UNAS: Differentiable Architecture Search Meets Reinforcement Learning

Arash Vahdat, Arun Mallya, Ming-Yu Liu, Jan Kautz
NVIDIA
{avahdat, amallya, mingyul, jkautz}@nvidia.com

Abstract

Neural architecture search (NAS) aims to discover network architectures with desired properties such as high accuracy or low latency. Recently, differentiable NAS (DNAS) has demonstrated promising results while maintaining a search cost orders of magnitude lower than reinforcement learning (RL) based NAS. However, DNAS models can only optimize differentiable loss functions in search, and they require an accurate differentiable approximation of non-differentiable criteria. In this work, we present UNAS, a unified framework for NAS, that encapsulates recent DNAS and RL-based approaches under one framework. Our framework brings the best of both worlds, and it enables us to search for architectures with both differentiable and non-differentiable criteria in one unified framework while maintaining a low search cost. Further, we introduce a new objective function for search based on the generalization gap that prevents the selection of architectures prone to overfitting. We present extensive experiments on the CIFAR-10, CIFAR-100 and ImageNet datasets and we perform search in two fundamentally different search spaces. We show that UNAS obtains the state-of-the-art average accuracy on all three datasets when compared to the architectures searched in the DARTS [18] space. Moreover, we show that UNAS can find an efficient and accurate architecture in the ProxylessNAS [28] search space, that outperforms existing MobileNetV2 [28] based architectures. The source code is available at https://github.com/NVlabs/unas.

1. Introduction

Since the success of deep learning, designing neural network architectures with desirable performance criteria (*e.g.* high accuracy, low latency, *etc*.) for a given task has been a challenging problem. Some call it alchemy and some refer to it as intuition, but the task of discovering a novel architecture often involves a tedious and costly process of trial-and-error for searching in an exponentially large space of hyper-parameters. The goal of neural architecture search (NAS) [6] is to find novel networks for new problem domains and criteria automatically and efficiently.

Early work on NAS used reinforcement learning [1, 3, 24, 42, 43], or evolutionary algorithms [17, 27, 26, 37] to obtain state-of-the-art performance on a variety of tasks. Although, these methods are generic and can search for architecture with a broad range of criteria, they are often computationally demanding. For example, the RL-based approach [43], and evolutionary method [26] each requires over 2000 GPU days.

Recently, several differentiable neural architecture search (DNAS) frameworks [18, 39, 36, 4] have shown promising results while reducing the search cost to a few GPU days. However, these approaches assume that the objective function is differentiable with respect to the architecture parameters and cannot directly optimize non-differentiable criteria like network latency, power consumption, memory usage, *etc*. To tackle this problem, DNAS methods [36, 4, 40] approximate network latency using differentiable functions. However, these approximations may fail when the underlying criteria cannot be accurately modeled. For example, if compiler optimizations are used, methods such as layer fusion, mixed-precision inference, and kernel auto-tuning can dramatically change latency, making it challenging to approximate it accurately. In addition to the loss approximation, DNAS relies on the continuous approximation of discrete variables in search, introducing additional mismatch in network performance between discovered architecture and the corresponding continuous relaxations.

In this paper, we introduce UNAS, a unified framework for NAS that bridges the gap between DNAS and RL-based architecture search. **(i)** UNAS offers the best of both worlds and enables us to search for architectures using both differentiable objective functions (*e.g.*, cross-entropy loss) and non-differentiable functions (*e.g.*, network latency). UNAS keeps the search time low similar to other DNAS models, but it also eliminates the need for accurate approximation of non-differentiable criteria. **(ii)** UNAS training does not introduce any additional biases due to the continuous relaxation of architecture parameters. We show that the gradient estimation in UNAS is equal to the estimations obtained by RL-based frameworks that operate on discrete variables. Finally, **(iii)** UNAS proposes a new objective function based

on the generalization gap which is empirically shown to find architectures less prone to overfitting.

We perform extensive experiments in both DARTS [18] and ProxylessNAS [4] search spaces. We show that UNAS achieves the state-of-the-art average performance on all three datasets in comparison to the recent gradient-based NAS models in the DARTS space. Moreover, UNAS can find architectures that are faster and more accurate than architectures, searched in the ProxylessNAS space.

1.1. Related Work

Zoph and Le [42] introduced the paradigm of NAS, where a controller recurrent neural network (RNN) was trained to output the specification of a network (filter sizes, number of channels, *etc*.). The controller was trained using REINFORCE [35] to maximize the expected accuracy of the output network on the target validation set, after training on the target task. Requiring the method to specify every layer of the network made it challenging to deepen or transfer an obtained network to other tasks. Based on the observation that popular manually-designed convolutional neural networks (CNNs) such as ResNet [10] or Inception [30] contained repeated generic blocks with the same structure, Zoph *et al*. [43] trained the RNN to output stackable '*cells*'. The task of NAS was thus reduced to learning two types of cells, the *Normal Cell* - convolutional cells that preserve the spatial dimensions, and the *Reduce Cell* - convolutional cells that reduce spatial dimensions while increasing feature maps.

Recently, DARTS [18] relaxed the architecture search space to be continuous by using a weighted mixture-of-operations and optimized the candidate architecture through gradient descent. Using weight-sharing [2, 24], they brought search down to a few GPU days. As the final architecture is required to be discrete, DARTS only retained the top two operations based on the weight assigned to each operation. Building upon DARTS, SNAS [39] used weights sampled from a trainable Gumbel-Softmax distribution instead of continuous weights. Both DARTS and SNAS assume that the objective function for search is differentiable. We extend these frameworks by introducing unbiased gradient estimators that can work for both differentiable and non-differentiable objective functions.

Recent works [4, 36, 40, 11, 31] consider latency in architectures search. ProxylessNAS [4], FBNet [36] and NetAdapt [40] convert the non-differentiable latency objective to a differentiable function by learning an accurate latency approximation. However, these approximations may fail when latency cannot be predicted by a trainable function. MnasNet [31] does not require a differentiable approximation of the latency as it relies on an RL-objective, however, it requires ∼300 TPU-days for each architecture search. Our framework bridges the gap between differentiable and RL-based NAS; it can search with differentiable and non-differentiable functions and it does not require an accurate approximation of non-differentiable terms in the objective. Our work is compared against previous works in Table. 1.

	DARTS [18]	SNAS [39]	P-DARTS [16]	ProxylessNAS [4]	FBNet [36]	MnasNet [31]	UNAS (ours)
Differentiable loss	✓	✓	✓	✓	✓	✓	✓
Non-differentiable loss	✗	✗	✗	✗	✗	✓	✓
Latency optimization	✗	✗	✗	✓	✓	✓	✓
Low search cost	✓	✓	✓	✓	✓	✗	✓

Table 1: Comparison with differentiable NAS methods.

