

FreeREA: Training-Free Evolution-based Architecture Search

Niccolò Cavagnero, Luca Robbiano, Barbara Caputo, Giuseppe Averta

Politecnico di Torino, Italy

{niccolo.cavagnero, luca.robbiano, barbara.caputo, giuseppe.averta}@polito.it

Abstract

In the last decade, most research in Machine Learning contributed to the improvement of existing models, with the aim of increasing the performance of neural networks for the solution of a variety of different tasks. However, such advancements often come at the cost of an increase of model memory and computational requirements. This represents a significant limitation for the deployability of research output in realistic settings, where the cost, the energy consumption, and the complexity of the framework play a crucial role. To solve this issue, the designer should search for models that maximise the performance while limiting its footprint. Typical approaches to reach this goal rely either on manual procedures, which cannot guarantee the optimality of the final design, or upon Neural Architecture Search algorithms to automatise the process, at the expenses of extremely high computational time. This paper provides a solution for the fast identification of a neural network that maximises the model accuracy while preserving size and computational constraints typical of tiny devices. Our approach, named FreeREA, is a custom cell-based evolution NAS algorithm that exploits an optimised combination of training-free metrics to rank architectures during the search, thus without need of model training. Our experiments, carried out on the common benchmarks NAS-Bench-101 and NATS-Bench, demonstrate that i) FreeREA is a fast, efficient, and effective search method for models automatic design; ii) it outperforms State of the Art training-based and training-free techniques in all the datasets and benchmarks considered, and iii) it can easily generalise to constrained scenarios, representing a competitive solution for fast Neural Architecture Search in generic constrained applications. The code is available at <https://github.com/NiccoloCavagnero/FreeREA>.

1. Introduction

In recent years, we observed a dramatic expansion of the impact of Machine Learning in several fields, such as education [23], industry [14], and healthcare [24]. However, this groundbreaking technological advancement suffers of a major flaw related to the significant impact on the energy consumption and, therefore, to climate change. As an example, a recent study demonstrated how the carbon footprint associated to a single training of common Natural Language Processing models is more than twice the impact of a standard American citizen's daily life for one year [28, 3]. When Neural Architecture Search (NAS) algorithms are used to search for the best model, this amount is increased of a factor ten [28, 3]. These observations strongly call for a reduction of the impact that Machine Learning models have in terms of required resources. For this reason, it is likely that current and future trends of research will encompass the development of models that can run on very tiny devices. This represents a significant paradigmatic shift for Machine Learning research. Indeed, when the hardware resources are limited, the designer should pay much more attention on how these resources are used by the model. To do this, Neural Architecture Search [9] appears to be the best option to obtain models that can fully exploit the available budget, with minimum waste of time for human experts.

However, standard NAS approaches are based on algorithms that - during the search - need to train and test all the candidates to assess their performance. Therefore, NAS pipelines are usually very expensive in terms of computational time and resources, making them hardly exploitable by standard Machine Learning users, with no access to super-computing resources. As an additional drawback, it is worth mentioning that the energy consumption of a standard training-based NAS algorithm would make the effort of designing tiny models vain [28, 3].

For these reasons, an effective development of tiny neural networks on large scale requires the implementation of efficient NAS algorithms that do not need train all the can-

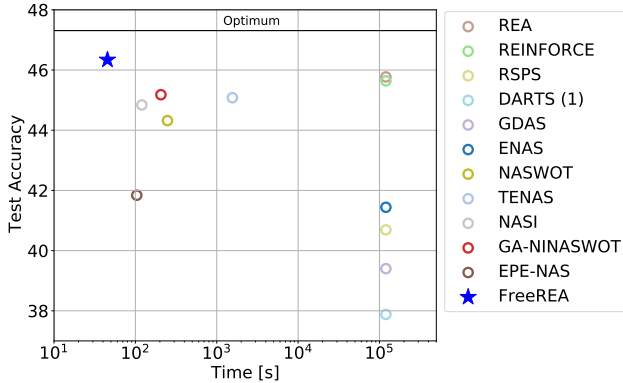


Figure 1: Average test accuracy vs time [s] on ImageNet16-120 of NATS-Bench [5]. X-axis is in log-scale. Accuracy and times for training-based methods are from the original NATS-Bench paper [5]. Similar plots can be drawn for the other datasets, and are here omitted for sake of space.

didates. To solve this problem, the research community has recently been focusing on the adoption of proxies to score the architectures at initialisation [21, 2, 1], thus avoiding the heavy and slow training phase and achieving consistent speedups even with limited hardware resources. However, training-free NAS algorithms proposed so far are not competitive with training-based approaches and typically also do not consider the opportunity to include hardware-driven constraints such as model footprint and FLOPs.

