

# Neural Architecture Search: A Survey

Thomas Elsken<sup>1,2</sup>, Jan Hendrik Metzen<sup>1</sup>, and Frank Hutter<sup>2</sup>

<sup>1</sup> Bosch Center for Artificial Intelligence, Robert Bosch GmbH

<sup>2</sup> University of Freiburg

**Abstract.** Deep Learning has enabled remarkable progress over the last years on a variety of tasks, such as image recognition, speech recognition, and machine translation. One crucial aspect for this progress are novel neural architectures. Currently employed architectures have mostly been developed manually by human experts, which is a time-consuming and error-prone process. Because of this, there is growing interest in automated *neural architecture search* methods. We provide an overview of existing work in this field of research and categorize them according to three dimensions: search space, search strategy, and performance estimation strategy.

**Keywords:** Neural Architecture Search, Search Space Design, Search Strategy, Performance Estimation Strategy


## 1 Introduction

The success of deep learning in perceptual tasks is largely due to its automation of the feature engineering process: hierarchical feature extractors are learned in an end-to-end fashion from data rather than manually designed. This success has been accompanied, however, by a rising demand for *architecture engineering*, where increasingly more complex neural architectures are designed manually. *Neural Architecture Search* (NAS), the process of automating architecture engineering, is thus a logical next step in automating machine learning. NAS can be seen as subfield of AutoML and has significant overlap with hyperparameter optimization and meta-learning. We categorize methods for NAS according to three dimensions: search space, search strategy, and performance estimation strategy:

- **Search Space.** The search space defines which architectures can be represented in principle. Incorporating prior knowledge about properties well-suited for a task can reduce the size of the search space and simplify the search. However, this also introduces a human bias, which may prevent finding novel architectural building blocks that go beyond the current human knowledge.
- **Search Strategy.** The search strategy details how to explore the search space. It encompasses the classical exploration-exploitation trade-off since, on the one hand, it is desirable to find well-performing architectures quickly, while on the other hand, premature convergence to a region of suboptimal architectures should be avoided.

- **Performance Estimation Strategy.** The objective of NAS is typically to find architectures that achieve high predictive performance on unseen data. *Performance Estimation* refers to the process of estimating this performance: the simplest option is to perform a standard training and validation of the architecture on data, but this is unfortunately computationally expensive and limits the number of architectures that can be explored. Much recent research therefore focuses on developing methods that reduce the cost of these performance estimations.

We refer to Figure 1 for an illustration. The chapter is also structured according to these three dimensions: we start with discussing search spaces in Section 2, cover search strategies in Section 3, and outline approaches to performance estimation in Section 4. We conclude with an outlook on future directions in Section 5.




**Fig. 1.** Abstract illustration of Neural Architecture Search methods. A search strategy selects an architecture  $A$  from a predefined search space  $\mathcal{A}$ . The architecture is passed to a performance estimation strategy, which returns the estimated performance of  $A$  to the search strategy.

## 2 Search Space

The search space defines which neural architectures a NAS approach might discover in principle. We now discuss common search spaces from recent works.

A relatively simple search space is the space of *chain-structured neural networks*, as illustrated in Figure 2 (left). A chain-structured neural network architecture  $A$  can be written as a sequence of  $n$  layers, where the  $i$ 'th layer  $L_i$  receives its input from layer  $i - 1$  and its output serves as the input for layer  $i + 1$ , i.e.,  $A = L_n \circ \dots \circ L_1 \circ L_0$ . The search space is then parametrized by: (i) the (maximum) number of layers  $n$  (possibly unbounded); (ii) the type of operation every layer can execute, e.g., pooling, convolution, or more advanced layer types like depthwise separable convolutions (Chollet, 2016) or dilated convolutions (Yu and Koltun, 2016); and (iii) hyperparameters associated with the operation, e.g., number of filters, kernel size and strides for a convolutional layer (Baker et al, 2017a; Suganuma et al, 2017; Cai et al, 2018a), or simply number of units for fully-connected networks (Mendoza et al, 2016). Note that the parameters from

(iii) are conditioned on (ii), hence the parametrization of the search space is not fixed-length but rather a conditional space.




**Fig. 2.** An illustration of different architecture spaces. Each node in the graphs corresponds to a layer in a neural network, e.g., a convolutional or pooling layer. Different layer types are visualized by different colors. An edge from layer  $L_i$  to layer  $L_j$  denotes that  $L_i$  receives the output of  $L_j$  as input. Left: an element of a chain-structured space. Right: an element of a more complex search space with additional layer types and multiple branches and skip connections.

