This ECCV 2018 paper, provided here by the Computer Vision Foundation, is the author-created version.
The content of this paper is identical to the content of the officially published ECCV 2018
LNCS version of the paper as available on SpringerLink: https://link.springer.com/conference/eccv

Progressive Neural Architecture Search

Chenxi Liu[1⋆], Barret Zoph[2], Maxim Neumann[2], Jonathon Shlens[2], Wei Hua[2], Li-Jia Li[2], Li Fei-Fei[2,3], Alan Yuille[1], Jonathan Huang[2], and Kevin Murphy[2]

[1] Johns Hopkins University
[2] Google AI
[3] Stanford University

Abstract. We propose a new method for learning the structure of convolutional neural networks (CNNs) that is more efficient than recent state-of-the-art methods based on reinforcement learning and evolutionary algorithms. Our approach uses a sequential model-based optimization (SMBO) strategy, in which we search for structures in order of increasing complexity, while simultaneously learning a surrogate model to guide the search through structure space. Direct comparison under the same search space shows that our method is up to 5 times more efficient than the RL method of Zoph et al. (2018) in terms of number of models evaluated, and 8 times faster in terms of total compute. The structures we discover in this way achieve state of the art classification accuracies on CIFAR-10 and ImageNet.

1 Introduction

There has been a lot of recent interest in automatically learning good neural net architectures. Some of this work is summarized in Section 2, but at a high level, current techniques usually fall into one of two categories: evolutionary algorithms (see e.g. [28,24,35]) or reinforcement learning (see e.g., [40,41,39,5,2]). When using evolutionary algorithms (EA), each neural network structure is encoded as a string, and random mutations and recombinations of the strings are performed during the search process; each string (model) is then trained and evaluated on a validation set, and the top performing models generate "children". When using reinforcement learning (RL), the agent performs a sequence of actions, which specifies the structure of the model; this model is then trained and its validation performance is returned as the reward, which is used to update the RNN controller. Although both EA and RL methods have been able to learn network structures that outperform manually designed architectures, they require significant computational resources. For example, the RL method in [41] trains and evaluates 20,000 neural networks across 500 P100 GPUs over 4 days.

In this paper, we describe a method that is able to learn a CNN which matches previous state of the art in terms of accuracy, while requiring 5 times fewer model evaluations during the architecture search. Our starting point is the

⋆ Work done while an intern at Google.

structured search space proposed by [41], in which the search algorithm is tasked with searching for a good convolutional "cell", as opposed to a full CNN. A cell contains B "blocks", where a block is a combination operator (such as addition) applied to two inputs (tensors), each of which can be transformed (e.g., using convolution) before being combined. This cell structure is then stacked a certain number of times, depending on the size of the training set, and the desired running time of the final CNN (see Section 3 for details). This modular design also allows easy architecture transfer from one dataset to another, as we will show in experimental results.

We propose to use heuristic search to search the space of cell structures, starting with simple (shallow) models and progressing to complex ones, pruning out unpromising structures as we go. At iteration b of the algorithm, we have a set of K candidate cells (each of size b blocks), which we train and evaluate on a dataset of interest. Since this process is expensive, we also learn a model or surrogate function which can predict the performance of a structure without needing to training it. We expand the K candidates of size b into $K' \gg K$ children, each of size $b+1$. We apply our surrogate function to rank all of the K' children, pick the top K, and then train and evaluate them. We continue in this way until $b = B$, which is the maximum number of blocks we want to use in our cell. See Section 4 for details.

Our progressive (simple to complex) approach has several advantages over other techniques that directly search in the space of fully-specified structures. First, the simple structures train faster, so we get some initial results to train the surrogate quickly. Second, we only ask the surrogate to predict the quality of structures that are slightly different (larger) from the ones it has seen (c.f., trust-region methods). Third, we factorize the search space into a product of smaller search spaces, allowing us to potentially search models with many more blocks. In Section 5 we show that our approach is 5 times more efficient than the RL method of [41] in terms of number of models evaluated, and 8 times faster in terms of total compute. We also show that the structures we discover achieve state of the art classification accuracies on CIFAR-10 and ImageNet.[4]

2 Related Work

Our paper is based on the "neural architecture search" (NAS) method proposed in [40,41]. In the original paper [40], they use the REINFORCE algorithm [34] to estimate the parameters of a recurrent neural network (RNN), which represents a policy to generate a sequence of symbols (actions) specifying the structure of the CNN; the reward function is the classification accuracy on the validation set of a CNN generated from this sequence. [41] extended this by using a more structured search space, in which the CNN was defined in terms of a series of stacked

[4] The code and checkpoint for the PNAS model trained on ImageNet can be downloaded from the TensorFlow models repository at `http://github.com/tensorflow/models/`. Also see `https://github.com/chenxi116/PNASNet.TF` and `https://github.com/chenxi116/PNASNet.pytorch` for author's reimplementation.

"cells". (They also replaced REINFORCE with proximal policy optimization (PPO) [29].) This method was able to learn CNNs which outperformed almost all previous methods in terms of accuracy vs speed on image classification (using CIFAR-10 [19] and ImageNet [8]) and object detection (using COCO [20]).

There are several other papers that use RL to learn network structures. [39] use the same model search space as NAS, but replace policy gradient with Q-learning. [2] also use Q-learning, but without exploiting cell structure. [5] use policy gradient to train an RNN, but the actions are now to widen an existing layer, or to deepen the network by adding an extra layer. This requires specifying an initial model and then gradually learning how to transform it. The same approach, of applying "network morphisms" to modify a network, was used in [12], but in the context of hill climbing search, rather than RL. [26] use parameter sharing among child models to substantially accelerate the search process.

An alternative to RL is to use evolutionary algorithms (EA; "neuro-evolution" [32]). Early work (e.g., [33]) used EA to learn both the structure and the parameters of the network, but more recent methods, such as [28,24,35,21,27], just use EA to search the structures, and use SGD to estimate the parameters.