Recently, P-DARTS [5] proposes a progressive version of DARTS and shows that by gradually increasing the depth of the network during the search, deeper cells can be discovered. UNAS explores an orthogonal direction to P-DARTS and it proposes generic gradient estimators that work with both differentiable and non-differentiable losses and new generalization-based search objective functions.

2. Background

In differentiable architecture search (DARTS) [18], a network is represented by a directed acyclic graph, where each node in the graph denotes a hidden representation (*e.g*., feature maps in CNNs) and each directed edge represents an operation transforming the state of the input node. The nth node x_n is connected to its predecessors (*i.e*., P_n) and its content is computed by applying a set of operations to the predeceasing nodes, represented by $x_n = \sum_{x_m \in P_n} O_{m,n}(x_m)$, where $O_{m,n}$ is the operation applied to x_m. The goal of architecture search is then to find the operation $O_{m,n}$ for each edge (m, n). Representing the set of all possible operations that can be applied to the edge $e := (m, n)$ using $\{O_e^{(1)}, O_e^{(2)}, \dots, O_e^{(K)}\}$ where K is the number of operations, this discrete assignment problem can be formulated as a *mixed operation* denoted by $O_e(x_m) = \sum_{k=1}^{K} z_e^{(k)} O_e^{(k)}(x_m)$, where $\boldsymbol{z}_e = [z_e^{(1)}, z_e^{(2)}, \dots, z_e^{(K)}]$ is a one-hot binary vector (*i.e*., $z_e^{(k)} \in \{0, 1\}$) with a single one indicating the selected operation. Typically, it is assumed that the set of operations also includes a *zero* operation that enables omitting edges in the network, and thus, learning the connectivity as well.

We can construct a network architecture given the set of all operation assignments for all edges denoted by $\boldsymbol{z} = \{\boldsymbol{z}_e\}$. Therefore, the objective of the architecture search is to find a distribution over architecture parameters, $\boldsymbol{z}$ such that it minimizes the expected loss $\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}(\boldsymbol{z})]$ where $p_{\boldsymbol{\phi}}$ is a $\boldsymbol{\phi}$-parameterized distribution over $\boldsymbol{z}$ and $\mathcal{L}(\boldsymbol{z})$ is a loss function measuring the performance of the architecture specified by $\boldsymbol{z}$ using a performance measure such as classification loss.

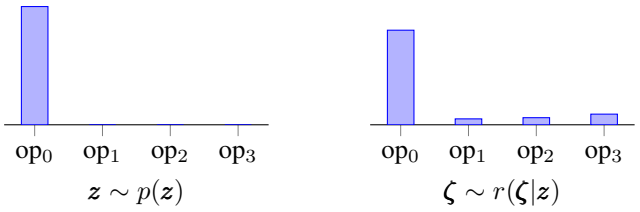

(a) categorical sample (b) smoothed sample

Figure 1: (a) Operation selection corresponds to sampling from a categorical distribution that selects an operation. (b) Sampling from the conditional Gumbel-Softmax distribution $r(\boldsymbol{\zeta}|\boldsymbol{z})$ acts as a smoothing distribution that yields continuous samples ($\boldsymbol{\zeta}$), correlated with the discrete samples ($\boldsymbol{z}$).

We assume that the architecture distribution is a factorial distribution with the form $p_{\boldsymbol{\phi}}(\boldsymbol{z}) = \prod_e p_{\boldsymbol{\phi}_e}(\boldsymbol{z}_e)$ where $p_{\boldsymbol{\phi}_e}(\boldsymbol{z}_e)$ is a $\boldsymbol{\phi}_e$-parameterized categorical distribution defined over the one-hot vector $\boldsymbol{z}_e$. Recently, SNAS [39] proposed using the Gumbel-Softmax relaxation [20, 14] for optimizing the expected loss. In this case, the categorical distribution $p_{\boldsymbol{\phi}}(\boldsymbol{z})$ is replaced with a Gumbel-Softmax distribution $p_{\boldsymbol{\phi}}(\boldsymbol{\zeta})$ where $\boldsymbol{\zeta}$ denotes the continuous relaxation of the architecture parameter $\boldsymbol{z}$. SNAS assumes that the loss $\mathcal{L}(\boldsymbol{z})$ is differentiable with respect to $\boldsymbol{z}$ and it uses the reparameterization trick to minimize the expectation of the relaxed loss $\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{\zeta})}[\mathcal{L}(\boldsymbol{\zeta})]$ instead of $\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}(\boldsymbol{z})]$.

3. Method

As discussed above, the problem of NAS can be formulated as optimizing the expected loss $\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}(\boldsymbol{z})]$. In this section, we present our framework in two parts. In Sec. 3.1, we start by presenting a general framework for computing $\frac{\partial}{\partial \boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}(\boldsymbol{z})]$ which is required for optimizing the expected loss. Then, we present our formulation of the loss function $\mathcal{L}(\boldsymbol{z})$ in Sec. 3.2.

3.1. Gradient Estimation

The most generic approach for optimizing the expected loss is the REINFORCE gradient estimator

$$\frac{\partial}{\partial \boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}(\boldsymbol{z})] = \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}\left[\mathcal{L}(\boldsymbol{z}) \partial_{\boldsymbol{\phi}} \log p_{\boldsymbol{\phi}}(\boldsymbol{z})\right], \quad (1)$$

where $\partial \log p_{\boldsymbol{\phi}}(\boldsymbol{z})$ is known as the score function and $\mathcal{L}(\boldsymbol{z})$ is a loss function. As we can see, the gradient estimator in Eq. 1 only requires computing the loss function $\mathcal{L}(\boldsymbol{z})$ (not the gradient $\partial_{\boldsymbol{z}} \mathcal{L}(\boldsymbol{z})$), so it can be applied to any differentiable and non-differentiable loss function. However, this estimator is known to suffer from high variance and therefore a large number of trained architecture samples are required to reduce its variance, making it extremely compute intensive. The REINFORCE estimator in Eq. 1 can be also rewritten as

$$\frac{\partial}{\partial \boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}(\boldsymbol{z})] = \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[(\mathcal{L}(\boldsymbol{z}) - c(\boldsymbol{z}))\, \partial_{\boldsymbol{\phi}} \log p_{\boldsymbol{\phi}}(\boldsymbol{z})] + \partial_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[c(\boldsymbol{z})], \quad (2)$$

