

GeNAS: Neural Architecture Search with Better Generalization

Joonhyun Jeong^{1,2}, Joonsang Yu^{1,3}, Geondo Park², Dongyoon Han³ and YoungJoon Yoo¹

¹NAVER Cloud, ImageVision

²KAIST

³NAVER AI Lab

{joonhyun.jeong, joonsang.yu}@navercorp.com, geondopark@kaist.ac.kr,
{dongyoon.han, youngjoon.yoo}@navercorp.com

Abstract

Neural Architecture Search (NAS) aims to automatically excavate the optimal network architecture with superior test performance. Recent neural architecture search (NAS) approaches rely on validation loss or accuracy to find the superior network for the target data. In this paper, we investigate a new neural architecture search measure for excavating architectures with better generalization. We demonstrate that the flatness of the loss surface can be a promising proxy for predicting the generalization capability of neural network architectures. We evaluate our proposed method on various search spaces, showing similar or even better performance compared to the state-of-the-art NAS methods. Notably, the resultant architecture found by flatness measure generalizes robustly to various shifts in data distribution (e.g. ImageNet-V2,-A,-O), as well as various tasks such as object detection and semantic segmentation.

1 Introduction

Recently, Neural Architecture Search (NAS) [Liu *et al.*, 2018b; Liu *et al.*, 2018a; Hong *et al.*, 2020] has evolved to achieve remarkable accuracy along with the development of human-designed networks [He *et al.*, 2016; Dosovitskiy *et al.*, 2020] on the image recognition task. Several NAS methods [Zoph *et al.*, 2018; Chu *et al.*, 2020; Zhang *et al.*, 2021; Liu *et al.*, 2018b; Xu *et al.*, 2019; Hong *et al.*, 2020] further have demonstrated generalization ability (generalizability) of automatically designed networks with test accuracy and transfer performance onto the other datasets. For the widespread leverage of architectures found by NAS on the various tasks such as object detection [Lin *et al.*, 2014] and segmentation [Cordts *et al.*, 2016] (task-generalizability), investigating generalizability of each architecture candidate is prerequisite and indispensable. Despite its importance, quantitative measurement of generalizability during the architecture search process is still under-explored. In this paper, we aim to find an optimal proxy measurement to discriminate generaliz-

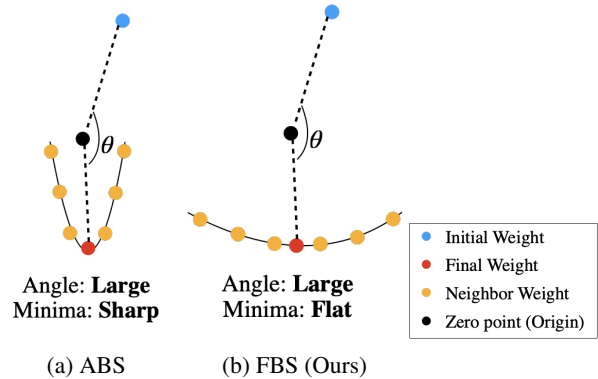


Figure 1: Shape of local loss minima found by angle-based search (ABS) and flatness-based search (FBS). (a) The architecture found by ABS can not guarantee to be located within flat local minima. (b) FBS searches for architectures with flat local minima by inspecting loss values of local neighborhood weights.

Kendall's Tau		
CIFAR-10	CIFAR-100	ImageNet16-120
0.4302	0.4724	0.4097

Table 1: Low correlation of angle measure with flatness measure on NAS-Bench-201 [Dong and Yang, 2020] search space. We evaluated the angle and flatness of all architectures and compared Kendall's Tau [Kendall, 1938] rank correlation between these search metrics on CIFAR-10, CIFAR-100, and ImageNet16-120 [Chrabaszcz *et al.*, 2017] dataset.

able architectures during the search process.^{1 2}

Previous NAS algorithms including the pioneering differentiable search method, DARTS [Liu *et al.*, 2018b] and evolutionary search method, SPOS [Guo *et al.*, 2020] use validation performance as a proxy measure for the generalizability as follows:

$$a^* = \operatorname{argmax}_{a \in A} S(a), \tag{1}$$

where a and A denote an architecture candidate and the entire search space, respectively, and $S(\cdot)$ represents a measurement

¹Code is available at <https://github.com/clovaai/GeNAS>.

²Extended paper (including the appendix) is available at <https://arxiv.org/abs/2305.08611>.

function which is broadly defined by accuracy [Guo *et al.*, 2020] or negative of loss value [Liu *et al.*, 2018b] on a validation dataset. Although these performance-based search (PBS) methods find the optimal architecture for generalization on the validation set, they show poor generalizability on the test set and other tasks, caused by overfitting on validation set [Zela *et al.*, 2019; Oymak *et al.*, 2021]. In addition, PBS methods represent a large discrepancy between the validation accuracy and ground truth test accuracy provided by NAS benchmark [Dong and Yang, 2020] as shown in [Guo *et al.*, 2020; Zhang *et al.*, 2021].

To search generalizable architectures, several literatures [Shu *et al.*, 2019; Zhang *et al.*, 2021] empirically observe that architectures with fast convergence during training have a tendency to show better generalizability on test set. Based on the empirical connection between convergence speed and generalization, RLNAS [Zhang *et al.*, 2021] proposed an Angle-Based Search (ABS) method, which estimates angle between initial and final network parameters after convergence of the model (i.e. convergence speed) as a proxy performance measure during the search process. However, we argue that ABS still has a large headroom for better generalization in terms of flat (wide) local minima, which has been considered as one of the key signals for inspecting generalizability of a trained network [Keskar *et al.*, 2016; Zhang *et al.*, 2018; Pereyra *et al.*, 2017; Cha *et al.*, 2020; He *et al.*, 2019]. Intuitively, since the architecture with flat loss minima has widely low loss values around the minimum, it can achieve a low generalization error even if the loss surface is shifted due to the distribution gap from the test dataset.