RL and EA are local search methods that search through the space of fully-specified graph structures. An alternative approach, which we adopt, is to use heuristic search, in which we search through the space of structures in a progressive way, from simple to complex. There are several pieces of prior work that explore this approach. [25] use Monte Carlo Tree Search (MCTS), but at each node in the search tree, it uses random selection to choose which branch to expand, which is very inefficient. Sequential Model Based Optimization (SMBO) [17] improves on MCTS by learning a predictive model, which can be used to decide which nodes to expand. This technique has been applied to neural net structure search in [25], but they used a flat CNN search space, rather than our hierarchical cell-based space. Consequently, their resulting CNNs do not perform very well. Other related works include [23], who focus on MLP rather than CNNs; [33], who used an incremental approach in the context of evolutionary algorithms; [40] who used a schedule of increasing number of layers; and [13] who search through the space of latent factor models specified by a grammar. Finally, [7,16] grow CNNs sequentially using boosting.

Several other papers learn a surrogate function to predict the performance of a candidate structure, either "zero shot" (without training it) (see e.g., [4]), or after training it for a small number of epochs and extrapolating the learning curve (see e.g., [10,3]). However, most of these methods have been applied to fixed sized structures, and would not work with our progressive search approach.

3 Architecture Search Space

In this section we describe the neural network architecture search space used in our work. We build on the hierarchical approach proposed in [41], in which we first learn a cell structure, and then stack this cell a desired number of times, in order to create the final CNN.

3.1 Cell Topologies

A cell is a fully convolutional network that maps an $H \times W \times F$ tensor to another $H' \times W' \times F'$ tensor. If we use stride 1 convolution, then $H' = H$ and $W' = W$; if we use stride 2, then $H' = H/2$ and $W' = W/2$. We employ a common heuristic to double the number of filters (feature maps) whenever the spatial activation is halved, so $F' = F$ for stride 1, and $F' = 2F$ for stride 2.

The cell can be represented by a DAG consisting of B blocks. Each block is a mapping from 2 input tensors to 1 output tensor. We can specify a block b in a cell c as a 5-tuple, (I_1, I_2, O_1, O_2, C), where $I_1, I_2 \in \mathcal{I}_b$ specifies the inputs to the block, $O_1, O_2 \in \mathcal{O}$ specifies the operation to apply to input I_i, and $C \in \mathcal{C}$ specifies how to combine O_1 and O_2 to generate the feature map (tensor) corresponding to the output of this block, which we denote by H_b^c.

The set of possible inputs, $\mathcal{I}_b$, is the set of all previous blocks in this cell, $\{H_1^c, \ldots, H_{b-1}^c\}$, plus the output of the previous cell, H_B^{c-1}, plus the output of the previous-previous cell, H_B^{c-2}.

The operator space $\mathcal{O}$ is the following set of 8 functions, each of which operates on a single tensor[5]:

- 3x3 depthwise-separable convolution
- 5x5 depthwise-separable convolution
- 7x7 depthwise-separable convolution
- 1x7 followed by 7x1 convolution
- identity
- 3x3 average pooling
- 3x3 max pooling
- 3x3 dilated convolution

This is less than the 13 operators used in [41], since we removed the ones that their RL method discovered were never used.

For the space of possible combination operators $\mathcal{C}$, [41] considerd both elementwise addition and concatenation. However, they discovered that the RL method never chose to use concatenation, so to reduce our search space, we always use addition as the combination operator. Thus in our work, a block can be specified by a 4-tuple.

We now quantify the size of the search space to highlight the magnitude of the search problem. Let the space of possible structures for the b'th block be $\mathcal{B}_b$; this has size $|\mathcal{B}_b| = |\mathcal{I}_b|^2 \times |\mathcal{O}|^2 \times |\mathcal{C}|$, where $|\mathcal{I}_b| = (2 + b - 1)$, $|\mathcal{O}| = 8$ and $|\mathcal{C}| = 1$. For $b = 1$, we have $\mathcal{I}_1 = \{H_B^{c-1}, H_B^{c-2}\}$, which are the final outputs of the previous two cells, so there are $|\mathcal{B}_1| = 256$ possible block structures.

If we allow cells of up to $B = 5$ blocks, the total number of cell structures is given by $|\mathcal{B}_{1:5}| = 2^2 \times 8^2 \times 3^2 \times 8^2 \times 4^2 \times 8^2 \times 5^2 \times 8^2 \times 6^2 \times 8^2 = 5.6 \times 10^{14}$. However, there are certain symmetries in this space that allow us to prune it to a more reasonable size. For example, there are only 136 unique cells composed of 1 block. The total number of unique cells is $\sim 10^{12}$. This is much smaller than the search space used in [41], which has size 10^{28}, but it is still an extremely large space to search, and requires efficient optimization methods.

[5] The depthwise-separable convolutions are in fact two repetitions of ReLU-SepConv-BatchNorm; 1x1 convolutions are also inserted when tensor sizes mismatch.

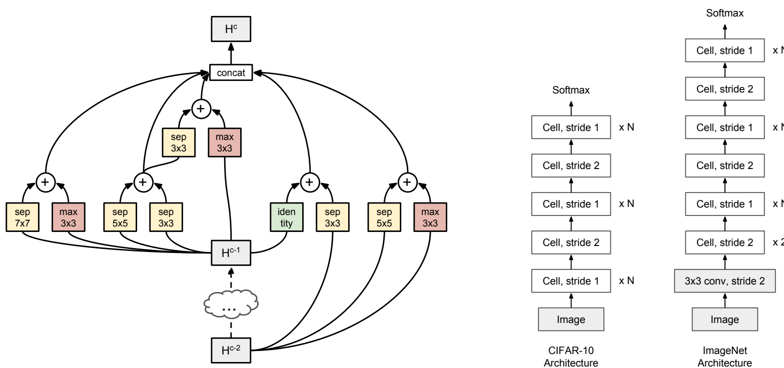

Fig. 1. *Left*: The best cell structure found by our Progressive Neural Architecture Search, consisting of 5 blocks. *Right*: We employ a similar strategy as [41] when constructing CNNs from cells on CIFAR-10 and ImageNet. Note that we learn a single cell type instead of distinguishing between Normal and Reduction cell.

