

# Approximate Neural Architecture Search via Operation Distribution Learning

Xingchen Wan<sup>1,2</sup> Binxin Ru<sup>1,2</sup> Pedro M. Esparana<sup>1</sup> Fabio M. Carlucci<sup>1</sup>

<sup>1</sup>Huawei Noah’s Ark Lab, London, UK <sup>2</sup>Machine Learning Research Group, University of Oxford, Oxford, UK

{xwan, robin}@robots.ox.ac.uk pedro.esperanca@huawei.com fabiom.carlucci@gmail.com

## Abstract

The standard paradigm in Neural Architecture Search (NAS) is to search for a fully deterministic architecture with specific operations and connections. In this work, we instead propose to search for the optimal **operation distribution**, thus providing a stochastic and approximate solution, which can be used to sample architectures of arbitrary length. We propose and show, that given an architectural cell, its performance largely depends on the ratio of used operations, rather than any specific connection pattern in typical search spaces; that is, small changes in the ordering of the operations are often irrelevant. This intuition is orthogonal to any specific search strategy and can be applied to a diverse set of NAS algorithms. Through extensive validation on 4 data-sets and 4 NAS techniques (Bayesian optimisation, differentiable search, local search and random search), we show that the operation distribution (1) holds enough discriminating power to reliably identify a solution and (2) is significantly easier to optimise than traditional encodings, leading to large speed-ups at little to no cost in performance. Indeed, this simple intuition significantly reduces the cost of current approaches and potentially enable NAS to be used in a broader range of applications.

## 1. Introduction

Neural architecture search (NAS) is an extremely challenging problem and in the last few years, a huge number of solutions have been proposed in the literature [12, 34, 28, 55, 32]. From genetic algorithms, to Bayesian optimisation, many approaches have been evaluated, each presenting different strengths and weaknesses, with no clear winner overall. At its core, NAS is a combinatorial problem for which no algorithm can guarantee to find the global optimum in polynomial time. For example, the commonly used DARTS [28] search space, even with simplifying constraints, contains  $8^{14} \approx 4.3 \times 10^{12}$  architectures, an unreasonably large number to explore. Given the sheer size of the search space, it is highly unlikely that NAS algorithms are able to reliably find the global optimum and in fact recent research has shown

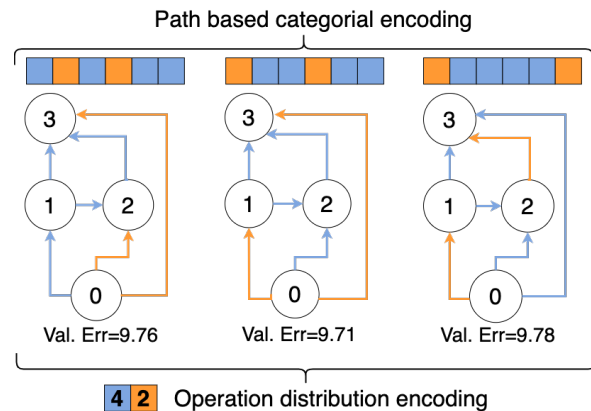


Figure 1: Three architecture cells from NAS-Bench-201 (NB201) [11] giving very similar validation performance. Each edge represents an operation: either  $\text{conv}1 \times 1$  or  $\text{conv}3 \times 3$ . From the viewpoint of existing encodings, they represent three distinct architectures which all need to be evaluated. The ANASOD framework assumes that they all belong to the same operation distribution ( $4 \times \text{conv}3 \times 3$  and  $2 \times \text{conv}1 \times 1$  in 6 operations) and does not repeatedly re-evaluate each one.

that local search algorithms can perform comparably to the state of the art [45, 31]. Furthermore, thanks to a number of benchmarks, we know that many architectures perform similarly one another [50, 11, 38, 49].

Knowing that most architectures perform similarly and that the true global optimum is very unlikely to be found, one can ask whether the current paradigm of searching for a specific architecture is the way forward. Differently from prior work, we argue that tackling NAS with an approximate algorithm that learns a distribution rather than a specific architecture allows resources to be used more efficiently. An intuitive understanding can be had from Figure 1. Indeed, the accuracy distribution for each encoding is shown to have small standard deviation in Figure 2, which motivates the optimization over encodings rather than architectures. The key issue is that comparing every single architecture is intractable and unnecessary: intractable due to the sheer size of the considered search spaces, and unnecessary because

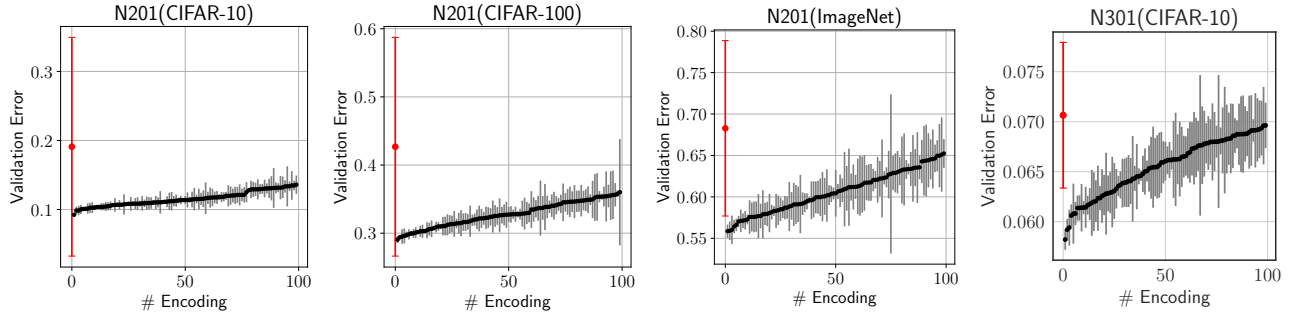


Figure 2: To what extent does the ANASOD encoding determine performance? We randomly draw 200 ANASOD encodings in 4 tasks. Within each, we draw 5 architectures *for each encoding* and show the mean  $\pm 1$  standard deviation (black and gray, respectively) of the top-100 encodings vs those of *all 1,000 sampled architectures* (red). Architectures sampled from the same encoding usually perform similarly and encodings that on average perform better also have smaller variability.

we experimentally know that small differences in the architecture have little to no effect on the final result [46, 34, 49].

Rather than searching for a specific architecture, we propose to find an approximate solution by instead searching for the optimal operation distribution, defined as the relative ratio between operations (e.g. `conv3x3`, `conv5x5`, `maxpool`, ...) in the whole architecture. Once the operation distribution is defined, we can sample from it, seamlessly generating architectures of variable length. As we experimentally show, this simple re-framing of the NAS problem enables us to significantly improve the sample efficiency of all exploitative methods. Indeed, ANASOD is orthogonal to existing NAS solutions and we experimentally show how it improves over Bayesian optimisation, local search and random search. Not only is this approach more sample-efficient, but it has less potential for overfitting, as we show by successfully transferring to new datasets.

To summarise, we propose searching for an approximate rather than exact solution to the NAS problem. This is easily applicable to most existing approaches and enables a substantial speed-up (from sample-efficiency) without a sacrifice in performance. We empirically show this both on the NAS benchmarks [11, 38] and with open-domain experiments on CIFAR-10 and CIFAR-100.