Since ABS only concerns the angle between initial model weights and trained ones in terms of convergence speed, found architectures can not be guaranteed to have flat local minima, as shown in Figure 1. Specifically, architectures not chosen by ABS (i.e. small angle) might have better generalizability based on the flat property of loss minima. Table 1 corroborates that angle is indeed in short of correlation with flatness of local minima.

To explicitly design a search proxy measure that has a high correlation with the generalizability of the found model, we propose a flatness-based search method, namely FBS, which excavates a well-generalizable architecture by measuring the flatness of loss surface. FBS can find out robust architecture with low generalization error on shifted data distribution (e.g. test data, out-of-distribution datasets, downstream tasks) by inspecting both the depth and flatness of loss curvature near local minima through injecting random noise. In addition, FBS can be either replaced or incorporated with other state-of-the-art search measures to enhance performance as well as generalizability.

Consequently, building upon our search method FBS, we propose a novel flatness-based NAS framework, namely GeNAS, to exactly discriminate generalizability of architectures during searching. We show the effectiveness of the proposed GeNAS for both cases when using FBS solely or integrated into the conventional architecture score measurements such as PBS and ABS. Specifically, our GeNAS achieves comparable or even better performances on several NAS benchmarks compared to PBS- and ABS-based search

methods [Liu *et al.*, 2018b; Zhang *et al.*, 2021; Xu *et al.*, 2019; Guo *et al.*, 2020; Chu *et al.*, 2020; Chen *et al.*, 2019; Hong *et al.*, 2020]. Furthermore, we also show that the proposed FBS can be combined with conventional search metrics (e.g. PBS, ABS), inducing significant performance gain. Finally, we also demonstrate that our FBS can well-generalize on various data distribution shifts, as well as on multiple downstream tasks such as object detection and semantic segmentation.

Our contributions can be summarized as follows:

- We firstly demonstrate that the flatness of local minima can be employed to quantify generalizability of architecture in NAS domain, which only had been a means of confirming the generalizability after training a neural network.
- We propose a new architecture search proxy measure, flatness of local minima, well-suited for finding architectures with better generalization, which can replace or even significantly enhance the search performance of the existing search proxy measures.
- The found architecture induced by our FBS demonstrates the state-of-the-art performance on various search spaces and datasets, even showing great robustness on data distribution shift and better generalization on various downstream tasks.

2 Related Work

2.1 Neural Architecture Search

Early NAS methods are based on the reinforcement learning (RL) [Baker *et al.*, 2016; Zoph *et al.*, 2018], which train the agent network to choose better architecture. The RL-based methods require the test accuracy of each candidate network for reward value, so training every candidate network from scratch is also required to measure that. For this reason, it is not feasible on a large-scale dataset such as ImageNet [Krizhevsky *et al.*, 2012]. To solve this problem, the weight-sharing NAS methods are introduced [Liu *et al.*, 2018b; Xu *et al.*, 2019; Xie *et al.*, 2018; Guo *et al.*, 2020; Zhang *et al.*, 2021]. The weight-sharing NAS generally uses the *SuperNet*, which contains all operations in objective search space, and chooses several operations from the *SuperNet* to decide the final architecture, which is called *SubNet*. Among these weight-sharing NAS frameworks, [Liu *et al.*, 2018b; Xu *et al.*, 2019] introduced a gradient-based architecture search method, where they jointly train the architecture parameters with weight parameters using gradient descent. After training, the final architecture is decided according to the architecture parameters. Meanwhile, the one-shot NAS methods [Guo *et al.*, 2020; Bender *et al.*, 2018; Brock *et al.*, 2017] pointed out the critical drawback of these gradient-based search methods as there exists a strongly coupled and biased connection between *SuperNet* weight parameters and its architecture parameters; only a small subset of *SuperNet* weight parameters with large architecture parameter value will be intensely optimized, leaving the others trained insufficiently. Therefore, [Guo *et al.*, 2020; Bender *et al.*, 2018; Brock *et al.*, 2017] sequentially decoupled

the optimization process for *SuperNet* and architecture parameters, showing superior search performance over the gradient-based search methods. Inspired by these breakthroughs and its flexibility of introducing various search proxy measures, we construct our GeNAS based on the one-shot NAS framework.

2.2 Architecture Search Proxy Measure

During the search time, it is hard to check the actual test performance of each architecture candidate when it is trained from scratch, so the proxy measure has to be employed for the candidate evaluation. Several approaches proposed to predictively discriminate well-trained neural networks without any training by inspecting either the correlation between the linear maps of variously augmented image [Mellor *et al.*, 2021] or spectrum of Neural Tangent Kernel (NTK) [Chen *et al.*, 2021]. Although these training-free search proxy measures significantly reduced the search costs within even four GPU hours, actual test performance was inferior to that of training-involved search proxy measures such as validation accuracy and loss. Meanwhile, ABS methods [Zhang *et al.*, 2021; Hu *et al.*, 2020] introduced a new search proxy measure, angle, for indicating the generalizability of a neural network architecture, showing search accuracy improvement [Zhang *et al.*, 2021] over conventional search proxy measures such as validation accuracy [Guo *et al.*, 2020]. Since ABS method only investigates the convergence speed of an architecture, [Zhang *et al.*, 2021] successfully searched a well-trainable architecture under ground truth label absent during *SuperNet* training. However, searching with randomly-distributed label still shows large performance gap (about 0.15 Kendall’s Tau score gap on NAS-Bench-201) to that of searching with the ground-truth label. Therefore, in order to fulfill higher test generalization of a searched architecture, we train *SuperNet* and searched architecture under ground-truth label setting.

