

# Hyper-parameter optimization with REINFORCE and Masked Attention Auto-regressive Density Estimators

Chepuri Shri Krishna  
Walmart Global Tech India  
Bengaluru, India  
Chepurishri.Krishna@walmartlabs.com

Ashish Gupta  
Walmart Global Tech India  
Bengaluru, India  
Ashish.Gupta@walmartlabs.com

Swarnim Narayan  
Walmart Global Tech India  
Bengaluru, India  
Swarnim.Narayan@walmartlabs.com

Himanshu Rai  
Walmart Global Tech India  
Bengaluru, India  
Himanshu.Rai@walmartlabs.com

Diksha Manchanda  
Walmart Global Tech India  
Bengaluru, India  
Diksha.Manchanda@walmartlabs.com

**Abstract**—Reinforcement Learning has yielded promising results for Neural Architecture Search (NAS). In this paper, we demonstrate how its performance can be improved by using a simplified Transformer block to model the policy network. The simplified Transformer uses a 2-stream attention-based mechanism to model hyper-parameter dependencies while avoiding layer normalization and position encoding. We posit that this parsimonious design balances model complexity against expressiveness, making it suitable for discovering optimal architectures in high-dimensional search spaces with limited exploration budgets. We demonstrate how the algorithm’s performance can be further improved by a) using an actor-critic style algorithm instead of plain vanilla policy gradient and b) ensembling Transformer blocks with shared parameters, each block conditioned on a different auto-regressive factorization order. Our algorithm works well as both a NAS and generic hyper-parameter optimization (HPO) algorithm: it outperformed most algorithms on NAS-Bench-101 [1], a public data-set for benchmarking NAS algorithms. In particular, it outperformed RL based methods that use alternate architectures to model the policy network, underlining the value of using attention-based networks in this setting. As a generic HPO algorithm, it outperformed Random Search in discovering more accurate multi-layer perceptron model architectures across 2 regression tasks. We have adhered to guidelines listed in Lindauer and Hutter [2] while designing experiments and reporting results.

**Index Terms**—Neural Architecture Search, Hyperparameter optimization, Transformer, Reinforcement Learning

## I. INTRODUCTION

Hyper-parameter optimization (HPO) is a key component of the ML model development cycle. Recent work has shown that older deep learning architectures such as the LSTM [3] and original GAN [4] outperformed contemporary architectures on benchmark tasks after being extensively tuned for regularization and other hyper-parameters [5,6].

HPO can be framed as the optimization of an unknown, possibly stochastic, objective function mapping from the hyper-parameter search space to a real valued scalar, the ML model’s

accuracy or any other performance metric on the validation data-set. The search-space can extend beyond algorithm or architecture specific elements to encompass the space of data pre-processing and data-augmentation techniques, feature selections, as well as choice of algorithms. This is sometimes referred to as the CASH (Combined Algorithm Search and Hyperparameter tuning) problem [7].

Neural Architecture Search (NAS) is a special type of HPO problem where the focus is on algorithm driven design of neural network architecture components or cells [8]. Here, the search space is usually discrete and of variable dimensionality. Deep learning architectures designed via NAS algorithms have surpassed their hand-crafted counterparts for tasks such as image recognition and language modeling [9,10], underlining the practical importance of this field of research.

We present a new HPO algorithm, REINFORCE with Masked Attention Auto-regressive Density Estimators (ReMAADE), that uses a Transformer like architecture [11] to specify the policy network and employs Policy Gradient to tune its parameters. ReMAADE works well as both a NAS and generic HPO algorithm. For NAS, it outperforms most NAS algorithms on NASBench-101 [1]. It is able to discover better multi-layer perceptron models for regression tasks relative to Random Search on the Boston Housing and Naval Propulsion data-sets.

Our contributions can be summarized as:

- We present a 2-stream attention based architecture for capturing dependencies between hyper-parameters. This architecture’s parametric complexity is invariant to the dimensionality of the search-space, yet the architecture is expressive enough to capture long range dependencies.
- We present an actor-critic style algorithm for stabilizing training and improving performance. This approach is useful in low computational budget settings.

- We also investigate the effect of ensembling models with shared parameters, each conditioned on a different autoregressive factorization order. We posit that this approach can be useful in large computational budget settings for discovering optimal architectures.

## II. RELATED WORK

HPO algorithms can be categorized under Bayesian Optimization, evolutionary learning, and gradient based methods, with random search widely regarded as a competitive baseline [12,13].

In methods based on Bayesian optimization [14]–[17], the function mapping from the hyper-parameter search space to the validation score is modelled as a Gaussian Process (GP) [18]. This permits easy computation of the posterior distribution of the validation error for a new architecture. The performance of the method hinges on designing an appropriate kernel function for the GP and an acquisition function for fetching the next set of architectures for evaluation. A limitation of these approaches is that performance degrades in high-dimensional search spaces, because larger samples are required to update the posterior distribution [19].

In evolutionary learning models, inspired by genetic evolution, the architectures are modelled as gene strings. The search proceeds by mutating and combining strings so as to hone in on promising architectures [20]–[22]. Gradient based methods specify the objective function as a parametric model and proceed to optimize it with respect to the hyper-parameters via gradient-descent [9,23,24].

Model free RL based techniques [10,25]–[28] specify a policy network that learns to output desirable architectures. Search proceeds by training the policy network using Q-learning or policy gradient. These techniques are flexible in that they can search over variable length architectures, and have shown very promising results for neural architecture search.

NAS was initially formulated as a Reinforcement Learning (RL) problem, where a policy network was trained to sample more efficient architectures [25]. In [10], a cell-based search in a search space of 13 operations is performed to find *reduction* and *normal* cells. These cells are then stacked to form larger networks. In [26], the authors use Q-learning to discover promising architectures. Using a similar approach, [29] performs NAS by sampling blocks of operations instead of cells, which can then be stacked to form networks.

DARTS [30] deploys gradient-based search based on a continuous relaxation of an otherwise non-differentiable search space to discover high quality architectures. [31]–[34] improve on this by using regularization to bridge the gap between validation and test scores.

Evolutionary learning algorithms have been applied to NAS as well. [21,35] discover high quality architectures for image classification using evolutionary learning techniques.

