our method at the same cost level than the ResNet variants. When dealing with the B-CNF model, we have chosen C to be a fraction of the cost of the (full) CNF model, allowing us to measure the quality of our model to discover low-cost architectures.

For each experiment, multiple versions of the different models are evaluated over the validation set during learning. Since our evaluation now involves both a cost and an accuracy, we select the best models using the pareto front on the cost/accuracy curve on the validation set. The reported performance are then obtained by evaluating these selected models over the test set. The detailed values of the hyperparameters and of the model selection procedure are given in Appendix and the source code of the model is publicly available⁷. The learning algorithm is a classical gradient-descent optimization technique, SGD with a momentum of 0.9 in our case.

For each experiment, we give the performance of both reference models (ResNet [11] and CNF [25]), and of related existing models i.e Low Cost Collaborative Layer

Model	FLOPs (millions)	Accuracy
ResNet [11]		<i>our/original</i>
ResNet-110	253.70	94.09/93.57
ResNet-56	126.01	93.61/93.03
ResNet-44	97.64	93.21/92.83
ResNet-32	69.27	92.91/92.49
ResNet-20	40.90	92.19/91.25
Low Cost Collaborative Layer [7]		
LCCL (ResNet-110)	166	93.44
LCCL (ResNet-44)	65	92.71
LCCL (ResNet-32)	49	92.56
LCCL (ResNet-20)	26	91.68
Multi Scale DenseNet [14] (values read on plot)		
MSDNet	≈ 255	94.1
	≈ 225	94.0
	≈ 205	94.0
	≈ 180	94.0
	≈ 145	93.8
	≈ 119	93.7
	≈ 97	93.3
	≈ 80	92.8
	≈ 45	91.8
	≈ 25	90.0
Budgeted ResNet		
B-ResNet	407.51	94.29
	258.20	94.15
	152.60	94.01
	120.20	93.71
	56.47	92.92
	42.69	92.48
	39.25	92.39
Convolutional Neural Fabrics [25]		<i>our/original</i>
CNF W=8	2,219.00	94.83/90.58
CNF W=4	1,010.00	93.75/87.91
CNF W=2	406.00	92.54/86.21
CNF W=1	54.00	89.91
Budgeted CNF		
B-CNF	2,150.00	94.92
	1,407.00	94.85
	1,144.00	94.69
	103.00	93.14
	85.00	92.17

Table 1: Accuracy/speed trade-off on CIFAR-10 using B-ResNet and B-CNF. Values reported as *our* corresponds to results we obtained when training a reproduction of the models, *original* corresponds to values from the original article.

(LCCL)[7] and MSDNet [14] (under the anytime classification settings). Note that the baselines methods have been

⁵The number of Multi-Add operations required to fully evaluate a network.

⁶Expressing constraint directly in *milliseconds* has been also investigated, with results similar to the ones obtain using the computation cost, and are not presented here.

⁷Blindreview

designed to reduce exclusively the *computation cost*, while our technique is able to deal with any type of cost. We provide the performance of our budgeted version of ResNet (B-ResNet) and of our budgeted version of CNF (B-CNF). Note that, for a fair comparison, we present the published results of ResNet and CNF, but also the ones that we have obtained by training these models by ourselves, **our results being of better quality** than the previously published performance.

4.2.1 Experimental results

Model	FLOPs (millions)	Accuracy (%)
ResNet-110	253.7	71.85
ResNet-56	126	70.57
ResNet-44	97.64	70.28
ResNet-32	69.27	69.28
ResNet-20	40.9	67.14
MSDNet [14]	215	76
	180	75
	150	74
	109	72.5
	80	71
	45	67.5
B-ResNet	15	62.5
	349.5	73.28
	115.09	71.46
	69.84	70.27
	64.96	70.12
	46.29	69.02
	39.22	68.45

Table 2: Accuracy/speed trade-off on Cifar-100 using ResNet Fabrics.

Reducing the computation cost: We first consider the *computation cost*. Figure 2 and Table 1 show the performance of different models over CIFAR-10. Each point corresponds to a model evaluated both in term of accuracy and *computation cost*. When considering the B-ResNet model, and by fixing the value of C to the computation costs of the different ResNet architectures, we obtain budgeted models that have approximatively the same costs than the ResNets, but with a higher accuracy. For example, ResNet-20 obtains an accuracy of 92.19% at a cost of 40.9×10^6 flop, while B-ResNet is able to discover an architecture with 92.39% accuracy at a slightly lower cost (39.25×10^6 flop). Moreover, the B-ResNet model also outperforms existing approaches like MSDNet or LCCL, particularly when the *computation cost* is low i.e for architectures that can be computed at a high speed. When comparing CNF to B-CNF, one can see that our approach is able to considerably reduce the computation cost while keeping a high accuracy. For example,

one of our learned models obtained an accuracy of 93.14% with a cost of 103×10^6 flop while CNF has an accuracy of 92.54% for a cost of 406×10^6 flop. Note that the same observations can be drawn for CIFAR-100 (Table 2).