Recent work on NAS (Brock et al, 2017; Elsken et al, 2017; Zoph et al, 2018; Elsken et al, 2018; Real et al, 2018; Cai et al, 2018b) incorporate modern design elements known from hand-crafted architectures such as skip connections, which allow to build complex, *multi-branch networks*, as illustrated in Figure 2 (right). In this case the input of layer  $i$  can be formally described as a function  $g_i(L_{i-1}^{out}, \dots, L_0^{out})$  combining previous layer outputs. Employing such a function results in significantly more degrees of freedom. Special cases of these multi-branch architectures are (i) the chain-structured networks (by setting  $g_i(L_{i-1}^{out}, \dots, L_0^{out}) = L_{i-1}^{out}$ ), (ii) Residual Networks (He et al, 2016), where previous layer outputs are summed ( $g_i(L_{i-1}^{out}, \dots, L_0^{out}) = L_{i-1}^{out} + L_j^{out}, j < i$ ) and (iii) DenseNets (Huang et al, 2017), where previous layer outputs are concatenated ( $g_i(L_{i-1}^{out}, \dots, L_0^{out}) = \text{concat}(L_{i-1}^{out}, \dots, L_0^{out})$ ).

Motivated by hand-crafted architectures consisting of repeated motifs (Szegedy et al, 2016; He et al, 2016; Huang et al, 2017), Zoph et al (2018) propose to search for such motifs, dubbed *cells*, rather than for whole architectures. They optimize two different kind of cells: a *normal cell* that preserves the dimensionality of the input and a *reduction cell* which reduces the spatial dimension. The final architecture is then built by stacking these cells in a predefined manner, as il-

illustrated in Figure 3. This search space has two major advantages compared to the ones discussed above:

1. The size of the search space is drastically reduced since cells can be comparably small. For example, Zoph et al (2018) estimate a seven-times speed-up compared to their previous work (Zoph and Le, 2017) while achieving better performance.
2. Cells can more easily be transferred to other datasets by adapting the number of cells used within a model. Indeed, Zoph et al (2018) transfer cells optimized on CIFAR-10 to ImageNet and achieve state-of-the-art performance.



**Fig. 3.** Illustration of the cell search space. Left: Two different cells, e.g., a normal cell (top) and a reduction cell (bottom) (Zoph et al, 2018). Right: an architecture built by stacking the cells sequentially. Note that cells can also be combined in a more complex manner, such as in multi-branch spaces, by simply replacing layers with cells.

Consequently, this cell-based search space was also successfully employed by many later works (Real et al, 2018; Liu et al, 2017; Pham et al, 2018; Elsen et al, 2018; Cai et al, 2018b; Liu et al, 2018b). However, a new design-choice arises when using a cell-based search space, namely how to choose the *meta-architecture*: how many cells shall be used and how should they be connected to build the actual model? For example, Zoph et al (2018) build a sequential model from cells, in which each cell receives the outputs of the two preceding cells as input, while Cai et al (2018b) employ the high-level structure of well-known manually designed architectures, such as DenseNet (Huang et al, 2017), and use their cells within these models. In principle, cells can be combined arbitrarily,

e.g., within the multi-branch space described above by simply replacing layers with cells. Ideally, the meta-architecture should be optimized automatically as part of NAS; otherwise one easily ends up doing meta-architecture engineering and the search for the cell becomes overly simple if most of the complexity is already accounted for by the meta-architecture.

One step in the direction of optimizing meta-architectures is the hierarchical search space introduced by Liu et al (2018a), which consists of several levels of motifs. The first level consists of the set of primitive operations, the second level of different motifs that connect primitive operations via a direct acyclic graphs, the third level of motifs that encode how to connect second-level motifs, and so on. The cell-based search space can be seen as a special case of this hierarchical search space where the number of levels is three, the second level motifs corresponds to the cells, and the third level is the hard-coded meta-architecture.

The choice of the search space largely determines the difficulty of the optimization problem: even for the case of the search space based on a single cell with fixed meta-architecture, the optimization problem remains (i) non-continuous and (ii) relatively high-dimensional (since more complex models tend to perform better, resulting in more design choices). We note that the architectures in many search spaces can be written as fixed-length vectors; e.g., the search space for each of the two cells by Zoph et al (2018) can be written as a 40-dimensional search space with categorical dimensions, each of which chooses between a small number of different building blocks and inputs. Similarly, unbounded search spaces can be constrained to have a maximal depth, giving rise to fixed-size search spaces with (potentially many) conditional dimensions.

In the next section, we discuss Search Strategies that are well-suited for these kinds of search spaces.

### 3 Search Strategy

Many different search strategies can be used to explore the space of neural architectures, including random search, Bayesian optimization, evolutionary methods, reinforcement learning (RL), and gradient-based methods. Historically, evolutionary algorithms were already used by many researchers to evolve neural architectures (and often also their weights) decades ago (see, e.g., Angeline et al, 1994; Stanley and Miikkulainen, 2002; Floreano et al, 2008; Stanley et al, 2009).

