

Demystifying the Neural Tangent Kernel from a Practical Perspective: Can it be trusted for Neural Architecture Search without training?

Jisoo Mok^{1*} Byunggook Na¹ Ji-Hoon Kim^{2,3†} Dongyoon Han^{2†} Sungroh Yoon^{1,4†}
¹ Department of ECE, Seoul National University ² NAVER AI Lab ³ NAVER CLOVA
⁴ AIIS, ASRI, INMC, ISRC, and Interdisciplinary Program in AI, Seoul National University

Abstract

In Neural Architecture Search (NAS), reducing the cost of architecture evaluation remains one of the most crucial challenges. Among a plethora of efforts to bypass training of each candidate architecture to convergence for evaluation, the Neural Tangent Kernel (NTK) is emerging as a promising theoretical framework that can be utilized to estimate the performance of a neural architecture at initialization. In this work, we revisit several at-initialization metrics that can be derived from the NTK and reveal their key shortcomings. Then, through the empirical analysis of the time evolution of NTK, we deduce that modern neural architectures exhibit highly non-linear characteristics, making the NTK-based metrics incapable of reliably estimating the performance of an architecture without some amount of training. To take such non-linear characteristics into account, we introduce Label-Gradient Alignment (LGA), a novel NTK-based metric whose inherent formulation allows it to capture the large amount of non-linear advantage present in modern neural architectures. With minimal amount of training, LGA obtains a meaningful level of rank correlation with the final test accuracy of an architecture. Lastly, we demonstrate that LGA, complemented with few epochs of training, successfully guides existing search algorithms to achieve competitive search performances with significantly less search cost. The code is available at: <https://github.com/nutellamok/DemystifyingNTK>.

1. Introduction

Deep Neural Networks (DNNs) continue to produce impressive results in a wide variety of domains and applications. The remarkable success of DNNs is due in no small part to the development of novel neural architectures, all of which used to be designed manually by machine learning engineers by testing a number of architectural design

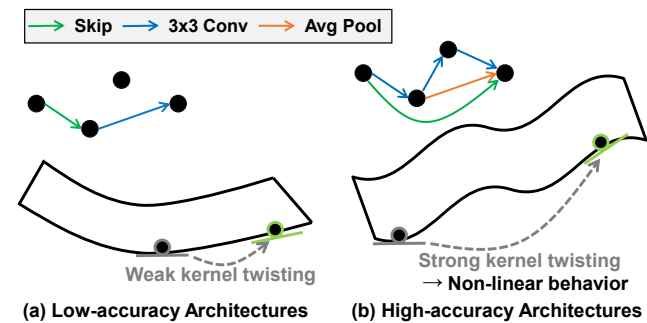


Figure 1. A conceptualized view of how (a) the NTK of a low-accuracy architecture and (b) that of a high-accuracy architecture evolve during training (gray \rightarrow green). Black planes denote the function space realization of weight parameters. On the top left corners of (a) and (b), an example of a low- and high-accuracy architecture is provided. Unlike a low-accuracy architecture, a high-accuracy architecture equipped with a large amount of non-linear advantage experiences strong kernel twisting, such that the principal components of the NTK become more aligned with target labels. In Figure 2, we illustrate that LGA captures this difference in the architectures residing in two polar accuracy regimes.

choices. To remedy this issue, Neural Architecture Search (NAS), a sub-field of automated machine learning, has emerged as a feasible alternative to hand-designing neural architectures [13].

Although the architectures derived by NAS are beginning to outperform hand-designed architectures, the tremendous computational cost required to execute NAS makes its immediate deployment rather challenging [39, 51, 52]. The majority of the search cost in NAS is induced by the need to train each candidate architecture to convergence for evaluation [35]. In more recently proposed NAS algorithms, the individual training of candidate architectures is circumvented by a weight-sharing strategy [5, 8, 12, 29, 34, 35, 46, 48, 50]. With weight sharing, the computational cost of NAS is reduced by orders of magnitude: from tens of thousands of GPU hours to < 1 GPU day. Unfortunately, NAS algorithms that rely on weight-sharing are known to experience an optimization gap between the performance of an architecture approximated through weight-sharing and

*Work done while interning at NAVER (magicshop1118@snu.ac.kr)

†Corresponding Authors

its stand-alone performance [43].

Another line of research that aims to accelerate the architecture evaluation process focuses on developing a performance predictor with as few architecture-accuracy pairs as possible. Minimizing the mean squared error between the predicted and ground-truth accuracy is the most straightforward way of training such a performance predictor because the problem of performance prediction can naturally be interpreted as a regression task [9, 26, 40]. The family of architecture comparators replaces the deterministic evaluation of neural architectures with a relativistic approach that compares two architectures and determines which one yields better performance [6, 45]. Apart from weight sharing and performance prediction, some works propose more general proxies for architecture evaluation [1, 27, 30, 31, 49]. White *et al.* [41] offer a comprehensive survey of performance predictors in NAS, and in the Appendix, we discuss related works in more detail.

The need to explicitly measure the test accuracy of architectures or train a performance predictor arises from our lack of theoretical understanding regarding how and what DNNs learn. Among diverse deep learning theories that claim to offer a quantifiable bound on the learning capacity of DNNs, the Neural Tangent Kernel (NTK) framework [21] is garnering a particular amount of attention. Based on the observation that DNNs of infinite widths are equivalent to Gaussian processes, the NTK framework proposes to characterize DNNs as kernel machines [21]. At the core of the NTK framework lies an assumption that the NTK computed from an infinitely wide DNN parameterized with randomly initialized weights remains unchanged throughout training. Thus, the NTK framework suggests that the training dynamics of such a DNN can be fully characterized by the NTK at initialization. Motivated by the solid theoretical ground on which the NTK framework is built, in the field of NAS, NTK-based metrics [4, 44], measured at initialization, have been proposed as an attractive alternative to computing the test accuracy directly.