This paper bridges this gap and provides a custom constrained evolutionist method, named FreeREA, to exploit an optimal combination of metrics that serve as proxy of model accuracy. We demonstrate how the metrics we use work as a valid replacement for model training, enabling the possibility to identify automatically highly performing models in just a few minutes of search. We tested our solution on two popular benchmarks for NAS of tiny models, and we demonstrated how our approach outperforms State of the Art methods for the constraint-free case. We also performed, for the first time for a training-free method, experiments in a constrained scenario, setting a very competitive baseline for future works.

In conclusion, this paper advances the State of the Art with the following contributions:

1. an optimised combination of training-free metrics for models ranking, that can completely substitute the training phase in NAS algorithms;
2. an improved evolutionist search algorithm that fully exploits the ranking strategy to identify very accurate models in few minutes of search;
3. a number of experiments to demonstrate how training-free methods can effectively replace training-based

NAS approaches, even in hardware-constrained scenarios.

2. Related Work

Neural Architecture Search was firstly introduced in [37], where a Reinforcement Learning approach, REINFORCE [33], with an LSTM [11] controller was employed to generate high-performance neural networks. The whole search exploited more than 800 GPUs for 28 days, for a total of 22400 GPU-hours, requiring the partial (35 epochs) training of more than 12k different architectures. After this groundbreaking work, the research community focused on the development of more efficient methods [22, 6, 25, 17] with the aim of mitigating the huge computational requirements of a NAS algorithm.

The first attempts in this direction focused on Parameter Sharing [22, 6, 17], where newly found architectures inherit their weights from previously discovered ones. While [22] still adopts a Reinforcement Learning approach, two other search paradigms emerged in NAS literature: differentiable search and evolution. The first ones [17, 6] move from the purpose to make the whole search differentiable so that it can be optimised by gradient descent algorithms, leading to significant speedups w.r.t. original methods [37]. On the other hand, evolution-based strategies are very simple to implement, and naturally allows parameter inheritance from the parents, but suffered of lower performance [18, 26]. In REA [25], the authors implemented a regularised Tournament Selection search algorithm, which proven to be the first evolution-based NAS able to outperform human-crafted architectures. Still, these methods were not able to completely solve the core issue, and thousands of different candidates still needed to be trained and evaluated.

For this reason, the research community has recently focused on the adoption of metrics to score the different architectures at initialisation, completely avoiding the training and evaluation phases, which constitute the true bottleneck in NAS algorithms. The first metric proposed was Linear Regions (NASWOT) [21], which measures the expressivity power of an architecture and allows the search of competitive candidates in few seconds. However, these networks were still less accurate than the ones found by standard NAS algorithms. Therefore, the use of this approach was limited to a mere initialisation for classic search. Later, GA-NINASWOT [34] adopted Linear Regions together with a genetic algorithm, while EPE-NAS [20] proposed an improvement over the base metrics. Yet, none of these demonstrated to be competitive with training-based methods. Taking inspiration from NASWOT [21], in [2] the authors combined a variation of Linear Regions with the Neural Tangent Kernel (NTK) metric [12], to encode in the search also the trainability of the candidates. Due to the complexity of NTK computation, NASI [27] proposed an

approximation of the metric obtaining significant speedups. For a detailed analysis on these and other training-free metrics, the interested reader can refer to [1].

3. Method

In this paper, we introduce FreeREA, an innovative NAS method that leverages on training-free metrics to deliver high performance architectures satisfying constraints on FLOPs and number of parameters. Indeed, in previous works the training-free phase could not compete with standard search approaches, and was therefore usually exploited only to initialise the population of a classic NAS algorithm [21]. In our implementation, instead, we show that a proper selection of metrics, combined with an optimised evolutionist search, can yield accurate models with a significantly lower search time (up to four orders of magnitude lower than standard NAS).

3.1. Metrics

As anticipated in the previous section, training-free algorithms heavily rely on metrics that serve as accurate proxy of the model performance. Unfortunately, these do not exactly correlate with the test accuracy, and therefore it is important to consider a combination of multiple metrics, possibly accounting for different properties such as trainability and expressivity. With this aim, we performed preliminary experiments on the NATS-Bench dataset to assess the reliability of a selection of metrics as surrogate of the test accuracy. Details on the correlation of different metrics are given in Tab. 1. This analysis demonstrates that the most performing metrics are a modified version of Synflow [30], which we named LogSynflow, and the Linear Regions [21]. In addition, we complemented these measures with the number of skipped layers to favour more trainable architectures. To combine the different metrics, we sum the normalised scores of the single metrics, which leads to better empirical results w.r.t. a standard cumulative ranking score [2]. The rationale behind this choice is that a model that is both trainable and expressive is more likely to show high performances [2]. In our implementation, the fitness function f for a given model i can be expressed as:

$$f_i = \frac{LS_i}{\max_{j \in J} LS_j} + \frac{LR_i}{\max_{j \in J} LR_j} + \frac{Skip_i}{\max_{j \in J} Skip_j} \quad (1)$$

where LS stands for LogSynflow, LR for Linear Regions, $Skip$ for Skipped Layers and J is the set of the explored networks. An ablation study on the three terms composing the function can be found in Tab. 1 of the supplementary material. In addition, to increase the robustness of the measure, we compute both *Linear Regions* and *LogSynflow*

three times, with different initialisation, and take the average over the three runs as the final metric. In the following sections we discuss the implementation of the metrics we considered in this work, while we refer the interested reader also to [1] for further details.

3.1.1 Linear Regions

Linear Regions was introduced in [21] with the purpose of measuring the expressiveness of an architecture at initialisation. Indeed, when one training sample is forwarded into a ReLU network, the activation values divide the tensors in active (positive values), and inactive (negative values) regions, generating binary masks. Given an input space of fixed dimension, the number of continuous regions of the input space that are mapped to different network activations provides a measure of the model capabilities to discriminate between different input values [21]. Indeed, when two different input items result in similar masks, the model cannot easily distinguish between them. As a consequence, an expressive model is capable to map relatively close values in input space to different activation tensors. To quantitatively evaluate this measure, following the implementation of [21] we build and forward a mini-batch of data composed of 64 samples into the network, and collect the resulting binary masks. Then, we evaluate the differences in the activation patterns via Hamming distance. Once a Kernel matrix K_H is built from the Hamming distances, the score s is computed as:

$$s = \log |\det(K_H)| \quad (2)$$

The higher the score, the higher the expressiveness of the network and its ability to distinguish the samples. As depicted in Fig. 2-A, models with high values of Linear Regions score tend to present high test accuracy, as also confirmed by a positive correlation between the two entities (see Tab. 1). Therefore, maximising the Linear Regions score yields models that are more expressive and, likely, capable of higher test accuracy.

3.1.2 LogSynflow

Synflow was initially introduced in [30] as a pruning metric for the selection of weights to be removed to compress an architecture while maximally preserving the performances. Its original implementation locally evaluated the relevance of single weights, and was then extended by [1] to score a whole architecture by summing up the contributions of the single weights. Formally, the global metric can be defined as the scalar product between the weight vector θ and the gradient vector $\frac{\partial \mathcal{R}}{\partial \theta}$:

$$S(\Theta) = \theta \cdot \frac{\partial \mathcal{R}}{\partial \theta} \quad (3)$$

Table 1: Kendall and Spearman correlation between training-free metrics and the test accuracy, evaluated on the three datasets of NATS-Bench [5]. Each metric has been computed three times with different initialisations and the average is taken as final score.

Metric	CIFAR10		CIFAR100		ImageNet16-120	
	Kendall	Spearman	Kendall	Spearman	Kendall	Spearman
NTK [12]	-0.33	-0.49	-0.30	-0.45	-0.39	-0.56
Snip [15]	0.45	0.61	0.47	0.62	0.41	0.55
Fisher [31]	0.39	0.54	0.40	0.55	0.36	0.48
Grasp [32]	0.28	0.41	0.35	0.50	0.35	0.49
PathNorm [13]	0.41	0.59	0.42	0.60	0.45	0.63
LinearRegions [21]	0.61	0.79	0.62	0.81	0.60	0.78
Synflow [30]	0.57	0.77	0.56	0.76	0.56	0.75
LogSynflow (ours)	0.61	0.81	0.60	0.79	0.59	0.78




Figure 2: Relationship between test accuracy and the two metrics considered in this work for all the models available in NATS-Bench: Linear Regions in A, *LogSynflow* in B. Accuracy is provided for the ImageNet16-120 dataset. Similar plots can be drawn for the other datasets and are here omitted for sake of space.

To compute the metric, first the network is initialised with the absolute value of its original weights, so that $\theta_i \geq 0$. Then an all-ones tensor is forwarded into the network, and the output is backwarded to compute the gradients. It is worth noticing that, to compute Synflow, Batch Normalisation layers must be suppressed since they interfere with the gradient flow. However, this is likely to result in gradient explosion, even in relatively small architectures. As a consequence, the metric neglects the significance of the weight values because the term associated to the gradients may be several order of magnitudes higher than the corresponding weight. To address this issue, we propose *LogSynflow* which scales down the gradients with a logarithmic function before summing up the contributions of each network weight:

$$S(\Theta) = \theta \cdot \log\left(\frac{\partial \mathcal{R}}{\partial \theta} + 1\right) \quad (4)$$

To verify that our modification provides a more expressive metric, we computed the numerical values of Synflow and *LogSynflow* over both the benchmark datasets we considered [36, 5] and calculated the correlation of the result with

Table 2: Kendall and Spearman correlation between the Synflow/*LogSynflow* and the test accuracy, evaluated on NAS-Bench-101 [36]. Numerical values suggest that *LogSynflow* is consistently a more accurate proxy of the test accuracy w.r.t. the original definition.