3.2 From Cell to CNN

To evaluate a cell, we have to convert it into a CNN. To do this, we stack a predefined number of copies of the basic cell (with the same structure, but untied weights), using either stride 1 or stride 2, as shown in Figure 1 (right). The number of stride-1 cells between stride-2 cells is then adjusted accordingly with up to N number of repeats. At the top of the network, we use global average pooling, followed by a softmax classification layer. We then train the stacked model on the relevant dataset.

In the case of CIFAR-10, we use 32×32 images. In the case of ImageNet, we consider two settings, one with high resolution images of size 331×331, and one with smaller images of size 224×224. The latter results in less accurate models, but they are faster. For ImageNet, we also add an initial 3×3 convolutional filter layer with stride 2 at the start of the network, to further reduce the cost.

The overall CNN construction process is identical to [41], except we only use one cell type (we do not distinguish between Normal and Reduction cells, but instead emulate a Reduction cell by using a Normal cell with stride 2), and the cell search space is slightly smaller (since we use fewer operators and combiners).

4 Method

4.1 Progressive Neural Architecture Search

Many previous approaches directly search in the space of full cells, or worse, full CNNs. For example, NAS uses a 50-step RNN[6] as a controller to generate cell specifications. In [35] a fixed-length binary string encoding of CNN architecture

[6] 5 symbols per block, times 5 blocks, times 2 for Normal and Reduction cells.

Algorithm 1 Progressive Neural Architecture Search (PNAS).

Inputs: B (max num blocks), E (max num epochs), F (num filters in first layer), K (beam size), N (num times to unroll cell), trainSet, valSet.
$\mathcal{S}_1 = \mathcal{B}_1$ // *Set of candidate structures with one block*
$\mathcal{M}_1$ = cell-to-CNN($\mathcal{S}_1$, N, F) // *Construct CNNs from cell specifications*
$\mathcal{C}_1$ = train-CNN($\mathcal{M}_1$, E, trainSet) // *Train proxy CNNs*
$\mathcal{A}_1$ = eval-CNN($\mathcal{C}_1$, valSet) // *Validation accuracies*
π = fit($\mathcal{S}_1$, $\mathcal{A}_1$) // *Train the reward predictor from scratch*
for $b = 2 : B$ **do**
  $\mathcal{S}'_b$ = expand-cell($\mathcal{S}_{b-1}$) // *Expand current candidate cells by one more block*
  $\hat{\mathcal{A}}'_b$ = predict($\mathcal{S}'_b$, π) // *Predict accuracies using reward predictor*
  $\mathcal{S}_b$ = top-K($\mathcal{S}'_b$, $\hat{\mathcal{A}}'_b$, K) // *Most promising cells according to prediction*
  $\mathcal{M}_b$ = cell-to-CNN($\mathcal{S}_b$, N, F)
  $\mathcal{C}_b$ = train-CNN($\mathcal{M}_b$, E, trainSet)
  $\mathcal{A}_b$ = eval-CNN($\mathcal{C}_b$, valSet)
  π = update-predictor($\mathcal{S}_b$, $\mathcal{A}_b$, π) // *Finetune reward predictor with new data*
end for
Return top-K($\mathcal{S}_B$, $\mathcal{A}_B$, 1)

is defined and used in model evolution/mutation. While this is a more direct approach, we argue that it is difficult to directly navigate in an exponentially large search space, especially at the beginning where there is no knowledge of what makes a good model.

As an alternative, we propose to search the space in a progressive order, simplest models first. In particular, we start by constructing all possible cell structures from $\mathcal{B}_1$ (i.e., composed of 1 block), and add them to a queue. We train and evaluate all the models in the queue (in parallel), and then expand each one by adding all of the possible block structures from $\mathcal{B}_2$; this gives us a set of $|\mathcal{B}_1| \times |\mathcal{B}_2| = 256 \times 576 = 147,456$ candidate cells of depth 2. Since we cannot afford to train and evaluate all of these child networks, we refer to a learned predictor function (described in Section 4.2); it is trained based on the measured performance of the cells we have visited so far. (Our predictor takes negligible time to train and apply.) We then use the predictor to evaluate all the candidate cells, and pick the K most promising ones. We add these to the queue, and repeat the process, until we find cells with a sufficient number B of blocks. See Algorithm 1 for the pseudocode, and Figure 2 for an illustration.

4.2 Performance Prediction with Surrogate Model

As explained above, we need a mechanism to predict the final performance of a cell before we actually train it. There are at least three desired properties of such a predictor:

- *Handle variable-sized inputs*: We need the predictor to work for variable-length input strings. In particular, it should be able to predict the performance of any cell with $b+1$ blocks, even if it has only been trained on cells with up to b blocks.

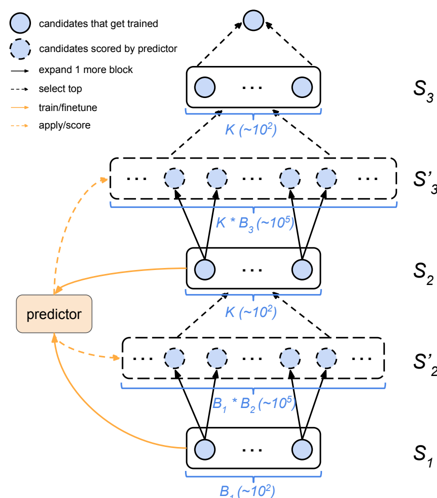

Fig. 2. Illustration of the PNAS search procedure when the maximum number of blocks is $B = 3$. Here $\mathcal{S}_b$ represents the set of candidate cells with b blocks. We start by considering all cells with 1 block, $\mathcal{S}_1 = \mathcal{B}_1$; we train and evaluate all of these cells, and update the predictor. At iteration 2, we expand each of the cells in $\mathcal{S}_1$ to get all cells with 2 blocks, $\mathcal{S}'_2 = \mathcal{B}_{1:2}$; we predict their scores, pick the top K to get $\mathcal{S}_2$, train and evaluate them, and update the predictor. At iteration 3, we expand each of the cells in $\mathcal{S}_2$, to get a subset of cells with 3 blocks, $\mathcal{S}'_3 \subseteq \mathcal{B}_{1:3}$; we predict their scores, pick the top K to get $\mathcal{S}_3$, train and evaluate them, and return the winner. $B_b = |\mathcal{B}_b|$ is the number of possible blocks at level b and K is the beam size (number of models we train and evaluate per level of the search tree).