Figure 3a and 3b illustrate two architectures discovered by B-ResNet and B-CNF with a low computation cost. One can see that B-ResNet has converged to an architecture which is a little bit different than the standard ResNet architecture, explaining why its accuracy is better. On the CNF side, our technique is able to extract a model that has a minimum of high-resolution convolutions operations, resulting in a high speedup.

Model	# params (millions)	Accuracy (%)
ResNet [11]		our/original
ResNet-110	1.73	94.09/93.57
ResNet-56	0.86	93.61/93.03
ResNet-44	0.66	93.21/92.83
ResNet-32	0.47	92.91/92.49
ResNet-20	0.27	92.19/91.25
Budgeted ResNet		
B-ResNet	4.38	94.35
	2.27	94.2
	1.29	93.85
	0.48	93.42
	0.34	92.72
	0.3	92.52
0.29	92.17	
Convolutional Neural Fabrics [25]		our/original
CNF W=8	18.04	94.83/90.58
CNF W=4	8.58	93.75/87.91
CNF W=2	3.85	92.54/86.21
CNF W=1	0.74	89.91
Budgeted CNF		
B-CNF	7.56	94.88
	4.98	94.58
	4.28	94.55
	3.67	94.42
	2.65	94.00
	1.19	93.53

Table 3: Accuracy/memory trade-off on Cifar-10 using B-ResNet and B-CNF.

Reducing the memory consumption: Similar experiments have been made considering the *memory consumption cost* that measures the number of parameters of the learned architectures. We want to demonstrate here the ability of our technique to be used with a large variety of costs, and not only to reduce the computation time. Table 3 illustrates the results obtained on CIFAR-10. As with the *computation cost*, one can see that our approach is able to

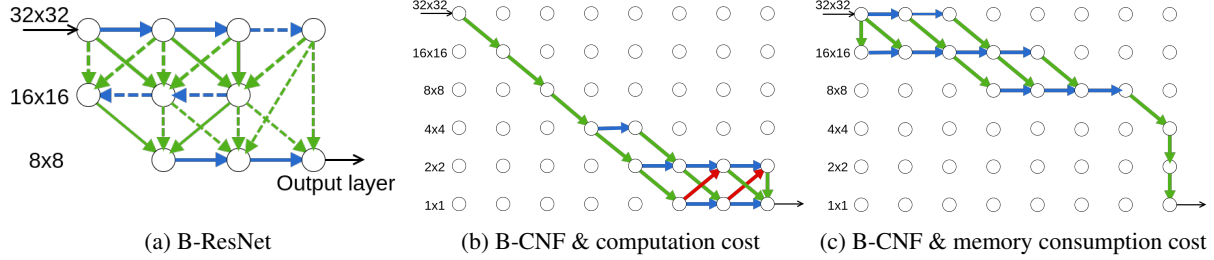


Figure 3: Discovered architectures: **(Left)** is a low *computation cost* B-ResNet where dashed edges correspond to connections in which the two convolution layers have been removed (only shortcuts connections are kept). **(Center)** is a low *computation cost* B-CNF where high-resolution operations have been removed. **(Right)** is a low memory consumption cost B-CNF: the algorithm has mostly kept all high resolution convolutions since they allow fine-grained feature maps and have the same number of parameters than lower-resolution convolutions. It is interesting to note that our algorithm, constrained with two different costs, automatically learned two different efficient architectures.

discover architectures that, given a particular memory cost, obtain a better accuracy. For example, for a model which size is ≈ 0.47 millions parameters, ResNet-32 has a classification error of 7.81% while B-ResNet only makes 6.58% error with ≈ 0.48 million parameters.

Image Segmentation: We also perform experiments on the image segmentation task using the Part Label dataset with CNF and B-CNF since CNF can also be used for this task – Table 4. In that case, the model computes a map of pixel probabilities, the output layer being now located at the top-right position of the CNF matrix. It is thus more difficult to reduce the overall computation cost. On the Part Label dataset, we are able to learn a BSN model with a computation gain of 40%. Forcing the model to reduce further the computation cost by decreasing the value of C results in inconsistent models i.e models where the input and output layers are not connected. At a Flop gain level of 40%, BSN obtains an error rate of 4.57%, which can be compared with the error of 4.94% for the full model. The B-CNF best learned architecture is given in appendix.