where $c(\boldsymbol{z})$ is a control variate [22]. The gradient estimator in Eq. 2 has lower variance than Eq. 1, if $c(\boldsymbol{z})$ is correlated with $\mathcal{L}(\boldsymbol{z})$, and $\partial_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[c(\boldsymbol{z})]$ has a low-variance gradient estimator [21, 23, 25].[1] Without loss of generality, we assume that the loss function is decomposed into $\mathcal{L}(\boldsymbol{z}) = \mathcal{L}_d(\boldsymbol{z}) + \mathcal{L}_n(\boldsymbol{z})$ where $\mathcal{L}_d(\boldsymbol{z})$ contains the terms that are differentiable with respect to $\boldsymbol{z}$ and $\mathcal{L}_n(\boldsymbol{z})$ includes the non-differentiable terms. We present a baseline function $c(\boldsymbol{z}) = c_d(\boldsymbol{z}) + c_n(\boldsymbol{z})$, where $c_d(\boldsymbol{z})$ and $c_n(\boldsymbol{z})$ are for $\mathcal{L}_d(\boldsymbol{z})$ and $\mathcal{L}_n(\boldsymbol{z})$ respectively. Intuitively, the baseline is designed such that the term $\partial_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[c(\boldsymbol{z})]$ in Eq. 2 is approximated using the low-variance reparameterization trick.

Gradient Estimation for Differentiable Loss $\mathcal{L}_d$: Following REBAR [33], in order to construct $c_d(\boldsymbol{z})$, a control variate for $\mathcal{L}_d$, we use stochastic continuous relaxation $r_{\boldsymbol{\phi}}(\boldsymbol{\zeta}|\boldsymbol{z})$ that samples from a conditional Gumbel-Softmax distribution given the architecture sample $\boldsymbol{z}$. Here, $\boldsymbol{\zeta}$ can be considered as a smooth architecture defined based on $\boldsymbol{z}$ as shown in Fig. 1. Hence, it is highly correlated with $\boldsymbol{z}$ (see REBAR [33] for details). With the definition $c_d(\boldsymbol{z}) := \mathbb{E}_{r_{\boldsymbol{\phi}}(\boldsymbol{\zeta}|\boldsymbol{z})}[\mathcal{L}_d(\boldsymbol{\zeta})]$, the gradient in Eq. 2 can be written as

$$\frac{\partial}{\partial \boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}_d(\boldsymbol{z})] = \underbrace{\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}\left[(\mathcal{L}_d(\boldsymbol{z}) - c_d(\boldsymbol{z}))\, \partial_{\boldsymbol{\phi}} \log p_{\boldsymbol{\phi}}(\boldsymbol{z})\right]}_{\text{(i) reinforce}} - \underbrace{\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}\left[\partial_{\boldsymbol{\phi}} c_d(\boldsymbol{z})\right]}_{\text{(ii) correction}} + \underbrace{\partial_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[c_d(\boldsymbol{z})]}_{\text{(iii) Gumbel-Softmax}}. \quad (3)$$

The gradient estimator in Eq. 3 consists of three terms: (i) is the reinforce term, which is estimated using the Monte Carlo method by sampling $\boldsymbol{z} \sim p_{\boldsymbol{\phi}}(\boldsymbol{z})$ and $\boldsymbol{\zeta} \sim r_{\boldsymbol{\phi}}(\boldsymbol{\zeta}|\boldsymbol{z})$. (ii) is the correction term due to the dependency of $c_d(\boldsymbol{z})$ on $\boldsymbol{\phi}$. This term is approximated using the reparameterization trick applied to the conditional Gumbel-Softmax $r_{\boldsymbol{\phi}}(\boldsymbol{\zeta}|\boldsymbol{z})$. (iii) is the Gumbel-Softmax term that can be written as

$$\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[c_d(\boldsymbol{z})] = \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}\left[\mathbb{E}_{r_{\boldsymbol{\phi}}(\boldsymbol{\zeta}|\boldsymbol{z})}[\mathcal{L}_d(\boldsymbol{\zeta})]\right] = \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{\zeta})}[\mathcal{L}_d(\boldsymbol{\zeta})], \quad (4)$$

which is the expected value of loss evaluated under the Gumbel-Softmax distribution $p_{\boldsymbol{\phi}}(\boldsymbol{\zeta})$. Thus, its gradient can be computed also using the low-variance reparameterization trick. In practice, we only need two function evaluations for estimating the gradient in Eq. 3, one for computing $\mathcal{L}_d(\boldsymbol{z})$, and one for $\mathcal{L}_d(\boldsymbol{\zeta})$. The gradients are computed using an automatic differentiation library.

[1] The low variance of Eq. 2 comes from fact that $\mathrm{Var}(X - Y) = \mathrm{Var}(X) + \mathrm{Var}(Y) - 2\mathrm{Cov}(X, Y)$ for any random variable X and Y. If X and Y are highly correlated the negative contribution from $-2\mathrm{Cov}(X, Y)$ reduces the overall variance of $X - Y$.

Eq. 3 unifies the differentiable architecture search with policy gradient-based NAS methods [42, 39, 31]. This estimator does not introduce any bias due to the continuous relaxation, as in expectation the gradient is equal to the REINFORCE estimator that operates on discrete variables. Moreover, this estimator uses the Gumbel-Softmax estimation of the differentiable loss for reducing the variance of the estimate. Under this framework, it is easy to see that SNAS [39] is a biased estimation of the policy gradient as it only uses (iii) for search, ignoring other terms. On the other hand, policy gradient-based NAS [24, 42, 43] assumes a constant control variate ($c_d(\boldsymbol{z}) = C$) which only requires computing (i) as $\partial_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[C] = 0$.