In this paper, we aim to rigorously evaluate how trustworthy of a theory the NTK framework is in the context of NAS by conducting a series of empirical investigations. To begin with, we revisit previously-proposed metrics that spawn out of the NTK framework: Frobenius Norm (F-Norm), Mean [44], and Negative Condition Number (NCN) [4]. In order to assess whether the NTK-based metrics, computed with randomly initialized weights, are truly applicable to NAS, we test them on various NAS benchmarks by measuring Kendall’s Tau rank correlation with the test accuracy at convergence. Our experimental results show that the predictive performance of the NTK-based metrics obtained at initialization fluctuates significantly from one benchmark to another. A more up-close study of how their predictive ability changes according to the evaluated archi-

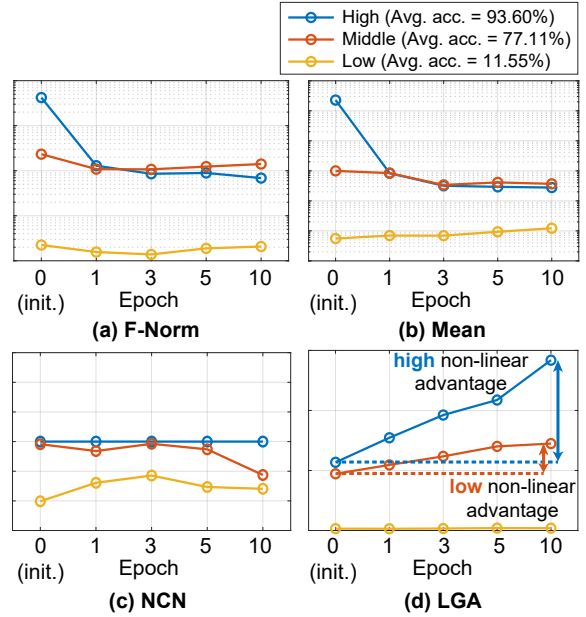


Figure 2. How the NTK-based metrics change in the early epochs for high-, mid-, and low-accuracy architectures. For each accuracy range, 200 architectures are randomly sampled from the NAS-Bench-201 search space, and the averaged test accuracy per architecture set is included in the legend.

itecture pool and the weight initialization scheme uncovers additional pitfalls of the NTK framework. Comprehensively, our results seem to indicate that the NTK at initialization does not exhibit a substantial level of reliability for architecture selection.

Empirically analyzing the time evolution of the NTK reveals that in modern neural architectures that constitute NAS search spaces, the NTK evolves in a highly non-linear manner. As a result, modern neural architectures tend to exhibit a large amount of non-linear advantage [15, 17, 32]. Figure 1 depicts on a high-level how the NTK rotates and evolves during the training process. Inspired by this observation, we introduce Label-Gradient Alignment (LGA), a novel NTK-based metric whose mathematical formulation allows it to coherently capture the non-linear characteristics of modern neural architectures. After only few epochs of training, LGA shows a considerable level of rank correlation with the test accuracy at convergence. As illustrated in Figure 2, delving deeper into how each metric changes throughout training consolidates that LGA is the only metric that can accurately reflect the non-linear behavior of modern neural architectures. Lastly, we conduct random [25] and evolutionary search [38] algorithms solely by using post-training LGA to demonstrate that it can be used to accelerate existing search algorithms.

Our main contributions can be summarized as follows:

- We rigorously assess the predictive ability of previous NTK-based metrics on various NAS benchmarks

and under different hyperparameter settings. Our results imply that the NTK at initialization may be insufficient for architecture selection in NAS.

- In order to understand the cause of the aforementioned limitation of the current NTK framework, we analyze the time evolution of the NTK and reveal that a considerable amount of non-linear advantage is present in modern neural architectures considered in NAS.
- We introduce LGA, a novel NTK-based metric that can reflect the change in NTK with respect to the target function. Integrating LGA after minimal amount of training with existing search algorithms yields a competitive search performance to state-of-the-art NAS algorithms, while noticeably reducing the search cost.

2. Neural Tangent Kernel

This section provides an overview of the NTK framework and the NTK-metrics that will be subject to a series of investigations in the later sections. In Section 2.1, we introduce the concept of the NTK, and in Section 2.2, we briefly review previously proposed NTK-based metrics.

2.1. Preliminaries

Let us define a DNN as a function $f_\theta : \mathbb{R}^d \rightarrow \mathbb{R}$, where θ is the set of trainable weight parameters. Given the target dataset, $\mathcal{D} = \{(x_i, y_i)\}_{i=1}^N$, without loss of generality, the NTK framework focuses on a binary classification problem, whose objective is to minimize the squared loss, $L_{\mathcal{D}}(\theta) = \sum_{i=1}^N \|y_i - f_\theta(x_i)\|_2^2$. Here, $x_i \in \mathcal{X}$ and $y_i \in \mathcal{Y}$ denote image samples and the corresponding class labels, respectively. In a small neighborhood region around the randomly initialized weights θ_0 , a DNN can be linearly approximated through a first-order Taylor expansion [21]:

$$f_\theta(x) \approx \hat{f}_\theta(x; \theta) = f_{\theta_0}(x) + (\theta - \theta_0)^\top \nabla_\theta f_{\theta_0}(x), \quad (1)$$

where $\nabla_\theta f_{\theta_0}$ corresponds to the Jacobian of a DNN’s prediction, computed with respect to θ_0 . The obtained approximation \hat{f}_θ can be regarded as a linearized network that maps weight vectors to functions residing in a reproducible kernel Hilbert space (RKHS) $\mathcal{H} \subseteq L_2(\mathbb{R}^d)$, determined by the NTK at θ_0 [32]:

$$\Theta_{\theta_0}(x, x') = \langle \nabla_\theta f_{\theta_0}(x), \nabla_\theta f_{\theta_0}(x')^\top \rangle. \quad (2)$$

Note that the NTK is essentially the dot product of two gradient vectors, and is thus equivalent to the Gram matrix of per-sample gradients. Intuitively speaking, the NTK can be interpreted as a condensed representation of gradient values and gradient correlations. From a geometric perspective, gradient values influence the extent of gradient descent at each step, and gradient correlations determine the stochasticity of gradient directions [44].

It has recently been discovered that under an infinitesimal learning rate and certain types of initialization, in a DNN of infinite width, the approximation in Eq. (1) is exact, and the NTK remains constant throughout training. Therefore, provided that the above assumptions hold, several aspects of the NTK at initialization can be used to fully characterize the training dynamics of a DNN and estimate its generalization performance. Following this theoretical discovery, in NAS, Xu *et al.* [44] and Chen *et al.* [4] have proposed to score DNNs at initialization based on the metrics that spawn out of the NTK.