Metric	Kendall	Spearman
Synflow [30]	0.25	0.37
LogSynflow (ours)	0.31	0.45

the test accuracy (see Tab. 1 and 2). We observed that our implementation consistently increases the correlation with the test accuracy w.r.t. the standard implementation and, hereinafter we will therefore employ *LogSynflow* in place of the original definition. Furthermore, the positive correlation between test accuracy and *LogSynflow* score, also evident from Fig. 2-B, suggests that searching for high values of *LogSynflow* scores provides architectures that are likely more accurate in test.

3.1.3 # Skipped Layers

In recent literature for training-free NAS, Neural Tangent Kernel (NTK) [12] is usually employed as a proxy for the trainability of models [2, 4]. However, NTK computation is consistently heavier than other training-free metrics [12] and is characterised by a low correlation with the test accuracy (see Tab. 1). In order to favour more trainable architectures, we instead adopted as score the number of skipped layers divided by the total number of skip connections in a cell, i.e. $s = \frac{\# \text{Skipped Layers}}{\# \text{Skip Connections}}$. Indeed, since their introduction in [10], which laid the foundation of modern Deep Learning, all the State of the Art networks (e.g. EfficientNet [29], ResNeXt [35], ViT [8], and ConvNext [19] among the others) present skip connections around convolutional cells in their architecture. These have proven to be crucial for

the training of really deep networks and single skip connection around the cell effectively mitigates the vanishing of the gradients and allows the information to be backpropagated more easily through the network. The rationale behind this metric is to favour such configuration by encouraging few connections with a large number of skipped layers, i.e. long range connections within a cell. Indeed, the maximum possible value is achieved when there is a single skip connection from input node to output node.

3.2. Search Algorithm

To search for the optimal model, we exploited an evolution approach. In particular, an improved Tournament Selection with Ageing (REA) [25] was implemented. Tournament Selection is initialised with a random population of size N . Each individual is characterised by a genotype that encodes the structure of the network. For each iteration, a subset of the surviving population of size n is randomly sampled, and the best individual in the sample is selected as parent for reproduction. The exemplar is then mutated to generate a child which is included into the existing population. Classic Tournament Selection keeps the population size constant by discarding the less fitting individuals, while REA [25] removes the oldest model, thus biasing the tournament in favour of younger exemplars and therefore enhancing the exploration of the search space. With the aim of further improving the search capability of REA, we implemented a variation of the original algorithm which samples two parents at each step. The two parents are independently mutated to generate two children, while a third child is generated through a crossover operation. Crossover is implemented by uniformly sampling a gene from one of the parents. This increases the exploration capability of the search algorithm since, while mutation is inherently local, the combination of two different genotypes can lead to more varied individuals which significantly differ from their parents [18]. We empirically prove the effectiveness of this strategy w.r.t. the base algorithm in the following experiments. To maintain the population size constant, after killing the oldest individual we only keep the top N ones. All gene choices in both mutation and crossover are uniformly sampled.

Constrained search is performed by imposing limits to the FLOPs and to the Number of Parameters for the models considered during the search, because these are not dependent on the specific deployment architecture. In case of constraints, only feasible individuals are added to the population and we keep generating offspring until a feasible child is discovered. In all the experiments of this work, N and n are set to 25 and 5 respectively, if not explicitly reported differently. The results for other values of N and n are reported in Tab. 2 of the supplementary material.

4. Benchmarks

To help the comparison between our method and State of the Art approaches, and to speed-up the experiments, we leverage on two different benchmarks that include pre-trained models. This enables an easy test of the search algorithms without the need of consistent computational resources, and eliminates the intrinsic variability due to different training pipelines, thus increasing the reproducibility of results.

4.1. NAS-Bench-101

NAS-Bench-101 [36] is the first public dataset built for this purpose. It contains 423k unique convolutional architectures trained and evaluated on CIFAR10. Each network is characterised by a cell, which is repeated with interleaved downsampling operators to form the full structure. Cells are represented by a Direct Acyclic Graph (DAG), with each node corresponding to an operation. There are three possible operators, 1x1 convolution, 3x3 convolution and 3x3 max pooling. The number of nodes in a cell is at most seven and the number of edges is at most nine.