2.3 Flatness of Local Minima

The flatness of loss landscape near local minima has been considered as a key signal for representing generalizability. [Keskar *et al.*, 2016; Jastrzebski *et al.*, 2017; Hoffer *et al.*, 2017] empirically observed that appropriate training hyper-parameters such as batch size, learning rate, and the number of training iterations can implicitly enable a model to have wide and flat minima, enhancing test generalization performance. [Chaudhari *et al.*, 2019] further explicitly drives a neural network model to the flat minima through an entropy-regularized SGD. Several works also promoted the flat local minima in terms of regularization during training using *Label Smoothing* [Pereyra *et al.*, 2017] and *Knowledge Distillation* [Zhang *et al.*, 2018], enjoying test performance gain. Based on these empirical connections between test generalizability and flatness of local minima of a neural network, we investigate the role of flatness of minima on the architecture search process. [Zela *et al.*, 2019] has in common with our work in that they also employed a flatness of local minima during the architecture search process, but in an indirect manner. They proposed an early-stopping search process to prevent overfitting on the validation set when the approximated sharpness of local minima exceeds the threshold. Similarly, [Chen and Hsieh, 2020; Wang *et al.*, 2021] tackled to alleviate fluctuating loss surface

and accuracy caused by the discretization of architecture parameters in DARTS [Liu *et al.*, 2018b]. We point out these previous similar works lack general usage on various NAS frameworks since they heavily depend on the DARTS [Liu *et al.*, 2018b] framework. Meanwhile, our method can be applied to any architecture search framework without dependence on the architecture parameters of DARTS, such as evolutionary-based search algorithm [Guo *et al.*, 2020].

3 Method

3.1 GeNAS: Generalization-Aware NAS With Flatness of Local Minima

GeNAS is aimed to search for network architectures with better generalization performance. To this end, we introduce a procedure for quantitatively estimating the flatness of an architecture’s converged minima as a search proxy measure $F_{val}(\cdot)$ as follows:

$$a^* = \operatorname{argmax}_{a \in A} F_{val}(W_A^*(a)). \quad (2)$$

Namely, we select the maximal flat architecture a^* by evaluating flatness of every *SubNet* extracted from the pre-trained *SuperNet* W_A^* . From the previous studies [Zhang *et al.*, 2018; Cha *et al.*, 2020] that empirically investigated the landscape of converged local minima, the neural networks having flat local minima where the changes of the validation loss around the local minima are relatively small show better generalization performance at the test phase. Based on these simple but effective empirical connections, we introduce a novel method that searches for the architecture with maximal loss flatness around converged minima which can be formulated as below, following [Zhang *et al.*, 2018]:

$$F_{val}(\theta) = \left(\sum_{i=1}^{t-1} \frac{L(\theta + N(\sigma_{i+1})) - L(\theta + N(\sigma_i))}{\sigma_{i+1} - \sigma_i} \right)^{-1}, \quad (3)$$

where $L(\theta)$ denotes validation loss value given weight parameter θ abbreviating $W_A^*(a)$, and $N(\sigma_i)$ denotes random Gaussian noise with its mean and standard deviation being 0 and σ_i , respectively. Namely, we inspect the flatness of minima near converged weight parameter space by injecting random Gaussian noise. The hyper-parameters σ denote the range for inspection of flatness and t denotes the number of perturbations. To perturb the weight parameters, we use unidirectional random noise, much more cost-efficient than recent flatness measuring approaches using Hessian [Yao *et al.*, 2019] and bidirectional random noise [He *et al.*, 2019] which can induce a considerable amount of computational cost. We observe that our choice is sufficient to discriminate architectures with high generalization performance.

Eq (2) and (3) would find architecture a^* having the flattest local minima in the entire search space, but a^* might have sub-optimal local minima far from the global minimum. In Figure 2, bottom-K architectures with the lowest ground test accuracy given by NAS-Bench-201 show the flattest local minima with relatively large loss values compared to the middle-K and top-K architectures. Therefore, naive investigation of the flatness of an architecture comes to achieve such sub-optimal

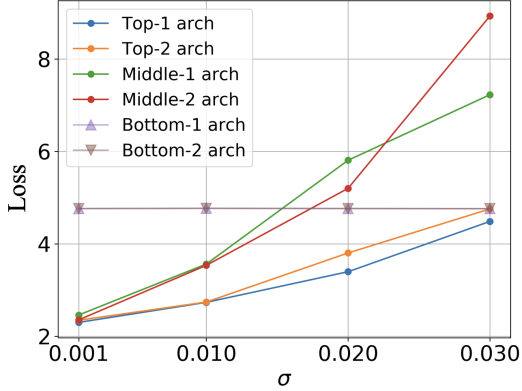


Figure 2: Validation loss curvatures of top-k, middle-k, bottom-k architectures sorted by the ground-truth test accuracy which is given by NAS-Bench-201 [Dong and Yang, 2020] on CIFAR-100.

architecture in terms of loss value. Note that the top-K architectures have the lowest loss values compared to middle and bottom architectures, equipped with flatness near converged minima. Correspondingly, considering both the flatness of loss landscape and the depth of minima is essential for excavating a generalizable architecture. For this reason, we add an additional term on Eq (3) to search for architectures with deep minima, along with flatness as follows:

$$F_{val}(\theta) = \left(\sum_{i=1}^{t-1} \left| \frac{L(\theta + N(\sigma_{i+1})) - L(\theta + N(\sigma_i))}{\sigma_{i+1} - \sigma_i} \right| \right) + \alpha \left| \frac{L(\theta + N(\sigma_1))}{\sigma_1} \right|^{-1} \quad (4)$$

Here, σ_1 denotes the smallest perturbation degree among σ , hence the second term inspects how low the loss value nearest converged minima is. The term α denotes the balancing coefficient term between flat and deep minima, which is set to 1 unless specified.