- *Correlated with true performance*: we do not necessarily need to achieve low mean squared error, but we do want the predictor to rank models in roughly the same order as their true performance values.
- *Sample efficiency*: We want to train and evaluate as few cells as possible, which means the training data for the predictor will be scarce.

The requirement that the predictor be able to handle variable-sized strings immediately suggests the use of an RNN, and indeed this is one of the methods we try. In particular, we use an LSTM that reads a sequence of length $4b$ (representing I_1, I_2, O_1 and O_2 for each block), and the input at each step is a one-hot vector of size $|\mathcal{I}_b|$ or $|\mathcal{O}|$, followed by embedding lookup. We use a shared embedding of dimension D for the tokens $I_1, I_2 \in \mathcal{I}$, and another shared embedding for $O_1, O_2 \in \mathcal{O}$. The final LSTM hidden state goes through a fully-connected layer and sigmoid to regress the validation accuracy. We also try a simpler MLP baseline in which we convert the cell to a fixed length vector as follows: we embed each token into an D-dimensional vector, concatenate the embeddings for each block to get an $4D$-dimensional vector, and then average over blocks. Both models are trained using L_1 loss.

When training the predictor, one approach is to update the parameters of the predictor using the new data using a few steps of SGD. However, since the sample size is very small, we fit an ensemble of 5 predictors, each fit (from scratch) to 4/5 of all the data available at each step of the search process. We observed empirically that this reduced the variance of the predictions.

In the future, we plan to investigate other kinds of predictors, such as Gaussian processes with string kernels (see e.g., [1]), which may be more sample efficient to train and produce predictions with uncertainty estimates.

5 Experiments and Results

5.1 Experimental Details

Our experimental setting follows [41]. In particular, we conduct most of our experiments on CIFAR-10 [19]. CIFAR-10 has 50,000 training images and 10,000 test images. We use 5000 images from the training set as a validation set. All images are whitened, and 32×32 patches are cropped from images upsampled to 40×40. Random horizontal flip is also used. After finding a good model on CIFAR-10, we evaluate its quality on ImageNet classification in Section 5.5.

For the MLP accuracy predictor, the embedding size is 100, and we use 2 fully connected layers, each with 100 hidden units. For the RNN accuracy predictor, we use an LSTM, and the hidden state size and embedding size are both 100. The embeddings use uniform initialization in range [-0.1, 0.1]. The bias term in the final fully connected layer is initialized to 1.8 (0.86 after sigmoid) to account for the mean observed accuracy of all $b = 1$ models. We use the Adam optimizer [18] with learning rate 0.01 for the $b = 1$ level and 0.002 for all following levels.

Our training procedure for the CNNs follows the one used in [41]. During the search we evaluate $K = 256$ networks at each stage (136 for stage 1, since there are only 136 unique cells with 1 block), we use a maximum cell depth of $B = 5$ blocks, we use $F = 24$ filters in the first convolutional cell, we unroll the cells for $N = 2$ times, and each child network is trained for 20 epochs using initial learning rate of 0.01 with cosine decay [22].

5.2 Performance of the Surrogate Predictors

In this section, we compare the performance of different surrogate predictors. Note that at step b of PNAS, we train the predictor on the observed performance of cells with up to b blocks, but we apply it to cells with $b+1$ blocks. We therefore consider predictive accuracy both for cells with sizes that have been seen before (but which have not been trained on), and for cells which are one block larger than the training data.

More precisely, let $\mathcal{U}_{b,1:R}$ be a set of randomly chosen cells with b blocks, where $R = 10,000$. (For $b = 1$, there are only 136 unique cells.) We convert each of these to CNNs, and train them for $E = 20$ epochs. (Thus in total we train

Algorithm 2 Evaluating performance of a predictor on a random dataset.

for $b = 1 : B - 1$ **do**
  for $t = 1 : T$ **do**
    $\mathcal{S}_{b,t,1:K}$ = random sample of K models from $\mathcal{U}_{b,1:R}$
    $\pi_{b,t} = \text{fit}(\mathcal{S}_{b,t,1:K},\, A(\mathcal{S}_{b,t,1:K}))$ // *Train or finetune predictor*
    $\hat{A}_{b,t,1:K} = \text{predict}(\pi_{b,t},\, \mathcal{S}_{b,t,1:K})$ // *Predict on same b*
    $\tilde{A}_{b+1,t,1:R} = \text{predict}(\pi_{b,t},\, \mathcal{U}_{b+1,1:R})$ // *Predict on next b*
  end for
end for

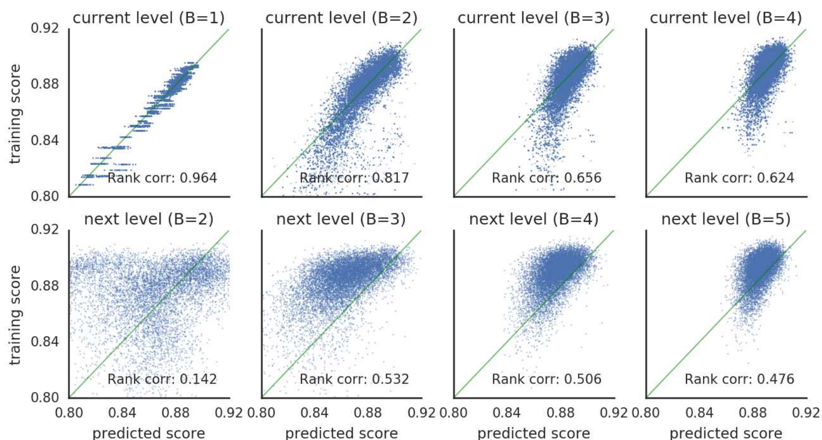

Fig. 3. Accuracy of MLP-ensemble predictor. Top row: true vs predicted accuracies on models from the training set over different trials. Bottom row: true vs predicted accuracies on models from the set of all unseen larger models. Denoted is the mean rank correlation from individual trials.