## III. PRELIMINARIES

Let the search space comprise  $N$  hyper-parameters, indexed by  $i \in \{1, \dots, N\}$ . RL methods based on policy gradient

Keyword	Description
NAS	Neural Architecture Search
ReMAADE	REINFORCE and Masked Attention Auto-Regressive Density Estimators
BANANAS	Bayesian Optimization with Neural Architectures for Neural Architecture Search
GP	Gaussian Process
MADE	Masked Autoregressive Density Estimator
MAADE	Masked Attention Autoregressive Density Estimator
NADE	Neural Autoregressive Density
PPO	Proximal Policy Optimization
NASBOT	Neural Architecture Search with Bayesian Optimization and optimal Transport
MCTS	Monte Carlo Tree Search

TABLE I: List of abbreviations

[10,25] specify a policy network, parametrized by  $\theta$ , to learn a desired probability distribution over values of the hyperparameters,  $P(a_1, a_2, \dots, a_N; \theta)$ , where  $a_i$  denotes the value of the  $i^{th}$  hyper-parameter.

We set up the training regime such that the policy network learns to assign higher probabilities to those sequences of hyperparameter values (henceforth, referred to as strings) that yield a higher accuracy on the cross-validation dataset. Accordingly, we maximize

$$J(\theta) = \mathbb{E}_{a_{1:N} \sim P(a_{1:N}; \theta)} (f(a_{1:N}) - b) \quad (1)$$

where  $f: \mathcal{H} \rightarrow \mathbb{R}$  is an unknown function that maps strings from the hyperparameter search space to the ML model’s accuracy on the validation dataset.  $b$  is a baseline function to reduce the variance of the estimate of the gradient of (1).

We optimize the objective via gradient ascent where the gradient can be estimated using Reinforce [36]:

$$\nabla J(\theta) \approx \frac{1}{m} \sum_{k=1}^m \nabla_{\theta} \log P(a_{1:n}^k; \theta) [f(a_{1:n}^k) - b] \quad (2)$$

where  $a_{1:n}^k \sim P(a_{1:n}; \theta)$ .  $f(a_{1:n}) - b$  is referred to as the advantage function, denoted by  $A(a_{1:n})$ .

The optimization procedure alternates between two steps until we exhaust the exploration budget:

- 1) sample a batch of action strings based on the current state of the policy network and fetch the corresponding rewards from the environment
- 2) update the policy network’s parameters using policy gradient

The exploration budget can be quantified in units of computation such as number of GPU/TPU hours for training architectures or, alternately, as the number of times the policy network can query the environment to fetch the architecture’s score. In case of the latter, it is assumed that all architectures consume identical compute for getting trained.

## IV. AUTOREGRESSIVE MODELS FOR DENSITY ESTIMATION

The policy network can be set up as an auto-regressive model, an approach that has been successfully applied to language models [37,38], generative models of images [39]–[41] and speech [42].

$$P(a_{1:N}; \theta) = \prod_{i=1}^N P(a_i | a_{1:i-1}; \theta) \quad (3)$$

The choice of a parametric architecture for modelling terms in (3) becomes crucial as it needs to balance expressiveness against model complexity. The former is important to learn dependencies between hyper-parameters over increasing string lengths, while the latter needs to be economized for discovering optimal strings within an exploration budget. RNN based networks struggle to learn adequate context representations over longer sequence lengths because of the vanishing gradient/exploding gradient problem [43]. On the other hand, the parametric complexity of masked multi-layer perceptron based models, such as MADE [44] and NADE [45], increases with string length  $N$  (table II).

### A. Masked Attention Autoregressive Density Estimators (MAADE)

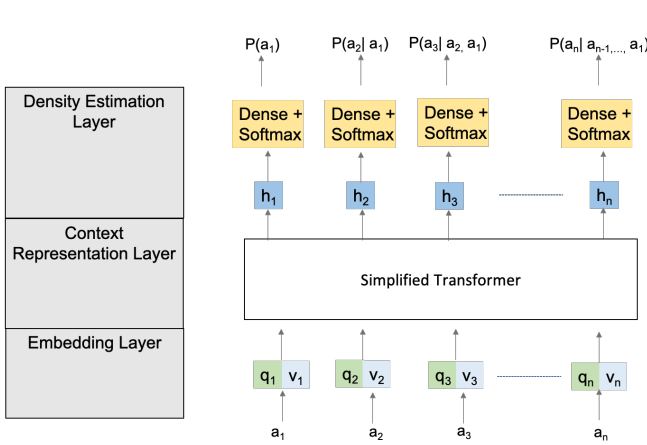


Fig. 1: Overall architecture layout

To strike a balance between expressiveness and complexity in designing the policy network, we propose the Masked Attention Autoregressive Density Estimation (MAADE) architecture, comprising three layers:

- 1) **Embedding layer** maps hyper-parameters and hyper-parameter values to a  $d$  dimensional vector space
- 2) **Context Representation layer** models dependencies between hyper-parameters as specified by the autoregressive factorization
- 3) **Density Estimation Layer** computes the probability density for a string

1) **Embedding Layer:** Let  $q_i \in \mathbb{R}^d, \forall i = 1, \dots, N$ , be query vectors for each of the  $N$  hyper-parameters. We also maintain value vectors for the values that each hyper-parameter can take. We assume without loss of generality that all hyper-parameters assume categorical values in the same space and dimensionality  $D$ . We therefore share the embedding layer  $V \in \mathbb{R}^{d \times D}$  across hyper-parameters. Given the value of the  $i^{\text{th}}$  hyper-parameter,  $a_i \in \{1, \dots, D\}$ , the corresponding value vector is  $V_{:,a_i} \in \mathbb{R}^d$  which we denote as  $v_i$  with slight abuse of notation. Note that this framework can easily be extended to deal with hyper-parameters operating in different spaces as well, such that we maintain a separate embedding layer for each hyper-parameter family.

2) **Context Representation Layer:** This layer produces  $d$  dimensional contextual vectors,  $h_i \in \mathbb{R}^d$  for each hyper-parameter as a function of the hyper-parameter being predicted and previously seen hyper-parameters:

$$h_i \leftarrow H_\theta(q_i, q_{1:i-1}, v_{1:i-1}) \quad (4)$$

Inspired by XLNet [46], we use a two-stream masked attention based architecture comprising query and key vectors to compose  $H_\theta$ . A notable departure from XLNet is that since we are not predicting probabilities for a position but for a given hyper-parameter, we let the query vector of the target hyper-parameter attend to preceding key vectors. Each key vector attends to preceding key vectors as well as itself.