## 2. Related Work

**Architecture encodings** Several encoding schemes have been proposed to embed the directed acyclic graph of an cell-based architecture into a real-valued tensor to facilitate optimisation [43]. These encodings either represent the architecture via a flattened adjacency matrix and a list of operations [50] or characterise the architecture as a set of paths from input to output [44]. While multiple architectures may be projected to the same encoding and thus a particular encoding can also define a distribution of architectures, such performance standard deviation is rather small [43] and the search dimension remains high. On the other hand, some prior works use graph neural networks [37], graph kernels

[35] or variational autoencoders [29] to implicitly encode an architecture. These methods all require surrogate training and focus on comparing individual architectures rather than distributions of architectures.

**Stochastic methods** Recently, a number of methods based on probabilistic formulations have been proposed to overcome the memory requirements of differentiable approaches [6, 48, 5]. These works learn a distribution over architectures and, at a first glance, might appear similar to what is being proposed in this work. The differences are significant and two-fold: (1) first and foremost the output of these methods is a deterministic architecture and the distribution is learned for each possible placement of an operation—which does not reduce the search space size; and (2), their focus is strictly applied to the one-shot methods only. In other words, these works [6, 48] propose a different search strategy rather than a different encoding, and as such are orthogonal to ANASOD and could be combined with it. Another related class of works propose the concept of stochastic architecture generators [46, 34]. Specifically, they employ random graph generators to define the wiring pattern within an architecture and thus recast the problem of searching for an optimal architecture to finding the optimal generator hyper-parameters [34]. This significantly reduces the dimension of the search space and improves the search efficiency. However, these prior works focus on the graph architecture topology, while keeping the operation type fixed. On the other hand, the focus of our work is to show that, regardless of any specific cell types, the operation distribution is an efficient and informative representation.

## 3. Our Method: ANASOD

### 3.1. Overview

Searching for the optimal configuration of an architecture cell can be both extremely difficult (due to the formulation of NAS as a combinatorial optimisation problem) and ineffective (due to the common practice of using different pipelines

during search and evaluation and the imperfect correlation of performance between the two). Instead, we propose to learn the optimal *distribution of operations* within a cell. Formally, given an architecture cell  $C$  with  $N$  operation blocks  $\{n_1, \dots, n_N\}$  where each operation is chosen from a pool of  $\{o_1, \dots, o_k\}$  candidate operations (e.g. in the NB201 cell, we have  $N = 6$  and  $k = 5$ ), its normalised distributional encoding (which we term the **ANASOD** - **A**pproximate **NAS** via **O**peration **D**istribution - encoding)  $\tilde{\mathbf{p}}$  is a  $k$ -dimensional vector defined over the  $k$ -dimensional probability simplex  $\{\tilde{\mathbf{p}} \in \mathbb{R}^k; \sum_{i=1}^k \tilde{p}_i = 1, \tilde{p}_i \geq 0 \forall i = \{1, \dots, k\}\}$  where  $\tilde{p}_i = \frac{n(o_i)}{N}$  and  $n(o_i)$  is the number of blocks in the cell  $o_i$ . The unnormalised encoding is  $\mathbf{p} = N\tilde{\mathbf{p}}$ .

Because the encoding now only incorporates information about the operations in the cell, a single encoding maps to a *distribution* of architectures sharing the same distribution of operations but possibly with different wiring. A key feature of ANASOD is the easy reversibility in *decoding*  $\tilde{\mathbf{p}}$  back to a distribution of architectures defined by it  $p(\alpha|\tilde{\mathbf{p}})$ : Here we outline two possible ways to achieve so which are asymptotically equivalent when either the size of the cell becomes large, or when are sampling a large number of architectures. First, we may directly interpret the  $i$ -th element of the unnormalised encoding vector  $\mathbf{p} = N\tilde{\mathbf{p}}$  as the number of operation  $i$  in the cell  $C$ ; with the operations fixed, we may then randomly shuffle their ordering and the wiring amongst them to obtain a set of architectures  $\{\alpha_1, \alpha_2, \dots\}$ . This, however, requires  $\sum_{i=1}^k p_i = N$  and  $\mathbf{p} \in \mathbb{N}^k$ .

Correspondingly,  $\tilde{\mathbf{p}}$  is only defined on a regular grid on the  $k$ -simplex, which could be problematic if we would like to apply standard continuous optimisation on the encoding space. To overcome this, we utilise the following rounding rule proposed in [4] to “snap” an arbitrary point on the simplex to a valid, all-integer configuration: denoting the standard  $k$ -simplex on which the *unnormalised* encoding lies as  $\Delta^k$  and an arbitrary point as  $\mathbf{m}$ . Define the fractional part of  $\mathbf{m}$  as  $s(\mathbf{m}) = [m_i - \lfloor m_i \rfloor]_{i \in [1, k]}$ , we have  $g(\mathbf{m}) := \sum_{i=1}^k s(m_i) = N - \sum_{i=1}^k \lfloor m_i \rfloor \in [0, k - 1]$  as a non-negative integer. We then round  $g(\mathbf{m})$  largest elements of  $s(m_i)$  to 1 and the rest to 0 to obtain a rounded integer vector  $\mathbf{m}_r$ :

$$\mathbf{m}_r := [\lfloor m_1 \rfloor + 1, \dots, \lfloor m_{g(\mathbf{m})} \rfloor + 1, \lfloor m_{g(\mathbf{m})+1} \rfloor, \dots, \lfloor m_k \rfloor]^\top \quad (1)$$

which, as [4] prove, is the closest valid configuration on  $\Delta^k$  to  $\mathbf{m}$  with respect to all  $\ell^p$ -norms for  $p \geq 1$ .

The second decoding method uses a probabilistic interpretation, and does not constrain  $\tilde{\mathbf{p}}$  to a finite set of points. At each operation block  $n_j \forall j \in [1, N]$ , we draw a sample from the categorical distribution parameterised by  $\tilde{\mathbf{p}}$ :

$$n_j \sim \text{Cat}(\tilde{p}_1, \dots, \tilde{p}_k) \forall j \in [1, N] \quad (2)$$

While the cells sampled via this decoding method only have

Table 1: Overall standard deviation (SD), median SD of *different* architectures sampled from the *same encoding* and the median SD of the *same* architecture trained with *different seeds*. All SD are w.r.t validation error in percentage.

Benchmark	NB201		NB301	
Task	C10	C100	ImageNet16	C10
Overall	9.5	14	10	0.77
Median (same encoding)				
- All encodings	1.2	2.4	2.7	0.25
- Top 50%-performing	0.84	1.4	1.9	0.22
Median (different seeds)	0.19	0.35	0.36	0.17

operation distribution to be  $\tilde{\mathbf{p}}$  on expectation, in this case  $\tilde{\mathbf{p}}$  is not constrained to the grid like before, enabling gradient-based continuous optimisation to be used effectively and as we show in Sec. 3.2, is key for applying ANASOD in the differentiable NAS setting.