3.2 Searching With Combined Metrics

Recent works [Hosseini *et al.*, 2021; Mellor *et al.*, 2021] adopted a combined search metric for enhancing the performance of the resultant architecture. [Hosseini *et al.*, 2021] employed an integrated search metric where the conventional cross-entropy loss over a clean image is combined with approximately measured adversarial robustness lower bound to enhance test accuracy of both clean images and adversarially attacked images. Inspired by the weak correlation between the existing search metrics (e.g. angle) and flatness (Table 1), we target to explicitly fulfill the large headroom of conventional search metrics to find better generalizable architectures in terms of our proposed flatness-based search measure (Eq (4)). Formally, we combine the existing metrics with flatness as a search proxy measure as follows:

$$a^* = \operatorname{argmax}_{a \in A} S(W_A^*(a)) + \gamma \beta F_{val}(W_A^*(a)) \quad (5)$$

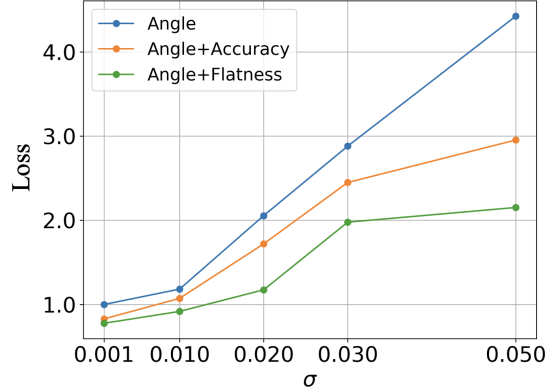


Figure 3: Test loss curvatures of architectures found by *Angle*, *Angle+Accuracy*, *Angle+Flatness*.

where S denotes conventional search metrics such as angle and validation accuracy, γ is a balancing parameter between the existing metric and flatness, and β is a normalization term, which is fixed as σ_1^{-1} , for matching scale of flatness term with the existing search metric.

4 Experiments

We first evaluate our proposed GeNAS framework on widely used benchmark dataset, ImageNet with DARTS [Liu *et al.*, 2018b] search space. Furthermore, we thoroughly conduct ablation studies with regard to the components of GeNAS on NAS-Bench-201 [Dong and Yang, 2020] benchmark. We refer the reader to the appendix for more experimental details. For better confirming robust generalization effect with regard to data distribution shift, we evaluate the found architectures on ImageNet variants (ImageNetV2 [Recht *et al.*, 2019],-A,-O [Hendrycks *et al.*, 2021]). Furthermore, we test the transferability of our excavated architectures onto other task domains, object detection, and semantic segmentation, with MS-COCO [Lin *et al.*, 2014] and Cityscapes [Cordts *et al.*, 2016] dataset.

4.1 ImageNet

Searching on CIFAR-10

We analyze the transferability of architectures found on small datasets such as CIFAR-10 and CIFAR-100 onto ImageNet. Specifically, we search architectures with 8 normal cells (i.e., *stride* = 1) and 2 reduction cells (i.e., *stride* = 2) on CIFAR-10/100, and transfer these normal / reduction cell architectures onto ImageNet by training from scratch and evaluating top-1 accuracy on ImageNet validation set. We compare our proposed FBS with other search metrics on CIFAR-10 in the upper part of Table 2. As a stand-alone search metric, flatness measure shows the best search performance among the other metrics including accuracy and angle with comparable FLOPs (\approx 0.6G) and parameters, when transferring searched architecture from CIFAR-10 onto ImageNet. Furthermore, when the angle is combined with flatness, loss landscape of found

Search Dataset	Search Metric	Params (M)	FLOPs (G)	Top-1 Acc (%)	Top-5 Acc (%)
CIFAR-10	Angle	5.3	0.59	75.70	92.45
	Accuracy	5.4	0.60	75.32	92.20
	Flatness	5.6	0.61	75.95	92.74
	Angle + Flatness	5.3 (+0.0)	0.59 (+0.00)	76.06 (+0.36)	92.77 (+0.32)
	Accuracy + Flatness	5.6 (+0.2)	0.61 (+0.01)	75.72 (+0.40)	92.59 (+0.39)
CIFAR-100	Angle	5.4	0.61	75.00	92.31
	Accuracy	5.4	0.60	75.37	92.23
	Flatness	5.2	0.58	76.05	92.64
	Angle + Flatness	5.4 (+0.0)	0.60 (-0.01)	75.72 (+0.72)	92.46 (+0.15)
	Accuracy + Flatness	5.4 (+0.0)	0.60 (+0.00)	75.85 (+0.48)	92.74 (+0.51)
ImageNet	Angle	5.4	0.60	75.09	92.30
	Accuracy	5.3	0.58	74.78	92.11
	Flatness	5.3	0.59	75.49	92.38
	Angle + Flatness	5.5 (+0.1)	0.60 (+0.00)	75.66 (+0.57)	92.62 (+0.32)
	Accuracy + Flatness	5.3 (+0.0)	0.59 (+0.01)	75.33 (+0.55)	92.41 (+0.30)

Table 2: Performance of various search metrics on ImageNet. The amount of change from adding *Flatness* term is denoted with blue color.

architecture becomes to be flatter and deeper as shown in Figure 3. As a result, search performance is further improved by 0.36% top-1 accuracy without any increase of either FLOPs or parameters. Also, the accuracy-based proxy measure also achieves performance gain when flatness is combined. The results show that our proposed flatness search metric indeed serves as a powerful search proxy measure for finding well-transferable architectures and also enhances the other search metrics to have a stronger ability to find architectures with better test generalization performance.