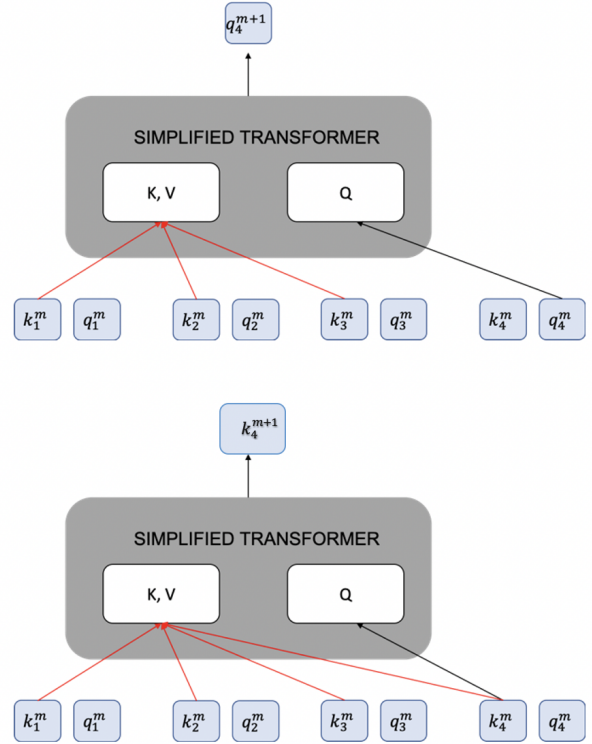


Fig. 2: 2-stream attention

Stream 1 is initialized as  $q_i^{(0)} \leftarrow q_i$  stream 2 is initialized as  $k_i^{(0)} \leftarrow v_i + q_i$ . Then we update the streams as:

$$q_i^{(m+1)} \leftarrow \text{Tran}(q_i^m, k_{1:i-1}^{(m)}) \quad (5)$$

$$k_i^{(m+1)} \leftarrow \text{Tran}(k_i^m, k_{1:i}^{(m)}) \quad (6)$$

where  $\text{Tran}$  is the transformer block referred in Figure 3. Finally, we get the contextual representations as:  $h_i \leftarrow q_i^{(M)}$

To specify the Simplified Transformer block, we use a simplified version of Transformer [11]. In the attention layer, we eschew dot production attention in favour of additive attention [47] to model interactions between query and key/value vectors as it was found to marginally improve the policy

network’s performance. We also found the policy network’s performance to deteriorate when  $M > 2$ , obviating the need for both residual connections and layer-normalization. Finally, we do away with positional encoding since the sequence in which preceding hyper-parameters in the auto-regressive order were encountered doesn’t matter.

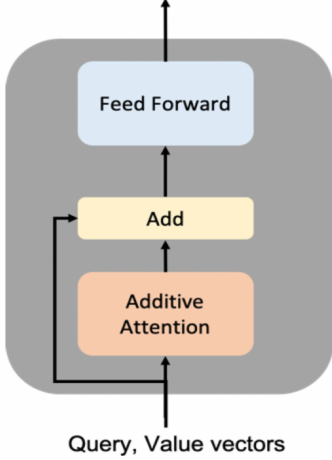


Fig. 3: Simplified Transformer block

We provide details of the computation steps in the Simplified Transformer block:

$$q_i^{(m+1)} \leftarrow \text{PosFF}(q_i^{(m)} + \text{Masked\_Attention}(Q = q_i^{(m)}, KV = k_{1:i-1}^{(m)})) \quad \forall i = 1, \dots, n \quad (7)$$

$$k_i^{(m+1)} \leftarrow \text{PosFF}(k_i^{(m)} + \text{Masked\_Attention}(Q = k_i^{(m)}, KV = k_{1:i}^{(m)})) \quad \forall i = 1, \dots, n \quad (8)$$

where *Masked Attention* is the additive attention operation, and *PosFF*, the position-wise feed-forward operation with *relu* non-linearity replaced by *tanh*:

$$\text{PosFF}(x) = W_2(\tanh(W_1x + b_1)) + b_2 \quad (9)$$

Note that both the streams share parameters of the masked attention and feed-forward operations.

3) **Density Estimation Layer**: In this layer, we pass the context representations through an affine transformation specific to the target hyper-parameter followed by a softmax,

$$P(a_i | a_{1:i-1}; \theta) = \frac{\exp(h_{a_i}^T W_{a_i}^i + b_{a_i}^i)}{\sum_j \exp(h_j^T W_j^i + b_j^i)} \quad (10)$$

**Complexity** A key advantage of attention-based networks is that model complexity of the context representation layer doesn’t change with string length  $N$ . This is crucial for the NAS problem where the policy network needs to discover the

best architecture within a limited exploration budget.

Table II defines the parametric complexity of each architecture.

Architecture	Complexity of Context Representation layer
MAADE	$O(d^2)$
NADE	$O(Nd^2)$
MADE	$O(N^2d^2)$
LSTM based controller	$O(d^2)$

TABLE II: Parametric Complexity

## V. TRAINING

For training the policy network, we can optimize either (1) or a PPO [48] objective as follows:

$$J'(\theta) = \mathbb{E}_{a_{1:N} \sim P(a_{1:N}; \theta)} [A(a_{1:N}) \min(r(a_{1:N}; \theta, \theta'), \text{clip}(r(a_{1:N}; \theta, \theta'), 1 - \epsilon, 1 + \epsilon))] \quad (11)$$

where,

$$r(a_{1:N}; \theta, \theta') = \frac{P(a_{1:N}; \theta)}{P(a_{1:N}; \theta')} \quad (12)$$

An unbiased estimate of the gradient of (11) is:

$$\nabla J'(\theta) \approx \frac{1}{m} \sum_{k=1}^m \nabla_{\theta} [A(a_{1:n}^k) \min(r(a_{1:n}^k; \theta, \theta'), \text{clip}(r(a_{1:n}^k; \theta, \theta'), 1 - \epsilon, 1 + \epsilon))] \quad (13)$$

where  $a_{1:n}^k \sim P(a_{1:n}; \theta)$ . We update the policy network parameters via gradient ascent:

$$\theta \leftarrow \theta + \alpha \nabla J'(\theta) \quad (14)$$

$\epsilon$ ,  $B$  constitute ReMAADE’s hyper-parameters, which, along with the learning rate,  $\alpha$ , can be tuned via cross-validation. or the term in (1). We can optionally add an entropy term to the objective to encourage exploration if we have a large exploration budget [49].

### A. ReMAADE algorithm

We now describe the REINFORCE with Masked Attention Auto-regressive Density Estimators (ReMAADE) algorithm.