For the proposed encoding to provide a reasonable approximation of the NAS solution, it is imperative for the architectures sampled from the same encoding to perform similarly, which we empirically show to be the case in the popular cell-based search spaces represented by the NB201 [11] and NAS-Bench-301 (NB301) [38] datasets. With reference to Figure 2 and Table 1, in each of the 4 tasks considered, we draw 200 random encoding  $\tilde{\mathbf{p}}$  and within each encoding, we draw 5 unique architectures sampled from that encoding and query the validation performance of all 1,000 architectures. To investigate the extent to which ANASOD encoding determines the validation performance, we compute the overall SD across all sampled architectures which acts as an unbiased estimate of the performance variability over the entire search space. We also consider the median SD amongst architectures sharing the same encoding, with particular attention to the better-performing encodings in which NAS is particularly interested. Additionally, we also report the median SD of the same architectures trained with different seeds, representing the amount of irreducible aleatory uncertainty due to evaluation noise. It is evident that fixing the ANASOD encoding massively reduces the validation error variability and hence is highly predictive of the performance, and comparatively, better encodings have lower variability.

Having established the *validity* of ANASOD, we now argue there are 3 key *advantages* of ANASOD. First, it requires no learning and massively compresses the search space, yet remains predictive enough about the performance of the architectures. While as discussed, the original architecture space explodes exponentially as  $k^N$ , the number of unique ANASOD encoding is simply  $\binom{N+k-1}{k-1}$ . For instance, in the DARTS-like cell with  $k = 8$  and  $N = 14$ , the number of architectures is  $8^{14} \approx 4.3 \times 10^{12}$ , but the number of ANASOD encodings is only  $\binom{21}{7} \approx 1.2 \times 10^5$ . This enables us to reduce the difficult combinatorial NAS optimisation into a much more tractable, easier-to-solve problem and thus to

achieve superior sample efficiency – this is a result of ANASOD being an approximate encoding; exact encoding such as the adjacency or path encoding, while representing the architecture in a way that is amenable for predictor learning, usually do not reduce the optimisation difficulty *per se*. Second, as mentioned, we may generatively sample architectures from the encoding easily. This allows us to directly search for an optimal encoding  $\tilde{\mathbf{p}}^*$  first and to map back to the architectures thereafter for effective (approximate) global optimisation over the entire search space; this is in contrast with many previous approaches which, although using advanced encoding schemes, still rely on random sampling and/or mutation algorithm in the *architecture* space for architecture optimisation [35, 37, 44]. Lastly, unlike most existing encodings, ANASOD is defined on a well-studied simplex space with well-defined distance metrics and is of a modest dimensionality; all of which are key for effective deployment of model-based methods such as Gaussian Process (GP)-based Bayesian optimisers, which are often difficult to apply NAS due to the high-dimensional search spaces and/or the lack of natural distance metrics.

### 3.2. Applying ANASOD

In this section, we give some examples that orthogonally combine various methods with ANASOD to achieve improved efficiency, performance, or both. It is worth emphasising that the methods proposed here are exemplary, and in no way exhaust the applicability of the proposed encoding.

**Random Search (RS)** While simplistic in implementation, RS is shown to be a remarkably strong baseline in NAS where even the state-of-the-art optimisers often fail to significantly outperform RS [50, 11, 49]. Implementing random sampling in the proposed encoding space is trivial: to sample uniformly from the entire search space, we simply draw encoding from the uniform Dirichlet distribution  $\tilde{\mathbf{p}}_i \sim \text{Dir}(1, \dots, 1)$ . More importantly, by virtue of the unique property that our encoding lies on a probability simplex on which well-principled distributions can be easily defined, we also propose a simple modification that mitigates the key shortfall of RS, the fact that it is purely exploratory and does not exploit the information gathered as the optimisation progresses: at iteration  $t$ , instead of randomly drawing the next encoding  $\tilde{\mathbf{p}}_{t+1}$  from the uniform distribution, we instead draw randomly the next sample from the following distribution whose mode is  $\tilde{\mathbf{p}}^*$ , the best-performing encoding seen so far:

$$\tilde{\mathbf{p}}_{t+1} \sim \text{Dir}(\alpha_1, \dots, \alpha_k) \text{ where } \alpha_i = k\beta_t\tilde{p}_i^* + 1 \forall i, \quad (3)$$

where  $\beta_t$  is a parameter similar to the temperature parameter in simulated annealing that monotonously increases from 0 (the uniform Dirichlet) to some positive  $\beta_T$  at termination of

the search that controls the scale of the Dirichlet distribution. Essentially, in this modified RS (which we term *biased* RS), we add exploitation by progressively biasing the generating distribution of the encoding to the locality around the best encoding seen so far with the hope of finding even better ones, with  $\beta_t$  controlling the amount of exploitation-exploration trade-off. This modification is simple, poses almost no additional computing cost (apart from the need to track  $\tilde{\mathbf{p}}^*$  as the search progresses), and is orthogonal to other modifications of RS (e.g. RS with successive halving [23] and RS with parameter sharing [25]), but as we will show in Sec. 4 it significantly improves on the vanilla RS on a wide range of search spaces to a level that often matches or even exceeds more sophisticated methods, and thus could serve as a strong baseline for the NAS community.

**Differentiable NAS** Another popular paradigm in NAS is that of *Differentiable Neural Architecture Search* (DNAS) pioneered by [28], which, generally speaking, utilises *continuous relaxation* of the discrete architecture space (followed by *discretisation* at the end of the pipeline) with weight-sharing supernet so that standard first-order optimisers commonly used in deep learning might be applied in this context. Differentiable NAS finds reasonable solutions, and often offers a massive speedup compared to earlier query-based methods, but it is also subjected various criticisms, such as the so-called *rank disorder* between the found cells during search and evaluation and the unstable training with a propensity to collapse into cells with parameterless operations that corresponds to poorly-performing architectures [49, 54] In this section, we describe a prototypical application of ANASOD in the context of DNAS, and shows that it performs competitively and leads to further saving in computation over the already efficient DNAS pipeline while remaining much simpler than many of the existing methods.

To apply ANASOD in the differentiable setting, we first notice that the stochastic interpretation of the ANASOD encoding (described in Section 3.1) as the parameters governing the categorical distribution at *each* operation block  $j \in [1, N]$  of a cell. As described, this relaxes the constraint that ANASOD encoding is only defined on a regular grid (the encoding only needs to obey  $\sum_i \tilde{p}_i = 1$ ). With this, we might compute the derivative of the supernet validation loss  $\mathcal{L}_{\text{val}}(\tilde{\mathbf{p}}, w)$  w.r.t the *encoding*,  $\nabla_{\tilde{\mathbf{p}}}(\mathcal{L}_{\text{val}}(\tilde{\mathbf{p}}, w))$  where  $w$  is the supernet parameter weights. We formulate the procedure as a bi-level optimisation problem like [28], and include an algorithmic description in Algorithm 1.