2.2. Previous NTK-based Metrics

Metric I: Frobenius Norm Suppose Θ_{θ_t} is the NTK at the t -th epoch. According to Xu *et al.* [44], for any $t > 0$, the following inequality holds:

$$\|y_i - f_{\theta_t}(x_i)\|_2^2 \leq \exp(-\lambda_{\min} t) \|y_i - f_{\theta_0}(x_i)\|_2^2, \quad (3)$$

where λ_{\min} is the minimum eigenvalue of the NTK matrix Θ_{θ_t} . From Eq. (3), we see that the upper bound on the loss term is determined by λ_{\min} ; the larger λ_{\min} is, the tighter the upper bound becomes, thereby yielding a smaller training loss.

Since Θ_{θ_t} is always symmetrical by definition, λ_{\min} can be bounded by the Frobenius norm of Θ_{θ_t} :

$$\lambda_{\min} \leq \sqrt{\sum_k |\lambda_k|^2} = \|\Theta_{\theta_t}\|_F, \quad (4)$$

where λ_k denotes the k -th eigenvalue of Θ_{θ_t} , ordered by $\lambda_{\min} \leq \dots \leq \lambda_{\max}$. Utilizing the Frobenius norm as a metric to score DNNs allows us to circumvent the eigen-decomposition of Θ_{θ_t} with the time complexity of $\mathcal{O}(n^3)$. Provided that the NTK does remain constant regardless of training as mentioned in Section 2.1, for any value of t , $\|\Theta_{\theta_t}\|_F$ can be replaced with $\|\Theta_{\theta_0}\|_F$. For the remainder of this paper, we use the abbreviation *F-Norm* to refer to this metric, which must be *positively correlated* with the final test accuracy of a DNN.

Metric II: Mean Although Xu *et al.* [44] show that the $\|\Theta_{\theta_0}\|_F$ can be leveraged to evaluate randomly-initialized DNNs, they do not directly use F-norm as a metric. Instead, the mean of Θ_{θ_0} is proposed as a metric for evaluating DNNs at initialization. The mean of the NTK matrix, denoted by $\mu(\Theta_{\theta_0})$, can be expressed as follows:

$$\mu(\Theta_{\theta_0}) = \frac{1}{N^2} \sum_{i=1}^N \sum_{j=1}^N \left(\frac{\partial f_{\theta_0}(x_i)}{\partial \theta_0} \right) \left(\frac{\partial f_{\theta_0}(x_j)}{\partial \theta_0} \right)^\top \quad (5)$$

Like F-norm, the Mean metric must also be *positively correlated* with the final test accuracy.

Metric III: Negative Condition Number Lee *et al.* [23] prove that the training dynamics of infinitely wide DNNs

are controlled by ordinary differential equations that can be solved as:

$$f_{\theta_t}(\mathcal{X}) = (\mathbf{I} - \exp(-\eta\Theta_{\theta_t}t))\mathcal{Y}, \quad (6)$$

where η and \mathbf{I} represent the learning rate and the Identity matrix, respectively. Lee *et al.* also hypothesize that the maximum feasible learning rate is given by: $\eta \sim 2/\lambda_{\max}$. A further study into the relationship between Θ_{θ_t} and the trainability of DNNs leads Xiao *et al.* [42] to conclude that Eq. (6) can be re-written in terms of the eigenspectrum of Θ_{θ_t} as follows:

$$f_{\theta_t}(\mathcal{X}) = (\mathbf{I} - \exp(-\eta\lambda_k t))\mathcal{Y}, \quad (7)$$

where λ_k denotes the k -th eigenvalue of Θ_{θ_t} . Plugging the maximum feasible learning rate $2/\lambda_{\min}$ into Eq. (7), Chen *et al.* [4] see that λ_{\min} converges exponentially at the rate of $1/c$, where $c = \lambda_{\max}/\lambda_{\min}$ is the condition number (CN) of Θ_{θ_t} . As CN grows larger, the output of a DNN $f_{\theta_t}(\mathcal{X})$ will fail to converge to the target label \mathcal{Y} . Thus, CN must exhibit a negative correlation with the final test accuracy. In this paper, to keep the trend in rank correlation consistent with the rest of the investigated metrics, we use the Negative Condition Number (NCN) instead; hence, NCN must be *positively correlated* with the final test accuracy.

3. Limitations of the NTK at Initialization

Here, we test the universal applicability of previous NTK-based metrics, measured at initialization, to diverse search spaces offered by NAS benchmarks. Even though these at-initialization metrics have been believed to highly correlated with the final accuracy, empirical demonstrations of their predictive abilities have been limited to a single search space: NAS-Bench-201. Therefore, extending the scope of evaluation to a far more diverse set of search spaces that contain different candidate operations and connectivity patterns is crucial for rigorously verifying the reliability of the NTK-based metrics. In Section 3.1, we provide a summary of the NAS benchmarks that are used for evaluation; more details on the construction of these benchmarks, as well as the image datasets they utilize, can be found in the Appendix. In Section 3.2, we present the evaluation results and report key findings regarding the practicality of the NTK-based metrics. Lastly, in Sections 3.3 and 3.4, we discuss additional pitfalls of the NTK identified from more up-close analyses of the NTK-based metrics.

3.1. Benchmarks for Neural Architecture Search

NAS-Bench-101 [47] contains 423,000 computationally unique neural architectures evaluated on CIFAR-10 [22]. All of the architectures in NAS-Bench-101 adopt the cell topology, a smaller feedforward module that is stacked repeatedly to construct the final architecture. The maximum

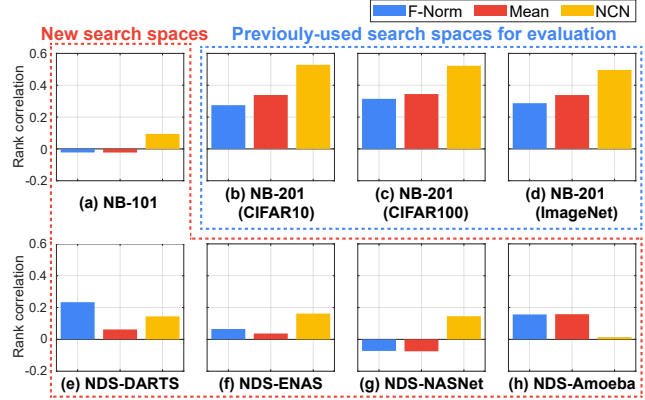


Figure 3. Rank correlation evaluation results on various NAS benchmarks. We compute the three metrics using Train and Eval Mode BNs. For simplicity, the higher correlation coefficient obtained from the two settings is reported here. The scale and the range of y-axes are set to be the same across all search spaces.

depth of each cell and the number of possible connections are set to be 7 and 9, respectively, and the following are the available candidate operations: 3×3 convolution, 1×1 convolution, and 3×3 max pooling.