Method	$b=1$		$b=2$		$b=3$		$b=4$	
	$\hat{\rho}_1$	$\tilde{\rho}_2$	$\hat{\rho}_2$	$\tilde{\rho}_3$	$\hat{\rho}_3$	$\tilde{\rho}_4$	$\hat{\rho}_4$	$\tilde{\rho}_5$
MLP	0.938	0.113	0.857	0.450	0.714	0.469	0.641	0.444
RNN	0.970	**0.198**	**0.996**	0.424	0.693	0.401	**0.787**	0.413
MLP-ensemble	**0.975**	0.164	0.786	**0.532**	0.634	**0.504**	0.645	**0.468**
RNN-ensemble	0.972	0.164	0.906	0.418	**0.801**	0.465	0.579	0.424

Table 1. Spearman rank correlations of different predictors on the training set, $\hat{\rho}_b$, and when extrapolating to unseen larger models, $\tilde{\rho}_{b+1}$. See text for details.

$\sim (B-1) \times R = 40,000$ models for 20 epochs each.) We now use this random dataset to evaluate the performance of the predictors using the pseudocode in Algorithm 2, where $A(\mathcal{H})$ returns the true validation set accuracies of the models in some set $\mathcal{H}$. In particular, for each size $b = 1 : B$, and for each trial $t = 1 : T$ (we use $T = 20$), we do the following: randomly select $K = 256$ models (each of size b) from $\mathcal{U}_{b,1:R}$ to generate a training set $\mathcal{S}_{b,t,1:K}$; fit the predictor on the training set; evaluate the predictor on the training set; and finally evaluate the predictor on the set of all unseen random models of size $b + 1$.

The top row of Figure 3 shows a scatterplot of the true accuracies of the models in the training sets, $A(\mathcal{S}_{b,1:T,1:K})$, vs the predicted accuracies, $\hat{A}_{b,1:T,1:K}$ (so there are $T \times K = 20 \times 256 = 5120$ points in each plot, at least for $b > 1$). The bottom row plots the true accuracies on the set of larger models, $A(\mathcal{U}_{b+1,1:R})$, vs the predicted accuracies $\tilde{A}_{b+1,1:R}$ (so there are $R = 10$K points in each plot). We see that the predictor performs well on models from the training set, but

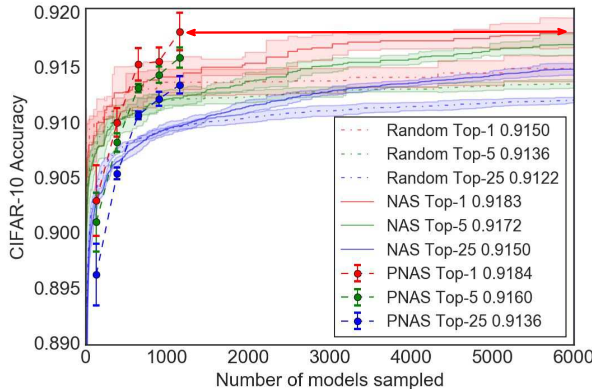

Fig. 4. Comparing the relative efficiency of NAS, PNAS and random search under the same search space. We plot mean accuracy (across 5 trials) on CIFAR-10 validation set of the top M models, for $M \in \{1, 5, 25\}$, found by each method vs number of models which are trained and evaluated. Each model is trained for 20 epochs. Error bars and the colored regions denote standard deviation of the mean.

not so well when predicting larger models. However, performance does increase as the predictor is trained on more (and larger) cells.

Figure 3 shows the results using an ensemble of MLPs. The scatter plots for the other predictors look similar. We can summarize each scatterplot using the Spearman rank correlation coefficient. Let $\hat{\rho}_b$ = rank-correlation($\hat{A}_{b,1:T,1:K}$, $A(\mathcal{S}_{b,1:T,1:K})$) and $\tilde{\rho}_{b+1}$ = rank-correlation($\tilde{A}_{b+1,1:R}$, $A(\mathcal{U}_{b+1,1:R})$). Table 1 summarizes these statistics across different levels. We see that for predicting the training set, the RNN does better than the MLP, but for predicting the performance on unseen larger models (which is the setting we care about in practice), the MLP seems to do slightly better. This will be corroborated by our end-to-end test in Section 5.3, and is likely due to overfitting. We also see that for the extrapolation task, ensembling seems to help.

5.3 Search Efficiency

In this section, we compare the efficiency of PNAS to two other methods: random search and the NAS method. To perform the comparison, we run PNAS for $B = 5$, and at each iteration b, we record the set $\mathcal{S}_b$ of $K = 256$ models of size b that it picks, and evaluate them on the CIFAR-10 validation set (after training for 20 epochs each). We then compute the validation accuracy of the top M models for $M \in \{1, 5, 25\}$. To capture the variance in performance of a given model due to randomness of the parameter initialization and optimization procedure, we repeat this process 5 times. We plot the mean and standard error of this statistic in Figure 4. We see that the mean performance of the top $M \in \{1, 5, 25\}$ models steadily increases, as we search for larger models. Furthermore, performance is

B	Top	Accuracy	# PNAS	# NAS	Speedup (# models)	Speedup (# examples)
5	1	0.9183	1160	5808	5.0	8.2
5	5	0.9161	1160	4100	3.5	6.8
5	25	0.9136	1160	3654	3.2	6.4

Table 2. Relative efficiency of PNAS (using MLP-ensemble predictor) and NAS under the same search space. B is the size of the cell, "Top" is the number of top models we pick, "Accuracy" is their average validation accuracy, "# PNAS" is the number of models evaluated by PNAS, "# NAS" is the number of models evaluated by NAS to achieve the desired accuracy. Speedup measured by number of examples is greater than speedup in terms of number of models, because NAS has an additional reranking stage, that trains the top 250 models for 300 epochs each before picking the best one.

better when using an MLP-ensemble (shown in Figure 4) instead of an RNN-ensemble (see supplementary material), which is consistent with Table 1.