While we inherit some key elements from the existing DNAS algorithms such as the idea of bi-level optimisation and most hyperparameter choices (e.g. the learning rate of the architecture optimiser and the network architecture), there are two key differences: first, inspired by our fundamental observation that often determining the operation mix

---

**Algorithm 1** ANASOD-DNAS. Key differences from existing DNAS algorithms marked **blue**.

- 1: Create a mixed operation  $\tilde{\sigma}^j \sim \text{Cat}(\tilde{\mathbf{p}})$  for each operation block. **Note that the parameters of the categorical distribution are shared.**
  - 2: **while** not converged **do**
  - 3:   Update the **encoding**  $\tilde{\mathbf{p}}$  by descending  $\nabla_{\tilde{\mathbf{p}}} (\mathcal{L}_{\text{val}}(\tilde{\mathbf{p}}, w))$ , and keep  $w$  constant
  - 4:   **Enforce the simplex constraint**  $\tilde{\mathbf{p}} \leftarrow \frac{\tilde{\mathbf{p}}}{\sum_i \tilde{p}_i}$  (i.e. mirror descent)
  - 5:   Update the supernet weights  $w$  by descending  $\nabla_w \mathcal{L}_{\text{train}}(\tilde{\mathbf{p}}, w)$ , and keep  $\tilde{\mathbf{p}}$  constant.
  - 6: **end while**
  - 7: **Sample cells from the optimised encoding**  $\alpha \sim p(\alpha|\tilde{\mathbf{p}}^*)$  **and stack them into a final neural architecture.**
- 

is already sufficiently predictive of the architecture performance, we use the  $k$ -dimensional ANASOD encoding of the cell as the parameters of the single categorical distribution that governs the all edges in the cell; on the other hand, existing methods uses a  $k$ -dimensional vector at each edge, leading to a matrix parameterisation of the cell of dimensionality  $N \times k$ ; as we will show later, we argue that our simpler representation implicitly regularises the architecture learning and avoids the catastrophic collapse of many DNAS methods due to over-fitting; at the same time, the reduced dimensionality allows fast convergence as we will demonstrate. Second, instead of taking  $\arg \max$  of the categorical distribution parameters on each edge at the end of the search, which is hypothesised as a key reason leading to rank disorder (as the discretised architecture might differ a lot from the continuously relaxed architecture in the parameter space [49, 51, 54]), we simply sample architecture cells from the ANASOD-encoding at the end to build a final network.

**Local Search (LS)** At the other end of the exploration-exploitation continuum, LS is purely exploitative, and has recently been shown to be an extremely powerful optimiser, achieving state-of-the-art performance in small search spaces but fails to perform on larger ones [45]. We show that combining ANASOD with LS universally further improves performance on LS but particularly in the latter case, addressing the previous failure mode of LS. The readers are referred to App. B for a detailed discussion.

**Sequential Model-based Optimisation (SMBO)** SMBO algorithms, in particular Bayesian Optimisation (BO), have seen great success in hyperparameter optimisation due to their exploitation-exploration balancing, sample efficiency, and principled treatment of and robustness against noises that inevitably arise in real life. Nonetheless, their application in NAS has been somewhat constrained, partly due to the discrete search space, extreme dimensionality of the existing encoding and the common lack of well-defined distances in such space (the presence of which is crucial for kernel-based methods, such as GP). As discussed, ANASOD addresses all the three aforementioned issues and enables GP-based BO to

---

**Algorithm 2** ANASOD-BO. Key differences from conventional BO are marked **blue**.

- 1: **Input:** Objective function (default: validation error)  $y$ , number of initialising random samples  $n_{\text{init}}$
  - 2: Initialise the *encoding generating distribution* to the uniform Dirichlet distribution  $p(\tilde{\mathbf{p}}) = \text{Dir}(1, \dots, 1)$
  - 3: Sample  $n_{\text{init}}$  random encodings  $\tilde{\mathbf{p}}_{[1:n_{\text{init}}]} \sim p(\tilde{\mathbf{p}})$  and evaluate to obtain  $y(\tilde{\mathbf{p}})$  to initialise the surrogate GP.
  - 4: **for**  $i=n_{\text{init}}, \dots, T$  **do**
  - 5:   Sample a pool of  $B$  candidate encodings from  $p(\tilde{\mathbf{p}})$
  - 6:   Select the next query point(s) by identifying the encoding that maximises the acquisition function  $\tilde{\mathbf{p}}_i = \arg \max (\text{acq}(\tilde{\mathbf{p}}))$ .
  - 7:   **Evaluate a single architecture**  $\alpha_i$  **from the encoding**  $\tilde{\mathbf{p}}_i$  **to approximate the performance of all architectures parameterised by**  $\tilde{\mathbf{p}}_i$ .
  - 8:   Augment the surrogate GP with new encoding-observation pair(s)  $\mathcal{D}_i \leftarrow \mathcal{D}_{i-1} \cup \{\tilde{\mathbf{p}}_i, y(\tilde{\mathbf{p}}_i)\}$  and optimise the GP hyperparameters via the marginal log-likelihood maximisation.
  - 9:   Update the **encoding generating distribution**  $p(\tilde{\mathbf{p}})$ .
  - 10: **end for**
- 

be effectively used. ANASOD-BO is described in Algorithm 2. Nonetheless, we emphasise that ANASOD-BO is not restricted to the GP surrogate; as we show in Sec. 4, it may also be used with Neural Ensemble Predictor (NEP) proposed in [44] to deliver further performance improvements.

For easy parallelism which is key in NAS, we opt for a sampling-based optimisation of the acquisition function which allows trivial modifications to be able to recommend batches of points simultaneously, although we are nonetheless compatible with, for example, the common gradient-based optimiser. On this, in addition to the usual BO components, inspired by the recent *trust region*-based innovations in BO literature [13, 39] we also adopt dynamic adjustment of the *encoding generating distribution* (Lines 2 and 9 in Algorithm 2). Specifically, similar to ANASOD-RS, we use the knowledge gained previously (i.e. the location of the best encoding so far  $\tilde{\mathbf{p}}_t^*$  as a prior to generate samples for the next iteration by biasing the Dirichlet distribution. However, instead of using a simple annealing schedule for the  $\beta$  parameter, here we adopt a probabilistic trust-region setting, where we halve  $\beta$  (i.e. reducing the concentration of the Dirichlet density around  $\tilde{\mathbf{p}}^*$  to allow candidate encodings to be drawn from a larger region) upon successive successes in reducing the best function value seen so far  $y_t^*$  up to 0.<sup>1</sup> Conversely, we double  $\beta$ , which is equivalent to increasing the concentration of the Dirichlet p.d.f around the mode and thus to shrink the trust region size upon successive failures, up to a certain  $\beta_{\text{max}}$ .

Since we conduct BO on the encoding space directly, at termination ANASOD-BO returns the optimised encoding  $\tilde{\mathbf{p}}^*$  instead of a single architecture  $\alpha^*$ . This may be of independent benefits: for example, the ANASOD encoding, which has much fewer parameters than an exact architecture cell and its exact encoding thereof, should be much less prone to

---

<sup>1</sup>In [13], the trust regions are set as “hard” hyper-rectangular bounding boxes around  $\tilde{\mathbf{p}}_t^*$ . Here, we place a “soft”, high-probability trust region by adjusting  $\beta$  to constrain most of the probability mass within the locality.