Bayesian optimization celebrated several early successes in NAS since 2013, leading to state-of-the-art vision architectures (Bergstra et al, 2013), state-of-the-art performance for CIFAR-10 without data augmentation (Domhan et al, 2015), and the first automatically-tuned neural networks to win competition datasets against human experts (Mendoza et al, 2016). NAS became a mainstream research topic in the machine learning community after Zoph and Le (2017) obtained competitive performance on the CIFAR-10 and Penn Treebank benchmarks with a search strategy based on reinforcement learning. As Zoph and Le (2017) use vast computational resources to achieve this result (800 GPUs for

three to four weeks), after their work, a wide variety of methods have been published in quick succession to reduce the computational costs and achieve further improvements in performance.

To frame NAS as a *reinforcement learning* (RL) problem (Baker et al, 2017a; Zoph and Le, 2017; Zhong et al, 2018; Zoph et al, 2018), the generation of a neural architecture can be considered to be the agent’s action, with the action space identical to the search space. The agent’s reward is based on an estimate of the performance of the trained architecture on unseen data (see Section 4). Different RL approaches differ in how they represent the agent’s policy and how they optimize it: Zoph and Le (2017) use a recurrent neural network (RNN) policy to sequentially sample a string that in turn encodes the neural architecture. They initially trained this network with the REINFORCE policy gradient algorithm, but in follow-up work use Proximal Policy Optimization (PPO) instead (Zoph et al, 2018). Baker et al (2017a) use Q-learning to train a policy which sequentially chooses a layer’s type and corresponding hyperparameters. An alternative view of these approaches is as sequential decision processes in which the policy samples actions to generate the architecture sequentially, the environment’s “state” contains a summary of the actions sampled so far, and the (undiscounted) reward is obtained only after the final action. However, since no interaction with an environment occurs during this sequential process (no external state is observed, and there are no intermediate rewards), we find it more intuitive to interpret the architecture sampling process as the sequential generation of a single action; this simplifies the RL problem to a stateless multi-armed bandit problem.

A related approach was proposed by Cai et al (2018a), who frame NAS as a sequential decision process: in their approach the state is the current (partially trained) architecture, the reward is an estimate of the architecture’s performance, and the action corresponds to an application of function-preserving mutations, dubbed network morphisms (Chen et al, 2016; Wei et al, 2017), see also Section 4, followed by a phase of training the network. In order to deal with variable-length network architectures, they use a bi-directional LSTM to encode architectures into a fixed-length representation. Based on this encoded representation, actor networks decide on the sampled action. The combination of these two components constitute the policy, which is trained end-to-end with the REINFORCE policy gradient algorithm. We note that this approach will not visit the same state (architecture) twice so that strong generalization over the architecture space is required from the policy.

An alternative to using RL are *neuro-evolutionary* approaches that use evolutionary algorithms for optimizing the neural architecture. Early neuro-evolutionary approaches (Angeline et al, 1994; Stanley and Miikkulainen, 2002; Floreano et al, 2008; Stanley et al, 2009) use genetic algorithms to optimize both the neural architecture and its weights; however, this is problematic when scaling to contemporary neural architectures with millions of weights. More recent neuro-evolutionary approaches (Real et al, 2017; Suganuma et al, 2017; Liu et al, 2018a; Real et al, 2018; Miikkulainen et al, 2017; Xie and Yuille, 2017; Elsen

et al, 2018) use gradient-based methods for optimizing weights and evolutionary algorithms are used solely for optimizing the neural architecture itself. Evolutionary algorithms evolve a population of models, i.e., a set of (possibly trained) networks; in every evolution step, at least one model from the population is sampled and serves as a parent to generate offsprings by applying mutations to it. In the context of NAS, mutations are local operations, such as adding or removing a layer, altering the hyperparameters of a layer, adding skip connections, as well as altering training hyperparameters. After training the offsprings, their fitness (e.g., performance on a validation set) is evaluated and they are added to the population.

Neuro-evolutionary methods differ in how they sample parents, update populations, and generate offsprings. For example, Real et al (2017), Real et al (2018), and Liu et al (2018a) use tournament selection (Goldberg and Deb, 1991) to sample parents, whereas Elsken et al (2018) sample parents from a multi-objective Pareto front using an inverse density. Real et al (2017) remove the worst individual from a population, while Real et al (2018) found it beneficial to remove the oldest individual (which decreases greediness), and Liu et al (2018a) do not remove individuals at all. To generate offspring, most approaches initialize child networks randomly, while Elsken et al (2018) employ Lamarckian inheritance, i.e, knowledge (in the form of learned weights) is passed on from a parent network to its children by using network morphisms. Real et al (2017) also let an offspring inherit all parameters of its parent that are not affected by the applied mutation; while this inheritance is not strictly function-preserving it might also speed up learning compared to a random initialization. Moreover, they also allow mutating the learning rate which can be seen as a way for optimizing the learning rate schedule during NAS.