NAS-Bench-201 [11] contains 15,625 architectures, all of which are evaluated on CIFAR-10, CIFAR-100 [22], and ImageNet-16-120 [7]. Similar to NAS-Bench-101, NAS-Bench-201 architectures are also based on the cell topology. Each one of the cells in NAS-Bench-201 has the fixed depth of 4, and the following candidate operations are included in the search space: zeroize, skip connection, 1×1 convolution, 3×3 convolution, and 3×3 average pooling.

NDS [36] offers a comprehensive analysis of commonly-adopted search spaces in NAS. The search spaces supported by the NDS benchmark include: DARTS [29], ENAS [35], NASNet [52], AmoebaNet [38], and PNAS [28]. Although all of these search spaces adopt the cell topology, the design of the cell structure differs from one another; please refer to the Appendix for the summary of differences. How the cells are stacked to generate the final neural architecture also varies among papers, but NDS standardizes this aspect of the search space by utilizing the DARTS architecture configuration. For each search space, NDS trains and evaluates $\sim 1K$ architectures on CIFAR-10.

3.2. Benchmark Evaluation Results

By using Kendall’s Tau as the measure of rank correlation, we evaluate how reliably the NTK-based metrics computed at initialization can predict the final test accuracy on various search spaces. For the sake of computational efficiency, we randomly sample 1,000 architectures from each search space for evaluation. It also appears that there exists no consensus on which batch statistics must be used for the batch normalization (BN) layer [20] when computing

the NTK. We thus test out both Train and Eval mode BN available in PyTorch [33]. Please refer to the Appendix for detailed experimental settings used in this section.

In Figure 3, we report the abbreviated evaluation results that only include the highest rank correlation coefficient obtained for each metric; a comprehensive visualization of all rank correlation measurement results is provided in the Appendix. Due to the page constraint, the results on NDS-PDARTS have also been moved to the Appendix. On NAS-Bench-201, we have successfully reproduced the rank correlation measure for Mean [44] and NCN [4] reported in their original papers. In NAS-Bench-101 and NDS search spaces, the degree of rank correlation decreases noticeably for all three metrics. In NAS-Bench-101 and NDS-NASNet, in particular, F-Norm and Mean seem to be negatively correlated with the final test accuracy, which goes against their theoretical motivation. Considering that the search spaces of NAS-Bench-101 and NDS are more complicated than that of NAS-Bench-201, such results may call into question whether the NTK framework can be deployed universally to more complex search spaces.

We also note that no single BN usage seems compatible with all three metrics. For instance, on the one hand, using fixed batch statistics (*i.e.* Eval mode BN) generally yields high rank correlation for NCN in NAS-Bench-201. On the other hand, in the same search space, using per-sample batch statistics (*i.e.* Train mode BN) improves rank correlation for F-Norm and Mean. This finding makes it evident that the NTK framework as is may lack consideration of the effect of BN on modern neural architectures.

3.3. Fine-grained Rank Correlation Evaluation

In the previous section, 1,000 architectures were randomly sampled from each benchmark to uniformly represent the entire architecture set. We now design a more challenging experiment, where we rank architectures in descending order and divide them into deciles, denoted by P; P1 contains Top-10% of architectures, P2 contains Top-10 ~ 20% of architectures, and so on. From each decile, 100 architectures are sampled for evaluation. This experiment allows us to determine whether the NTK-based metrics can stably guide the search process by gradually searching for a better architecture. Such a fine-grained experiment is no longer valid in search spaces that contradict the theoretical motivation. Therefore, the experiments in this section are conducted only on NAS-Bench-201.

We repeat this experiment with 20 different seeds for architecture sampling and visualize the results in the form of box-and-whisker plots. Please refer to Section A6 and Figures A2, A3, and A4 in the Appendix for the evaluation results. They suggest that in most deciles, the predictive ability of NTK-based metrics fluctuates significantly according to the choice of architectures used for evaluation.




Figure 4. Rank correlation evaluation results on NAS-Bench-201 obtained from different initialization schemes. NCN appears relatively robust to change in initialization schemes, but F-Norm and Mean are destroyed with the Gaussian initialization.

These results imply that guiding a search algorithm with the NTK-based metrics may not be able to escape from a locally optimal architecture and thus may often lead to unstable search results. Also, progressively shrinking the initial search space based on the error distribution of the architectures within it has become a commonly adopted technique in NAS or general architecture design [5, 19, 24, 37]. In such a refined search space that consists only of high-accuracy architectures, NTK-based metrics may fail to identify a particularly better architecture.

3.4. Sensitivity to Weight Initialization

Considering that previous NTK-based metrics have always been computed at initialization, the choice of weight initialization can be expected to have a non-negligible influence over the NTK computation result. We test how the NTK-based metrics are affected by Xavier [16], Kaiming [18] and Gaussian initializations. The experiments in this section are conducted only on NAS-Bench-201 as well. Figure 4 shows the change in rank correlation according to different initialization schemes. All three metrics show some degree of fluctuation when using Xavier and Kaiming initializations, but when the Gaussian initialization is used, to our surprise, the rank correlation for F-Norm and Mean plummet close to zero. This is an unexpected result because the NTK framework assumes that the parameters in a DNN are initialized as iid Gaussians, and thus their function realizations asymptotically converge to a Gaussian distribution in the infinite width limit [21].

4. Methodology

We conjecture that the unreliability of the NTK-based metrics obtained at initialization occurs because the underlying theoretical assumptions in the NTK framework are violated in modern DNNs. As a result, the NTK derived from a modern DNN is likely to evolve in a non-linear manner as training progresses, diverging away from the NTK at initialization [15, 17, 32]. In Section 4.1, we witness that the architectures considered in NAS indeed exhibit highly non-linear characteristics. Then, in Section 4.2, we present Label-Gradient Alignment, a novel NTK-based metric that has yet to be studied in NAS, and show how it can capture

the evolution of the NTK with respect to target labels. Afterwards, in Section 4.3, we corroborate the theoretical motivation behind LGA by demonstrating that after little amount of training, LGA shows a meaningful level of rank correlation with the test accuracy. Please refer to the Appendix for detailed experimental settings used in this section.