4.2. NATS-Bench

NATS-Bench [5] is an extension of NAS-Bench-201 [7] containing architectures trained and evaluated on three different Computer Vision datasets, CIFAR10, CIFAR100 and ImageNet16-120 (listed from the simplest to the more complex). These are complemented with additional information such as FLOPs and latency, which is particularly relevant for constrained scenarios. However, the number of architectures in this benchmark is limited to 15625, with about 6k unique cells. The possible operators are 1x1 convolution, 3x3 convolution, 3x3 average pooling, skip connection and zeroise. In this case, edges in the DAG represent the operators, while each node stands for the sum of all the feature maps transformed by the edges pointing to the node. The number of edges is six, while the number of nodes is set to four.

5. Experiments

Given the lack of experiments testing the efficiency of existing constrained NAS algorithms on the benchmarks mentioned above, we evaluate, as a first step, FreeREA against State of the Art training-based and training-free techniques, without imposing any constraint on both NAS-Bench-101 and NATS-Bench. Then, to also prove its effectiveness in a constrained scenario, we directly compare the architectures discovered by our approach on NATS-Bench against the optimal architectures within the imposed limits. We stop FreeREA after 45 seconds of execution time on NATS-Bench and 12 minutes on NAS-Bench-101. All experiments were performed on a single RTX 3080 GPU. The




Figure 3: Mean and standard deviation of test accuracy achieved on the ImageNet16-120 dataset of NATS-Bench [5]. Orange shaded area collects training-based methods, while green shaded area collects training-free methods. Similar plots can be drawn for the other datasets, and are here omitted for sake of space.

results for training-based approaches on NATS-Bench are extracted from the original NATS-Bench paper [5], while the ones for training-free methods are taken from the respective papers.

5.1. Unconstrained Scenario

As a first analysis, we considered the NATS-Bench benchmark. As reported in Tab. 3 and depicted in Fig. 3, our experiments verified that among the State of the Art approaches REA [25] is the algorithm that discovers the best architectures. Our solution, however, is able to outperform the previous best result for all the datasets considered in the benchmark, with a gain that increases with the complexity of the task. More specifically, we obtained that FreeREA yields a model with 94.36% of test accuracy for CIFAR10, 73.51% on CIFAR100 and 46.34% for ImageNet16-120, only 0.01%, and 0.97% less than the optimum for CIFAR10 and ImageNet16-120, while we hit the optimum in CIFAR100.

Notably, the performance demonstrated by our approach significantly outperforms all the methods (training-based and training-free) considered in this paper, which is especially surprising considering the comparison with training-based approaches. Indeed, while these provide a solution in a (capped) time of $2e4$, $4e4$ and $1.2e5$ seconds, FreeREA iterates for approximately 45 seconds. It is also relevant to notice that our approach is characterised by a significantly smaller variation across different runs (see Tab. 3 and Fig. 3). Indeed, the largest variance of the test accuracy we observed (for the CIFAR100 case) is as low as 0.05, while REA shows significantly higher values (0.31 for CIFAR-10, 0.84 for CIFAR100 and 0.80 for ImageNet16-120). Similar considerations can be made for all the other methods and are here omitted to improve readability of the

manuscript. This suggests that our convergence to the optimal model is affected by a smaller uncertainty w.r.t. all the other approaches, thus ideally vouching for a lower number of search runs to obtain the final model.

It is worth remarking that, while for training-free methods this potentially represents a gain in the order of minutes, for training-based approaches the amount is of several hours or even days. We argue that the increase of average test accuracy and the reduction of variance are a consequence of both the optimal combination of metrics we considered, and of the modifications to the search algorithm we implemented. To verify this claim, we performed experiments in which we used the vanilla REA search algorithm and ranked the models considering our metrics rather than through the validation accuracy after training. We named this implementation FreeREA⁻. Interestingly, the results achieved with FreeREA⁻ (see Tab. 3) are pretty similar to FreeREA, albeit characterised by a larger variability in certain conditions. Such similarity vouches for the effectiveness of the metrics selected, while it does not provide evidence on the role played by the search algorithm. It is reasonable to expect that this is due to the relative small dimension of the benchmark, while differences should be more evident for benchmarks of higher complexity (such as NAS-Bench-101).

In Fig. 1 we summarise the average test accuracy achieved by our approach and all the competitors versus the execution time. Our findings demonstrate that FreeREA represents the best trade-off between computation time and performance. Of note, all the training-based methods, reported on the right of the plot, are limited by the same search time (see Tab. 3). To verify the performances of our method on a more challenging scenario, we also tested FreeREA on the more complex search space of NAS-Bench-101 against the other training-free methods and the best-in-class for the training-based (i.e. REA). In this experiment, we selected as population and tournament size for REA 50 and 5 respectively. These values are decreased w.r.t. the original implementation provided in NAS-Bench-101 [36] because we allow ≈ 24 hours of search, and only the initialisation of 100 candidates requires ≈ 18 hours. We adopted the same configuration for FreeREA⁻ for comparison purposes, while the population size is instead kept as default equal to 25 for FreeREA. Numerical results, collected in Tab. 4, demonstrate how FreeREA still outperforms all the competitors while searching for approximately twelve minutes.