**Inputs:**

- Search space of architectures :  $\mathcal{H}$
- Environment function that maps an architecture from the search-space to the corresponding validation accuracy :  $f : a_{1:N} \in \mathcal{H} \rightarrow \mathbb{R}$
- Exploration budget:  $E$
- ReMAADE hyperparameters:  $B, S, \alpha, \epsilon$  where  $B$  is the batch-size,  $S$  is the set size of auto-regressive orderings to sample from,  $\alpha$  is the learning rate,  $\epsilon$  is the PPO coefficient.
- MAADE hyperparameters:  $d, M$  where  $d$  is the embedding dimension for the transformer block and  $M$  is the number of transformer blocks stacked.

The algorithm is described as follows:

Algorithm	Source	Test Error (in %)	Std-Deviation (in %)
TPE	Bergstra et al. [16]	6.43	0.16
BOHB	Falkner et al. [15]	6.40	0.12
Random Search	Bergstra et al. [12]	6.36	0.12
NASBOT	Kandasamy et al. [14]	6.35	0.10
Alpha X	Wang et al. [28]	6.31	0.13
Reg Evolution	Real et al. [21]	6.20	0.13
ReMAADE	Ours	6.16	0.25
ReACTS	Ours	6.13	0.25
BANANAS	White et al. [17]	5.77	0.31

TABLE III: Performance of different search algorithms on NASBench-101 for short term run

---

**Algorithm 1** *ReMAADE*


---

```

1:  $e \leftarrow 0$ 
2:  $f(a^*) \leftarrow -\infty$ 
3: Initialize policy network and  $\theta$ 
4: while  $e < E$  do
5:   Sample  $B$  valid hyper-parameter strings,
    $\{a_{1:N}^1, \dots, a_{1:N}^B\}$ , using the policy network
6:   Fetch corresponding rewards,  $\{f(a_{1:N}^1), \dots, f(a_{1:N}^B)\}$ 
7:    $a^* \leftarrow \operatorname{argmax}(f(a^*), \{f(a_{1:N}^1), \dots, f(a_{1:N}^B)\})$ 
8:   Update  $\theta$  as  $\theta \leftarrow \theta + \alpha \nabla J'(\theta)$ 
9:    $e \leftarrow e + B$ 
10: end while
11: Return best architecture found  $a^*$ 

```

---

## VI. REACTS: REDUCING VARIANCE USING A CRITIC

Equation (2) provides an unbiased but high variance estimate of the gradient. The variance can be reduced by recourse to actor-critic algorithms [50,51]. We adopt a similar procedure, as follows. Define the value function as:

$$V_\theta(a_{\leq i}) = \mathbb{E}_{a_{i+1:N} \sim P(a_{i+1:N} | a_{\leq i}; \theta)} [f(a_{1:N})] \quad (15)$$

In other words, the value function is the expected reward if we sample actions given the first  $i$  actions,  $a_1, a_2, \dots, a_i$ . We can use Monte-Carlo policy evaluation to obtain an unbiased estimate of the value function. However, that would require us to query the environment further. Instead we take recourse to a simulator to estimate the value function as:

$$\widehat{V}_\theta(a_{\leq i}) = \frac{1}{L} \sum_{l=1}^L S_\phi(a_{\leq i}, a_{>i}^l | a_{\leq i}) \approx V_\theta(a_{\leq i}) \quad (16)$$

where  $a_{>i}^l | a_{\leq i} \sim P(a_{i+1:N} | a_{\leq i}; \theta)$ .  $S_\phi$  is a simulator, such as the meta-network in BANANAS [17], that predicts the validation set reward associated with an architecture. Then, an unbiased, lower variance (relative to equation (2)) estimate of  $\nabla J(\theta)$  can be computed as:

$$\nabla J(\theta) \approx \frac{1}{B} \sum_{k=1}^B \sum_{i=1}^N \nabla_\theta \log P(a_i^k | a_{<i}^k; \theta) [f(a_{1:N}^k) - \widehat{V}_\theta(a_{<i}^k)] \quad (17)$$

where  $a_{1:N}^k \sim P(a_{1:N}; \theta)$ . Search for the optimal architecture proceeds in discrete steps. At a given step  $t$ , we sample  $B$  valid architectures using the policy network and fetch the corresponding rewards by querying the environment. Let the assembled data-set be denoted by

$$B_{\theta^t} := \{(f(a_{1:N}^1), a_{1:N}^1; \theta^t), \dots, (f(a_{1:N}^B), a_{1:N}^B; \theta^t)\} \quad (18)$$

We update  $\phi$  such that it can predict rewards for architectures drawn from the policy network at step  $t$ :

$$\phi = \operatorname{argmax}(\mathbb{E}_{a_{1:N} \sim P(a_{1:N}; \theta^t)} [-\operatorname{loss}(S_\phi, f(a_{1:N}))]) \quad (19)$$

To make use of samples accumulated from earlier states of the policy network,  $\{B_{\theta^0}, B_{\theta^1}, \dots, B_{\theta^t}\}$  we need to adjust for co-variate shift [52]. Accordingly, we update  $\phi$  as:

$$\phi = \operatorname{argmax}(\mathbb{E}_{a_{1:N} \sim P(a_{1:N}; \theta^{0:t})} [-\operatorname{loss}(S_\phi; f(a_{1:N})) \frac{(t+1)P(a_{1:N}; \theta^t)}{P(a_{1:N}; \theta^{0:t})}]) \quad (20)$$

The policy network's state is then updated using equation (17).

### A. ReACTS Algorithm

We now have the machinery to describe ReACTS.

---

**Algorithm 2** *ReACTS*


---

```

1:  $f(a^*) \leftarrow -\infty$ 
2: Initialize policy network parameters,  $\theta^0$ 
3: for  $t = 0 \dots T$  do
4:   Sample  $B$  valid hyper-parameter strings using policy
   network, Fetch corresponding rewards to assemble  $B_{\theta^t}$ 
5:    $a^* \leftarrow \operatorname{argmax}(f(a^*), \{f(a_{1:N}^1), \dots, f(a_{1:N}^B)\})$ 
6:   Compute  $\phi^t$  using  $\{B_{\theta^0}, B_{\theta^1}, \dots, B_{\theta^t}\}$ , equation (20)
7:    $\theta^{t+1} \leftarrow \operatorname{ProcedureII}(B_{\theta^t}, \phi^t)$ 
8: end for
9: Return best architecture found  $a^*$ 

```

---

---

**Algorithm 3** *ProcedureII*

---

```
1: Inputs:  $B_{\theta^t}, S(\phi^t), P(a_{1:N}; \theta^t)$ 
2: for  $k = 1, \dots, B$  do
3:   for  $i = 1, \dots, N$  do
4:     Compute  $\widehat{V}_{\theta}(a_{\leq i}^k)$  as per equation (16)
5:   end for
6: end for
7: Compute  $\nabla J(\theta^t)$  as per equation (17)
8:  $\theta^{t+1} \leftarrow \theta^t + \alpha \nabla J(\theta^t)$ 
9: Return  $\theta^{t+1}$ 
```

---

The **ReACTS** algorithm should outperform naive policy gradient-based methods if we have recourse to a good simulator. The Actor-Critic formula in equation (17) fully exploits the Markovian nature of the implicit MDP of the policy network. Further, the transition dynamics are deterministic, eliminating another source of variance. Therefore, we expect this way of computing the gradient to reduce the variance of the gradient estimate and stabilize the training regime.