Model	FLOPs(billions)	Accuracy
CNF	35.614	95.06
CNF W=8 [25]	35.614	95.39
B-CNF	28.49	95.21
B-CNF	21.37	95.43

Table 4: Accuracy/Speed trade-off on Part Label using CNF.

Learning Dynamics: Figure 4 illustrates the learning dynamics of B-CNF and CNF. First, one can see (entropy curve) that the model becomes deterministic at the end of the learning iteration, and thus converges to a unique archi-

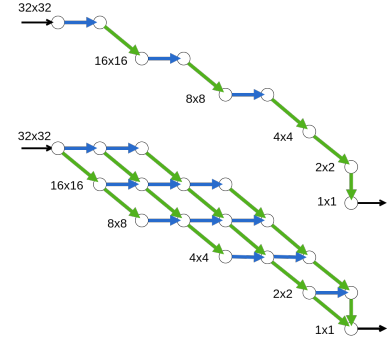
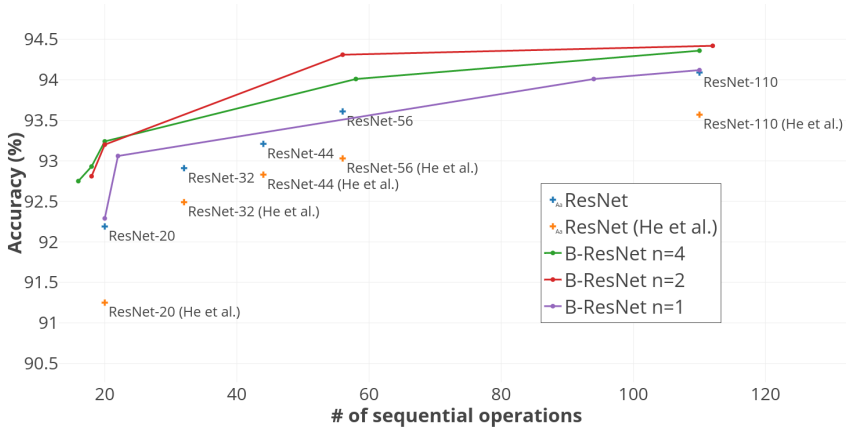


Figure 4: The figure illustrates the evolution of the loss function during learning iterations for CNF and B-CNF, and also the entropy of Γ . the period between iteration 0 and 50 is the burn-in phase, at iteration 150, the learning rate has been decreased to increase the convergence speed.

ture. Moreover, the training speed of B-CNF and CNF are comparable showing that our method does not result in a slower training procedure. Note that the figure illustrates the fact that during a burn-in period, we don’t update the edges probabilities, which allows us to obtain a faster convergence speed (see appendix).

4.3. Learning Distributed Architectures

At last, we perform a third set of experiments focused on distributed computing where different edges can be computed simultaneously on different computers/processors of a distributed platform. We thus evaluate the quality of an architecture by its ability to be efficiently parallelized. The *distributed computation cost* corresponds to the number of steps needed to compute the output of the network e.g on an architecture with $n = 2$ computers, depending on the structure of the network, two edges could be computed simultaneously. If the architecture is a sequence of layers,



(b) Architectures discovered with B-CNF for different number of cores: $n = 1$ (top) and $n = 4$ (bottom)

Figure 5: Architectures discovered on CIFAR-10 for different number of distributed cores n .

then this parallelization becomes impossible. These experiments allow us to measure the ability of BSN to handle complex costs that cannot be decomposed as a sum of individual modules costs as it is usually done in related works.

Results and corresponding architectures are illustrated in Figure 5 for the CIFAR-10 dataset and for both the B-ResNet and the B-CNF architectures. Note that ResNet is typically an architecture that cannot be efficiently distributed since it is a sequences of modules. One can see that our approach is able to find efficient architectures for $n = 2$ and $n = 4$. Surprisingly, when $n = 4$ the discovered architectures are less efficient, which is mainly due to an over fitting on the training set, the cost constraint becoming too large and stopping to act as a regularizer on the network architecture. On Figure 5b, one can see two examples of architectures discovered when $n = 1$ and $n = 4$. The shape of the architecture when $n = 4$ clearly confirm that BSN is able to discover parallelized architectures, and to ‘understand’ the structure of this complex cost.