Gradient Estimation for Non-Differentiable Loss $\mathcal{L}_n$: The gradient estimator in Eq. 3 cannot be applied to non-differentiable loss $\mathcal{L}_n(\boldsymbol{z})$ as the reparameterization trick is only applicable to differentiable functions. For $\mathcal{L}_n(\boldsymbol{z})$, we use RELAX [8] that lifts this limitation by defining the baseline function $c_n(\boldsymbol{z}) := \mathbb{E}_{r_{\boldsymbol{\phi}}(\boldsymbol{\zeta}|\boldsymbol{z})}[g(\boldsymbol{\zeta})]$, where $g(.)$ is a surrogate function (*e.g.*, a neural network) trained to be correlated with $\mathcal{L}_n(\boldsymbol{z})$. The gradient estimator for $\mathcal{L}_n$ is obtained by replacing c_d in Eq. 3 with c_n:

$$\frac{\partial}{\partial \boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}_n(\boldsymbol{z})] = \underbrace{\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}\left[(\mathcal{L}_d(\boldsymbol{z}) - c_n(\boldsymbol{z}))\, \partial_{\boldsymbol{\phi}} \log p_{\boldsymbol{\phi}}(\boldsymbol{z})\right]}_{\text{(i) reinforce}} - \underbrace{\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}\left[\partial_{\boldsymbol{\phi}} c_n(\boldsymbol{z})\right]}_{\text{(ii) correction}} + \underbrace{\partial_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{\zeta})}[g(\boldsymbol{\zeta})]}_{\text{(iii) Gumbel-Softmax}}, \quad (5)$$

However, the main difference is that here the reparameterization trick is applied to $\mathbb{E}_{r_{\boldsymbol{\phi}}(\boldsymbol{\zeta}|\boldsymbol{z})}[g(\boldsymbol{\zeta})]$ in (ii) and similarly to $\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{\zeta})}[g(\boldsymbol{\zeta})]$ in (iii). Here, to make $g(\boldsymbol{z})$ be correlated with $\mathcal{L}_n(\boldsymbol{z})$, we train g by minimizing $||g(\boldsymbol{z}) - \mathcal{L}_n(\boldsymbol{z})||_2^2$. In the case of latency, this corresponds to training g to predict latency on a set of randomly generated architectures before search. Similar to FBNet [36] and ProxylessNAS [4], we use a simple linear function to represent $g(\boldsymbol{z})$.

It is worth noting that the Gumbel-Softmax term, (iii) in Eq. 5, minimizes the expectation of the approximation of the non-differentiable loss (e.g., latency) using the Gumbel-Softmax relaxation. This gradient estimator was used in FBNet [36] for optimizing latency. In Eq. 5, we can see that if g cannot predict latency correctly, $\mathcal{L}_d(\boldsymbol{z}) - c_n(\boldsymbol{z})$ will be large, thus, optimizing only (iii) will suffer from additional bias due to the approximation error. However, even if $g(\boldsymbol{z})$ cannot approximate $\mathcal{L}_n(\boldsymbol{z})$ accurately, for example in the case of compile-time performance optimizations, our gradient estimator is equal to the REINFORCE estimator, and it optimizes the true expected latency. Hence, UNAS does not suffer from any bias introduced due to the approximation of non-differentiable criteria.

3.2. Training Objective

Several recent works on differentiable NAS have proposed bi-level training of architecture parameters and network parameters. In the architecture update, either training loss [39], or validation loss [18] given the current network parameters $\boldsymbol{w}$, are used to update architecture parameters using

$$\min_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}_{\text{train}}(\boldsymbol{z}, \boldsymbol{w})], \text{ or } \min_{\boldsymbol{\phi}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}_{\text{val}}(\boldsymbol{z}, \boldsymbol{w})]. \quad (6)$$

Then, the network parameters $\boldsymbol{w}$ are updated given samples from the architecture by minimizing

$$\min_{\boldsymbol{w}} \mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}_{\text{train}}(\boldsymbol{z}, \boldsymbol{w})]. \quad (7)$$

The parameters $\boldsymbol{\phi}$ and $\boldsymbol{w}$ are updated iteratively by taking a single gradient step in Eq. 6 and Eq. 7. It has been shown that by sharing network parameters among all the architecture instances, we gain several orders of magnitude speedup in search [18, 24]. However, this comes with the cost of updating architecture parameters at suboptimial $\boldsymbol{w}$. Intuitively, this translates to making decision on architecture without considering its optimal performance.

To avoid overfitting, we base our objective function on the generalization gap of an architecture. The rationale behind this is that the selected architecture not only should perform well on the training set, but also, should generalize equally well to the examples in the validation set, even if network weights are suboptimal. This prevents search from choosing architectures that do not generalize well. Formally, we define the generalization loss in search $\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}_{\text{gen}}(\boldsymbol{z}, \boldsymbol{w})]$ by:

$$\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\mathcal{L}_{\text{train}}(\boldsymbol{z}, \boldsymbol{w}) + \lambda |\mathcal{L}_{\text{val}}(\boldsymbol{z}, \boldsymbol{w}) - \mathcal{L}_{\text{train}}(\boldsymbol{z}, \boldsymbol{w})|], \quad (8)$$

where λ is a scalar balancing the training loss and generalization gap. We observe that $\lambda = 0.5$ often works well in our experiments.[2] For training, we iterate between updating $\boldsymbol{\phi}$ using Eq. 8 and updating $\boldsymbol{w}$ using Eq. 7. In each parameter update, we perform a simple gradient descent update.

Latency Loss: In resource-constrained applications, we might be interested in finding an architecture that has a low latency as well as high accuracy. In this case, we can measure the latency of the network specified by $\boldsymbol{z}$ in each parameter update[3]. Representing the latency of the network using $\mathcal{L}_{lat}(\boldsymbol{z})$, we augment the objective function in Eq. 8 with $\mathbb{E}_{p_{\boldsymbol{\phi}}(\boldsymbol{z})}[\lambda_{lat} \mathcal{L}_{lat}(\boldsymbol{z})]$, where λ_{lat} is a scalar balancing the trade-off between the architecture loss and the latency loss. Although $\mathcal{L}_{lat}(\boldsymbol{z})$ is not differentiable w.r.t. $\boldsymbol{z}$, we construct a low-variance gradient estimator using Eq. 5 for optimizing this term.

4. Experiments in DARTS Search Space

In this section, we apply the proposed UNAS framework to the problem of architecture search for image classifica-

[2] We also explored with the objective function without the absolute value, *i.e.*, $\mathcal{L}_{\text{train}}(\boldsymbol{z}, \boldsymbol{w}) + \lambda(\mathcal{L}_{\text{val}}(\boldsymbol{z}, \boldsymbol{w}) - \mathcal{L}_{\text{train}}(\boldsymbol{z}, \boldsymbol{w}))$. We observed that this variants does not perform as good as Eq. 8.

[3] We measure latency on the same hardware that the model is being trained.