4.1. Time Evolution of the NTK

Kernel Correlation measures the Pearson’s correlation coefficient between Θ_{θ_0} and Θ_{θ_t} : $\text{Cov}(\Theta_{\theta_0}, \Theta_{\theta_t}) / (\sigma(\Theta_{\theta_0}) \sigma(\Theta_{\theta_t}))$. The correlation measurement results are presented on the left panel of Figure 5. For all three datasets, the correlation between Θ_{θ_0} and Θ_{θ_t} decreases rapidly in the initial epochs and start to stabilize after some amount of training, and such a trend becomes more conspicuous with the growth in data complexity.

Relative Kernel Difference measures the relative change in the NTK from Θ_{θ_0} to Θ_{θ_t} : $|\Theta_{\theta_t} - \Theta_{\theta_0}| / |\Theta_{\theta_0}|$. The kernel difference measurement results are visualized on the right panel of Figure 5. We once again observe that the NTK deviates noticeably from Θ_{θ_0} in the initial epochs, but the relative difference starts to saturate mid-training.

Based on both the correlation and the distance measurement results, a singular conclusion can be drawn: modern neural architectures primarily studied in NAS exhibit a highly non-linear behavior during training, and thus, the NTK in such architectures experiences a large amount of kernel twisting. Consequently, the NTK framework, whose core theoretical results are built on the assumption that the NTK remains constant throughout training, loses its credibility, and the characteristics of the NTK at initialization become incapable of accurately representing the final test accuracy of a neural architecture. This finding can be interpreted as being consistent with recent discovery that the non-linear advantage in DNNs is what allows them to outperform their linear kernel counterparts [15, 32]. Therefore, we introduce Label-Gradient Alignment, a novel NTK-based metric that can capture the non-linear characteristics of neural architectures with only few epochs of training.

4.2. Label-Gradient Alignment

While a neural architecture’s generalization performance must be one of the most prioritized factors in NAS, obtaining a closed-form characterization of generalization error on test data that cannot be accessed is impossible. However, provided that the approximation in Eq. (1) is exact, it is possible to formulate generalization guarantees for DNNs by transferring the generalization bounds computed from their linear kernel equivalents. In Bartlett *et al.* [2], it is shown that with high probability, the following relationship holds:

$$\mathcal{R}(f^*) \leq \hat{\mathcal{R}}(f^*) + \mathcal{O} \left(\sqrt{\frac{\|f\|_{\Theta_{\theta_0}}^2 \mathbb{E}_x[\Theta_{\theta_0}(x, x')]}{m}} \right), \quad (8)$$

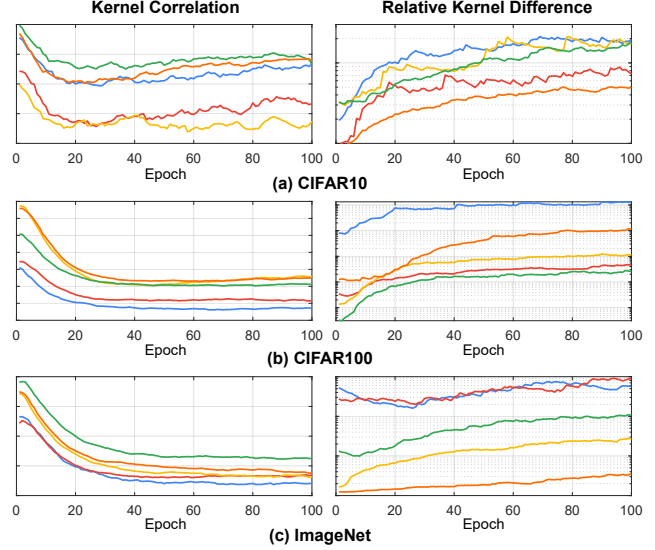


Figure 5. Analyzing the time evolution of the NTK as training progresses. Five unique architectures, represented by lines of different colors, are randomly sampled. For all three datasets, the kernel correlation decreases, and the kernel difference increases.

where \mathcal{R} and $\hat{\mathcal{R}}$ denote the expected and the empirical risks, respectively, of $f^* = \text{argmin}_{h \in \mathcal{H}} \hat{\mathcal{R}}(h) + r \|h\|_{\Theta_{\theta_0}}^2$, with $r > 0$ as a regularization constant. $f : \mathbb{R}^d \rightarrow \pm 1$ corresponds to the target function that the DNN is trying to learn, $\|f\|_{\Theta_{\theta_0}}$ is the RKHS norm of this target function. One can assume that a DNN generalizes well whenever it achieves a low expected risk, and Eq. (8) indicates that the difference between the expected and the empirical risks decreases as the term $\|f\|_{\Theta_{\theta_0}}^2$ becomes smaller. Through eigendecomposition, $\|f\|_{\Theta_{\theta_0}}^2$ can be re-written as:

$$\|f\|_{\Theta_{\theta_0}}^2 = \sum_k \frac{1}{\lambda_k} (\mathbb{E}_{x \sim \mathcal{D}} [v_k(x) f(x)])^2, \quad (9)$$

where $\{\lambda_k, v_k\}$ denotes k -th eigenvalue-eigenvector pair of Θ_{θ_0} . We can now see that better generalization performance may be expected when targets (or labels) align well with top eigenvectors of the NTK matrix, *i.e.* the first principal components of the variability of the per-sample gradients.

Instead of directly computing Eq. (9), Ortiz *et al.* [32] offer a more tractable bound on $\|f\|_{\Theta_{\theta_0}}^2$:

$$\|f\|_2^4 / \|f\|_{\Theta_{\theta_0}}^2 \leq \alpha(f), \quad (10)$$

$$\text{where } \alpha(f) = \mathbb{E}_{x, x' \sim \mathcal{D}} [f(x) \Theta_{\theta_0}(x, x') f(x')],$$

where $\|\cdot\|_{\Theta_{\theta_0}}$ and $\|\cdot\|_2$ denote the RKHS norm and the l_2 norm, respectively. In a fully supervised setting, where the target function is defined in the form of class labels, target function f in $\alpha(f)$ can be replaced with target labels \mathcal{Y} , thereby yielding:

$$\alpha(\mathcal{Y}) = \mathcal{Y}^T \Theta_{\theta_0} \mathcal{Y}. \quad (11)$$

From here on, we refer to $\mathcal{Y}^\top \Theta_{\theta_0} \mathcal{Y}$ as *LGA*, a shorthand for label-gradient alignment. Replacing $\alpha(f)$ in Eq. (10) with $\alpha(\mathcal{Y})$, we see that smaller $\|f\|_{\Theta_{\theta_0}}^2$ will increase $\alpha(\mathcal{Y})$. From Eq. (8), it is evident that a small value of $\|f\|_{\Theta_{\theta_0}}^2$ is preferred to minimize the gap between the expected and the empirical risks. Therefore, LGA must be *positively correlated* with the generalization performance of an architecture and thus with its final test accuracy. The noteworthy difference between LGA and previously-proposed metrics is that LGA takes both the NTK and the target labels into consideration. Such a mathematical formulation of LGA allows it to accurately follow the orientation of the NTK with respect to the target functions that the neural architecture is trying to estimate.