5. Related Work

Learning cost-efficient models: One of the first approaches to learn efficient models is to *a posteriori* compress the learned network, typically by pruning some connections. The oldest work is certainly the Optimal Brain Surgeon [10] which removes weights in a classical neural network. The problem of network compression can also be seen as a way to speed up a particular architecture, for example by using quantization of the weights of the network [29], or by combining pruning and quantization [9]. Other algorithms include the use of hardware efficient operations that allow a high speedup [6].

Efficient architectures: Architecture improvements have been widely used in convolutional neural networks to improve cost efficiency of network components, some examples are the bottleneck units in the ResNet model [11], the use of depthwise separable convolution in Xception [3] and the lightweight MobileNets[13] or the combination of pointwise group convolution and channel shuffle in ShuffleNet[30].

End-to-end approaches: A first example of end-to-end approaches is the usage of quantization at training time: different authors trained models using binary weight quantization coupled with full precision arithmetic operations [4],[17]. Recently, [19] proposed an method using half precision floating numbers during training. Another technique proposed by [12], [24] and used in [31, 21] is the distillation of knowledge, which consists of training a smaller network to imitate the outputs of a larger network. Other approaches are dynamic networks which conditionally select the modules to respect a budget objective.[2, 22, 14, 1, 18].

Architecture Search: Different authors have proposed to provide networks with the ability to learn to select the computations that will be applied i.e choosing the right architecture for a particular task. This is the case for example for classification in [5, 32] based on Reinforcement learning techniques, in [26] based on gating mechanisms, in [23] based on evolutionary algorithms or even in [8] based on both RL and evolutionary techniques. With respect to existing methods, the strongest different is that we do not make any assumption concerning the nature of the cost which can be a computation cost, a memory cost, etc... Our model is thus more generic than existing techniques and allow to handle a large variety of problems.

6. Conclusion and Perspectives

We proposed a new model called *Budgeted Super Network* able to automatically discover cost-constrained neural network architectures by specifying a maximum authorized cost. The experiments in the computer vision domain show the effectiveness of our approach. Its main advantage is that BSN can be used for any costs (computation cost, memory cost, etc.) without any assumption on the shape of this cost. A promising research direction is now to study whether this model could be adapted in order to reduce the training time (instead of the test computation time). This could for example be done using meta-learning approaches.

Acknowledgments

This work has been funded in part by grant ANR-16-CE23-0016 “PAMELA” and grant ANR-16-CE23-0006 “Deep in France”.