For our random search baseline, we uniformly sample 6000 cells of size $B = 5$ blocks from the random set of models $\mathcal{U}_{5,1:R}$ described in Section 5.2. Figure 4 shows that PNAS significantly outperforms this baseline.

Finally, we compare to NAS. Each trial sequentially searches 6000 cells of size $B = 5$ blocks. At each iteration t, we define H_t to be the set of all cells visited so far by the RL agent. We compute the validation accuracy of the top M models in H_t, and plot the mean and standard error of this statistic in Figure 4. We see that the mean performance steadily increases, but at a slower rate than PNAS.

To quantify the speedup factor compared to NAS, we compute the number of models that are trained and evaluated until the mean performance of PNAS and NAS are equal (note that PNAS produces models of size B after evaluating $|\mathcal{B}_1| + (B - 1) \times K$ models, which is 1160 for $B = 5$). The results are shown in Table 2. We see that PNAS is up to 5 times faster in terms of the number of models it trains and evaluates.

Comparing the number of models explored during architecture search is one measure of efficiency. However, some methods, such as NAS, employ a secondary reranking stage to determine the best model; PNAS does not perform a reranking stage but uses the top model from the search directly. A more fair comparison is therefore to count the total number of examples processed through SGD throughout the search. Let M_1 be the number of models trained during search, and let E_1 be the number of examples used to train each model.[7] The total number of examples is therefore $M_1 E_1$. However, for methods with the additional reranking stage, the top M_2 models from the search procedure are trained using E_2 examples each, before returning the best. This results in a total cost of

[7] The number of examples is equal to the number of SGD steps times the batch size. Alternatively, it can be measured in terms of number of epoch (passes through the data), but since different papers use different sized training sets, we avoid this measure. In either case, we assume the number of examples is the same for every model, since none of the methods we evaluate use early stopping.

Model	B	N	F	Error	Params	M_1	E_1	M_2	E_2	Cost
NASNet-A [41]	5	6	32	3.41	3.3M	20000	0.9M	250	13.5M	21.4-29.3B
NASNet-B [41]	5	4	N/A	3.73	2.6M	20000	0.9M	250	13.5M	21.4-29.3B
NASNet-C [41]	5	4	N/A	3.59	3.1M	20000	0.9M	250	13.5M	21.4-29.3B
Hier-EA [21]	5	2	64	3.75±0.12	15.7M	7000	5.12M	0	0	35.8B[9]
AmoebaNet-B [27]	5	6	36	3.37±0.04	2.8M	27000	2.25M	100	27M	63.5B[10]
AmoebaNet-A [27]	5	6	36	3.34±0.06	3.2M	20000	1.13M	100	27M	25.2B[11]
PNASNet-5	5	3	48	3.41±0.09	3.2M	1160	0.9M	0	0	1.0B

Table 3. Performance of different CNNs on CIFAR test set. All model comparisons employ a comparable number of parameters and exclude cutout data augmentation [9]. "Error" is the top-1 misclassification rate on the CIFAR-10 test set. (Error rates have the form $\mu \pm \sigma$, where μ is the average over multiple trials and σ is the standard deviation. In PNAS we use 15 trials.) "Params" is the number of model parameters. "Cost" is the total number of examples processed through SGD ($M_1E_1 + M_2E_2$) before the architecture search terminates. The number of filters F for NASNet-{B, C} cannot be determined (hence N/A), and the actual E_1, E_2 may be larger than the values in this table (hence the range in cost), according to the original authors.

$M_1E_1 + M_2E_2$. For NAS and PNAS, $E_1 = 900$K for NAS and PNAS since they use 20 epochs on a training set of size 45K. The number of models searched to achieve equal top-1 accuracy is $M_1 = 1160$ for PNAS and $M_1 = 5808$ for NAS. For the second stage, NAS trains the top $M_2 = 250$ models for $E_2 = 300$ epochs before picking the best.[8] Thus we see that PNAS is about 8 times faster than NAS when taking into account the total cost.

5.4 Results on CIFAR-10 Image Classification

We now discuss the performance of our final model, and compare it to the results of other methods in the literature. Let PNASNet-5 denote the best CNN we discovered on CIFAR using PNAS, also visualized in Figure 1 (left). After

[8] This additional stage is quite important for NAS, as the NASNet-A cell was originally ranked 70th among the top 250.

[9] In Hierarchical EA, the search phase trains 7K models (each for 4 times to reduce variance) for 5000 steps of batch size 256. Thus, the total computational cost is 7K $\times$ 5000 $\times$ 256 $\times$ 4 = 35.8B.

[10] The total computational cost for AmoebaNet consists of an architecture search and a reranking phase. The architecture search phase trains over 27K models each for 50 epochs. Each epoch consists of 45K examples. The reranking phase searches over 100 models each trained for 600 epochs. Thus, the architecture search is 27K $\times$ 50 $\times$ 45K = 60.8B examples. The reranking phase consists of 100 $\times$ 600 $\times$ 45K = 2.7B examples. The total computational cost is 60.8B + 2.7B = 63.5B.

[11] The search phase trains 20K models each for 25 epochs. The rest of the computation is the same as AmoebaNet-B.

we have selected the cell structure, we try various N and F values such that the number of model parameters is around 3M, train them each for 300 epochs using initial learning rate of 0.025 with cosine decay, and pick the best combination based on the validation set. Using this best combination of N and F, we train it for 600 epochs on the union of training set and validation set. During training we also used auxiliary classifier located at 2/3 of the maximum depth weighted by 0.4, and drop each path with probability 0.4 for regularization.