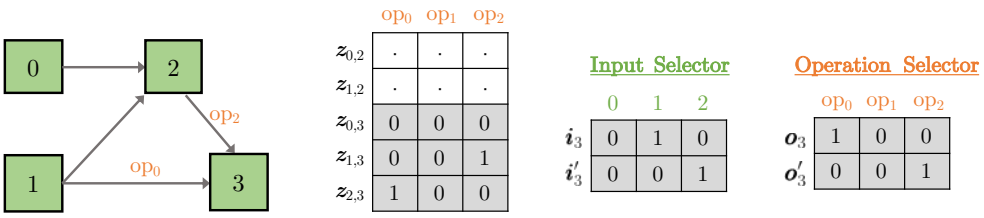

Figure 2: The factorized cell structure ensures that each node depends on two previous nodes. On the left, a small graph with 4 nodes is visualized. In the middle, $\boldsymbol{z} = \{\boldsymbol{z}_e\}$, the operation assignment for the incoming edges to node 3 is shown. On the right, the input and operation selectors for these edges are shown. The shaded matrix on $\boldsymbol{z}$ is parameterized by the outer product $\boldsymbol{i}_3 \otimes \boldsymbol{o}_3 + \boldsymbol{i}'_3 \otimes \boldsymbol{o}'_3$.

tion using DARTS [18] search space, which was also used in [43, 24, 39, 5]. We closely follow the experimental setup introduced DARTS [18]. In the search phase, we search for a normal and reduction cell using a network with a small number of feature maps and/or layers. Given the stochastic representation of the architecture, the final cells are obtained by taking the configuration that has highest probability for each node as discussed below. Then, in the evaluation phase, the cells are stacked into a larger network which is retrained from scratch. Sec. 4.1 discuses a simple approach for factorizing cells that eliminates the necessity of post-search heuristics. Sec. 4.2 provides comparisons to previous work on three datasets.

4.1. Factorized Cell Structure

Training the cell structure introduced in DARTS [18] may result in a densely connected cell where each node depends on the output of all the previous nodes. In order to induce sparsity on the connectivity, prior work [5, 18, 43] heavily relies on post-search heuristics to limit the number of incoming edges for each node. DARTS [18] uses a heuristic to prune the number of input edges to two by choosing operations with the largest weights. P-DARTS [5] uses an iterative optimization to limit the number of skip-connections and the number of incoming edges to two. The main issue with such post-search methods is that they create inconsistency between search and evaluation by constructing a cell structure without directly measuring its performance [39].

In order to explicitly induce sparsity, we factorize the operation assignment problem on the edges using two selection problems: i) an *input selector* that selects two nodes out of the previous nodes and ii) an *operation selector* that selects two operations that are applied to each selected input. We name this structure a *factorized cell* as it enables us to ensure that the content of each node depends only on two previous nodes without relying on any post-search heuristic. Formally, we introduce $\boldsymbol{i}_n$ and $\boldsymbol{i}'_n$, two one-hot vectors for the nth node representing the input selectors as well as two one-hot vectors $\boldsymbol{o}_n$ and $\boldsymbol{o}'_n$ denoting the operation selectors. The architecture is specified by the sets $\{\boldsymbol{i}_n, \boldsymbol{i}'_n\}_{n=1}^{N}$ and $\{\boldsymbol{o}_n, \boldsymbol{o}'_n\}_{n=1}^{N}$, where N is the number of nodes in a cell. This formulation is easily converted to the operation assignment problem on edges (*i.e.* $\{\boldsymbol{z}_e\}$) in Sec. 2 using the outer product $\boldsymbol{i}_n \otimes \boldsymbol{o}_n + \boldsymbol{i}'_n \otimes \boldsymbol{o}'_n$, as shown in Fig. 2. We use the product of categorical distributions in the form $\prod_n p(\boldsymbol{i}_n)p(\boldsymbol{i}'_n)p(\boldsymbol{o}_n)p(\boldsymbol{o}'_n)$ to represent the distribution over architecture parameters.

4.2. Comparison with the Previous Work

The current literature on NAS often reports the final performance obtained by the best discovered cell. Unfortunately, such qualitative metric fails to capture i) the number of searches conducted before finding the best cell, ii) the performance variation resulted from different searches, iii) the effect of each model component on the final performance, and iv) the effect of post-search heuristics used for creating the best architecture. To better provide insights into our framework, we conduct extensive ablation experiments on the CIFAR-10, CIFAR-100 and ImageNet datasets. We run the search and evaluation phases end-to-end four times on each dataset and we report mean and standard deviation of the final test error as well as the best cell out of the four searches. We do not use any post-search heuristic, as our factorized cell structure always yields two-incoming edges per node in the cell. This stands in a stark contrast to DARTS [18] and P-DARTS [5] that use post-search heuristics to sparsify the discovered cell.

Here, we only consider the differentiable cross-entropy loss functions as the search objective function (*i.e.*, we do not optimize for latency). Since the direct search on ImageNet is computationally expensive, we reduce the search space on this dataset to five operations including skip connection, depthwise-separable 3×3 convolution, max pooling, dilated depthwise-separable 3×3 convolution, and depthwise-separable 3×3 convolution. Prior work on ResNets [10], DenseNets [13], as well as the recent RandWire [38] suggest that it should be possible to achieve high accuracy by using only these three operations.

Below, we discuss the different baselines summarized in Table 2. Additional details of search and evaluation can be found in Appendix A, and Appendix B respectively.

The state-of-the-art: The previous works closest to our

Table 2: Comparison against the state-of-the art methods. Different objective functions for updating architecture parameters and different gradient estimators are examined for UNAS. We run UNAS and the original publicly-available source code for DARTS [18] and P-DARTS [5] end-to-end four times with different initialization seeds. Mean±standard deviation of all four discovered architectures as well as the best architecture at the end of the evaluation phase are reported. For other techniques, the original best results are reported. The search cost is reported on CIFAR-10. UNAS with $\mathcal{L}_{\text{gen}}$ and REBAR significantly outperforms gradient-based methods on all three datasets.