References

- [1] E. Bengio, P. Bacon, J. Pineau, and D. Precup. Conditional computation in neural networks for faster models. *CoRR*, abs/1511.06297, 2015.
- [2] T. Bolukbasi, J. Wang, O. Dekel, and V. Saligrama. Adaptive neural networks for fast test-time prediction. *CoRR*, abs/1702.07811, 2017.
- [3] F. Chollet. Xception: Deep learning with depthwise separable convolutions. *CoRR*, abs/1610.02357, 2016.
- [4] M. Courbariaux, Y. Bengio, and J. David. Binaryconnect: Training deep neural networks with binary weights during propagations. *CoRR*, abs/1511.00363, 2015.
- [5] L. Denoyer and P. Gallinari. Deep sequential neural network. *CoRR*, abs/1410.0510, 2014.
- [6] J. Devlin. Sharp Models on Dull Hardware: Fast and Accurate Neural Machine Translation Decoding on the CPU, 2017.
- [7] X. Dong, J. Huang, Y. Yang, and S. Yan. More is less: A more complicated network with less inference complexity. *CoRR*, abs/1703.08651, 2017.
- [8] C. Fernando, D. Banarse, C. Blundell, Y. Zwols, D. Ha, A. A. Rusu, A. Pritzel, and D. Wierstra. Pathnet: Evolution channels gradient descent in super neural networks. *CoRR*, abs/1701.08734, 2017.
- [9] S. Han, H. Mao, and W. J. Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. In *International Conference on Learning Representations (ICLR’16 best paper award)*, 2015.
- [10] B. Hassibi and D. G. Stork. Second order derivatives for network pruning: Optimal brain surgeon. In *Advances in Neural Information Processing Systems 5, [NIPS Conference]*, pages 164–171, San Francisco, CA, USA, 1993. Morgan Kaufmann Publishers Inc.
- [11] K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. *CoRR*, abs/1512.03385, 2015.
- [12] G. E. Hinton, O. Vinyals, and J. Dean. Distilling the knowledge in a neural network. *CoRR*, abs/1503.02531, 2015.
- [13] A. G. Howard, M. Zhu, B. Chen, D. Kalenichenko, W. Wang, T. Weyand, M. Andreetto, and H. Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *CoRR*, abs/1704.04861, 2017.
- [14] G. Huang, D. Chen, T. Li, F. Wu, L. van der Maaten, and K. Q. Weinberger. Multi-scale dense convolutional networks for efficient prediction. *CoRR*, abs/1703.09844, 2017.
- [15] A. Kae, K. Sohn, H. Lee, and E. Learned-Miller. Augmenting CRFs with Boltzmann machine shape priors for image labeling. In *CVPR*, 2013.
- [16] A. Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009.
- [17] L. Lu. Toward computation and memory efficient neural network acoustic models with binary weights and activations. *CoRR*, abs/1706.09453, 2017.
- [18] M. McGill and P. Perona. Deciding how to decide: Dynamic routing in artificial neural networks. *CoRR*, abs/1703.06217, 2017.
- [19] P. Micikevicius, S. Narang, J. Alben, G. F. Diamos, E. Elsen, D. Garcia, B. Ginsburg, M. Houston, O. Kuchaiev, G. Venkatesh, and H. Wu. Mixed precision training. *CoRR*, abs/1710.03740, 2017.
- [20] R. Miiikkulainen, J. Z. Liang, E. Meyerson, A. Rawal, D. Fink, O. Francon, B. Raju, H. Shahrzad, A. Navruzryan, N. Duffy, and B. Hodjat. Evolving deep neural networks. *CoRR*, abs/1703.00548, 2017.
- [21] F. Nan and V. Saligrama. Adaptive Classification for Prediction Under a Budget. *ArXiv e-prints*, May 2017.
- [22] A. Odena, D. Lawson, and C. Olah. Changing model behavior at test-time using reinforcement learning. *CoRR*, abs/1702.07780, 2017.
- [23] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, Q. V. Le, and A. Kurakin. Large-scale evolution of image classifiers. *CoRR*, abs/1703.01041, 2017.
- [24] A. Romero, N. Ballas, S. E. Kahou, A. Chassang, C. Gatta, and Y. Bengio. Fitnets: Hints for thin deep nets. *CoRR*, abs/1412.6550, 2014.
- [25] S. Saxena and J. Verbeek. Convolutional neural fabrics. *CoRR*, abs/1606.02492, 2016.
- [26] R. K. Srivastava, K. Greff, and J. Schmidhuber. Highway networks. *CoRR*, abs/1505.00387, 2015.
- [27] K. O. Stanley and R. Miiikkulainen. Efficient reinforcement learning through evolving neural network topologies. In *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference, New York, USA, 9-13 July 2002*, pages 569–577, 2002.
- [28] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. E. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich. Going deeper with convolutions. *CoRR*, abs/1409.4842, 2014.
- [29] V. Vanhoucke, A. Senior, and M. Z. Mao. Improving the speed of neural networks on cpus. In *Deep Learning and Unsupervised Feature Learning Workshop, NIPS 2011*, 2011.
- [30] X. Zhang, X. Zhou, M. Lin, and J. Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. *CoRR*, abs/1707.01083, 2017.

- [31] H. Zhu, F. Nan, I. Paschalidis, and V. Saligrama. Sequential Dynamic Decision Making with Deep Neural Nets on a Test-Time Budget. *ArXiv e-prints*, May 2017.
- [32] B. Zoph and Q. V. Le. Neural architecture search with reinforcement learning. *CoRR*, abs/1611.01578, 2016.

Supplementary Material

Demonstration of Proposition 1

Let us consider the stochastic optimization problem defined in Equation 4. The schema of the proof is the following:

- First, we lower bound the value of Equation 4 by the optimal value of Equation 3
- Then we show that this lower bound can be reached by some particular values of Γ and θ in Equation 4. Said otherwise, the solution of Equation 4 is equivalent to the solution of 3.

Let us denote:

$$B(H \odot E, \theta, \lambda) = \frac{1}{\ell} \sum_i \Delta(f(x^i, H \odot E, \theta), y^i) + \lambda \max(0, C(H \odot E) - C) \quad (8)$$

Given a value of Γ , let us denote $\text{supp}(\Gamma)$ all the H matrices that can be sampled following Γ . The objective function of Equation 4 can be written as:

$$\begin{aligned} E_{H \sim \Gamma}[B(H, E, \theta, \lambda)] &= \sum_{H \in \text{supp}(\Gamma)} B(H \odot E, \theta, \lambda) P(H|\Gamma) \\ &\geq \sum_{H \in \text{supp}(\Gamma)} B((H \odot E)^*, \theta^*, \lambda) P(H|\Gamma) \\ &= B((H \odot E)^*, \theta^*, \lambda) \end{aligned} \quad (9)$$

where $(H \odot E)^*$ and θ^* correspond to the solution of:

$$(H \odot E)^*, \theta^* = \arg \min_{H, \theta} B(H \odot E, \theta, \lambda) \quad (10)$$

Now, it is easy to show that this lower bound can be reached by considering a value of Γ^* such that $\forall H \in \text{supp}(\Gamma), H \odot E = (H \odot E)^*$. This corresponds to a value of Γ where all the probabilities associated to edges in E are equal to 0 or to 1.