## VII. RESULTS ON NASBENCH-101

**NASBench-101 search space** The NASBench-101 dataset [46] is a public architecture dataset to facilitate NAS research and compare NAS algorithms. The search space comprises the elements of small-feed forward structures called cells. These cells are assembled together in a predefined manner to form an overall convolutional neural network architecture that is trained on the CIFAR-10 dataset.

A cell comprises 7 nodes, of which the first node is the input node and the last node is the output. The remaining 5 nodes need to be assigned one of 3 operations: 1x1 convolution, 3x3 convolution, or 3x3 max pooling. The nodes then need to be connected to form a valid directed acyclic graph (DAG). The NAS algorithm therefore needs to specify the operations for each of the 5 nodes, and then specify the edges to form a valid DAG. To limit the search space, NASBench-101 imposes additional constraints: the total number of edges cannot exceed 9, and there needs to be a path from the input node to the output node. This results in 423K valid and unique ways to specify a cell. NASBench-101 has pre-computed the validation and test errors for all the neural network architectures that can be designed from these 423K cell configurations.

To benchmark ReMAADE on NASBench-101, we investigate short term performance (exploration budget of 150 architectures). We include random search [13], which is regarded as a competitive baseline, regularized evolution [21], and AlphaX, an RL algorithm that uses MCTS [28]. We also compare with several algorithms based on Bayesian optimization with GP priors: BOHB [53], tree-structured Parzen estimator (TPE) [16], BANANAS [17], and NASBOT [14]. For all NAS algorithms, during a trial, we track the best random validation error achieved after  $t$  explorations and the corresponding random test error. We report metrics averaged over 500 trials for each NAS algorithm. For ReMAADE, in

all experiments, we set  $M = 1, \epsilon = 0.1$ , and used ADAM [54] for updating  $\theta$ .

### Short Term Performance

NAS algorithms need to discover good architectures within 150 explorations to be of practical use. In this setting, (table III), ReMAADE outperforms all algorithms with the exception of BANANAS. For ReMAADE, we set  $\alpha = 1e - 2, d = 36, S = 1, B = 30$ .

#### A. Ablation studies

Algorithm	Test Error (in %)
Random Search	6.36 +-0.12
Plain vanilla REINFORCE	6.26 +- 0.22
REINFORCE with MADE	6.25 +- 0.22
ReMAADE w/o PPO	6.16 +- 0.25
ReACTS w/o PPO	6.13 +- 0.25

TABLE IV: Ablation Study

We perform an ablation study to understand the importance of the autoregressive component and MAADE in designing the context representation layer. We can use REINFORCE without an autoregressive model, which we call plain vanilla REINFORCE. In other words, we assume that all actions are independent:

$$P(a_1, a_2, \dots, a_N; \theta) = \prod_i P(a_i) \quad (21)$$

This amounts to updating only the bias terms in (10) using policy gradient, and yields a baseline test set error of 6.26%. Interestingly, using MADE [44] to design the autoregressive model failed to improve upon plain vanilla REINFORCE, underscoring the importance of MAADE in capturing autoregressive dependencies. Using ReACTS marginally improved performance relative to ReMAADE (table IV).

#### B. Effect of auto-regressive ordering ensembles

Autoregressive density estimation models struggle with terms in (3) with  $i \gg 1$  since they use a fixed capacity in the context representation layer. This can be mitigated to some extent by picking an autoregressive factorization order that exploits spatial-temporal dependencies using an appropriate architecture. For instance, in generative modelling of images, the raster scan ordering is preferred as it is able to capture spatial dependencies in the immediate neighborhood [39].

In case of neural architecture search, however, it is not clear, *a priori*, what autoregressive order to fix. Therefore, training an ensemble of models, each with a different autoregressive factorization order, with parameters shared across all models, can potentially improve performance as shown in [45,46]. To do so, we explicitly condition the density on the autoregressive factorization order and share the policy network’s parameters across orders.

Accordingly, we set up the following framework: Let  $Z_{N!}$  denote the set of all possible permutations of length  $N$  index

sequences. Let  $Z_S \subseteq Z_{N!}$  with set size  $S$ . Let  $z_{s,t}$  denote the  $t^{\text{th}}$  element of a permutation  $z_S \in Z_S$ . We then define the joint probability over the autoregressive factorization order and action string,  $(z_s, a_{1:N})$ , as:

$$P(z_s, a_{1:N}; \theta) = \frac{1}{S} \prod_{i=1}^N P(a_{z_s,i} | a_{z_s,1:i-1}; \theta) \quad (22)$$

To sample from this distribution, we sample a permutation uniformly at random from  $Z_S$ . We then fix the autoregressive factorization order based on the sampled permutation and sample the action string based on the MAADE architecture.

Accordingly, we maximize the following PPO objective:

$$J'(\theta) = \mathbb{E}_{z_s, a_{1:N} \sim P(z_s, a_{1:N}; \theta, \theta')} [A(a_{1:N}; \theta) \min(r(z_s, a_{1:N}; \theta, \theta'), \text{clip}(r(z_s, a_{1:N}; \theta, \theta'), 1 - \epsilon, 1 + \epsilon))] \quad (23)$$

where,

$$r(z_s, a_{1:N}; \theta, \theta') = \frac{P(z_s, a_{1:N}; \theta)}{P(z_s, a_{1:N}; \theta')} \quad (24)$$

Training proceeds as per the ReMAADE algorithm. We expect that higher values of  $S$  will lead to improved performance as we increase the exploration budget. Also note that in the limit when  $S = N!$ , the training objective is identical to the bidirectional training objective in [46]. We empirically validated this and found performance to improve as the number of orderings was increased to 4 and then deteriorate before asymptotically improving again, when given an exploration budget of 3,200 architectures on NAS-Bench 101 (ref Table V).