	Objective Function	Gradient Estimator	CIFAR-10 mean	CIFAR-10 best	CIFAR-100 mean	CIFAR-100 best	ImageNet mean	ImageNet best	Search Cost (GPU days)
UNAS	$\mathcal{L}_{\text{val}}$	Gumbel-Soft.	$2.79_{\pm 0.10}$	2.68	$17.11_{\pm 0.38}$	16.80	$26.06_{\pm 0.51}$	25.41	-
	$\mathcal{L}_{\text{gen}}$	Gumbel-Soft.	$2.81_{\pm 0.01}$	2.74	$16.98_{\pm 0.34}$	16.59	$24.64_{\pm 0.13}$	**24.46**	-
	$\mathcal{L}_{\text{gen}}$	REBAR	**2.65**$_{\pm 0.07}$	**2.53**	**16.72**$_{\pm 0.76}$	**15.79**	**24.60**$_{\pm 0.06}$	24.49	4.3
Gradient	DARTS [18]		$3.03_{\pm 0.16}$	2.80	$27.83_{\pm 8.47}$	20.49	$25.27_{\pm 0.06}$	25.20	4
	P-DARTS [5]		$2.91_{\pm 0.14}$	2.75	$18.09_{\pm 0.49}$	17.36	$24.98_{\pm 0.44}$	24.49	0.3
	SNAS [39]		-	2.85	-	-	-	27.3	1.5
Reinforce	NASNet-A [43]		-	2.65	-	-	-	26.0	2000
	BlockQNN [41]		-	3.54	-	18.06	-	-	96
	ENAS [24]		-	2.89	-	-	-	-	0.45
Evolution	AmoebaNet-A [26]		-	3.12	-	-	-	25.5	3150
	AmoebaNet-B [26]		-	2.55	-	-	-	26.0	3150
	AmoebaNet-C [26]		-	-	-	-	-	24.3	3150
	Hierarchical. Evolution [17]		-	3.75	-	-	-	-	300

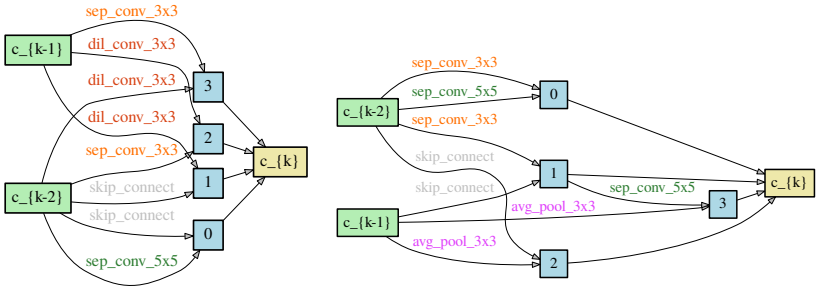

(a) Normal Cell (b) Reduce Cell

Figure 3: The best performing cell on CIFAR-10.

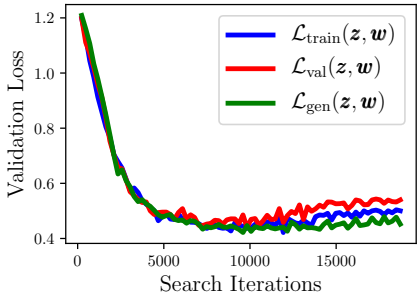

Figure 4: Validation loss in search.

work DARTS [18], P-DARTS [5] and SNAS [39] have unfortunately reported the performance for the best discovered cell. Since DARTS and P-DARTS implementations are publicly available, for a fair comparison, we run their original source code end-to-end four times similar to our model with different random initialization seeds using hyperparameters and commands released by the authors on CIFAR-10 and CIFAR-100.[4] For the ImageNet datasets, we transfer the discovered cells from CIFAR-10 to this dataset as described in [18, 5]. The implementation of SNAS [39] is not publicly available. So, we compare against this work using the original published results. Finally, in order to better contextualize our results, we compare UNAS against previous methods that use reinforcement learning or evolutionary search. On ImageNet, we only consider the mobile-setting (FLOPS $<$ 600M) which is often used to compare NAS models.

UNAS baselines: We also explore the different variants of UNAS. The objective function column in Table 2 represents the loss function used during search for updating $\boldsymbol{\phi}$. Here, $\mathcal{L}_{\text{val}}$ (Eq. 6) and $\mathcal{L}_{\text{gen}}$ (Eq. 8) are considered. The gradient estimator column represents the gradient estimator used for updating $\boldsymbol{\phi}$ during search. We examine Gumbel-Softmax and REBAR (Eq. 3).

Observations: From the first group of Table 2, we ob-

[4] We exactly followed the hyperparameters and commands using the search/eval code provided by the authors. We only set the initialization seed to a number in $\{0, 1, 2, 3\}$.

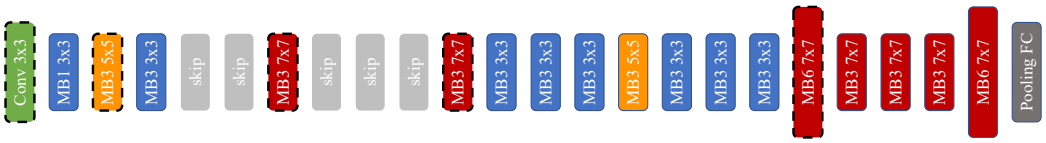

(a) Cell discovered by UNAS in the ProxylessNAS [4] search space with 9.8 ms GPU latency and 24.7% top-1 error

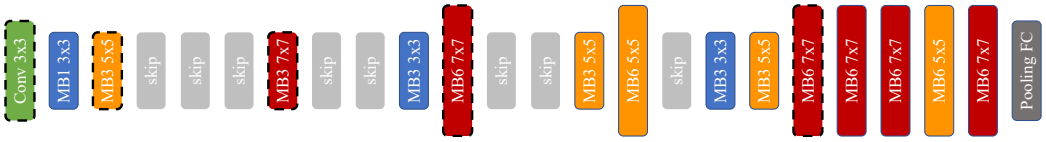

(b) Cell discovered by ProxylessNAS [4] with 10.1 ms GPU latency and 24.9% top-1 error

Figure 5: Visualization of the network discovered by UNAS in the ProxylessNAS [4] search space. MBe $K \times K$ denotes a mobile inverted residual block with expansion ratio e and kernel size K. UNAS, in contrast to ProxylessNAS, keeps the cells at the deeper layers (on the right side) computationally inexpensive by using a small expansion ratio, enabling more MBConv layers in the shallower layers. Although UNAS architecture is deeper, it has a lower latency with the same network width.

serve that architecture search with the generalization loss yields a better model often in terms of both average performance and best results. The improvement obtained by the generalization is especially profound in ImageNet as this loss function improves $\mathcal{L}_{\text{val}}$ by 1.4% in average. We can also see that our REBAR gradient estimator often improves the results across all datasets. From the second group of Table 2, we observe that our UNAS framework with REBAR estimator and the generalization loss significantly outperforms DARTS, P-DARTS, and SNAS on all three datasets.[5] Interestingly, our full model ($\mathcal{L}_{\text{gen}}$ with REBAR) exhibits a low variance on CIFAR-10 and ImageNet, showing the robustness of the framework in discovering high-performing architectures. Finally, comparing UNAS against the evolutionary and RL-based models shows that UNAS outperforms these models. The only exception is AmoebaNet-C [26] on ImageNet. However, note that this method requires 700x more GPU time to search.