A similar measure was used in Deshpande *et al.* [10] in the context of model selection for finetuning. Inspired by Deshpande *et al.*, we introduce additional procedures to effectively utilize LGA for NAS. To extend the binary classification setting of the NTK framework to multi-class classification in NAS, an $N \times N$ label matrix $L_{\mathcal{Y}}$, is introduced, in which $L_{\mathcal{Y}}[i, j] = 1$ if x_i and x_j belong in the same class, and $L_{\mathcal{Y}}[i, j] = -1$ otherwise. To induce the invariability in scale, LGA is normalized as follows:

$$LGA = \frac{(\Theta_{\theta_0} - \mu(\Theta_{\theta_0})) \cdot (L_{\mathcal{Y}} - \mu(L_{\mathcal{Y}}))}{\|\Theta_{\theta_0} - \mu(\Theta_{\theta_0})\|_2 \|L_{\mathcal{Y}} - \mu(L_{\mathcal{Y}})\|_2}, \quad (12)$$

where $\mu(L_{\mathcal{Y}})$ is the average of elements in $L_{\mathcal{Y}}$.

4.3. NTK-based Metrics after Training

We now repeat the rank correlation evaluation experiment conducted in Section 3 after training architectures for $t \in \{1, 3, 5, 10\}$ epochs. The post-training rank correlation evaluation results on NAS-Bench-201 are visualized in Figure 6. In terms of rank correlation, LGA is the only metric that exhibits a steady improvement as training progresses across all three datasets. In Table 1, we compare LGA after a single training epoch (LGA_1) with NTK-based metrics obtained at initialization on other NAS benchmarks. Independent of the choice of a benchmark, the rank correlation of LGA_1 exceeds that of previous NTK-based metrics.

To better understand this characteristic behavior of LGA, we analyze how the NTK-based metrics of high-, mid-, and low-accuracy architectures change during the training process; please refer back to Figure 2 for the results of this analysis. For high-accuracy architectures, we observe a surge in LGA, which indicates that the concentration of labels on to the principal components of the NTK matrix has significantly increased; LGA in mid-accuracy architectures behaves similarly in that it shows some amount of increase, albeit small. On the contrary, for low-accuracy architectures, LGA remains stationary. While the other metrics also change during training, they do so in a rather meaningless way that cannot distinguish different amounts of non-linear




Figure 6. Post-training rank correlation evaluation results on NAS-Bench-201. Regardless of the dataset, the predictive performance of LGA steadily improves from initialization.

Table 1. Comparison of LGA_1 with previous NTK-based metrics on various NAS benchmarks. In terms of rank correlation, LGA_1 outperforms other metrics across all benchmarks.

Benchmark	F-Norm	Mean	NCN	LGA_1
NAS-Bench-101	-0.022	-0.023	0.094	0.308
NDS-DARTS	0.234	0.062	0.145	0.408
NDS-ENAS	0.065	0.037	0.162	0.416
NDS-NASNet	-0.073	-0.075	0.146	0.357
NDS-Amoeba	0.157	0.158	0.013	0.396

advantage in high-, mid- and low-accuracy architectures. This experimental analysis justifies the need for target labels in LGA to understand how the NTK rotates and evolves with respect to the target functions. As stated in Section 4.2, with target labels as a type of an anchor point, LGA can distinguish between highly trainable architectures, in which target labels gradually become more aligned with the principal components of the NTK in the initial epochs, and less trainable architectures, which do not benefit much from kernel twisting. In the absence of target labels, the other metrics cannot determine the direction in which the NTK evolves, and hence, they do not gain any meaningful information from the training process.

5. Searching with LGA

To demonstrate that LGA can be utilized to improve the computational efficiency of NAS, we integrate LGA with random search and evolutionary search algorithms. Based on the evaluation results in Section 4.1, LGAs after 3 (LGA_3) and 5 (LGA_5) epochs of training are used for searching. In random search (RS), 100 architectures are sampled from the search space for evaluation, and the architecture with the maximum LGA is selected. For evolutionary search (REA), we adopt the regularized evolutionary search algorithm of Real *et al.* [38]. The regularized approach of Real *et al.* differs from a naïve evolutionary algorithm in that it prefers newer candidate architectures. We use a fixed search cost budget for evolutionary search on all three datasets. Please refer to the Appendix for the experimental settings used in this section and comprehensive pseudo codes of both search processes.

In Table 2, the search performances of RS and REA

Table 2. Comparison against state-of-the-art NAS algorithms on NAS-Bench-201. “Optimal” refers to the best test accuracy achievable in the NAS-Bench-201 search space. The search process is executed separately for each image dataset. The search cost is reported in GPU seconds. All of our search experiments are conducted on a single NVIDIA Tesla A40 GPU. †: search based on NTK-based metrics.

Model	CIFAR-10			CIFAR-100			ImageNet-16-120			Search Method
	Acc.	Cost	Speed-up	Acc.	Cost	Speed-up	Acc.	Cost	Speed-up	
ResNet	93.97	N/A	N/A	70.86	N/A	N/A	43.63	N/A	N/A	Manual
RS [3]	93.63	216K	1.0×	71.28	460K	1.0×	44.88	1M	1.0×	Random
RL [51]	93.72	216K	1.0×	70.71	460K	1.0×	44.10	1M	1.0×	RL
REA [38]	93.72	216K	1.0×	72.12	460K	1.0×	45.01	1M	1.0×	EA
BOHB [14]	93.49	216K	1.0×	70.84	460K	1.0×	44.33	1M	1.0×	HPO
RSPS [25]	91.67	10K	21.6×	57.99	46K	21.6×	36.87	104K	9.6×	RS+WS
DARTS [29]	88.32	23K	9.4×	67.34	80K	5.8×	33.04	110K	9.6×	Gradient
GDAS [12]	93.36	22K	12.0×	67.60	39K	11.7×	37.97	130K	7.7×	Gradient
NASWOT [30]	92.96	2.2K	100×	70.03	4.6K	100×	44.43	10K	100×	Random
†TE-NAS	93.90	2.2K	100×	71.24	4.6K	100×	42.38	10K	100×	Pruning-based
†KNAS ($k = 2$)	93.05	4.2K	50×	68.91	9.2K	50×	34.11	20K	50×	Random
†KNAS ($k = 5$)	93.42	10.8K	20×	71.42	23K	20×	45.35	50K	20×	Random
†RS + LGA ₃	93.64	3.6K	60×	69.77	5K	92×	45.03	10.1K	99×	Random
†RS + LGA ₅	94.03	5.4K	40×	71.56	7K	66×	46.30	15K	67×	Random
†REA + LGA ₃	94.30	3.6K	60×	71.18	3.6K	127×	45.30	3.6K	277×	EA
†REA + LGA ₅	93.94	5.4K	40×	72.42	5.4K	85×	45.17	5.4K	185×	EA
Optimal	94.37			73.51			47.31			N/A