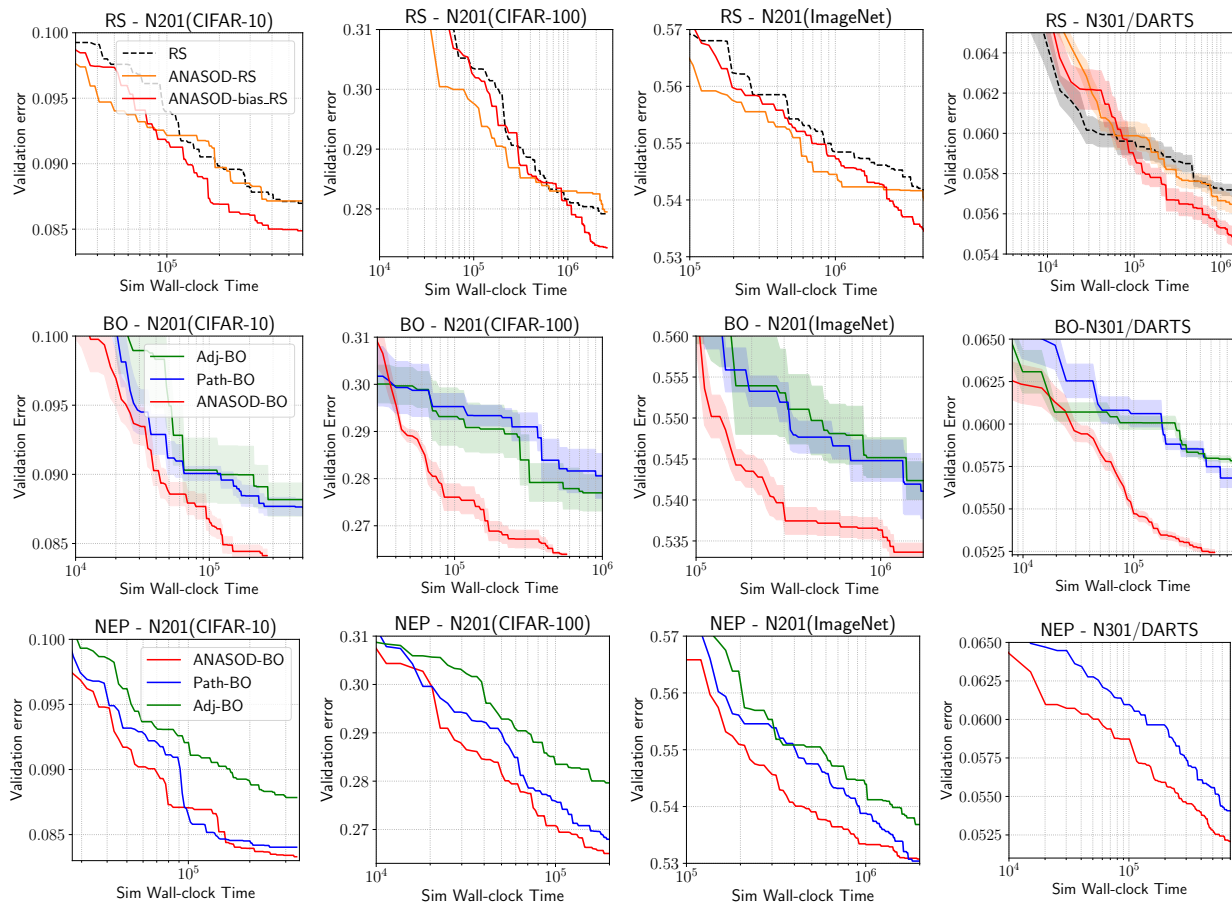


Figure 3: Performance of various methods on NB201 and NB301 with and without the ANASOD encoding: **(Top row)** random search, RS; **(Middle row)** GP-BO; **(Bottom row)** NEP-BO. Note the x-axis which shows the (simulated) GPU-seconds is in log-scale. For RS we set a budget of maximum 300 architecture queries; for BO, we set a more stringent budget of 150 queries for ANASOD-BO but allow the baselines to run for longer (to observe the amount of speedup). Adj-BO and Path-BO are the variants of BO that are otherwise identical to ANASOD-BO as outlined in Algorithm 2, but with the ANASOD encoding replaced by the adjacency [55, 50] and path [44, 42] encoding, respectively. Lines and shades denote mean  $\pm 1$  standard error, across 10 different random trials.

overfitting to a singular data-set which is cited as a problem many NAS methods have [26]. Furthermore, since multiple architectures are sampled from a single encoding, we might also obtain an *ensemble* of architectures which is recently shown to perform very competitively both in terms of accuracy and uncertainty quantification [2, 53]. We defer a thorough investigation to a future work.

## 4. Experiments

### 4.1. Experiments on the NAS-Bench Datasets

We first conduct optimisation experiments on the NB201 and NB301 datasets [11, 38]. On NB201, we consider all three tasks: CIFAR-10, CIFAR-100 and ImageNet16. NB301 is a surrogate-based benchmark that is trained on CIFAR-10 only, but it contains as many distinct architectures as the DARTS search space and is thus much more challenging. We do not conduct experiment on the earlier NAS-Bench-101

(NB101) [11], as similarly to NB301 it is also only trained on CIFAR-10, but its search space ( $\sim 4 \times 10^5$  architectures) is much smaller.

**Comparison with baseline strategies** To enable direct and fair comparison of algorithms within each class, we first compare RS, GP-BO and NEP-BO against their ANASOD counterpart, as described in Section 3.2.

In the top row of Figure 3, we observe that ANASOD-RS performs on par with RS—this is unsurprising, as both methods sample uniformly from the same, entire search space which is independent of the specific encoding chosen to represent it. This result also acts as a sanity check that there is no noticeable bias in sampling in the ANASOD space as contrasted to the original architecture space. On the other hand, ANASOD-RS with bias, which randomly samples from a non-uniform Dirichlet distribution, can be seen to significantly outperform them both as the optimisation progresses,

Table 2: Performance on NAS-Bench datasets. Unless otherwise specified, we report mean  $\pm 1$  standard error of the validation error across 10 random trials. For fair comparison, in RS and LS experiments, the numbers shown denote the best validation error seen after **300** architecture queries; in BO experiments, we show the best validation error seen for each method after **150** architecture queries (which is the budget we set for ANASOD-BO).

Benchmark Dataset	CIFAR-10	NB201 CIFAR-100	ImageNet16	NB301 CIFAR-10
RS	8.67 $\pm$ 0.03	27.91 $\pm$ 0.17	54.17 $\pm$ 0.10	5.71 $\pm$ 0.02
<b>ANASOD-RS</b>	8.71 $\pm$ 0.05	27.95 $\pm$ 0.20	53.85 $\pm$ 0.12	5.65 $\pm$ 0.04
<b>ANASOD-biasedRS</b>	<b>8.48</b> $\pm$ 0.06	<b>27.35</b> $\pm$ 0.20	<b>53.30</b> $\pm$ 0.10	<b>5.47</b> $\pm$ 0.04
RL [55]	8.91 $\pm$ 0.05	28.15 $\pm$ 0.18	54.45 $\pm$ 0.12	*
RE [33]	8.86 $\pm$ 0.05	28.40 $\pm$ 0.14	54.28 $\pm$ 0.10	5.62 $\pm$ 0.03
SMAC [18]	8.89 $\pm$ 0.05	27.80 $\pm$ 0.20	53.64 $\pm$ 0.13	5.45 $\pm$ 0.03
TPE [3]	8.57 $\pm$ 0.04	27.28 $\pm$ 0.14	53.54 $\pm$ 0.14	5.51 $\pm$ 0.02
GCNBO [37]	8.84 $\pm$ 0.01	27.93 $\pm$ 0.03	53.46 $\pm$ 0.06	5.54 $\pm$ 0.04
BANANAS [44]	8.51 $\pm$ 0.08	26.53 $\pm$ 0.02	53.41 $\pm$ 0.04	5.36 $\pm$ 0.05
NAS-BOWL [35]	8.50 $\pm$ 0.09	26.51 $\pm$ 0.00	<b>53.36</b> $\pm$ 0.04	5.31 $\pm$ 0.06
<b>ANASOD-BO</b>	<b>8.41</b> $\pm$ 0.05	<b>26.41</b> $\pm$ 0.02	<b>53.36</b> $\pm$ 0.10	<b>5.24</b> $\pm$ 0.02

\*: The original repo does not support the NB301 search space.

across all tasks, further highlighting its potential as a simple yet robust baseline in NAS research.