Searching on CIFAR-100

In middle part of Table 2, we analyze transferability of architectures found on CIFAR-100 onto ImageNet. The results show that flatness consistently reports significantly superior search performance even with fewer flops and parameters compared to ABS or PBS metrics, about 1.05% and 0.68% better top-1 accuracy, respectively. Furthermore, when flatness is appended onto angle and accuracy as a search proxy measure, top-1 accuracy drastically increases by 0.72% and 0.48%, respectively, which was consistently shown in CIFAR-10.

Searching on ImageNet

In the bottom part of Table 2, we directly search architectures on ImageNet and evaluate validation accuracy on ImageNet to compare in-domain search performance. Similar to the trend of the transfer experiments, our flatness metric achieves the best search performance compared to the existing search metrics and improves generalizability of them.

Comparison With SOTA NAS Methods

In Table 3, our GeNAS clearly represents large headroom compared to the other state-of-the-art NAS methods. Especially in comparison with SDARTS [Chen and Hsieh, 2020] which is a similar approach to GeNAS by using an implicit regularization for smoothing accuracy landscape, our GeNAS outperforms with a comparable number of FLOPs. Table 2 and 3 results show that our proposed flatness search metric indeed serves as

a powerful search proxy measure for finding well-transferable architectures and also enhances the other search metrics to have a stronger ability to find architectures with better test generalization performance.

4.2 Generalization Ability

For a more sophisticated investigation of generalization ability, we analyze GeNAS in terms of robustness towards data distribution shift and transferability onto various downstream tasks in Table 4.

Distribution Shift Robustness

To measure robustness towards data distribution shift, we evaluate our found architectures on ImageNet variants, ImageNet-V2 matched frequency [Recht *et al.*, 2019] and ImageNet-A [Hendrycks *et al.*, 2021], where the test-set is distinct from the original ImageNet validation set. The results demonstrate superior robustness compared to the other NAS methods. Our GeNAS widens the performance gap especially when the distribution shift is severe as in ImageNet-A, which has extremely confusing examples.

Task Generalization

Object Detection We evaluate the generalization capability of architectures found by GeNAS on the downstream task, specifically object detection. We firstly re-train architectures found on CIFAR-100 onto ImageNet, and finetune on MS-COCO [Lin *et al.*, 2014] dataset. For training, we adopt the default training strategy of RetinaNet [Lin *et al.*, 2017] from Detectron2 [Wu *et al.*, 2019]. We only replace the backbone network of RetinaNet for analyzing the sole impact of architectures found by each NAS method. The result shows that our GeNAS framework guided by the flatness measure clearly achieves the best AP scores. In case of RLNAS (angle) combined with flatness as a search metric, AP is enhanced by about 0.61%, without an increase of FLOPs or number of parameters.

Search Dataset	Method	Search Metric	Params (M)	FLOPs (G)	Top-1 Acc (%)	Top-5 Acc (%)
CIFAR-10	DARTS [Liu <i>et al.</i> , 2018b]	Val. loss	4.7	0.57	73.3	91.3
	PC-DARTS [Xu <i>et al.</i> , 2019]	Val. loss	5.3	0.59	74.9	92.2
	FairDARTS-B [Chu <i>et al.</i> , 2020]	Val. loss	4.8	0.54	75.1	92.5
	P-DARTS [Chen <i>et al.</i> , 2019]	Val. loss	4.9	0.56	75.6	92.6
	DropNAS [†] [Hong <i>et al.</i> , 2020]	Val. loss	5.4	0.60	76.0	92.8
	SANAS [Hosseini and Xie, 2022]	Val. loss	4.9	0.55	75.2	91.7
	SPOS [Guo <i>et al.</i> , 2020]	Val. acc	5.4	0.60	75.3	92.2
	MF-NAS [Xue <i>et al.</i> , 2022]	Val. acc	4.9	0.55	75.3	-
	Shapley-NAS [Xiao <i>et al.</i> , 2022]	Shapley value	5.1	0.57	75.7	-
	RLNAS [Zhang <i>et al.</i> , 2021]	Angle	5.3	0.59	75.7	92.5
	SDARTS-RS [Chen and Hsieh, 2020]	Flatness	5.5	0.61	75.5	92.7
	SDARTS-ADV [Chen and Hsieh, 2020]	Flatness	5.5	0.62	75.6	92.4
GeNAS (Ours)	Flatness	5.6	0.61	76.0	92.7	
GeNAS (Ours)	Angle + Flatness	5.3	0.59	76.1	92.8	
CIFAR-100	PC-DARTS [Xu <i>et al.</i> , 2019]	Val. loss	5.3	0.59	74.8	92.2
	DropNAS [†] [Hong <i>et al.</i> , 2020]	Val. loss	5.1	0.57	75.1	92.3
	P-DARTS [Chen <i>et al.</i> , 2019]	Val. loss	5.1	0.58	75.3	92.5
	SPOS [Guo <i>et al.</i> , 2020]	Val. acc	5.4	0.60	75.4	92.2
	RLNAS [Zhang <i>et al.</i> , 2021]	Angle	5.4	0.61	75.0	92.3
	GeNAS (Ours)	Flatness	5.2	0.58	76.1	92.6
	GeNAS (Ours)	Angle + Flatness	5.4	0.60	75.7	92.5

Table 3: ImageNet performance comparison of SOTA NAS methods searched with DARTS search space on CIFAR-10 and CIFAR-100 dataset. [†] denotes that SE [Hu *et al.*, 2018] module is excluded for fair comparison with other methods.