with LGA₃ and LGA₅ are compared against those of other state-of-the-art NAS algorithms. We would like to emphasize that for both RS and REA, no additional information besides LGA is used to evaluate architectures. RS with either LGA₃ or LGA₅ outperforms search algorithms based on other NTK-based metrics obtained at initialization; TE-NAS [4] and KNAS [44] utilize CN and Mean, respectively. This result is particularly impressive considering that TE-NAS and KNAS rely on some external signal during search; TE-NAS utilizes a more complicated search algorithm and another at-initialization metric, and in the final architecture derivation step of KNAS, the actual test accuracy is used. REA with LGA₃ or LGA₅ also achieves competitive test accuracy, and more importantly, it does so with far less search cost than other search algorithms on CIFAR-100 and ImageNet-16-120. Overall, despite introducing some amount of training, LGA₃ and LGA₅ appear to be highly competent and computationally efficient metrics. Lastly, we show that LGA can be applied more broadly to a variety of search spaces by conducting random search on various benchmarks other than NAS-Bench-201. The results and experimental details are provided in the Appendix.

6. Concluding Remarks

The technical and experimental contributions of this paper are largely three-fold. First, through a more extensive and fine-grained evaluation of the NTK-based metrics, we revealed that the current form of the NTK framework might

not be as reliable of a theoretical framework for NAS as previously believed to be. Second, through the empirical analysis of the time evolution of the NTK, we demonstrated that the aforementioned limitation occurs because in modern neural architectures, the NTK evolves in a highly non-linear manner during the training process, diverging significantly away from the NTK at initialization. Third, when complemented with little amount of training, LGA, first introduced in this work, rose as a strong predictor of test accuracy because its innate theoretical motivation could embody the non-linear characteristics of modern neural architectures, which the other NTK-based metrics were blind to. Integrating LGA into existing search algorithms provided a further empirical support for its effectiveness as a computationally effective predictor of test accuracy. We discuss the limitations and societal impact of our work in the Appendix.

Acknowledgements

This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government(MSIT) [NO.2021-0-01343, Artificial Intelligence Graduate School Program (Seoul National University)], the BK21 FOUR program of the Education and Research Program for Future ICT Pioneers, Seoul National University in 2022, AIRS Company in Hyundai Motor Company & Kia Corporation through HMC/KIA-SNU AI Consortium Fund, and SNU-Naver Hyperscale AI Center.

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
- [2] Peter L Bartlett and Shahar Mendelson. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3(Nov):463–482, 2002. 6
- [3] James Bergstra and Yoshua Bengio. Random search for hyper-parameter optimization. *Journal of machine learning research*, 13(2), 2012. 8
- [4] 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*, 2020. 2, 3, 4, 5, 8
- [5] Xiangning Chen, Ruochen Wang, Minhao Cheng, Xiaocheng Tang, and Cho-Jui Hsieh. Drnas: Dirichlet neural architecture search. In *International Conference on Learning Representations*, 2020. 1, 5
- [6] Yaofo Chen, Yong Guo, Qi Chen, Minli Li, Wei Zeng, Yaowei Wang, and Mingkui Tan. Contrastive neural architecture search with neural architecture comparators. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9502–9511, 2021. 2
- [7] Patryk Chrabaszcz, Ilya Loshchilov, and Frank Hutter. A downsampled variant of imagenet as an alternative to the cifar datasets. *arXiv preprint arXiv:1707.08819*, 2017. 4
- [8] Xiangxiang Chu, Tianbao Zhou, Bo Zhang, and Jixiang Li. Fair DARTS: Eliminating Unfair Advantages in Differentiable Architecture Search. In *European Conference On Computer Vision*, 2020. 1
- [9] Boyang Deng, Junjie Yan, and Dahua Lin. Peephole: Predicting network performance before training. *arXiv preprint arXiv:1712.03351*, 2017. 2
- [10] Aditya Deshpande, Alessandro Achille, Avinash Ravichandran, Hao Li, Luca Zancato, Charless Fowlkes, Rahul Bhotika, Stefano Soatto, and Pietro Perona. A linearized framework and a new benchmark for model selection for fine-tuning. *arXiv preprint arXiv:2102.00084*, 2021. 7
- [11] Xuanyi Dong and Yi Yang. Nas-bench-201: Extending the scope of reproducible neural architecture search. In *International Conference on Learning Representations*, 2019. 4
- [12] Xuanyi Dong and Yi Yang. Searching for a robust neural architecture in four gpu hours. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 1761–1770, 2019. 1, 8
- [13] Thomas Elsken, Jan Hendrik Metzen, and Frank Hutter. Neural architecture search: A survey. *The Journal of Machine Learning Research*, 20(1):1997–2017, 2019. 1
- [14] Stefan Falkner, Aaron Klein, and Frank Hutter. Bohb: Robust and efficient hyperparameter optimization at scale. In *International Conference on Machine Learning*, pages 1437–1446. PMLR, 2018. 8
- [15] Stanislav Fort, Gintare Karolina Dziugaite, Mansheej Paul, Sepideh Kharaghani, Daniel M Roy, and Surya Ganguli. Deep learning versus kernel learning: an empirical study of loss landscape geometry and the time evolution of the neural tangent kernel. *Advances in Neural Information Processing Systems*, 33, 2020. 2, 5, 6
- [16] Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings, 2010. 5
- [17] Micah Goldblum, Jonas Geiping, Avi Schwarzschild, Michael Moeller, and Tom Goldstein. Truth or backpropaganda? an empirical investigation of deep learning theory. In *International Conference on Learning Representations*, 2019. 2, 5
- [18] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015. 5
- [19] Yiming Hu, Yuding Liang, Zichao Guo, Ruosi Wan, Xiangyu Zhang, Yichen Wei, Qingyi Gu, and Jian Sun. Angle-based search space shrinking for neural architecture search. In *European Conference on Computer Vision*, pages 119–134. Springer, 2020. 5
- [20] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015. 4
- [21] Arthur Jacot, Franck Gabriel, and Clément Hongler. Neural tangent kernel: convergence and generalization in neural networks. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems*, pages 8580–8589, 2018. 2, 3, 5
- [22] Alex Krizhevsky et al. Learning multiple layers of features from tiny images. 2009. 4
- [23] Jaehoon Lee, Lechao Xiao, Samuel Schoenholz, Yasaman Bahri, Roman Novak, Jascha Sohl-Dickstein, and Jeffrey Pennington. Wide neural networks of any depth evolve as linear models under gradient descent. *Advances in neural information processing systems*, 32:8572–8583, 2019. 3
- [24] Guohao Li, Guocheng Qian, Itzel C Delgadillo, Matthias Muller, Ali Thabet, and Bernard Ghanem. Sgas: Sequential greedy architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1620–1630, 2020. 5
- [25] Liam Li and Ameet Talwalkar. Random search and reproducibility for neural architecture search. In *Uncertainty in artificial intelligence*, pages 367–377. PMLR, 2020. 2, 8
- [26] Zhihang Li, Teng Xi, Jiankang Deng, Gang Zhang, Shengzhao Wen, and Ran He. Gp-nas: Gaussian process based neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11933–11942, 2020. 2
- [27] Ming Lin, Pichao Wang, Zhenhong Sun, Heseng Chen, Xiuyu Sun, Qi Qian, Hao Li, and Rong Jin. Zen-nas: A zero-shot nas for high-performance image recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 347–356, 2021. 2