Real et al (2018) conduct a case study comparing RL, evolution, and random search (RS), concluding that RL and evolution perform equally well in terms of final test accuracy, with evolution having better anytime performance and finding smaller models. Both approaches consistently perform better than RS in their experiments, but with a rather small margin: RS achieved test errors of approximately 4% on CIFAR-10, while RL and evolution reached approximately 3.5% (after “model augmentation” where depth and number of filters was increased; the difference on the actual, non-augmented search space was approx. 2%). The difference was even smaller for Liu et al (2018a), who reported a test error of 3.9% on CIFAR-10 and a top-1 validation error of 21.0% on ImageNet for RS, compared to 3.75% and 20.3% for their evolution-based method, respectively.

*Bayesian Optimization (BO)* is one of the most popular methods for hyperparameter optimization, but it has not been applied to NAS by many groups since typical BO toolboxes are based on Gaussian processes and focus on low-dimensional continuous optimization problems. Swersky et al (2013) and Kandasmay et al (2018) derive kernel functions for architecture search spaces in order to use classic GP-based BO methods, but so far without achieving new state-of-the-art performance. In contrast, several works use tree-based models (in particular, treed Parzen estimators (Bergstra et al, 2011), or random forests (Hut-

ter et al, 2011)) to effectively search very high-dimensional conditional spaces and achieve state-of-the-art performance on a wide range of problems, optimizing both neural architectures and their hyperparameters jointly (Bergstra et al, 2013; Domhan et al, 2015; Mendoza et al, 2016; Zela et al, 2018). While a full comparison is lacking, there is preliminary evidence that these approaches can also outperform evolutionary algorithms (Klein et al, 2018).

Architectural search spaces have also been explored in a *hierarchical* manner, e.g., in combination with evolution (Liu et al, 2018a) or by sequential model-based optimization (Liu et al, 2017). Negrinho and Gordon (2017) and Wistuba (2017) exploit the tree-structure of their search space and use *Monte Carlo Tree Search*. Elsken et al (2017) propose a simple yet well performing *hill climbing* algorithm that discovers high-quality architectures by greedily moving in the direction of better performing architectures without requiring more sophisticated exploration mechanisms.

In contrast to the gradient-free optimization methods above, Liu et al (2018b) propose a continuous relaxation of the search space to enable *gradient-based optimization*: instead of fixing a single operation  $o_i$  (e.g., convolution or pooling) to be executed at a specific layer, the authors compute a convex combination from a set of operations  $\{o_1, \dots, o_m\}$ . More specifically, given a layer input  $x$ , the layer output  $y$  is computed as  $y = \sum_{i=1}^m \lambda_i o_i(x)$ ,  $\lambda_i \geq 0$ ,  $\sum_{i=1}^m \lambda_i = 1$ , where the convex coefficients  $\lambda_i$  effectively parameterize the network architecture. Liu et al (2018b) then optimize both the network weights and the network architecture by alternating gradient descent steps on training data for weights and on validation data for architectural parameters such as  $\lambda$ . Eventually, a discrete architecture is obtained by choosing the operation  $i$  with  $i = \arg \max_i \lambda_i$  for every layer. Shin et al (2018) and Ahmed and Torresani (2017) also employ gradient-based optimization of neural architectures, however they only consider optimizing layer hyperparameters or connectivity patterns, respectively.

## 4 Performance Estimation Strategy

The search strategies discussed in Section 3 aim at finding a neural architecture  $A$  that maximizes some performance measure, such as accuracy on unseen data. To guide their search process, these strategies need to estimate the performance of a given architecture  $A$  they consider. The simplest way of doing this is to train  $A$  on training data and evaluate its performance on validation data. However, training each architecture to be evaluated from scratch frequently yields computational demands in the order of thousands of GPU days for NAS (Zoph and Le, 2017; Real et al, 2017; Zoph et al, 2018; Real et al, 2018).

To reduce this computational burden, performance can be estimated based on *lower fidelities* of the actual performance after full training (also denoted as proxy metrics). Such lower fidelities include shorter training times (Zoph et al, 2018; Zela et al, 2018), training on a subset of the data (Klein et al, 2017a), on lower-resolution images (Chrabaszcz et al, 2017), or with less filters per layer (Zoph et al, 2018; Real et al, 2018). While these low-fidelity approximations



reduce the computational cost, they also introduce bias in the estimate as performance will typically be underestimated. This may not be problematic as long as the search strategy only relies on ranking different architectures and the relative ranking remains stable. However, recent results indicate that this relative ranking can change dramatically when the difference between the cheap approximations and the “full” evaluation is too big (Zela et al, 2018), arguing for a gradual increase in fidelities (Li et al, 2017; Falkner et al, 2018).

Another possible way of estimating an architecture’s performance builds upon learning curve extrapolation (Swersky et al, 2014; Domhan et al, 2015; Klein et al, 2017b; Baker et al, 2017b). Domhan et al (2015) propose to extrapolate initial learning curves and terminate those predicted to perform poorly to speed up the architecture search process. Swersky et al (2014), Klein et al (2017b), and Baker et al (2017b) also consider architectural hyperparameters for predicting which partial learning curves are most promising. Training a surrogate model for predicting the performance of novel architectures is also proposed by Liu et al (2017), who do not employ learning curve extrapolation but support predicting performance based on architectural/cell properties and extrapolate to architectures/cells with larger size than seen during training. The main challenge for predicting the performances of neural architectures is that, in order to speed up the search process, good predictions in a relatively large search space need to be made based on relatively few evaluations.