Interestingly, there is an appreciable difference between FreeREA and FreeREA⁻ in terms of average test accuracy, which supports our claim that the search algorithm of our implementation is more efficient than the vanilla REA. Indeed, the modifications we introduced, discussed in the previous section, play a fundamental role in giving better exploration capabilities to the search algorithm. In other

Table 3: Test accuracy and time [s] on NATS-Bench. Results for training-based methods are taken from [5].

Algorithm	CIFAR10		CIFAR100		ImageNet16-120	
	Accuracy	Time	Accuracy	Time	Accuracy	Time
Training-based						
REA [25]	94.02 ± 0.31	2e4	72.23 ± 0.84	4e4	45.77 ± 0.80	1.2e5
REINFORCE [37]	93.90 ± 0.26	2e4	71.86 ± 0.89	4e4	45.64 ± 0.78	1.2e5
RSPS [16]	91.05 ± 0.66	2e4	68.27 ± 0.72	4e4	40.69 ± 0.36	1.2e5
DARTS (1) [17]	59.84 ± 7.84	2e4	61.26 ± 4.43	4e4	37.88 ± 2.91	1.2e5
DARTS (2) [17]	65.38 ± 7.84	2e4	60.49 ± 4.95	4e4	36.79 ± 7.59	1.2e5
GDAS [6]	93.23 ± 0.58	2e4	68.17 ± 2.50	4e4	39.40 ± 0.00	1.2e5
ENAS [22]	93.76 ± 0.00	2e4	70.67 ± 0.62	4e4	41.44 ± 0.00	1.2e5
Training-free						
NASWOT (1000) [21]	93.10 ± 0.31	248	69.10 ± 1.61	248	45.08 ± 1.55	248
TENAS [2]	93.9 ± 0.47	1558	71.24 ± 0.56	1558	42.38 ± 0.46	1558
NASI [27]	93.55 ± 0.10	120	71.20 ± 0.14	120	44.84 ± 1.41	120
GA-NINASWOT [34]	93.70 ± 0.63	206	71.57 ± 1.37	206	45.18 ± 2.05	206
EPE-NAS [20]	91.31 ± 1.69	104	69.58 ± 0.83	104	41.84 ± 2.06	104
FreeREA ⁻ (ours)	94.30 ± 0.02	45	73.30 ± 0.31	45	46.34 ± 0.00	45
FreeREA (ours)	94.36 ± 0.00	45	73.51 ± 0.05	45	46.34 ± 0.00	45
Optimum						
-	94.37	-	73.51	-	47.31	-

Table 4: Test accuracy and time [s] of training-free methods and REA on NAS-Bench-101 [36].

Algorithm	Accuracy	Time
Training-based		
REA [25]	93.39 ± 0.09	86752
Training-free		
NASWOT [21]	92.23 ± 8.90	114
FreeREA ⁻ (ours)	93.13 ± 1.16	723
FreeREA (ours)	93.80 ± 0.02	724
Optimum		
	94.31	-

words, by mutating two parents we better inspect the neighbourhood of highly performing models, while by means of crossover we also explore more distant regions of the search space. As a consequence, we expect that the difference in terms of test accuracy between FreeREA and FreeREA⁻ increases with the complexity of the search space and the exploration time. To prove this, we collected the average test accuracy and its variance for FreeREA, FreeREA⁻ and REA versus the computation time (see Fig. 4). It is also interesting to observe that FreeREA shows a larger variability in the first iterations (i.e. between 10¹ and 10² s), while

the output shows a significantly smaller uncertainty towards the end of the search. This behaviour is not matched by FreeREA⁻ and REA, and is a consequence of the modifications made to the search algorithm, which suggests that our implementation not only provides higher exploration capabilities, but also lower search uncertainty.

5.2. Constrained Scenario

As we discussed extensively in the introduction, it is our opinion that future trends of ML research will devote an increasing effort towards the development of architectures that can fit on constrained devices. In this scenario, albeit the role of models optimisation is fundamental to maximise the exploitation of the available hardware, we believe that there should not be explicit modifications to the search algorithm that depend on the target hardware. For this reason, we tested our approach in a constrained scenario without explicitly accounting for the target constraints, besides the introduction of limits to the exploration phase.