- [28] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. Progressive neural architecture search. In *Proceedings of the European conference on computer vision (ECCV)*, pages 19–34, 2018. 4
- [29] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. In *International Conference on Learning Representations*, 2019. 1, 4, 8
- [30] 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. 2, 8
- [31] Byungsook Na, Jisoo Mok, Hyeokjun Choe, and Sungroh Yoon. Accelerating neural architecture search via proxy data. In *Proceedings of the 30th International Joint Conference on Artificial Intelligence*, 2021. 2
- [32] Guillermo Ortiz-Jiménez, Seyed-Mohsen Moosavi-Dezfooli, and Pascal Frossard. What can linearized neural networks actually say about generalization? *arXiv preprint arXiv:2106.06770*, 2021. 2, 3, 5, 6
- [33] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in neural information processing systems*, pages 8026–8037, 2019. 5
- [34] Houwen Peng, Hao Du, Hongyuan Yu, QI LI, Jing Liao, and Jianlong Fu. Cream of the crop: Distilling prioritized paths for one-shot neural architecture search. *Advances in Neural Information Processing Systems*, 33, 2020. 1
- [35] Hieu Pham, Melody Y. Guan, Barret Zoph, Quoc V. Le, and Jeff Dean. Efficient neural architecture search via parameter sharing. In *Proceedings of the 35th International Conference on Machine Learning*, pages 4095–4104, 2018. 1, 4
- [36] Ilija Radosavovic, Justin Johnson, Saining Xie, Wan-Yen Lo, and Piotr Dollár. On network design spaces for visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1882–1890, 2019. 4
- [37] Ilija Radosavovic, Raj Prateek Kosaraju, Ross Girshick, Kaiming He, and Piotr Dollár. Designing network design spaces. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10428–10436, 2020. 5
- [38] 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. 2, 4, 7, 8
- [39] 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 *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 2902–2911. JMLR.org, 2017. 1
- [40] Yehui Tang, Yunhe Wang, Yixing Xu, Hanting Chen, Boxin Shi, Chao Xu, Chunjing Xu, Qi Tian, and Chang Xu. A semi-supervised assessor of neural architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1810–1819, 2020. 2
- [41] Colin White, Arber Zela, Binxin Ru, Yang Liu, and Frank Hutter. How powerful are performance predictors in neural architecture search? *arXiv preprint arXiv:2104.01177*, 2021. 2
- [42] Lechao Xiao, Jeffrey Pennington, and Sam Schoenholz. Disentangling trainability and generalization in deep learning. *arXiv preprint arXiv:1912.13053*, 2019. 4
- [43] Lingxi Xie, Xin Chen, Kaifeng Bi, Longhui Wei, Yuhui Xu, Zhengsu Chen, Lanfei Wang, An Xiao, Jianlong Chang, Xiaopeng Zhang, et al. Weight-sharing neural architecture search: A battle to shrink the optimization gap. *arXiv preprint arXiv:2008.01475*, 2020. 2
- [44] Jingjing Xu, Liang Zhao, Junyang Lin, Rundong Gao, Xu Sun, and Hongxia Yang. Knas: Green neural architecture search. In *International Conference on Machine Learning*, pages 11613–11625. PMLR, 2021. 2, 3, 5, 8
- [45] Yixing Xu, Yunhe Wang, Kai Han, Yehui Tang, Shangling Jui, Chunjing Xu, and Chang Xu. Renas: Relativistic evaluation of neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4411–4420, 2021. 2
- [46] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. Pc-darts: Partial channel connections for memory-efficient architecture search. In *International Conference on Learning Representations*, 2019. 1
- [47] 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. 4
- [48] Miao Zhang, Huiqi Li, Shirui Pan, Xiaojun Chang, and Steven Su. Overcoming multi-model forgetting in one-shot nas with diversity maximization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 7809–7818, 2020. 1
- [49] Dongzhan Zhou, Xinchu Zhou, Wenwei Zhang, Chen Change Loy, Shuai Yi, Xuesen Zhang, and Wanli Ouyang. Econas: Finding proxies for economical neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11396–11404, 2020. 2
- [50] Pan Zhou, Caiming Xiong, Richard Socher, and Steven Hoi. Theory-inspired path-regularized differential network architecture search. In *Neural Information Processing Systems*, 2020. 1
- [51] Barret Zoph and Quoc V. Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations*, 2017. 1, 8
- [52] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 8697–8710, 2018. 1, 4