Segmentation architecture

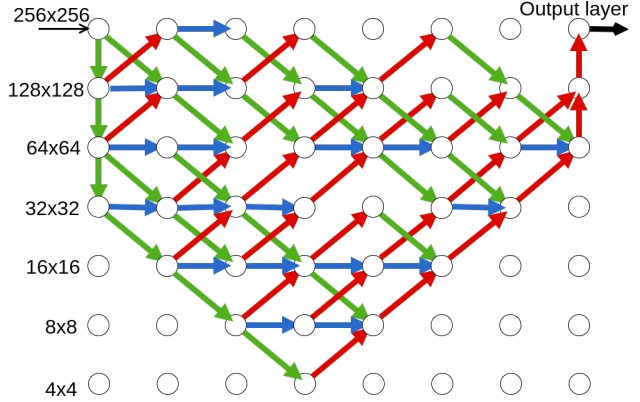


Figure 6: Segmentation architecture

Figure 6 is an example of segmentation architecture discovered on the Part Label dataset using the *flop cost*. It is interesting to note that only one layer with 256x256 input and output is kept and that most of the computations are done at lower less-expensive layers.

Model Selection Protocol

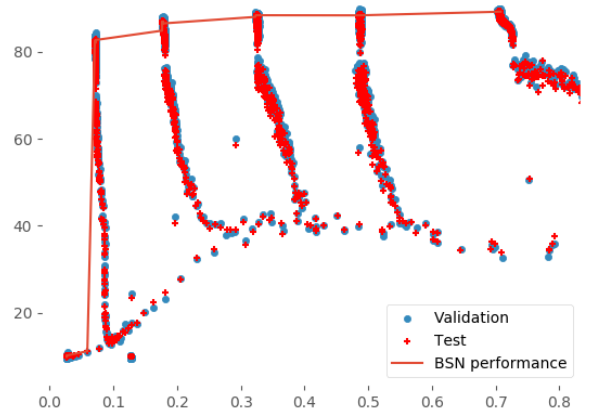


Figure 7: Model selection

The selection of reported models is obtained by learning many different models, computing the Pareto front of the accuracy/cost curve on the validation set, and reporting the performance obtained on the test set. This is illustrated in figure 7 where many different models are reported on the validation set (blue circles) with the corresponding performance on the test set (red crosses).

Considering non-differentiable costs

Stochastic costs in the REINFORCE algorithm:

As explained previously, the proposed algorithm can also be used when the cost $C(H \odot E)$ is a stochastic function that depends on the environment e.g the network latency, (or even on the input data x). Our algorithm is still able to learn with such stochastic costs since the only change in the learning objective is that the expectation is now made on both H and C (and x if needed). This property is interesting since it allows to discover efficient architecture on stochastic operational infrastructure.

Distributed computation cost Taking the real-life example of a network which will, once optimized, have to run on a given computing infrastructure, the *distributed computation cost* is a measure of how "parallelizable" an architecture is. This cost function takes the following three elements as inputs (i) A network architecture (represented as a graph for instance), (ii) An allocation algorithm and (iii) a maximum number of concurrent possible operations. The cost function then returns the number of computation cycles required to run the architecture given the allocation strategy.

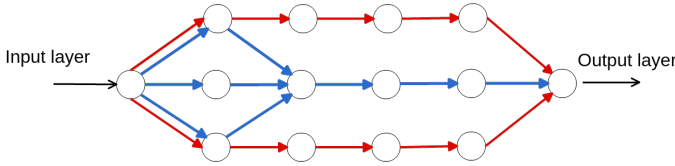


Figure 8: Two networks illustrating the need to have a cost function evaluating the global architecture of a network. Considering an environment with $n = 2$ machines performing computations in parallel, the blue network composed of 9 computational modules has a *distributed computation cost* of 6 while the red network, composed of 10 modules, has a smaller cost of 5.

Additional Architecture Details

ResNet Fabric

Based on the ResNet architecture, the structure of a ResNet Fabric is a stack of k groups of layers, each group being composed of $2n$ layers where n represents the width of the Fabric. The feature maps size and number of filters stay constant across the layers of each group and are modified between groups.