It is worth reporting that the introduction of those constraints is not straightforward, and good performances of the method in an unconstrained scenario do not necessarily correspond to analogous results in the constrained case. Indeed, FreeREA heavily relies on metrics that serve as proxy of the test accuracy, but the correlation between the two in the whole search space may be different than the correlation in a constrained one. Therefore, an efficient training-free




Figure 4: REA, FreeREA⁻ and FreeREA trajectories on NAS-Bench-101 [36]. Standard deviations are scaled down by a factor 10 for visualisation purposes. X-axis is in log-scale. Note that FreeREA shows a larger variability in the first iterations (i.e. between 10¹ and 10² s), and a lower uncertainty towards the end of the search. REA and FreeREA⁻, instead, present similar variances in all the phases of the search.

Table 5: Test accuracy of FreeREA on NATS-Bench in constrained scenarios. Regret is the difference between optimum and average test accuracy. The algorithm was stopped after 45 seconds.

FLOPs	Params	Accuracy	Optimum	Regret
CIFAR10				
1e8	8e5	94.12 ± 0.01	94.31	0.19
7e7	5e5	93.02 ± 0.13	93.75	0.73
4e7	3e5	91.30 ± 0.00	91.32	0.02
CIFAR100				
1e8	8e5	72.08 ± 0.01	72.43	0.35
7e7	5e5	71.07 ± 0.00	71.63	0.54
4e7	3e5	68.39 ± 0.00	68.48	0.09
ImageNet16-120				
1e8	8e5	46.08 ± 0.08	46.85	0.77
7e7	5e5	45.35 ± 0.04	45.95	0.65
4e7	3e5	40.73 ± 0.00	41.00	0.27

NAS algorithm should be able to search for a good model with and without limits imposed by the target hardware, ideally with no change to the metrics and the search algorithm used. Since, to the best of our knowledge, there are no experimental results in constrained scenarios for the considered benchmarks, to verify whether our implementation represents a good solution also in this case, we considered

three different thresholds for model FLOPs and number of parameters, namely (1e8, 8e5), (7e7, 5e5), and (4e7, 3e5), which stand for three levels of hardware constraints. We then tested FreeREA on the three datasets of NATS-Bench for all the constraints.

In absence of other competitors, in Tab. 5 we compare the results of our search with the best model available in the dataset that fits the constraints, which clearly represents an upper-bound for our problem. Interestingly, on the CIFAR datasets the average margin of improvement left (regret) is 0.32%, while on the more challenging ImageNet16-120, we under-perform the optimal model of 0.56% on average.

It is worth remarking that such results are coherent with the unconstrained scenario, where for the more challenging ImageNet16-120 the difference between the test accuracy of our best candidate and the optimal model is on average approximately 0.97%. Therefore, this suggests that our combination of metrics serves as a good surrogate of the test accuracy even when very hard constraints are imposed to the search, and that these results may serve as a strong baseline for future research on this topic.

6. Conclusions

In this paper, we presented FreeREA, a new fast, efficient and accurate training-free NAS for tiny models. Moving from State of the Art results, we selected a suitable combination of training-free metrics that serve as an accurate proxy of the model test accuracy, and we used this index to rank different networks during the search. The latter is implemented following an evolution-based policy, yielding very competitive results in a fraction of the time required by competitor approaches. We tested our method on two different benchmarks, NATS-Bench [5] (with three datasets and $\approx 15k$ available models as search space), and NAS-Bench-101 [36] (with one dataset and $\approx 500k$ available models as search space). Interestingly, our results demonstrate that FreeREA not only is able to effectively perform Neural Architecture Search without training any candidate, but it also is the first training-free method which competes and outperforms State of the Art training-based methods, resulting in the current most efficient and accurate NAS algorithm on the considered benchmarks. Such advancement comes also with the relevant benefit that our search time is as low as twelve minutes in the worst case, i.e. up to four order of magnitude lower than State of the Art algorithms [25]. Additional experiments on constrained settings suggest that FreeREA can deliver accurate models even when the target hardware imposes limits to the search space that exclude some of the models. Motivated by these results, our future efforts will be devoted towards the generalisation of our method to other search spaces, and considering more challenging tasks.