Method	Search Measure	Params (M)	FLOPs (G)	ImageNet-V2 <i>Acc</i>	ImageNet-A <i>Acc</i>	COCO <i>AP</i>	Cityscapes <i>mIoU</i>
PC-DARTS [Xu <i>et al.</i> , 2019]	Val. loss	5.3	0.59	62.53	3.85	35.56	70.68
DropNAS [Hong <i>et al.</i> , 2020]	Val. loss	5.1	0.57	63.14	4.28	36.39	71.16
SPOS [Guo <i>et al.</i> , 2020]	Val. acc	5.4	0.60	62.84	3.91	36.04	71.70
RLNAS [Zhang <i>et al.</i> , 2021]	Angle	5.4	0.61	62.95	3.81	35.98	70.84
SDARTS-ADV [Zhang <i>et al.</i> , 2021]	Flatness	5.5	0.62	62.88	4.24	36.36	71.77
GeNAS (Ours)	Flatness	5.2	0.58	63.38	5.65	37.05	72.58
GeNAS (Ours)	Angle + Flatness	5.4	0.60	63.32	4.37	36.59	72.05

Table 4: Comparison with SOTA NAS methods on various ImageNet variants and downstream tasks (object detection with COCO [Lin *et al.*, 2014] and segmentation with Cityscapes [Cordts *et al.*, 2016]).

Semantic Segmentation We also test the generalization of our GeNAS on Semantic Segmentation task with Cityscapes [Cordts *et al.*, 2016] dataset. Based on the DeepLab-v3 [Chen *et al.*, 2017], we only replaced the backbone network and trained with MMSegmentation [Contributors, 2020] framework. The results demonstrate the effectiveness of our flatness-guided architectures with a large performance margin. Consistently, our flatness guidance ensures a large performance gain, about 1.21%, when added onto angle-based search.

4.3 Ablation Study

To better analyze our proposed FBS-based GeNAS framework, we conduct an ablation study of each component and hyperparameters consisting of GeNAS.

Flatness Range

We analyze the effect of range of inspecting flatness near converged local minima in Table 5. The results demonstrate that searching flat architectures within too small area near

converged minima (1st row in Table 5) is not sufficient for discriminating generalizable architectures. When σ is increased to $\{2e - 3, 1e - 2, 2e - 2\}$, Kendall’s Tau is considerably improved, while further widening the flatness inspection range (4th row in Table 5) only significantly degrades the search performance on various datasets.

Deep and Low Minima

We further investigate the effect of searching architectures equipped with not only flatness but also the depth of loss landscape near converged minima. Specifically, we adjust α in Eq (4), where $\alpha = 0$ denotes searching with only flatness of local minima. Results on Table 6 demonstrate that as α value increases from zero to one, search performance is drastically enhanced, indicating the indispensability of searching with both flatness and depth of minima. Note that $\alpha = 0$ case can search out a sub-optimal architecture that has highly flat loss curvature but its loss values near local minima are too high, as shown in Figure 2. When α is further increased to $\alpha > 1$, Kendall’s Tau rank correlation starts to decrease, denoting that

σ	Kendall's Tau	
	CIFAR-10	IN16-120
$\{1e-6, 5e-6, 1e-5\}$	0.5756	0.5524
$\{5e-4, 1e-3, 2e-3\}$	0.5770	0.5531
$\{2e-3, 1e-2, 2e-2\}$	0.6047	0.5800
$\{2e-3, 2e-2, 4e-2\}$	0.5416	0.2364

Table 5: Kendall’s Tau on the NAS-Bench-201 search space according to the perturbation range σ , inspecting the effect of flatness range near local minima. IN16-120 denotes ImageNet16-120 dataset [Dong and Yang, 2020].

	$\alpha = 0$	$\alpha = 0.1$	$\alpha = 0.5$	$\alpha = 1$	$\alpha = 4$	$\alpha = 16$
Kendall’s Tau (CIFAR-10)	0.1777	0.4026	0.5890	0.6047	0.5898	0.5820

Table 6: Kendall’s Tau on CIFAR-10 with different α in Eq (4).

searching with largely depending on the depth of converged minima is not optimal for discriminating better generalizable architectures.

Perturbation Methodology

To quantitatively measure flatness of loss landscape, all the weight parameters of a given network are perturbed with random direction following Gaussian distribution as in Eq (4). Here, we investigate the effect of perturbation positions and directions. In Table 7, perturbing only weight parameters of target search cells (i.e. excluding stem *conv* layer and final *fully-connected* layer) only harms Kendall’s Tau. Moreover, with regard to the perturbation directions, strongly perturbing the given models’ parameters across the hessian eigenvectors [Yao *et al.*, 2019] suffers from a slight decrease of Kendall’s Tau (Table 7) with large computational overhead induced by approximation of hessian.

Effect of Flatness on ABS

We analyze the effect of integrating flatness on ABS. Specifically, we adjust γ in Eq (5), which balances the coefficient concerning the ratio of flatness to angle term. In Table 8, integrating flatness with a small proportion to angle mildly improves top-1 accuracy. As γ increases, top-1 accuracy of searched architecture gradually increases to reach 0.72% improvement over $\gamma = 0$ (ABS) case.

4.4 Search Cost Analysis

In Figure 4, we compare the required architecture search time of GeNAS with the other SOTA NAS frameworks. We measured the execution time spent for the *SuperNet* training and the search process, using a single NVIDIA V100 GPU. Our required search time is competitive to the other NAS methods while exhibiting shortened time compared to the other flatness-based search method (i.e., SDARTS-ADV).