Due to its linear structure, the standard ResNet architecture spans a limited number of possible (sub-)architectures. In order to increase the size of the search space, we add several connections between groups as shown in 1a: each block in the second to last groups receives two (for the first

and last block of each group) or three (for every other block) inputs from preceding groups. To stay consistent with the rest of the network, each connection is a *basic block* [11] composed of 2 convolutional layers and a shortcut connection.

In our experiments, we use stacks of $k = 3$ blocks and $n = \{3, 5, 7, 9, 18\}$ to respectively include the ResNet- $\{20, 32, 44, 56, 110\}$ in the Fabric. Between each block, the feature maps size is reduced by a factor of 2 and the number of feature maps is doubled.

Convolutional Neural Fabric

The second network we use in our experiments is based on the dense Convolutional Neural Fabrics, which can be seen as a multi-layer and multi-scale convolutional neural network. As shown in Figure 1b, this architecture has 2 axis: The first axis represents the different columns (or width) W of the network while the second axis corresponds to different scales (or height) H of output feature maps, the first scale being the size of the input images, each subsequent scale being of a size reduced by a factor of 2 up to the last scale corresponding to a single scalar.

Each layer (l, s) in this fabric takes its input from three different layers of the preceding column: (i) One with a finer scale $(l - 1, s - 1)$ on which a convolution with stride 2 is applied to obtain feature maps having half the size of the input, (ii) one with the same scale $(l - 1, s)$ on which a convolution with stride 1 is applied to obtain feature map of the same resolution as the input and (iii) one with a coarser scale $(l - 1, s + 1)$ on which convolution with stride 1 is applied after a factor 2 up-sampling to obtain feature maps having twice the size of the input. The three feature blocks are then added before passing through the ReLU activation function to obtain the final output of this layer (l, s) .

The first and last columns are the only two which have vertical connections within scales of the same layer (as can be seen in Figure 1b). This is made to allow the propagation of the information to all nodes in the first column and to aggregate the activations of the last column to compute the final prediction. A more detailed description of this architecture can be found in the CNF original article.

We used two different Convolutional Neural Fabrics in our experiments: One for the classification task (CIFAR-10 and CIFAR-100) with $W = 8$ columns, $H = 6$ scales and 128 filters per convolution and one for the segmentation task (Part Label) with $W = 8$ layers, $H = 9$ scales (from 256×256 to 1×1 feature map sizes) and 64 filters per convolution.

Additional Learning Details

Datasets

CIFAR-10. The CIFAR-10 dataset consists of 60k 32x32 images with 10 classes and 6000 images per class. The dataset is decomposed in 50k training and 10k testing images. We split the training set following the standard, i.e 45k training samples and 5k validation samples. We use two data augmentation techniques: padding the image to 36x36 pixels before extracting a random crop of size 32x32 and horizontally flipping. Images are then normalized in the range $[-1,1]$.

CIFAR100. The CIFAR-100 dataset is similar to CIFAR-10, with 100 classes and 600 images per class. We use the same train/validation split and data augmentation technique as with CIFAR-10.

Part Labels. The Part Labels dataset is a subset of the LFW dataset composed of 2927 250x250 face images in which each pixel is labeled as one of the Hair/Skin/Background classes. The standard split contains 1500 training samples, 500 validation samples and 927 test samples. Images are zero-padded from 250x250 to 256x256. We use horizontal flipping as data augmentation. Images are then normalized in the range $[-1,1]$.

Learning procedure

When training our budgeted models, we first train the network for 50 "warm-up" epochs during which no sampling is done (The whole super network is trained). After this warm-up phase, the probability of each edge is initialized and we start sampling architectures.

The real-valued distribution parameter associated with each layer (and used to generate the probability of sampling the edge) are all initialized to 3, resulting in a ≈ 0.95 initial probability once passed through the sigmoid activation function.

On CIFAR-10 and CIFAR-100 datasets we train all models for 300 epochs. We start with a learning rate of 10^{-1} and divide it by 10 after 150 and 225 epochs. On Part Label dataset all models are trained for 200 epochs with a learning rate initialized to 10^{-1} and divided by 10 after 130 epochs.

For all models and all cost functions, we select the λ hyper-parameter based on the order of magnitude m of the maximum authorized cost C . λ is determined using cross-validation on values logarithmically spaced between 10^{m-1} and 10^{m+1} .