The middle row of Figure 3, containing the Gaussian Process based BO experiments, shows that ANASOD achieves a massive speed-up compared to both the adjacency and path encodings, the two dominant types of encoding in NAS [43]. Perhaps owing to the high-dimensional, one-hot-encoded nature of the two alternative encoding, GP-BO struggles to learn meaningful relations and performs no more competitively than RS or LS, at least up to the query budget we consider (e.g. for a simple,  $N$ -node DAG, the adjacency encoding has  $\frac{N(N-1)}{2} + N$  dimensions whereas the path encoding, without truncation, has  $\sum_{j=1}^N k^j$  dimensions [43]. In contrast, ANASOD encoding only has  $k$  dimensions corresponding to the  $k$  operation choices). On the other hand, by rigorously compressing the search space, for which non-parametric models such as GPs are better suited, ANASOD enables the full power of GP-BO to be utilised.

We show the results on NEP-BO in the bottom row of Figure 3. The findings are consistent with those on GP-BO discussed above although Path-BO performs much stronger. This is unsurprising, as NEP-BO is originally proposed to be used with the path encoding [44], and that the hyperparameters we used are tuned for this purpose. Nonetheless, on NB201 search spaces, ANASOD encoding delivers significant speedup, whereas on the more difficult NB301 search space, it is also significantly better in terms of final performance.

**Comparison with SOTA** We also compare our ANASOD-based GP-BO with a wide range of competitive baselines, including reinforcement learning (RL [55]), regularised evolution (RE [33]), Sequential Model-based Algorithm Configuration (SMAC [18]), Tree-structured Parzen Esti-

Table 3: Comparison of one-shot NAS methods on NB201. To reflect real-world applications, we search on CIFAR-10 and transfer the search result to the other datasets.

Benchmark Dataset	Search epoch	CIFAR-10	NB201 CIFAR-100	ImageNet16
<i>Optimal</i>	-	5.70	26.50	52.70
RSPS* [25]	50	12.34 $\pm$ 1.7	41.67 $\pm$ 4.3	68.86 $\pm$ 3.9
DARTS <sup>†</sup> [28]	50	45.70 $\pm$ 0.0	84.39 $\pm$ 0.0	83.68 $\pm$ 0.0
SETN* [10]	50	12.36 $\pm$ 0.0	41.95 $\pm$ 0.2	67.48 $\pm$ 0.2
ENAS <sup>†</sup> [32]	50	45.70 $\pm$ 0.0	84.97 $\pm$ 0.0	83.68 $\pm$ 0.0
GAEA-DARTS* [24]	25	8.36 $\pm$ 2.6	31.61 $\pm$ 4.5	58.41 $\pm$ 4.2
<b>ANASOD-DNAS</b>	20	7.75 $\pm$ 1.2	31.33 $\pm$ 1.7	58.00 $\pm$ 2.9

\*: Results taken from [24]; <sup>†</sup>: Results taken from [11].

imator (TPE [3]), graph convolutional network based BO (GCNBO [37]), NEP-BO with path encodings (BANANAS [44]) and Weisfeiler–Lehman kernel-based BO (NAS-BOWL [35]). Similarly, we compare ANASOD-DNAS with a number of existing DNAS algorithms in Table 3. All additional details about the experimental setup can be found in App. A.

From Table 2, it is evident that ANASOD-BO clearly outperforms competing methods, with the only BANANAS and NAS-BOWL being relatively close competitors in terms of final validation error. For NAS-BOWL, we hypothesise that this could be because it relies on local mutation in the architecture space for candidate architecture generation (at each BO iteration, the algorithm generates a pool of candidates from which BO recommends the next evaluation), which might be overly greedy and exploitative and miss the better solution especially on larger and hence more complicated search spaces. For BANANAS, the performance gain mainly comes from the NEP (which is orthogonal to ANASOD), as we show previously, combining ANASOD with NEP also yields further improvements. With respect to Table 3, despite of the much simpler formulation, ANASOD-DNAS performs competitively, outperforming all baseline methods significantly with the exception of GAEA-DARTS where they perform on par (note that we consider GAEA-DARTS with bi-level formulation since our implementation of ANASOD-DNAS is also bi-level. While the single level version of GAEA-DARTS outperforms the bi-level counterpart, this modification is orthogonal to our contribution, and hence could be similarly exploited by ANASOD-DNAS for further performance gains). Furthermore, as we show the learning curves in Fig 5 in App. C, ANASOD-DNAS almost instantly converges to a location of good solutions, allowing us to shorten the search time – this is possible because ANASOD encoding compresses and smooths the search space into a comparatively low-dimensional vector, allowing more effective gradient-based optimisation. Furthermore, in contrast to the previously mentioned performance collapse observed in some DNAS methods [11], the performance of ANASOD-DNAS does not drop as optimisation proceeds.

Table 4: Performance on CIFAR-10 in the NASNET-style search space. ANASOD-BO experiment is conducted on 4× NVIDIA Tesla V100 GPUS using 0.6 wall-clock days.

Algorithm	Val. Err	#Params(M)	Method
Random-WS [46]	2.85±0.08	4.3	RS
NASNet-A [56]	2.65	3.3	RL
LaNAS [41]	2.53±0.05	3.2	MCTS
DARTS [28]	2.76±0.09	3.3	GD
DARTS+ <sup>†</sup> [27]	2.37±0.13	4.3	GD
P-DARTS [8]	2.50	3.4	GD
DropNAS [15]	2.58±0.14	4.1	GD
DropNAS <sup>†</sup> [15]	1.88	4.1	GD
BANANAS [44]	2.64	-	BO
BOGCN [37]	2.61	3.5	BO
NAS-BOWL [35]	2.61±0.08	3.7	BO
ANASOD-BO (Mixup)	2.63	3.5	BO
ANASOD-BO (CutMix)	2.41	3.5	BO
ANASOD-BO+	1.86	3.5	BO

<sup>†</sup>: Training protocol comparable to ANASOD-BO+  
MCTS: Monte Carlo tree search; GD: Gradient descent.