S	Test Error(in %)
1	5.95 +- 0.18
2	5.94 +- 0.19
4	5.91 +- 0.19
6	5.92 +- 0.19
8	5.93 +- 0.19
16	5.92 +- 0.19
256	5.92 +- 0.19

TABLE V: Effects of autoregressive factorization order ensembling

## VIII. RESULTS FOR HYPER-PARAMETER OPTIMIZATION

We evaluate the performance of ReMAADE in optimizing the hyper-parameters of a Multi-Layer Perceptron (MLP). The MLP model is trained on two real-world datasets for a regression task.

- 1) **Boston Housing** [55]: This data-set consists of 506 samples, each sample made of 13 scaled input variables and a scalar regression output, the housing price.

- 2) **Naval propulsion plants** [56]: The data-set consists of 11,934 samples, each sample made of 16 scaled input variables and a scalar regression output variable (turbine degradation coefficient).

ReMAADE is bench-marked against Random Search and TPE [57]. Each hyper-parameter optimization algorithm is given a budget of 100 model explorations per trial, and the Root Mean Square Error (RMSE) obtained on the test-set is reported at the end of the trial. The 3 HPO algorithms are evaluated on RMSE averaged over 100 trials.

### A. Case Study I: Boston Housing

In this case-study, we train an MLP model on the Boston Housing data-set to predict house-prices. The MLP model has 10 hyper-parameters, including learning rate, l1 and l2 regularization, size of the hidden layer, number of iterations and choice of activation function. All 3 algorithms performed identically for this data-set and search space (Figure 4).

### B. Case Study II: Naval propulsion plants [56]

In the second case study, we consider the task of improving condition based maintenance of naval propulsion plants. The data is generated using numerical simulator of a naval vessel which is characterized by a Gas Turbine (GT) propulsion plant. We trained a MLP model on this data-set with 10 hyper-parameters, including learning rate l1 and l2 penalties, size of hidden layer, number of iterations, activation function. In this setting, ReMAADE clearly outperforms Random Search and is competitive with TPE (Figure 5).

### C. Importance of Various Hyperparameters

We find that Boston housing dataset with numerical features gave RMSE of about 14.102 while Naval Propulsion dataset gave 1.98 as RMSE with numerical features. Other search algorithms shown as baselines performed worse than ReMAADE, highlighting the use of Reinforced and masked attention auto-regressive ordering.

Among other algorithms which are shown in Table III, we compared our results with Random and TPE(Tree-structured

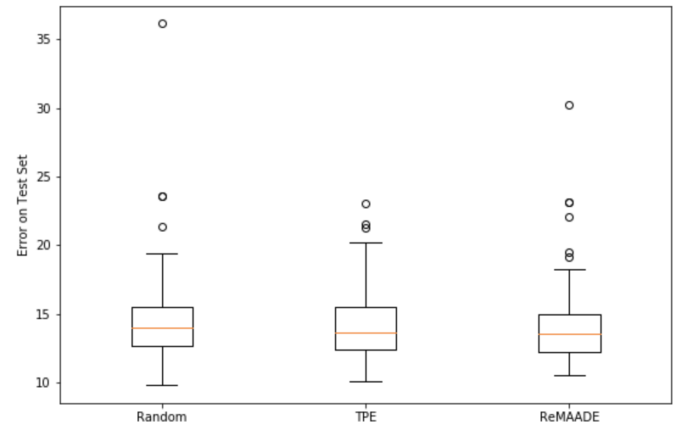


Fig. 4: Benchmarking on the Boston Dataset



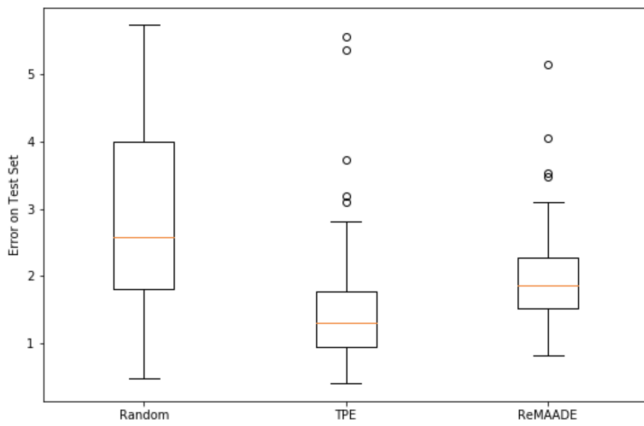


Fig. 5: Benchmarking on the Naval Dataset

Parzen Estimator) we got performance boost of 3.48% in RMSE for Boston Housing dataset and nearly 10% improvement in RMSE for Naval Propulsion dataset. Finally, we also observe that leveraging attention mechanism with auto-regressive ordering across different metadata helps the algorithm improve performance.

## IX. FUTURE WORK

We conclude that attention based auto-regressive models combined with policy gradient can be used as an effective hyper-parameter optimization problem. We need to further investigate the impact of conditioning on and ensembling across multiple autoregressive factorization orders on performance given large computational budgets. We also plan to investigate the performance of ReMAADE on other neural architecture search spaces such as DARTS [30]. Another interesting line of work would be to use deep generative graph models with Policy Gradient [58] to discover optimal NAS architectures. To model the graph context and capture node dependencies, we plan to investigate attention based mechanisms as was done in this paper.