Why does the generalization loss help in search: Recall that in differentiable architecture search, we often update the architecture distribution parameters $\boldsymbol{\phi}$ using suboptimal $\boldsymbol{w}$. We hypothesize that even if validation loss is used in search due to the suboptimality of $\boldsymbol{w}$, the architecture is not discovered using the the true generalization of the network to unseen examples. To illustrate this, the validation loss during architecture search visualized in Fig. 4 for different loss functions. We observe that even using the validation loss for updating architecture parameters does not prevent the network from overfitting.

One question is whether our generalization loss is required in the case of the original RL-based NAS [24, 42, 43], which updates architecture parameters using $\boldsymbol{w}$ closer to optimality. To answer this, we also examine with ENAS [24]-like training where network parameters $\boldsymbol{w}$ are updated for half epoch in every $\boldsymbol{\phi}$ update (*i.e.*, the network parameters $\boldsymbol{w}$ are brought closer to the optimum). In this case, the architectures found by generalization loss in average obtains test error 2.92% on CIFAR-10 compared with the validation loss based search that achieves 3.12%. This provides another evidence that architecture search can potentially benefit from considering generalization, opening up new research directions in NAS.

Cell visualization: The best cell discovered on the CIFAR-10 dataset is visualized in Fig. 3. See Appendix C for the visualization of best cells on other datasets.

More comparisons: We provide in-depth comparisons against the state-of-the-art techniques with more detailed information including the number of parameters, search cost, and the number of floating point operations in Appendix. D.

5. Experiments on Latency-based Search

In this section, we examine our proposed framework for searching architectures with low latency directly on the ImageNet dataset. Unfortunately, the DARTS search space results in high-latency networks due to the parallel branches and concatenation in each cell. So, here, we change the building blocks of our search space to the mobile inverted bottleneck convolution (MBConv) [28] that has been used in ProxylessNAS [4] and FBNet [36] for discovering low-latency networks. For this section, we closely follow the search space introduced in ProxylessNAS [4] for ImageNet in which a 21-layer network with seven choices of operations in each layer is searched. Specifically, for each layer, an MBConv is selected among various kernel sizes $\{3, 5, 7\}$ and expansion ratios $\{3, 6\}$. To allow layer removal, an additional skip-connection is used in ProxylessNAS yielding seven operations per layer. For search and evaluation

[5] The significance test between UNAS and any other approach passes on all the datasets with p-value < 0.05, except on ImageNet between UNAS and P-DARTS which yields p-value $= 0.18$.

Table 3: Latency-based architecture search. Models are sorted based on their top-1 error. For a better illustration, Fig. 6 compares the models visually.

Architecture	Val Error top-1	Val Error top-5	Latency (ms)
EfficientNet B0 [32]	23.7	6.8	14.5
MobileNetV3 Large [11]	24.7	7.6	11.0
MnasNet-A1 [31]	24.8	7.5	10.9
Single-Path NAS [29]	25.0	7.8	10.2
FBNet-C [36]	25.1	-	11.5
MobileNetV2 1.4x [28]	25.3	7.5	13.0
MnasNet-B1 [31]	25.5	-	9.4
ShuffleNet V2 2x [19]	26.3	-	9.16
MobileNetV2 1x [28]	28.0	9.	9.2
ProxylessNAS-GPU [4]	24.9	7.5	10.1
UNAS	24.7	7.6	9.8

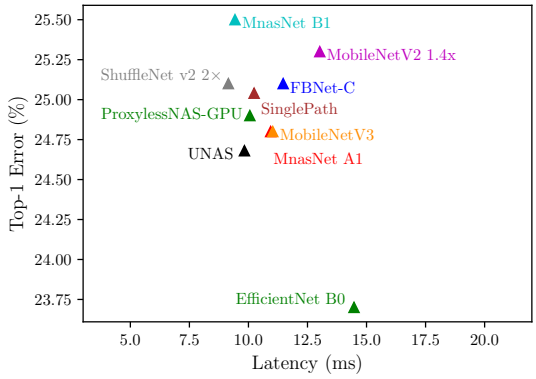

Figure 6: Latency-based architecture search. We seek architectures that are in the bottom-left side of the error vs. latency axes. UNAS discovers an architecture that is more accurate and has a low latency compared to the current state-of-the-art architectures based on MobileNetV2.

we closely follow the settings used in ProxylessNAS (see Appendix E for details).

For gradient estimation of the latency loss in UNAS, we use Eq. 5 with a simple linear function as the surrogate function, *i.e.*, $g(\boldsymbol{z}) = \sum_{i,j} l_{i,j} z_{i,j}$ where $z_{i,j} \in \{0, 1\}$ is a binary scalar indicating if operation i is used in layer j and $l_{i,j}$ is the approximate latency associated with the operation. Similar to ProxylessNAS, we randomly generate 10K network samples before search and we train the parameters of g (*i.e.*, all $l_{i,j}$) by minimizing an L_2 regression loss.

We search for architecture on V100 GPUs, as it allows us to measure the true latency on the device during search. These GPUs were also used in ProxylessNAS [4] which enables us to have a fair comparison against this method. We measure latency using a batch size of 32 images. We empirically observed that smaller batch sizes under-utilize GPUs, resulting in inaccurate latency measurements.

Table 3 and Fig. 6 report the latency and validation set error on ImageNet for our model in comparison to recent hardware-aware NAS frameworks that operate in a similar search space (*i.e.*, MobileNetV2 [28]) and have similar latency (∼10 ms on V100 GPUs). We can see that UNAS finds an architecture that is slightly faster but more accurate than the ProxylessNAS-GPU [4] architecture that uses exactly the same search space and the same target device. EfficientNet B0 [32] is the only architecture that is more accurate than UNAS but it is also 48% slower on the GPU. Although EfficientNet B0 has a low number of mathematical operations, it is not so efficient on TPU/GPU due to the heavy usage of depth-wise separable convolutions [9]. The architectures that are faster than UNAS including ShuffleNet v2 [19], MnasNet B1 [31] and MobileNetV2 1.0x [28] are also less accurate.[6]

In Fig. 5, the architecture discovered by UNAS is compared against ProxylessNAS-GPU that has been discovered for the same type of GPUs. Interestingly, UNAS discovers an architecture that is deeper, *i.e.*, it has 3 more MBConv layers. But, it also faster and more accurate than the architecture discovered by ProxylessNAS.