The results are shown in Table 3. We see that PNAS can find a model with the same accuracy as NAS, but using 21 times less compute. PNAS also outperforms the Hierarchical EA method of [21], while using 36 times less compute. Though the the EA method called "AmoebaNets" [27] currently give the highest accuracies (at the time of writing), it also requires the most compute, taking 63 times more resources than PNAS. However, these comparisons must be taken with a grain of salt, since the methods are searching through different spaces. By contrast, in Section 5.3, we fix the search space for NAS and PNAS, to make the speedup comparison fair.

5.5 Results on ImageNet Image Classification

We further demonstrate the usefulness of our learned cell by applying it to ImageNet classification. Our experiments reveal that CIFAR accuracy and ImageNet accuracy are strongly correlated ($\rho = 0.727$; see supplementary material).

To compare the performance of PNASNet-5 to the results in other papers, we conduct experiments under two settings:

- *Mobile*: Here we restrain the representation power of the CNN. Input image size is 224×224, and the number of multiply-add operations is under 600M.
- *Large*: Here we compare PNASNet-5 against the state-of-the-art models on ImageNet. Input image size is 331×331.

In both experiments we use RMSProp optimizer, label smoothing of 0.1, auxiliary classifier located at 2/3 of the maximum depth weighted by 0.4, weight decay of 4e-5, and dropout of 0.5 in the final softmax layer. In the *Mobile* setting, we use distributed synchronous SGD with 50 P100 workers. On each worker, batch size is 32, initial learning rate is 0.04, and is decayed every 2.2 epochs with rate 0.97. In the *Large* setting, we use 100 P100 workers. On each worker, batch size is 16, initial learning rate is 0.015, and is decayed every 2.4 epochs with rate 0.97. During training, we drop each path with probability 0.4.

The results of the *Mobile* setting are summarized in Table 4. PNASNet-5 achieves slightly better performance than NASNet-A (74.2% top-1 accuracy for PNAS vs 74.0% for NASNet-A). Both methods significantly surpass the previous state-of-the-art, which includes the manually designed MobileNet [14] (70.6%) and ShuffleNet [37] (70.9%). AmoebaNet-C performs the best, but note that this is a different model than their best-performing CIFAR-10 model. Table 5 shows that under the *Large* setting, PNASNet-5 achieves higher performance (82.9% top-1; 96.2% top-5) than previous state-of-the-art approaches, including SENet [15], NASNet-A, and AmoebaNets under the same model capacity.

Model	Params	Mult-Adds	Top-1	Top-5
MobileNet-224 [14]	4.2M	569M	70.6	89.5
ShuffleNet (2x) [37]	5M	524M	70.9	89.8
NASNet-A ($N = 4$, $F = 44$) [41]	5.3M	564M	74.0	91.6
AmoebaNet-B ($N = 3$, $F = 62$) [27]	5.3M	555M	74.0	91.5
AmoebaNet-A ($N = 4$, $F = 50$) [27]	5.1M	555M	74.5	92.0
AmoebaNet-C ($N = 4$, $F = 50$) [27]	6.4M	570M	75.7	92.4
PNASNet-5 ($N = 3$, $F = 54$)	5.1M	588M	74.2	91.9

Table 4. ImageNet classification results in the *Mobile* setting.

Model	Image Size	Params	Mult-Adds	Top-1	Top-5
ResNeXt-101 (64x4d) [36]	320×320	83.6M	31.5B	80.9	95.6
PolyNet [38]	331×331	92M	34.7B	81.3	95.8
Dual-Path-Net-131 [6]	320×320	79.5M	32.0B	81.5	95.8
Squeeze-Excite-Net [15]	320×320	145.8M	42.3B	82.7	96.2
NASNet-A ($N = 6$, $F = 168$) [41]	331×331	88.9M	23.8B	82.7	96.2
AmoebaNet-B ($N = 6$, $F = 190$) [27]	331×331	84.0M	22.3B	82.3	96.1
AmoebaNet-A ($N = 6$, $F = 190$) [27]	331×331	86.7M	23.1B	82.8	96.1
AmoebaNet-C ($N = 6$, $F = 228$) [27]	331×331	155.3M	41.1B	83.1	96.3
PNASNet-5 ($N = 4$, $F = 216$)	331×331	86.1M	25.0B	82.9	96.2

Table 5. ImageNet classification results in the *Large* setting.

6 Discussion and Future Work

The main contribution of this work is to show how we can accelerate the search for good CNN structures by using progressive search through the space of increasingly complex graphs, combined with a learned prediction function to efficiently identify the most promising models to explore. The resulting models achieve the same level of performance as previous work but with a fraction of the computational cost.

There are many possible directions for future work, including: the use of better surrogate predictors, such as Gaussian processes with string kernels; the use of model-based early stopping, such as [3], so we can stop the training of "unpromising" models before reaching E_1 epochs; the use of "warm starting", to initialize the training of a larger $b+1$-sized model from its smaller parent; the use of Bayesian optimization, in which we use an acquisition function, such as expected improvement or upper confidence bound, to rank the candidate models, rather than greedily picking the top K (see e.g., [31,30]); adaptively varying the number of models K evaluated at each step (e.g., reducing it over time); the automatic exploration of speed-accuracy tradeoffs (cf., [11]), etc.

Acknowledgements

We thank Quoc Le for inspiration, discussion and support; George Dahl for many fruitful discussions; Gabriel Bender, Vijay Vasudevan for the development of much of the critical infrastructure and the larger Google Brain team for the support and discussions. CL also thanks Lingxi Xie for support.