5 Conclusion

This paper demonstrates that the flatness of local minima can be directly employed as a proxy of discriminating and search-

Perturbation Position	Perturbation Direction	Kendall’s Tau
All	Random	0.6047
Search Cells	Random	0.5612 (-0.0435)
All	Hessian	0.5908 (-0.0139)

Table 7: Ablation study of perturbation position and direction on CIFAR-10 with NAS-BENCH-201 [Dong and Yang, 2020] search space. *All* denotes perturbing all the weight parameters of a given network, while *Search Cells* denotes perturbing only the weight parameters of search cells. The quantities in the parentheses denote the amount of change compared to the default case (first row).

γ	Flatness (%)	Top-1 Acc (%)	Top-5 Acc (%)
0	0	75.00	92.31
0.5	20	75.22 (+0.22)	92.39 (+0.08)
1.5	43	75.58 (+0.58)	92.44 (+0.13)
6	76	75.63 (+0.63)	92.54 (+0.23)
16	89	75.72 (+0.72)	92.46 (+0.15)

Table 8: Search performance of *Angle + Flatness* with different γ values, where searched on CIFAR-100 and transferred onto ImageNet. *Flatness (%)* denotes the average ratio of *Flatness* compared to *Angle* during evaluation of architectures on evolutionary search algorithm. The quantities in the parentheses denote the amount of change compared to the $\gamma = 0$ case.




Figure 4: Comparison of search cost with the SOTA NAS frameworks.

ing for generalizable architectures. Based on the quantitative benchmark experiments on various search spaces and datasets, we demonstrate the superior generalizability of our flatness-based search over conventional search metrics, while showing comparable or even better search performance compared to recent state-of-the-art NAS frameworks. We further analyze the insufficient generalizability of conventional search metrics in terms of the flatness of local minima. Consequently, integrating conventional search metrics with our proposed flatness measure can further lead to significantly boosting search performance. We also demonstrate superior generalization capability of GeNAS on the downstream object detection and semantic segmentation tasks while showing great robustness with regard to the data distribution shift.

References

- [Baker *et al.*, 2016] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. Designing neural network architectures using reinforcement learning. *arXiv preprint arXiv:1611.02167*, 2016.
- [Bender *et al.*, 2018] Gabriel Bender, Pieter-Jan Kindermans, Barret Zoph, Vijay Vasudevan, and Quoc Le. Understanding and simplifying one-shot architecture search. In *International Conference on Machine Learning*, pages 550–559. PMLR, 2018.
- [Brock *et al.*, 2017] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Smash: one-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344*, 2017.
- [Cha *et al.*, 2020] Sungmin Cha, Hsiang Hsu, Taebaek Hwang, Flavio P Calmon, and Taesup Moon. Cpr: Classifier-projection regularization for continual learning. *arXiv preprint arXiv:2006.07326*, 2020.
- [Chaudhari *et al.*, 2019] Pratik Chaudhari, Anna Choromanska, Stefano Soatto, Yann LeCun, Carlo Baldassi, Christian Borgs, Jennifer Chayes, Levent Sagun, and Riccardo Zecchina. Entropy-sgd: Biasing gradient descent into wide valleys. *Journal of Statistical Mechanics: Theory and Experiment*, 2019(12):124018, 2019.
- [Chen and Hsieh, 2020] Xiangning Chen and Cho-Jui Hsieh. Stabilizing differentiable architecture search via perturbation-based regularization. In *International conference on machine learning*, pages 1554–1565. PMLR, 2020.
- [Chen *et al.*, 2017] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017.
- [Chen *et al.*, 2019] Xin Chen, Lingxi Xie, Jun Wu, and Qi Tian. Progressive differentiable architecture search: Bridging the depth gap between search and evaluation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1294–1303, 2019.
- [Chen *et al.*, 2021] Wuyang Chen, Xinyu Gong, and Zhangyang Wang. Neural architecture search on imagenet in four gpu hours: A theoretically inspired perspective. *arXiv preprint arXiv:2102.11535*, 2021.
- [Chrabaszcz *et al.*, 2017] 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.
- [Chu *et al.*, 2020] Xiangxiang Chu, Tianbao Zhou, Bo Zhang, and Jixiang Li. Fair darts: Eliminating unfair advantages in differentiable architecture search. In *European conference on computer vision*, pages 465–480. Springer, 2020.
- [Contributors, 2020] MMSegmentation Contributors. MM-Segmentation: Openmmlab semantic segmentation toolbox and benchmark. <https://github.com/open-mmlab/mms Segmentation>, 2020. Accessed: 2022-06-24.
- [Cordts *et al.*, 2016] Marius Cordts, Mohamed Omran, Sebastian Ramos, Timo Rehfeld, Markus Enzweiler, Rodrigo Benenson, Uwe Franke, Stefan Roth, and Bernt Schiele. The cityscapes dataset for semantic urban scene understanding. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 3213–3223, 2016.
- [Dong and Yang, 2020] Xuanyi Dong and Yi Yang. Nasbench-201: Extending the scope of reproducible neural architecture search. *arXiv preprint arXiv:2001.00326*, 2020.
- [Dosovitskiy *et al.*, 2020] 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.
- [Guo *et al.*, 2020] Zichao Guo, Xiangyu Zhang, Haoyuan Mu, Wen Heng, Zechun Liu, Yichen Wei, and Jian Sun. Single path one-shot neural architecture search with uniform sampling. In *European Conference on Computer Vision*, pages 544–560. Springer, 2020.
- [He *et al.*, 2016] 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.
- [He *et al.*, 2019] Haowei He, Gao Huang, and Yang Yuan. Asymmetric valleys: Beyond sharp and flat local minima. *Advances in neural information processing systems*, 32, 2019.
- [Hendrycks *et al.*, 2021] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. *CVPR*, 2021.
- [Hoffer *et al.*, 2017] Elad Hoffer, Itay Hubara, and Daniel Soudry. Train longer, generalize better: closing the generalization gap in large batch training of neural networks. *Advances in neural information processing systems*, 30, 2017.
- [Hong *et al.*, 2020] Weijun Hong, Guilin Li, Weinan Zhang, Ruiming Tang, Yunhe Wang, Zhenguo Li, and Yong Yu. Dropnas: Grouped operation dropout for differentiable architecture search. In Christian Bessiere, editor, *Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI-20*, pages 2326–2332. International Joint Conferences on Artificial Intelligence Organization, 7 2020. Main track.
- [Hosseini and Xie, 2022] Ramtin Hosseini and Pengtao Xie. Saliency-aware neural architecture search. *Advances in Neural Information Processing Systems*, 35:14743–14757, 2022.
- [Hosseini *et al.*, 2021] Ramtin Hosseini, Xingyi Yang, and Pengtao Xie. Dsrna: Differentiable search of robust neural architectures. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 6196–6205, 2021.
- [Hu *et al.*, 2018] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE con-*