References

- [1] Mohamed S Abdelfattah, Abhinav Mehrotra, Łukasz Dudziak, and Nicholas Donald Lane. Zero-cost proxies for lightweight nas. In *International Conference on Learning Representations*, 2020.
- [2] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. In *International Conference on Learning Representations (ICLR)*, 2021.
- [3] Payal Dhar. The carbon impact of artificial intelligence. *Nature Machine Intelligence*, 2(8):423–425, 2020.
- [4] Tu Do and Ngoc Hoang Luong. Training-free multi-objective evolutionary neural architecture search via neural tangent kernel and number of linear regions. In *International Conference on Neural Information Processing*, pages 335–347. Springer, 2021.
- [5] Xuanyi Dong, Lu Liu, Katarzyna Musial, and Bogdan Gabrys. Nats-bench: Benchmarking nas algorithms for architecture topology and size. *IEEE transactions on pattern analysis and machine intelligence*, 2021.
- [6] Xuanyi Dong and Yi Yang. Searching for a robust neural architecture in four gpu hours. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1761–1770, 2019.
- [7] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *International Conference on Learning Representations*, 2020.
- [8] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- [9] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *The Journal of Machine Learning Research*, 20(1):1997–2017, 2019.
- [10] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [11] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [12] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: Convergence and generalization in neural networks. *Advances in neural information processing systems*, 31, 2018.
- [13] Yiding Jiang, Behnam Neyshabur, Hossein Mobahi, Dilip Krishnan, and Samy Bengio. Fantastic generalization measures and where to find them. In *International Conference on Learning Representations*, 2019.
- [14] Jay Lee, Hossein Davari, Jaskaran Singh, and Vibhor Pandhare. Industrial artificial intelligence for industry 4.0-based manufacturing systems. *Manufacturing letters*, 18:20–23, 2018.
- [15] Namhoon Lee, Thalaisyasingam Ajanthan, and Philip H. S. Torr. SNIP: single-shot network pruning based on connection sensitivity. *CoRR*, abs/1810.02340, 2018.
- [16] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *Uncertainty in artificial intelligence*, pages 367–377. PMLR, 2020.
- [17] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations*, 2018.
- [18] Yuqiao Liu, Yanan Sun, Bing Xue, Mengjie Zhang, Gary G Yen, and Kay Chen Tan. A survey on evolutionary neural architecture search. *IEEE transactions on neural networks and learning systems*, 2021.
- [19] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. *CoRR*, abs/2201.03545, 2022.
- [20] Vasco Lopes, Saeid Alirezazadeh, and Luís A Alexandre. Epe-nas: Efficient performance estimation without training for neural architecture search. In *International Conference on Artificial Neural Networks*, pages 552–563. Springer, 2021.
- [21] Joe Mellor, Jack Turner, Amos Storkey, and Elliot J Crowley. Neural architecture search without training. In *International Conference on Machine Learning*, pages 7588–7598. PMLR, 2021.
- [22] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *International conference on machine learning*, pages 4095–4104. PMLR, 2018.
- [23] Stefan AD Popenici and Sharon Kerr. Exploring the impact of artificial intelligence on teaching and learning in higher education. *Research and Practice in Technology Enhanced Learning*, 12(1):1–13, 2017.
- [24] Alvin Rajkomar, Jeffrey Dean, and Isaac Kohane. Machine learning in medicine. *New England Journal of Medicine*, 380(14):1347–1358, 2019.
- [25] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, volume 33, pages 4780–4789, 2019.
- [26] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V Le, and Alexey Kurakin. Large-scale evolution of image classifiers. In *International Conference on Machine Learning*, pages 2902–2911. PMLR, 2017.
- [27] Yao Shu, Shaofeng Cai, Zhongxiang Dai, Beng Chin Ooi, and Bryan Kian Hsiang Low. Nasi: Label-and data-agnostic neural architecture search at initialization. In *International Conference on Learning Representations*, 2021.
- [28] Emma Strubell, Ananya Ganesh, and Andrew McCallum. Energy and policy considerations for deep learning in nlp. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3645–3650, 2019.
- [29] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, pages 6105–6114. PMLR, 2019.
- [30] Hidenori Tanaka, Daniel Kunin, Daniel L Yamins, and Surya Ganguli. Pruning neural networks without any data by iteratively conserving synaptic flow. *Advances in Neural Information Processing Systems*, 33:6377–6389, 2020.

- [31] Jack Turner, Elliot J. Crowley, Gavin Gray, Amos J. Storkey, and Michael F. P. O’Boyle. Blockswap: Fisher-guided block substitution for network compression. *CoRR*, abs/1906.04113, 2019.
- [32] Chaoqi Wang, Guodong Zhang, and Roger Baker Grosse. Picking winning tickets before training by preserving gradient flow. *CoRR*, abs/2002.07376, 2020.
- [33] Ronald J Williams. Simple statistical gradient-following algorithms for connectionist reinforcement learning. *Machine learning*, 8(3):229–256, 1992.
- [34] Meng-Ting Wu, Hung-I Lin, and Chun-Wei Tsai. A training-free genetic neural architecture search. In *Proceedings of the 2021 ACM International Conference on Intelligent Computing and its Emerging Applications*, pages 65–70, 2021.
- [35] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1492–1500, 2017.
- [36] 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*, pages 7105–7114. PMLR, 2019.
- [37] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. *CoRR*, abs/1611.01578, 2016.