Another approach to speed up performance estimation is to initialize the weights of novel architectures based on weights of other architectures that have been trained before. One way of achieving this, dubbed *network morphisms* (Wei et al, 2016), allows modifying an architecture while leaving the function represented by the network unchanged (Cai et al, 2018a,b; Elsken et al, 2017; Elsken et al, 2018). This allows increasing capacity of networks successively and retaining high performance without requiring training from scratch. Continuing training for a few epochs can also make use of the additional capacity introduced by network morphisms. An advantage of these approaches is that they allow search spaces without an inherent upper bound on the architecture’s size (Elsken et al, 2017); on the other hand, strict network morphisms can only make architectures larger and may thus lead to overly complex architectures. This can be attenuated by employing approximate network morphisms that allow shrinking architectures (Elsken et al, 2018).

*One-Shot Architecture Search* is another promising approach for speeding up performance estimation, which treats all architectures as different subgraphs of a supergraph (the one-shot model) and shares weights between architectures that have edges of this supergraph in common (Saxena and Verbeek, 2016; Brock et al, 2017; Pham et al, 2018; Liu et al, 2018b; Bender et al, 2018). Only the weights of the one-shot model are trained, and architectures (which are just subgraphs of the one-shot model) can thereupon be evaluated without any further training by inheriting trained weights from the one-shot model. This greatly speeds up performance estimation of architectures, since no training is required (only evaluating performance on validation data). This approach typically incurs a large

bias as it underestimates the actual performance of architectures severely; nevertheless, it allows ranking architectures reliably, since the estimated performance correlates strongly with the actual performance (Bender et al, 2018). Different one-shot NAS methods differ in how the one-shot model is trained: ENAS (Pham et al, 2018) learns an RNN controller that samples architectures from the search space and trains the one-shot model based on approximate gradients obtained through REINFORCE. DARTS (Liu et al, 2018b) optimizes all weights of the one-shot model jointly with a continuous relaxation of the search space obtained by placing a mixture of candidate operations on each edge of the one-shot model. Bender et al (2018) only train the one-shot model once and show that this is sufficient when deactivating parts of this model stochastically during training using path dropout. While ENAS and DARTS optimize a distribution over architectures during training, the approach of Bender et al (2018) can be seen as using a fixed distribution. The high performance obtainable by the approach of Bender et al (2018) indicates that the combination of weight sharing and a fixed (carefully chosen) distribution might (perhaps surprisingly) be the only required ingredients for one-shot NAS. Related to these approaches is meta-learning of hypernetworks that generate weights for novel architectures and thus requires only training the hypernetwork but not the architectures themselves (Brock et al, 2017). The main difference here is that weights are not strictly shared but sampled from a distribution that is conditioned on the architecture and represented by the hypernetwork.

A general limitation of one-shot NAS is that the supergraph defined a-priori restricts the search space to its subgraphs. Moreover, approaches which require that the entire supergraph resides in GPU memory during architecture search will be restricted to relatively small supergraphs and search spaces accordingly and are thus typically used in combination with cell-based search spaces. While approaches based on weight-sharing have substantially reduced the computational resources required for NAS (from thousands to a few GPU days), it is currently not well understood which biases they introduce into the search if the sampling distribution of architectures is optimized along with the one-shot model. For instance, an initial bias in exploring certain parts of the search space more than others might lead to the weights of the one-shot model being better adapted for these architectures, which in turn would reinforce the bias of the search to these parts of the search space. This might result in premature convergence of NAS and might be one advantage of a fixed sampling distribution as used by Bender et al (2018). In general, a more systematic analysis of biases introduced by different performance estimators would be a desirable direction for future work.

## 5 Future Directions

In this section, we discuss several current and future directions for research on NAS. Most existing work has focused on NAS for image classification. On the one hand, this provides a challenging benchmark since a lot of manual engineering

has been devoted to finding architectures that perform well in this domain and are not easily outperformed by NAS. On the other hand, it is relatively easy to define a well-suited search space by utilizing knowledge from manual engineering. This in turn makes it unlikely that NAS will find architectures that substantially outperform existing ones considerably since the found architectures cannot differ fundamentally. We thus consider it important to go beyond image classification problems by applying NAS to less explored domains. Notable first steps in this direction are applying NAS to language modeling (Zoph and Le, 2017), music modeling (Rawal and Miikkulainen, 2018), and generative models (Suganuma et al, 2018); applications to reinforcement learning, generative adversarial networks, or sensor fusion could be further promising future directions.