- ference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [Hu *et al.*, 2020] 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.
- [Jastrzębski *et al.*, 2017] Stanisław Jastrzębski, Zachary Kenton, Devansh Arpit, Nicolas Ballas, Asja Fischer, Yoshua Bengio, and Amos Storkey. Three factors influencing minima in sgd. *arXiv preprint arXiv:1711.04623*, 2017.
- [Kendall, 1938] Maurice G Kendall. A new measure of rank correlation. *Biometrika*, 30(1/2):81–93, 1938.
- [Keskar *et al.*, 2016] Nitish Shirish Keskar, Dheevatsa Mudigere, Jorge Nocedal, Mikhail Smelyanskiy, and Ping Tak Peter Tang. On large-batch training for deep learning: Generalization gap and sharp minima. *arXiv preprint arXiv:1609.04836*, 2016.
- [Krizhevsky *et al.*, 2012] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012.
- [Lin *et al.*, 2014] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014.
- [Lin *et al.*, 2017] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *Proceedings of the IEEE international conference on computer vision*, pages 2980–2988, 2017.
- [Liu *et al.*, 2018a] 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.
- [Liu *et al.*, 2018b] Hanxiao Liu, Karen Simonyan, and Yiming Yang. Darts: Differentiable architecture search. *arXiv preprint arXiv:1806.09055*, 2018.
- [Mellor *et al.*, 2021] 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.
- [Oymak *et al.*, 2021] Samet Oymak, Mingchen Li, and Mahdi Soltanolkotabi. Generalization guarantees for neural architecture search with train-validation split, 2021.
- [Pereyra *et al.*, 2017] Gabriel Pereyra, George Tucker, Jan Chorowski, Łukasz Kaiser, and Geoffrey Hinton. Regularizing neural networks by penalizing confident output distributions. *arXiv preprint arXiv:1701.06548*, 2017.
- [Recht *et al.*, 2019] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar. Do imagenet classifiers generalize to imagenet? In *International Conference on Machine Learning*, pages 5389–5400. PMLR, 2019.
- [Shu *et al.*, 2019] Yao Shu, Wei Wang, and Shaofeng Cai. Understanding architectures learnt by cell-based neural architecture search. *arXiv preprint arXiv:1909.09569*, 2019.
- [Wang *et al.*, 2021] Xiaofang Wang, Shengcao Cao, Mengtian Li, and Kris M Kitani. Neighborhood-aware neural architecture search. *arXiv preprint arXiv:2105.06369*, 2021.
- [Wu *et al.*, 2019] Yuxin Wu, Alexander Kirillov, Francisco Massa, Wan-Yen Lo, and Ross Girshick. Detectron2. <https://github.com/facebookresearch/detectron2>, 2019.
- [Xiao *et al.*, 2022] Han Xiao, Ziwei Wang, Zheng Zhu, Jie Zhou, and Jiwen Lu. Shapley-nas: Discovering operation contribution for neural architecture search. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11892–11901, 2022.
- [Xie *et al.*, 2018] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. *arXiv preprint arXiv:1812.09926*, 2018.
- [Xu *et al.*, 2019] Yuhui Xu, Lingxi Xie, Xiaopeng Zhang, Xin Chen, Guo-Jun Qi, Qi Tian, and Hongkai Xiong. Pcdarts: Partial channel connections for memory-efficient architecture search. *arXiv preprint arXiv:1907.05737*, 2019.
- [Xue *et al.*, 2022] Chao Xue, Xiaoxing Wang, Junchi Yan, and Chun-Guang Li. A max-flow based approach for neural architecture search. In *Computer Vision—ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XX*, pages 685–701. Springer, 2022.
- [Yao *et al.*, 2019] Zhewei Yao, Amir Gholami, Kurt Keutzer, and Michael Mahoney. Pyhessian: Neural networks through the lens of the hessian. *arXiv preprint arXiv:1912.07145*, 2019.
- [Zela *et al.*, 2019] Arber Zela, Thomas Elsken, Tonmoy Saikia, Yassine Marrakchi, Thomas Brox, and Frank Hutter. Understanding and robustifying differentiable architecture search. *arXiv preprint arXiv:1909.09656*, 2019.
- [Zhang *et al.*, 2018] Ying Zhang, Tao Xiang, Timothy M Hospedales, and Huchuan Lu. Deep mutual learning. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4320–4328, 2018.
- [Zhang *et al.*, 2021] Xuanyang Zhang, Pengfei Hou, Xiangyu Zhang, and Jian Sun. Neural architecture search with random labels. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10907–10916, 2021.
- [Zoph *et al.*, 2018] 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.