6. Conclusions

In this paper, we presented UNAS that unifies differentiable and RL-based NAS. Our proposed framework uses the gradient of the objective function for search without introducing any bias due to continuous relaxation. In contrast to previous DNAS methods, UNAS search objective is not limited to differentiable loss functions as it can also search using non-differentiable loss functions. We also introduced a new objective function for search based on the generalization gap and we showed that it outperforms previously proposed training or validation loss functions.

In extensive experiments in both DARTS [18] and ProxylessNAS [4] search spaces, we showed that UNAS finds architectures that 1) are more accurate on CIFAR-10, CIFAR-100, and ImageNet and 2) are more efficient to run on GPUs. We will make our implementation publicly available to facilitate the research in this area.

References

[1] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. In *ICLR*, 2017. 1

[6] All models are examined in PyTorch. For MobileNetV2 and ShuffleNet V2 the official PyTorch implementations are used. For ProxylessNAS-GPU, the original code provided in [4] is used. Other networks implementations are obtained from EfficientNets repo [34] which are optimized for PyTorch.

[2] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *ICML*, 2018. 2

[3] Han Cai, Tianyao Chen, Weinan Zhang, Yong Yu, and Jun Wang. Efficient architecture search by network transformation. In *AAAI*, 2018. 1

[4] Han Cai, Ligeng Zhu, and Song Han. ProxylessNAS: Direct neural architecture search on target task and hardware. In *ICLR*, 2019. 1, 2, 4, 7, 8, 14

[5] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *ICCV*, 2019. 2, 5, 6, 14

[6] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *arXiv preprint arXiv:1808.05377*, 2018. 1

[7] Priya Goyal, Piotr Dollár, Ross Girshick, Pieter Noordhuis, Lukasz Wesolowski, Aapo Kyrola, Andrew Tulloch, Yangqing Jia, and Kaiming He. Accurate, large minibatch sgd: Training imagenet in 1 hour. *arXiv preprint arXiv:1706.02677*, 2017. 12

[8] Will Grathwohl, Dami Choi, Yuhuai Wu, Geoff Roeder, and David Duvenaud. Backpropagation through the void: Optimizing control variates for black-box gradient estimation. In *ICLR*, 2018. 4

[9] Suyog Gupta and Mingxing Tan. EfficientNet-EdgeTPU. `https://ai.googleblog.com/2019/08/efficientnet-edgetpu-creating.html`, Nov 2019. 8

[10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 2, 5

[11] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for MobileNetV3. In *ICCV*, 2019. 2, 8

[12] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017. 11, 14

[13] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *CVPR*, 2017. 5

[14] Eric Jang, Shixiang Gu, and Ben Poole. Categorical reparameterization with Gumbel-Softmax. In *ICLR*, 2017. 3

[15] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. *arXiv preprint arXiv:1902.07638*, 2019. 13

[16] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive Neural Architecture Search. In *ECCV*, 2018. 2, 13, 14

[17] Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. Hierarchical representations for efficient architecture search. In *ICLR*, 2018. 1, 6, 13

[18] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable Architecture Search. In *ICLR*, 2019. 1, 2, 4, 5, 6, 8, 11, 13, 14

[19] Ningning Ma, Xiangyu Zhang, Hai-Tao Zheng, and Jian Sun. Shufflenet v2: Practical guidelines for efficient cnn architecture design. In *ECCV*, 2018. 8, 14

[20] Chris J Maddison, Andriy Mnih, and Yee Whye Teh. The concrete distribution: A continuous relaxation of discrete random variables. In *ICLR*, 2017. 3

[21] Andriy Mnih and Karol Gregor. Neural variational inference and learning in belief networks. In *ICML*, 2014. 3

[22] Art B. Owen. *Monte Carlo theory, methods and examples*. 2013. 3

[23] John Paisley, David M Blei, and Michael I Jordan. Variational bayesian inference with stochastic search. In *ICML*, 2012. 3

[24] Hieu Pham, Melody Y Guan, Barret Zoph, Quoc V Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *ICML*, 2018. 1, 2, 4, 5, 6, 7, 13

[25] Rajesh Ranganath, Sean Gerrish, and David Blei. Black box variational inference. In *Artificial Intelligence and Statistics*, 2014. 3

[26] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. *arXiv preprint arXiv:1802.01548*, 2018. 1, 6, 7, 11, 13, 14

[27] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V Le, and Alexey Kurakin. Large-scale evolution of image classifiers. In *ICML*, 2017. 1

[28] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*, 2018. 1, 7, 8, 14

[29] Dimitrios Stamoulis, Ruizhou Ding, Di Wang, Dimitrios Lymberopoulos, Bodhi Priyantha, Jie Liu, and Diana Marculescu. Single-path NAS: Designing hardware-efficient convnets in less than 4 hours. In *arXiv preprint arXiv:1904.02877*, 2019. 8

[30] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *CVPR*, 2015. 2

[31] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. Mnasnet: Platform-aware neural architecture search for mobile. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 2, 4, 8

[32] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114, 2019. 8

[33] George Tucker, Andriy Mnih, Chris J Maddison, John Lawson, and Jascha Sohl-Dickstein. REBAR: Low-variance, unbiased gradient estimates for discrete latent variable models. In *NIPS*, 2017. 3

[34] Ross Wightman. Eficientnet for pytorch. `https://github.com/rwightman/gen-efficientnet-pytorch`, Nov 2019. 8

[35] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 1992. 2

[36] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. FBnet: Hardware-aware efficient convnet design via differentiable neural architecture search. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019. 1, 2, 4, 7, 8

[37] Lingxi Xie and Alan Yuille. Genetic CNN. In *ICCV*, 2017. 1

[38] Saining Xie, Alexander Kirillov, Ross Girshick, and Kaiming He. Exploring randomly wired neural networks for image recognition. *arXiv preprint arXiv:1904.01569*, 2019. 5, 11

[39] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. SNAS: Stochastic Neural Architecture Search. In *ICLR*, 2019. 1, 2, 3, 4, 5, 6, 11, 13, 14

[40] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, and Hartwig Adam. Netadapt: Platform-aware neural network adaptation for mobile applications. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 285–300, 2018. 1, 2

[41] Zhao Zhong, Junjie Yan, Wei Wu, Jing Shao, and Cheng-Lin Liu. Practical block-wise neural network architecture generation. In *CVPR*, 2018. 6, 13

[42] Barret Zoph and Quoc V Le. Neural architecture search with reinforcement learning. In *ICLR*, 2017. 1, 2, 4, 7

[43] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *CVPR*, 2018. 1, 2, 4, 5, 6, 7, 11, 13, 14