An alternative direction is developing NAS methods for multi-task problems (Liang et al, 2018; Meyerson and Miikkulainen, 2018) and for multi-objective problems (Elsken et al, 2018; Dong et al, 2018; Zhou et al, 2018), in which measures of resource efficiency are used as objectives along with the predictive performance on unseen data. Likewise, it would be interesting to extend RL/bandit approaches, such as those discussed in Section 3, to learn policies that are conditioned on a state that encodes task properties/resource requirements (i.e., turning the setting into a contextual bandit). A similar direction was followed by Ramachandran and Le (2018) in extending one-shot NAS such that it can generate different architectures depending on the task or instance on-the-fly. Moreover, applying NAS to searching for architectures that are more robust to adversarial examples (Cubuk et al, 2017) is an intriguing recent direction.

Related to this is research on defining more general and flexible search spaces. For instance, while the cell-based search space provides high transferability between different image classification tasks, it is largely based on human experience on image classification and does not generalize easily to other domains where the hard-coded hierarchical structure (repeating the same cells several times in a chain-like structure) does not apply (e.g., semantic segmentation or object detection). A search space which allows representing and identifying more general hierarchical structure would thus make NAS more broadly applicable, see Liu et al (2018a) for first work in this direction. Moreover, common search spaces are also based on predefined building blocks, such as different kinds of convolutions and pooling, but do not allow identifying novel building blocks on this level; going beyond this limitation might substantially increase the power of NAS.

The comparison of different methods for NAS is complicated by the fact that measurements of an architecture’s performance depend on many factors other than the architecture itself. While most authors report results on the CIFAR-10 dataset, experiments often differ with regard to search space, computational budget, data augmentation, training procedures, regularization, and other factors. For example, for CIFAR-10, performance substantially improves when using a cosine annealing learning rate schedule (Loshchilov and Hutter, 2017), data augmentation by CutOut (Devries and Taylor, 2017), by MixUp (Zhang et al, 2017) or by a combination of factors (Cubuk et al, 2018), and regularization by Shake-Shake regularization (Gastaldi, 2017) or scheduled drop-path (Zoph et al, 2018).

It is therefore conceivable that improvements in these ingredients have a larger impact on reported performance numbers than the better architectures found by NAS. We thus consider the definition of common benchmarks to be crucial for a fair comparison of different NAS methods. A first step in this direction is the definition of a benchmark for joint architecture and hyperparameter search for a fully connected neural network with two hidden layers (Klein et al, 2018). In this benchmark, nine discrete hyperparameters need to be optimized that control both architecture and optimization/regularization. All 62,208 possible hyperparameter combinations have been pre-evaluated such that different methods can be compared with low computational resources. However, the search space is still very simple compared to the spaces employed by most NAS methods. It would also be interesting to evaluate NAS methods not in isolation but as part of a full open-source AutoML system, where also hyperparameters (Mendoza et al, 2016; Real et al, 2017; Zela et al, 2018), and data augmentation pipeline (Cubuk et al, 2018) are optimized along with NAS.

While NAS has achieved impressive performance, so far it provides little insights into why specific architectures work well and how similar the architectures derived in independent runs would be. Identifying common motifs, providing an understanding why those motifs are important for high performance, and investigating if these motifs generalize over different problems would be desirable.

**Acknowledgements** We would like to thank Esteban Real, Arber Zela and Thomas Pfeil for feedback on earlier versions of this paper. This work has partly been supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme under grant no. 716721.

# Bibliography

- Ahmed K, Torresani L (2017) Connectivity learning in multi-branch networks. In: arXiv:1709.09582
- Angeline PJ, Saunders GM, Pollack JB (1994) An evolutionary algorithm that constructs recurrent neural networks. *IEEE transactions on neural networks* 5 1:54-65
- Baker B, Gupta O, Naik N, Raskar R (2017a) Designing neural network architectures using reinforcement learning. In: *International Conference on Learning Representations*
- Baker B, Gupta O, Raskar R, Naik N (2017b) Accelerating Neural Architecture Search using Performance Prediction. In: *NIPS Workshop on Meta-Learning*
- Bender G, Kindermans PJ, Zoph B, Vasudevan V, Le Q (2018) Understanding and simplifying one-shot architecture search. In: *International Conference on Machine Learning*
- Bergstra J, Yamins D, Cox DD (2013) Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In: *ICML*
- Bergstra JS, Bardenet R, Bengio Y, Kégl B (2011) Algorithms for hyperparameter optimization. In: Shawe-Taylor J, Zemel RS, Bartlett PL, Pereira F, Weinberger KQ (eds) *Advances in Neural Information Processing Systems* 24, pp 2546-2554
- Brock A, Lim T, Ritchie JM, Weston N (2017) SMASH: one-shot model architecture search through hypernetworks. In: *NIPS Workshop on Meta-Learning*
- Cai H, Chen T, Zhang W, Yu Y, Wang J (2018a) Efficient architecture search by network transformation. In: *Association for the Advancement of Artificial Intelligence*
- Cai H, Yang J, Zhang W, Han S, Yu Y (2018b) Path-Level Network Transformation for Efficient Architecture Search. In: *International Conference on Machine Learning*
- Chen T, Goodfellow IJ, Shlens J (2016) Net2net: Accelerating learning via knowledge transfer. In: *International Conference on Learning Representations*
- Chollet F (2016) Xception: Deep learning with depthwise separable convolutions. arXiv:161002357
- Chrabaszcz P, Loshchilov I, Hutter F (2017) A downsampled variant of imagenet as an alternative to the CIFAR datasets. CoRR abs/1707.08819
- Cubuk ED, Zoph B, Schoenholz SS, Le QV (2017) Intriguing Properties of Adversarial Examples. In: arXiv:1711.02846
- Cubuk ED, Zoph B, Mane D, Vasudevan V, Le QV (2018) AutoAugment: Learning Augmentation Policies from Data. In: arXiv:1805.09501
- Devries T, Taylor GW (2017) Improved regularization of convolutional neural networks with cutout. arXiv preprint abs/1708.04552
- Domhan T, Springenberg JT, Hutter F (2015) Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning

- curves. In: Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)
- Dong JD, Cheng AC, Juan DC, Wei W, Sun M (2018) Ppp-net: Platform-aware progressive search for pareto-optimal neural architectures. In: International Conference on Learning Representations Workshop
- Elsken T, Metzen JH, Hutter F (2017) Simple And Efficient Architecture Search for Convolutional Neural Networks. In: NIPS Workshop on Meta-Learning
- Elsken T, Metzen JH, Hutter F (2018) Efficient Multi-objective Neural Architecture Search via Lamarckian Evolution. ArXiv e-prints 1804.09081
- Falkner S, Klein A, Hutter F (2018) BOHB: Robust and efficient hyperparameter optimization at scale. In: Dy J, Krause A (eds) Proceedings of the 35th International Conference on Machine Learning, PMLR, Stockholmssan, Stockholm Sweden, Proceedings of Machine Learning Research, vol 80, pp 1436–1445
- Floreano D, Drr P, Mattiussi C (2008) Neuroevolution: from architectures to learning. *Evolutionary Intelligence* 1(1):47–62
- Gastaldi X (2017) Shake-shake regularization. In: International Conference on Learning Representations Workshop
- Goldberg DE, Deb K (1991) A comparative analysis of selection schemes used in genetic algorithms. In: Foundations of Genetic Algorithms, Morgan Kaufmann, pp 69–93
- He K, Zhang X, Ren S, Sun J (2016) Deep Residual Learning for Image Recognition. In: Conference on Computer Vision and Pattern Recognition
- Huang G, Liu Z, Weinberger KQ (2017) Densely Connected Convolutional Networks. In: Conference on Computer Vision and Pattern Recognition
- Hutter F, Hoos H, Leyton-Brown K (2011) Sequential model-based optimization for general algorithm configuration. In: LION, pp 507–523
- Kandasamy K, Neiswanger W, Schneider J, Poczos B, Xing E (2018) Neural Architecture Search with Bayesian Optimisation and Optimal Transport. arXiv:180207191
- Klein A, Falkner S, Bartels S, Hennig P, Hutter F (2017a) Fast Bayesian Optimization of Machine Learning Hyperparameters on Large Datasets. In: Singh A, Zhu J (eds) Proceedings of the 20th International Conference on Artificial Intelligence and Statistics, PMLR, Fort Lauderdale, FL, USA, Proceedings of Machine Learning Research, vol 54, pp 528–536
- Klein A, Falkner S, Springenberg JT, Hutter F (2017b) Learning curve prediction with Bayesian neural networks. In: International Conference on Learning Representations
- Klein A, Christiansen E, Murphy K, Hutter F (2018) Towards reproducible neural architecture and hyperparameter search. In: ICML 2018 Workshop on Reproducibility in ML (RML 2018)
- Li L, Jamieson K, DeSalvo G, Rostamizadeh A, Talwalkar A (2017) Hyperband: bandit-based configuration evaluation for hyperparameter optimization. In: International Conference on Learning Representations
- Liang J, Meyerson E, Miikkulainen R (2018) Evolutionary Architecture Search For Deep Multitask Networks. In: arXiv:1803.03745