## REFERENCES

- [1] C. Ying, A. Klein, E. Real, E. Christiansen, K. Murphy, and F. Hutter, "Nas-bench-101: Towards reproducible neural architecture search," *arXiv preprint arXiv:1902.09635*, 2019.
- [2] M. Lindauer and F. Hutter, "Best practices for scientific research on neural architecture search," *arXiv preprint arXiv:1909.02453*, 2019.
- [3] S. Hochreiter and J. Schmidhuber, "Long short-term memory," *Neural computation*, vol. 9, no. 8, pp. 1735–1780, 1997.
- [4] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in neural information processing systems*, 2014, pp. 2672–2680.
- [5] G. Melis, C. Dyer, and P. Blunsom, "On the state of the art of evaluation in neural language models," *arXiv preprint arXiv:1707.05589*, 2017.
- [6] M. Lucic, K. Kurach, M. Michalski, S. Gelly, and O. Bousquet, "Are gans created equal? a large-scale study," in *Advances in neural information processing systems*, 2018, pp. 700–709.
- [7] C. Thornton, F. Hutter, H. H. Hoos, and K. Leyton-Brown, "Auto-weka: Combined selection and hyperparameter optimization of classification algorithms," in *Proceedings of the 19th ACM SIGKDD international conference on Knowledge discovery and data mining*, 2013, pp. 847–855.
- [8] T. Elsken, J. H. Metzen, and F. Hutter, "Neural architecture search: A survey," *arXiv preprint arXiv:1808.05377*, 2018.
- [9] R. Luo, F. Tian, T. Qin, E. Chen, and T.-Y. Liu, "Neural architecture optimization," in *Advances in neural information processing systems*, 2018, pp. 7816–7827.
- [10] B. Zoph, V. Vasudevan, J. Shlens, and Q. V. Le, "Learning transferable architectures for scalable image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 8697–8710.
- [11] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin, "Attention is all you need," in *Advances in neural information processing systems*, 2017, pp. 5998–6008.
- [12] J. Bergstra and Y. Bengio, "Random search for hyper-parameter optimization," *Journal of machine learning research*, vol. 13, no. Feb, pp. 281–305, 2012.
- [13] L. Li and A. Talwalkar, "Random search and reproducibility for neural architecture search," *arXiv preprint arXiv:1902.07638*, 2019.
- [14] K. Kandasamy, W. Neiswanger, J. Schneider, B. Poczos, and E. P. Xing, "Neural architecture search with bayesian optimisation and optimal transport," in *Advances in Neural Information Processing Systems*, 2018, pp. 2016–2025.
- [15] S. Falkner, A. Klein, and F. Hutter, "Bohb: Robust and efficient hyperparameter optimization at scale," *arXiv preprint arXiv:1807.01774*, 2018.
- [16] J. S. Bergstra, R. Bardenet, Y. Bengio, and B. Kégl, "Algorithms for hyper-parameter optimization," in *Advances in neural information processing systems*, 2011, pp. 2546–2554.
- [17] C. White, W. Neiswanger, and Y. Savani, "Bananas: Bayesian optimization with neural architectures for neural architecture search," *arXiv preprint arXiv:1910.11858*, 2019.
- [18] C. K. Williams and C. E. Rasmussen, *Gaussian processes for machine learning*. MIT press Cambridge, MA, 2006, vol. 2, no. 3.
- [19] N. Fusi, R. Sheth, and M. Elibol, "Probabilistic matrix factorization for automated machine learning," in *Advances in neural information processing systems*, 2018, pp. 3348–3357.
- [20] E. Real, S. Moore, A. Selle, S. Saxena, Y. L. Suematsu, J. Tan, Q. V. Le, and A. Kurakin, "Large-scale evolution of image classifiers," in *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 2017, pp. 2902–2911.
- [21] E. Real, A. Aggarwal, Y. Huang, and Q. V. Le, "Regularized evolution for image classifier architecture search," in *Proceedings of the AAAI conference on artificial intelligence*, vol. 33, 2019, pp. 4780–4789.
- [22] L. Xie and A. Yuille, "Genetic cnn," in *Proceedings of the IEEE international conference on computer vision*, 2017, pp. 1379–1388.
- [23] D. Maclaurin, D. Duvenaud, and R. Adams, "Gradient-based hyperparameter optimization through reversible learning," in *International Conference on Machine Learning*, 2015, pp. 2113–2122.
- [24] C. Liu, B. Zoph, M. Neumann, J. Shlens, W. Hua, L.-J. Li, L. Fei-Fei, A. Yuille, J. Huang, and K. Murphy, "Progressive neural architecture search," in *Proceedings of the European Conference on Computer Vision (ECCV)*, 2018, pp. 19–34.
- [25] B. Zoph and Q. V. Le, "Neural architecture search with reinforcement learning," *arXiv preprint arXiv:1611.01578*, 2016.
- [26] B. Baker, O. Gupta, N. Naik, and R. Raskar, "Designing neural network architectures using reinforcement learning," *arXiv preprint arXiv:1611.02167*, 2016.
- [27] H. Pham, M. Y. Guan, B. Zoph, Q. V. Le, and J. Dean, "Efficient neural architecture search via parameter sharing," *arXiv preprint arXiv:1802.03268*, 2018.
- [28] L. Wang, Y. Zhao, Y. Jinnai, and R. Fonseca, "Alphax: exploring neural architectures with deep neural networks and monte carlo tree search," *arXiv preprint arXiv:1805.07440*, 2018.
- [29] Z. Zhong, J. Yan, W. Wu, J. Shao, and C.-L. Liu, "Practical block-wise neural network architecture generation," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2018, pp. 2423–2432.
- [30] H. Liu, K. Simonyan, and Y. Yang, "Darts: Differentiable architecture search," *arXiv preprint arXiv:1806.09055*, 2018.
- [31] A. Zela, T. Elsken, T. Saikia, Y. Marrakchi, T. Brox, and F. Hutter, "Understanding and robustifying differentiable architecture search," *arXiv preprint arXiv:1909.09656*, 2019.
- [32] X. Chen, L. Xie, J. Wu, and Q. Tian, "Progressive differentiable architecture search: Bridging the depth gap between search and evaluation," in