References

1. Baisero, A., Pokorny, F.T., Ek, C.H.: On a family of decomposable kernels on sequences. CoRR **abs/1501.06284** (2015)
2. Baker, B., Gupta, O., Naik, N., Raskar, R.: Designing neural network architectures using reinforcement learning. In: ICLR (2017)
3. Baker, B., Gupta, O., Raskar, R., Naik, N.: Accelerating neural architecture search using performance prediction. CoRR **abs/1705.10823** (2017)
4. Brock, A., Lim, T., Ritchie, J.M., Weston, N.: SMASH: one-shot model architecture search through hypernetworks. In: ICLR (2018)
5. Cai, H., Chen, T., Zhang, W., Yu, Y., Wang, J.: Efficient architecture search by network transformation. In: AAAI (2018)
6. Chen, Y., Li, J., Xiao, H., Jin, X., Yan, S., Feng, J.: Dual path networks. In: NIPS (2017)
7. Cortes, C., Gonzalvo, X., Kuznetsov, V., Mohri, M., Yang, S.: Adanet: Adaptive structural learning of artificial neural networks. In: ICML (2017)
8. Deng, J., Dong, W., Socher, R., Li, L., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: CVPR (2009)
9. Devries, T., Taylor, G.W.: Improved regularization of convolutional neural networks with cutout. CoRR **abs/1708.04552** (2017)
10. Domhan, T., Springenberg, J.T., Hutter, F.: Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves. IJCAI (2015)
11. Dong, J.D., Cheng, A.C., Juan, D.C., Wei, W., Sun, M.: PPP-Net: Platform-aware progressive search for pareto-optimal neural architectures. In: ICLR Workshop (2018)
12. Elsken, T., Metzen, J.H., Hutter, F.: Simple and efficient architecture search for convolutional neural networks. CoRR **abs/1711.04528** (2017)
13. Grosse, R.B., Salakhutdinov, R., Freeman, W.T., Tenenbaum, J.B.: Exploiting compositionality to explore a large space of model structures. In: UAI (2012)
14. Howard, A.G., Zhu, M., Chen, B., Kalenichenko, D., Wang, W., Weyand, T., Andreetto, M., Adam, H.: Mobilenets: Efficient convolutional neural networks for mobile vision applications. CoRR **abs/1704.04861** (2017)
15. Hu, J., Shen, L., Sun, G.: Squeeze-and-excitation networks. CoRR **abs/1709.01507** (2017)
16. Huang, F., Ash, J.T., Langford, J., Schapire, R.E.: Learning deep resnet blocks sequentially using boosting theory. CoRR **abs/1706.04964** (2017)
17. Hutter, F., Hoos, H.H., Leyton-Brown, K.: Sequential Model-Based optimization for general algorithm configuration. In: Intl. Conf. on Learning and Intelligent Optimization. pp. 507–523. Lecture Notes in Computer Science (2011)
18. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. In: ICLR (2015)

19. Krizhevsky, A., Hinton, G.: Learning multiple layers of features from tiny images. Technical report, University of Toronto (2009)
20. Lin, T., Maire, M., Belongie, S.J., Hays, J., Perona, P., Ramanan, D., Dollár, P., Zitnick, C.L.: Microsoft COCO: common objects in context. In: ECCV (2014)
21. Liu, H., Simonyan, K., Vinyals, O., Fernando, C., Kavukcuoglu, K.: Hierarchical representations for efficient architecture search. In: ICLR (2018)
22. Loshchilov, I., Hutter, F.: SGDR: stochastic gradient descent with restarts. In: ICLR (2017)
23. Mendoza, H., Klein, A., Feurer, M., Springenberg, J.T., Hutter, F.: Towards Automatically-Tuned neural networks. In: ICML Workshop on AutoML. pp. 58–65 (Dec 2016)
24. Miikkulainen, R., Liang, J.Z., Meyerson, E., Rawal, A., Fink, D., Francon, O., Raju, B., Shahrzad, H., Navruzyan, A., Duffy, N., Hodjat, B.: Evolving deep neural networks. CoRR **abs/1703.00548** (2017)
25. Negrinho, R., Gordon, G.J.: Deeparchitect: Automatically designing and training deep architectures. CoRR **abs/1704.08792** (2017)
26. Pham, H., Guan, M.Y., Zoph, B., Le, Q.V., Dean, J.: Efficient neural architecture search via parameter sharing. CoRR **abs/1802.03268** (2018)
27. Real, E., Aggarwal, A., Huang, Y., Le, Q.V.: Regularized evolution for image classifier architecture search. CoRR **abs/1802.01548** (2018)
28. Real, E., Moore, S., Selle, A., Saxena, S., Suematsu, Y.L., Tan, J., Le, Q.V., Kurakin, A.: Large-scale evolution of image classifiers. In: ICML (2017)
29. Schulman, J., Wolski, F., Dhariwal, P., Radford, A., Klimov, O.: Proximal policy optimization algorithms. CoRR **abs/1707.06347** (2017)
30. Shahriari, B., Swersky, K., Wang, Z., Adams, R.P., de Freitas, N.: Taking the human out of the loop: A review of bayesian optimization. Proceedings of the IEEE **104**(1), 148–175 (2016)
31. Snoek, J., Larochelle, H., Adams, R.P.: Practical bayesian optimization of machine learning algorithms. In: NIPS (2012)
32. Stanley, K.O.: Neuroevolution: A different kind of deep learning (Jul 2017)
33. Stanley, K.O., Miikkulainen, R.: Evolving neural networks through augmenting topologies. Evol. Comput. **10**(2), 99–127 (2002)
34. Williams, R.: Simple statistical gradient-following algorithms for connectionist reinforcement learning. Machine Learning **8**, 229–256 (1992)
35. Xie, L., Yuille, A.L.: Genetic CNN. In: ICCV (2017)
36. Xie, S., Girshick, R.B., Dollár, P., Tu, Z., He, K.: Aggregated residual transformations for deep neural networks. In: CVPR (2017)
37. Zhang, X., Zhou, X., Lin, M., Sun, J.: Shufflenet: An extremely efficient convolutional neural network for mobile devices. CoRR **abs/1707.01083** (2017)
38. Zhang, X., Li, Z., Loy, C.C., Lin, D.: Polynet: A pursuit of structural diversity in very deep networks. In: CVPR (2017)
39. Zhong, Z., Yan, J., Liu, C.L.: Practical network blocks design with Q-Learning. In: AAAI (2018)
40. Zoph, B., Le, Q.V.: Neural architecture search with reinforcement learning. In: ICLR (2017)
41. Zoph, B., Vasudevan, V., Shlens, J., Le, Q.V.: Learning transferable architectures for scalable image recognition. In: CVPR (2018)