## 4.2. Open Domain Experiments

For these experiments, we choose the popular NASNET-style search space [56], In our experiments, we use the setup of LANAS [41, 40] and due to the much higher computational cost, we only experiment with ANASOD GP-BO. It is worth noting that due to the distributional nature of the ANASOD encoding, differing from all previous NAS methods, we do *not* search for a single cell, but we instead search for an *encoding*: at training and evaluation time, we randomly draw concrete realisations of architecture *at each layer* by sampling the encoding according to the decoding scheme described in Section 3.1. To highlight the sample efficiency of our method, during search we only sample and train 50 architectures for 30 epochs each. At the end of the search phase, we re-train the architecture corresponding to the best encoding seen to evaluate its performance.

We report our results searched and evaluated on the CIFAR-10 dataset in Table 4, where we include two variants of ANASOD-BO, which only differ in the final evaluation: ANASOD-BO+ uses a range of additional enhancement techniques including CutMix [52], Auxiliary Towers, DropPath [22], AutoAugment [9] and trains for an extended 1,500 epochs; ANASOD-BO does not use AutoAugment and trains for the standard 600 epochs and we report ANASOD-BO using two regularisation schemes. Owing to both the shorter search budget and the smaller number of epochs per search (c.f. BANANAS [44], which trains 100 architectures for 50 epochs (~ 12 GPU days), and LANAS [41] which trains 800 architectures (≥ 150 GPU days)), ANASOD-BO is much more efficient than most other query-based NAS methods.

**Transferring to CIFAR-100** We then transfer the best encoding searched on CIFAR-10 to the CIFAR-100 data-set,

Table 5: Performance on CIFAR-100 in the NASNET-style search space. The ANASOD-BO result is transferred from CIFAR-10.

Algorithm	Val. Err	#Params(M)	Method
DARTS [28]	17.76	3.3	GD
DARTS+ <sup>†</sup> [27]	14.87	3.9	GD
P-DARTS* [8]	16.55	3.4	GD
DropNAS [15]	16.39	4.4	GD
DropNAS+ <sup>†</sup> [15]	14.10	4.4	GD
ANASOD-BO*	16.33	3.5	BO
ANASOD-BO+*	13.76	3.5	BO

\*: Transferred from the CIFAR-10 search.

<sup>†</sup>: Training protocol comparable to ANASOD-BO+.

and we present the results against a number of competitive baselines in Table 5 (note that some methods use advanced evaluation techniques similar to those used in ANASOD-BO+, and they are marked correspondingly). Remarkably, the performance obtained by us (16.33% and 13.76% for ANASOD-BO and ANASOD-BO+, respectively), outperforms some methods that are trained and searched on CIFAR-100 directly, empirically demonstrating the robustness of the ANASOD encoding across various datasets.

## 5. Conclusion

In this work, we have shown how ANASOD, an approximate solution to the NAS problem by effective compression of the search space, enables a speed-up in search with little to no deterioration in performance. Through combining with various search strategies, we show the versatility and effectiveness of our method on NAS benchmarks. We also perform *open domain* experiments CIFAR-10, showing competitive performances compared to the state-of-the-art and successfully transfer the encoding (and thus the architecture) to CIFAR-100. We believe the room for future works is ample: on one side, the operation distribution is a human-interpretable representation to communicate the *gist* of the architecture. On the other hand, by not restricting the algorithm to a single architecture, ANASOD can naturally generate diverse ensembles of related and high performing architectures. Ensembles have shown to provide benefits such as robustness [19] and future work should investigate this direction. Furthermore, we have currently only validated the effectiveness on computer vision tasks in popular cell-based NAS search spaces, and it would be similarly interesting to verify both alternative search spaces [16] and/or tasks [21, 30]. Finally, as opposed to most existing search strategies, ANASOD is orthogonal to a large number of different strategies and thus could be effectively combined with ever-improving new strategies for even further improvements.



## References

- [1] Emile Aarts, , and Jan Karel Lenstra. *Local search in combinatorial optimization*. Princeton University Press, 2003.
- [2] Randy Ardywibowo, Shahin Boluki, Xinyu Gong, Zhangyang Wang, and Xiaoning Qian. Nads: Neural architecture distribution search for uncertainty awareness. In *International Conference on Machine Learning*, pages 356–366. PMLR, 2020.
- [3] James S Bergstra, Rémi Bardenet, Yoshua Bengio, and Balázs Kégl. Algorithms for hyper-parameter optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 2546–2554, 2011.
- [4] Immanuel M Bomze, Stefan Gollowitz, and E Alper Yıldırım. Rounding on the standard simplex: Regular grids for global optimization. *Journal of Global Optimization*, 59(2-3):243–258, 2014.
- [5] Fabio Carlucci, Pedro M. Esperança, Marco Singh, Antoine Yang, Victor Gabillon, Hang Xu, Zewei Chen, and Jun Wang. MANAS: Multi-agent neural architecture search. *arxiv:1909.01051*, 2019.
- [6] Francesco Paolo Casale, Jonathan Gordon, and Nicolo Fusi. Probabilistic neural architecture search. *arXiv preprint arXiv:1902.05116*, 2019.
- [7] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- [8] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1294–1303, 2019.
- [9] Ekin D Cubuk, Barret Zoph, Dandelion Mane, Vijay Vasudevan, and Quoc V Le. Autoaugment: Learning augmentation strategies from data. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 113–123, 2019.
- [10] Xuanyi Dong and Yi Yang. One-shot neural architecture search via self-evaluated template network. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3681–3690, 2019.
- [11] Xuanyi Dong and Yi Yang. NAS-Bench-201: Extending the scope of reproducible neural architecture search. In *International Conference on Learning Representations (ICLR)*, 2020.
- [12] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *Journal of Machine Learning Research*, 20(55):1–21, 2019.
- [13] David Eriksson, Michael Pearce, Jacob R Gardner, Ryan Turner, and Matthias Poloczek. Scalable global optimization via local bayesian optimization. *arXiv preprint arXiv:1910.01739*, 2019.
- [14] Jacob R Gardner, Geoff Pleiss, David Bindel, Kilian Q Weinberger, and Andrew Gordon Wilson. Gpytorch: Blackbox matrix-matrix gaussian process inference with gpu acceleration. *arXiv preprint arXiv:1809.11165*, 2018.
- [15] Weijun Hong, Guilin Li, Weinan Zhang, Ruiming Tang, Yunhe Wang, Zhenguo Li, and Yong Yu. Dropnas: Grouped operation dropout for differentiable architecture search. In *International Joint Conference on Artificial Intelligence*, 2020.
- [16] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, et al. Searching for mobilenetv3. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1314–1324, 2019.
- [17] Juraj Hromkovič. *Algorithmics for hard problems: introduction to combinatorial optimization, randomization, approximation, and heuristics*. Springer Science & Business Media, 2013.
- [18] Frank Hutter, Holger H Hoos, and Kevin Leyton-Brown. Sequential model-based optimization for general algorithm configuration. In *Learning and Intelligent Optimization (LION)*, pages 507–523, 2011.
- [19] Sanjay Kariyappa and Moinuddin K Qureshi. Improving adversarial robustness of ensembles with diversity training. *arXiv preprint arXiv:1901.09981*, 2019.
- [20] Guolin Ke, Qi Meng, Thomas Finley, Taifeng Wang, Wei Chen, Weidong Ma, Qiwei Ye, and Tie-Yan Liu. Lightgbm: A highly efficient gradient boosting decision tree. *Advances in neural information processing systems*, 30:3146–3154, 2017.
- [21] Nikita Klyuchnikov, Ilya Trofimov, Ekaterina Artemova, Mikhail Salnikov, Maxim Fedorov, and Evgeny Burnaev. Nas-bench-nlp: neural architecture search benchmark for natural language processing. *arXiv preprint arXiv:2006.07116*, 2020.
- [22] Gustav Larsson, Michael Maire, and Gregory Shakhnarovich. Fractalnet: Ultra-deep neural networks without residuals. *arXiv preprint arXiv:1605.07648*, 2016.
- [23] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Ros-tamizadeh, and Ameet Talwalkar. Hyperband: A novel bandit-based approach to hyperparameter optimization. *arXiv:1603.06560*, 2016.
- [24] Liam Li, Mikhail Khodak, Maria-Florina Balcan, and Ameet Talwalkar. Geometry-aware gradient algorithms for neural architecture search. *arXiv preprint arXiv:2004.07802*, 2020.
- [25] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *Uncertainty in Artificial Intelligence (UAI)*, pages 367–377, 2020.
- [26] Dongze Lian, Yin Zheng, Yintao Xu, Yanxiong Lu, Leyu Lin, Peilin Zhao, Junzhou Huang, and Shenghua Gao. Towards fast adaptation of neural architectures with meta learning. In *International Conference on Learning Representations*, 2019.
- [27] Hanwen Liang, Shifeng Zhang, Jiacheng Sun, Xingqiu He, Weiran Huang, Kechen Zhuang, and Zhenguo Li. Darts+: Improved differentiable architecture search with early stopping. *arXiv preprint arXiv:1909.06035*, 2019.
- [28] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *International Conference on Learning Representations (ICLR)*, 2019.
- [29] Renqian Luo, Fei Tian, Tao Qin, Enhong Chen, and Tie-Yan Liu. Neural architecture optimization. In *International Conference on Neural Information Processing Systems (NeurIPS)*, pages 7827–7838, 2018.