- Proceedings of the IEEE International Conference on Computer Vision*, 2019, pp. 1294–1303.
- [33] H. Cai, L. Zhu, and S. Han, “Proxylessnas: Direct neural architecture search on target task and hardware,” *arXiv preprint arXiv:1812.00332*, 2018.
- [34] Y. Xu, L. Xie, X. Zhang, X. Chen, G.-J. Qi, Q. Tian, and H. Xiong, “Pc-darts: Partial channel connections for memory-efficient architecture search,” in *International Conference on Learning Representations*, 2019.
- [35] H. Liu, K. Simonyan, O. Vinyals, C. Fernando, and K. Kavukcuoglu, “Hierarchical representations for efficient architecture search,” *arXiv preprint arXiv:1711.00436*, 2017.
- [36] R. J. Williams, “Simple statistical gradient-following algorithms for connectionist reinforcement learning,” *Machine learning*, vol. 8, no. 3-4, pp. 229–256, 1992.
- [37] T. Mikolov, M. Karafiát, L. Burget, J. Černocký, and S. Khudanpur, “Recurrent neural network based language model,” in *Eleventh annual conference of the international speech communication association*, 2010.
- [38] Y. Kim, Y. Jernite, D. Sontag, and A. M. Rush, “Character-aware neural language models,” in *Thirtieth AAAI Conference on Artificial Intelligence*, 2016.
- [39] A. v. d. Oord, N. Kalchbrenner, and K. Kavukcuoglu, “Pixel recurrent neural networks,” *arXiv preprint arXiv:1601.06759*, 2016.
- [40] T. Salimans, A. Karpathy, X. Chen, and D. P. Kingma, “Pixelcnn++: Improving the pixelcnn with discretized logistic mixture likelihood and other modifications,” *arXiv preprint arXiv:1701.05517*, 2017.
- [41] X. Chen, N. Mishra, M. Rohaninejad, and P. Abbeel, “Pixel-snail: An improved autoregressive generative model,” *arXiv preprint arXiv:1712.09763*, 2017.
- [42] A. v. d. Oord, S. Dieleman, H. Zen, K. Simonyan, O. Vinyals, A. Graves, N. Kalchbrenner, A. Senior, and K. Kavukcuoglu, “Wavenet: A generative model for raw audio,” *arXiv preprint arXiv:1609.03499*, 2016.
- [43] R. Pascanu, T. Mikolov, and Y. Bengio, “On the difficulty of training recurrent neural networks,” in *International conference on machine learning*, 2013, pp. 1310–1318.
- [44] M. Germain, K. Gregor, I. Murray, and H. Larochelle, “Made: Masked autoencoder for distribution estimation,” in *International Conference on Machine Learning*, 2015, pp. 881–889.
- [45] B. Uribe, M.-A. Côté, K. Gregor, I. Murray, and H. Larochelle, “Neural autoregressive distribution estimation,” *The Journal of Machine Learning Research*, vol. 17, no. 1, pp. 7184–7220, 2016.
- [46] Z. Yang, Z. Dai, Y. Yang, J. Carbonell, R. R. Salakhutdinov, and Q. V. Le, “Xlnet: Generalized autoregressive pretraining for language understanding,” in *Advances in neural information processing systems*, 2019, pp. 5754–5764.
- [47] D. Bahdanau, K. Cho, and Y. Bengio, “Neural machine translation by jointly learning to align and translate,” *arXiv preprint arXiv:1409.0473*, 2014.
- [48] J. Schulman, F. Wolski, P. Dhariwal, A. Radford, and O. Klimov, “Proximal policy optimization algorithms,” *arXiv preprint arXiv:1707.06347*, 2017.
- [49] V. Mnih, A. P. Badia, M. Mirza, A. Graves, T. Lillicrap, T. Harley, D. Silver, and K. Kavukcuoglu, “Asynchronous methods for deep reinforcement learning,” in *International conference on machine learning*, 2016, pp. 1928–1937.
- [50] V. R. Konda and J. N. Tsitsiklis, “Actor-critic algorithms,” in *Advances in neural information processing systems*, 2000, pp. 1008–1014.
- [51] J. Schulman, P. Moritz, S. Levine, M. Jordan, and P. Abbeel, “High-dimensional continuous control using generalized advantage estimation,” *arXiv preprint arXiv:1506.02438*, 2015.
- [52] S. Bickel, M. Brückner, and T. Scheffer, “Discriminative learning under covariate shift,” *Journal of Machine Learning Research*, vol. 10, no. 9, 2009.
- [53] J. Snoek, H. Larochelle, and R. P. Adams, “Practical bayesian optimization of machine learning algorithms,” in *Advances in neural information processing systems*, 2012, pp. 2951–2959.
- [54] D. P. Kingma and J. L. Ba, “Adam: A method for stochastic optimization,” in *Proceedings of the International Conference on Learning Representations*, 2015.
- [55] D. Harrison and D. Rubinfeld, “Boston housing dataset,” 2015.
- [56] A. Coraddu, L. Oneto, A. Ghio, S. Savio, D. Anguita, and M. Figari, “Machine learning approaches for improving condition-based maintenance of naval propulsion plants,” *Proceedings of the Institution of Mechanical Engineers, Part M: Journal of Engineering for the Maritime Environment*, vol. 230, no. 1, pp. 136–153, 2016.
- [57] J. Bergstra, B. Komer, C. Eliasmith, D. Yamins, and D. D. Cox, “Hyperopt: a python library for model selection and hyperparameter optimization,” *Computational Science & Discovery*, vol. 8, no. 1, p. 014008, 2015.
- [58] J. You, B. Liu, Z. Ying, V. Pande, and J. Leskovec, “Graph convolutional policy network for goal-directed molecular graph generation,” in *Advances in neural information processing systems*, 2018, pp. 6410–6421.

## APPENDIX

In order for better evaluation and reproduction of our research we address all the items in the checklist as mentioned in Lindauer and Hutter [2].

- *Code for the training pipeline used to evaluate the final architectures* We used the search space of the architectures reported in NASBench-101 and thus the accuracy for all the architectures were precomputed. We have published code to train the policy network for ReMAADE algorithm, we also provide code for how different NAS algorithms were evaluated.
- *Code for the search space* The publicly available NASBench-101 dataset search space was used.
- *Hyperparameters used for the final evaluation pipeline, as well as random seeds* The hyperparameters were left unchanged.
- *For all NAS methods you compare, did you use exactly the same NAS benchmark, including the same dataset, search space, and code for training the architectures and hyperparameters for that code?* Yes, as NASBench-101 was used for evaluation, the search space was fixed accordingly. All the different NAS methods we surveyed were evaluated against the same NASBench-101 dataset.
- *Did you control for confounding factors?* Yes, all the experiments across all NAS algorithms were on the same NASBench-101 framework.
- *Did you run ablation studies?* Yes, the results for the ablation studies have been outlined in the paper.
- *Did you use the same evaluation protocol for the methods being compared?* Yes, the same evaluation protocol was used.
- *Did you compare performance over time?* Yes, the performance was evaluated both in the short term and the medium term with an exploration budget of 150 architectures in the short term and an exploration budget of 3200 architectures in the medium-term..
- *Did you compare to random search?* Yes.
- *Did you perform multiple runs of your experiments and report seeds?* Yes, we ran 500 trials of each experiment, with a different seed for each trial on NASBench-101. These results are completely reproducible.
- *Did you use tabular or surrogate benchmarks for indepth evaluations* Yes, all our experiments were evaluated against the NASBench-101 dataset.
- *Did you report how you tuned hyperparameters, and what time and resources this required?* We explored certain ranges of hyper-parameters get the best performance. These have been mentioned in the paper.
- *Did you report the time for the entire end-to-end NAS method?* Since all our experiments were run against the NASBench-101 framework for which we had pre-computed results, we were able to test the architectures generated by our network without training them from scratch.
- *Did you report all details of your experimental setup?*

Yes, all the details for the experimental setup have been reported.