- Liu C, Zoph B, Shlens J, Hua W, Li LJ, Fei-Fei L, Yuille A, Huang J, Murphy K (2017) Progressive Neural Architecture Search. In: arXiv:1712.00559
- Liu H, Simonyan K, Vinyals O, Fernando C, Kavukcuoglu K (2018a) Hierarchical Representations for Efficient Architecture Search. In: International Conference on Learning Representations
- Liu H, Simonyan K, Yang Y (2018b) Darts: Differentiable architecture search. In: arXiv:1806.09055
- Loshchilov I, Hutter F (2017) Sgdr: Stochastic gradient descent with warm restarts. In: International Conference on Learning Representations
- Mendoza H, Klein A, Feurer M, Springenberg J, Hutter F (2016) Towards Automatically-Tuned Neural Networks. In: International Conference on Machine Learning, AutoML Workshop
- Meyerson E, Miikkulainen R (2018) Pseudo-task Augmentation: From Deep Multitask Learning to Intratask Sharing and Back. In: arXiv:1803.03745
- Miikkulainen R, Liang J, Meyerson E, Rawal A, Fink D, Francon O, Raju B, Shahrzad H, Navruzyan A, Duffy N, Hodjat B (2017) Evolving Deep Neural Networks. In: arXiv:1703.00548
- Negrinho R, Gordon G (2017) DeepArchitect: Automatically Designing and Training Deep Architectures. arXiv:170408792
- Pham H, Guan MY, Zoph B, Le QV, Dean J (2018) Efficient neural architecture search via parameter sharing. In: International Conference on Machine Learning
- Ramachandran P, Le QV (2018) Dynamic Network Architectures. In: AutoML 2018 (ICML workshop)
- Rawal A, Miikkulainen R (2018) From Nodes to Networks: Evolving Recurrent Neural Networks. In: arXiv:1803.04439
- Real E, Moore S, Selle A, Saxena S, Suematsu YL, Le QV, Kurakin A (2017) Large-scale evolution of image classifiers. International Conference on Machine Learning
- Real E, Aggarwal A, Huang Y, Le QV (2018) Regularized Evolution for Image Classifier Architecture Search. In: arXiv:1802.01548
- Saxena S, Verbeek J (2016) Convolutional neural fabrics. In: Lee DD, Sugiyama M, Luxburg UV, Guyon I, Garnett R (eds) *Advances in Neural Information Processing Systems 29*, Curran Associates, Inc., pp 4053–4061
- Shin R, Packer C, Song D (2018) Differentiable neural network architecture search. In: International Conference on Learning Representations Workshop
- Stanley KO, Miikkulainen R (2002) Evolving neural networks through augmenting topologies. *Evolutionary Computation* 10:99–127
- Stanley KO, D’Ambrosio DB, Gauci J (2009) A hypercube-based encoding for evolving large-scale neural networks. *Artif Life* 15(2):185–212, DOI 10.1162/artl.2009.15.2.15202, URL <http://dx.doi.org/10.1162/artl.2009.15.2.15202>
- Suganuma M, Shirakawa S, Nagao T (2017) A genetic programming approach to designing convolutional neural network architectures. In: Genetic and Evolutionary Computation Conference
- Suganuma M, Ozay M, Okatani T (2018) Exploiting the potential of standard convolutional autoencoders for image restoration by evolutionary search. In:

- Dy J, Krause A (eds) Proceedings of the 35th International Conference on Machine Learning, PMLR, Stockholmsmssan, Stockholm Sweden, Proceedings of Machine Learning Research, vol 80, pp 4771–4780
- Swersky K, Duvenaud D, Snoek J, Hutter F, Osborne M (2013) Raiders of the lost architecture: Kernels for bayesian optimization in conditional parameter spaces. In: NIPS Workshop on Bayesian Optimization in Theory and Practice
- Swersky K, Snoek J, Adams RP (2014) Freeze-thaw bayesian optimization
- Szegedy C, Vanhoucke V, Ioffe S, Shlens J, Wojna Z (2016) Rethinking the Inception Architecture for Computer Vision. In: Conference on Computer Vision and Pattern Recognition
- Wei T, Wang C, Rui Y, Chen CW (2016) Network morphism. In: International Conference on Machine Learning
- Wei T, Wang C, Chen CW (2017) Modularized morphing of neural networks. arXiv:170103281
- Wistuba M (2017) Finding Competitive Network Architectures Within a Day Using UCT. In: arXiv:1712.07420
- Xie L, Yuille A (2017) Genetic CNN. In: International Conference on Computer Vision
- Yu F, Koltun V (2016) Multi-scale context aggregation by dilated convolutions
- Zela A, Klein A, Falkner S, Hutter F (2018) Towards automated deep learning: Efficient joint neural architecture and hyperparameter search. In: ICML 2018 Workshop on AutoML (AutoML 2018)
- Zhang H, Cissé M, Dauphin YN, Lopez-Paz D (2017) mixup: Beyond empirical risk minimization. arXiv preprint abs/1710.09412
- Zhong Z, Yan J, Liu CL (2018) Practical Network Blocks Design with Q-Learning. Association for the Advancement of Artificial Intelligence
- Zhou Y, Ebrahimi S, Ark S, Yu H, Liu H, Damos G (2018) Resource-efficient neural architect. In: arXiv:1806.07912
- Zoph B, Le QV (2017) Neural architecture search with reinforcement learning. In: International Conference on Learning Representations
- Zoph B, Vasudevan V, Shlens J, Le QV (2018) Learning transferable architectures for scalable image recognition. In: Conference on Computer Vision and Pattern Recognition