- [30] Abhinav Mehrotra, Alberto Gil CP Ramos, Sourav Bhat-tacharya, Łukasz Dudziak, Ravichander Vippera, Thomas Chau, Mohamed S Abdelfattah, Samin Ishtiaq, and Nicholas Donald Lane. Nas-bench-asr: Reproducible neural architecture search for speech recognition. In *International Conference on Learning Representations*, 2020.
- [31] T Den Ottelander, Arkadiy Dushatskiy, Marco Virgolin, and Peter AN Bosman. Local search is a remarkably strong baseline for neural architecture search. *arXiv preprint arXiv:2004.08996*, 2020.
- [32] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *International Conference on Machine Learning (ICML)*, pages 4092–4101, 2018.
- [33] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *AAAI Conference on Artificial Intelligence (AAAI)*, volume 33, pages 4780–4789, 2019.
- [34] Binxin Ru, Pedro Esperanca, and Fabio Carlucci. Neural architecture generator optimization. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [35] Binxin Ru, Xingchen Wan, Xiaowen Dong, and Michael Osborne. Interpretable neural architecture search via Bayesian optimisation with Weisfeiler–Lehman Kernels. In *International Conference on Learning Representations (ICLR)*, 2021.
- [36] Michael Ruchte, Arber Zela, Julien Siems, Josif Grabocka, and Frank Hutter. Naslib: a modular and flexible neural architecture search library, 2020.
- [37] Han Shi, Renjie Pi, Hang Xu, Zhenguo Li, James T Kwok, and Tong Zhang. Bridging the gap between sample-based and one-shot neural architecture search with BONAS. *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [38] Julien Siems, Lucas Zimmer, Arber Zela, Jovita Lukasik, Margret Keuper, and Frank Hutter. Nas-bench-301 and the case for surrogate benchmarks for neural architecture search. *arXiv preprint arXiv:2008.09777*, 2020.
- [39] Xingchen Wan, Vu Nguyen, Huong Ha, Binxin Ru, Cong Lu, and Michael A Osborne. Think global and act local: Bayesian optimisation over high-dimensional categorical and mixed search spaces. *International Conference on Machine Learning*, 2021.
- [40] Linnan Wang, Rodrigo Fonseca, and Yuandong Tian. Learning search space partition for black-box optimization using monte carlo tree search. *arXiv preprint arXiv:2007.00708*, 2020.
- [41] Linnan Wang, Saining Xie, Teng Li, Rodrigo Fonseca, and Yuandong Tian. Sample-efficient neural architecture search by learning action space. *arXiv preprint arXiv:1906.06832*, 2019.
- [42] Chen Wei, Chuang Niu, Yiping Tang, and Jimin Liang. Npe-nas: Neural predictor guided evolution for neural architecture search. *arXiv preprint arXiv:2003.12857*, 2020.
- [43] Colin White, Willie Neiswanger, Sam Nolen, and Yash Savani. A study on encodings for neural architecture search. *Neural Information Processing Systems (NeurIPS)*, 33, 2020.
- [44] Colin White, Willie Neiswanger, and Yash Savani. BANANAS: Bayesian optimization with neural architectures for neural architecture search. *AAAI Conference on Artificial Intelligence*, 2021.
- [45] Colin White, Sam Nolen, and Yash Savani. Local search is state of the art for neural architecture search benchmarks. *arXiv e-prints*, pages arXiv–2005, 2020.
- [46] Saining Xie, Alexander Kirillov, Ross Girshick, and Kaiming He. Exploring randomly wired neural networks for image recognition. In *International Conference on Computer Vision (ICCV)*, pages 1284–1293, 2019.
- [47] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.
- [48] Zhicheng Yan, Xiaoliang Dai, Peizhao Zhang, Yuandong Tian, Bichen Wu, and Matt Feiszli. Fp-nas: Fast probabilistic neural architecture search. *arXiv preprint arXiv:2011.10949*, 2020.
- [49] Antoine Yang, Pedro M. Esperança, and Fabio M. Carlucci. NAS evaluation is frustratingly hard. In *International Conference on Learning Representations (ICLR)*, 2020.
- [50] Chris Ying, Aaron Klein, Eric Christiansen, Esteban Real, Kevin Murphy, and Frank Hutter. NAS-Bench-101: Towards reproducible neural architecture search. In *International Conference on Machine Learning (ICML)*, pages 7105–7114, 2019.
- [51] Kaicheng Yu, Christian Sciuto, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the search phase of neural architecture search, 2019.
- [52] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6023–6032, 2019.
- [53] Sheheryar Zaidi, Arber Zela, Thomas Elsken, Chris Holmes, Frank Hutter, and Yee Whye Teh. Neural ensemble search for performant and calibrated predictions. *arXiv preprint arXiv:2006.08573*, 2020.
- [54] Arber Zela, Thomas Elsken, Tonmoy Saikia, Yassine Marakchi, Thomas Brox, and Frank Hutter. Understanding and robustifying differentiable architecture search. *arXiv preprint arXiv:1909.09656*, 2019.
- [55] Barret Zoph and Quoc Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations (ICLR)*, 2017.
- [56] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Computer Vision and Pattern Recognition (CVPR)*, pages 8697–8710, 2018.