Forward algorithm

Given the SS-Network (E, Γ, θ) and input x , the evaluation of $f(x, E, \Gamma, \theta)$ is done as follow :

Algorithm 1 Stochastic Super Network forward algorithm

```

1: procedure SSN-FORWARD( $x, E, \Gamma, \theta$ )
2:    $H \sim \Gamma$  ▷ as explained in Section 3.2
3:   for  $i \in [1..N]$  do
4:      $l_i \leftarrow \emptyset$ 
5:   end for
6:    $l_1 \leftarrow x$ 
7:   for  $i \in [2..N]$  do
8:      $l_i \leftarrow \sum_{k < i} e_{k,i} h_{k,i} f_{k,i}(l_k)$ 
9:   end for
10:  return  $l_N$ 
11: end procedure

```

Additional results

Model	# of sequential operations	Accuracy %
ResNet [11]		<i>our/original</i>
ResNet-110	110.00	94.09/93.57
ResNet-56	56.00	93.61/93.03
ResNet-44	44.00	93.21/92.83
ResNet-32	32.00	92.91/92.49
ResNet-20	20.00	92.19/91.25
Busgeted ResNet		
B-ResNet	110.00	94.36
	58.00	94.01
	20.00	93.24
	18.00	92.93
	16.00	92.75
Convolutional Neural Fabric [25]		<i>our/original</i>
CNF W=8	53.00	94.83/90.58
CNF W=4	31.00	93.75/87.91
CNF W=2	19.00	92.54/86.21
CNF W=1	12.00	89.91
Budgeted CNF		
B-ResNet	31.00	94.96
	25.00	94.72
	21.00	94.36
	18.00	93.86

Table 5: Results for *Distributed computation cost* on CIFAR-10 with $n = 4$

Model	# of sequential operations	Accuracy %
ResNet [11]		our/original
ResNet-110	112.00	94.09/93.57
ResNet-56	58.00	93.61/93.03
ResNet-44	46.00	93.21/92.83
ResNet-32	34.00	92.91/92.49
ResNet-20	22.00	92.19/91.25
Budgeted ResNet		
B-ResNet	184.00	94.42
	110.00	94.12
	94.00	94.01
	22.00	93.06
	20.00	92.29
Convolutional Neural Fabrics [25]		our/original
CNF W=8	171.00	94.83/90.58
CNF W=4	83.00	93.75/87.91
CNF W=2	39.00	92.54/86.21
CNF W=1	12.00	89.91
Budgeted CNF		
B-CNF	98.00	95.02
	50.00	94.62
	45.00	94.55
	39.00	94.35
	33.00	93.00
	26.00	92.91
	18.00	92.87

Table 6: Results for *Distributed computation cost* on CIFAR-10 with $n = 1$

Model	# of sequential operations	Accuracy %
ResNet [11]		our/original
ResNet-110	110.00	94.09/93.57
ResNet-56	56.00	93.61/93.03
ResNet-44	44.00	93.21/92.83
ResNet-32	32.00	92.91/92.49
ResNet-20	20.00	92.19/91.25
Budgeted ResNet		
B-ResNet	179.00	94.36
	112.00	94.42
	56.00	94.31
	20.00	93.20
	18.00	92.81
Convolutional Neural Fabric [25]		our/original
CNF W=8	90.00	94.83/90.58
CNF W=4	47.00	93.75/87.91
CNF W=2	26.00	92.54/86.21
CNF W=1	12.00	89.91
Budgeted CNF		
B-CNF	47.00	94.67
	30.00	94.68
	28.00	94.58
	24.00	94.41
	20.00	94.35
	18.00	92.86

Table 8: Results for *Distributed computation cost* on CIFAR-10 with $n = 2$

Model	# of sequential operations	Accuracy (%)
ResNet[11]		
ResNet-110	112.00	71.85
ResNet-56	58.00	70.57
ResNet-44	46.00	70.28
ResNet-32	34.00	69.28
ResNet-20	22.00	67.14
Budgeted ResNet		
B-ResNet	320.00	74.35
	184.00	73.85
	110.00	72.88
	67.00	72.02
	32.00	69.60
	20.00	68.48

Table 7: Results for *Distributed computation cost* on CIFAR-100 with $n = 1$

Model	# of sequential operations	Accuracy (%)
ResNe [11]		
ResNet110	110.00	71.85
ResNet56	56.00	70.57
ResNet44	44.00	70.28
ResNet32	34.00	69.28
ResNet20	20.00	67.14
Budgeted ResNet		
B-ResNet	179.00	74.35
	112.00	73.85
	49.00	71.84
	29.00	69.94
	22.00	69.09

Table 9: Results for *Distributed computation cost* on CIFAR-100 with $n = 2$