

LISSNAS: Locality-based Iterative Search Space Shrinkage for Neural Architecture Search

Bhavna Gopal¹, Arjun Sridhar¹, Tunhou Zhang¹ and Yiran Chen¹

¹Duke University

{bhavna.gopal, arjun.sridhar, tunhou.zhang, yiran.chen}@duke.edu

Abstract

Search spaces hallmark the advancement of Neural Architecture Search (NAS). Large and complex search spaces with versatile building operators and structures provide more opportunities to brew promising architectures, yet pose severe challenges on efficient exploration and exploitation. Subsequently, several search space shrinkage methods optimize by selecting a single sub-region that contains some well-performing networks. Small performance and efficiency gains are observed with these methods but such techniques leave room for significantly improved search performance and are ineffective at retaining architectural diversity. We propose LISSNAS, an automated algorithm that shrinks a large space into a diverse, small search space with SOTA search performance. Our approach leverages locality, the relationship between structural and performance similarity, to efficiently extract many pockets of well-performing networks. We showcase our method on an array of search spaces spanning various sizes and datasets. We accentuate the effectiveness of our shrunk spaces when used in one-shot search by achieving the best Top-1 accuracy in two different search spaces. Our method achieves a SOTA Top-1 accuracy of 77.6% in ImageNet under mobile constraints, best-in-class Kendal-Tau, architectural diversity, and search space size.

1 Introduction

Neural Architecture Search (NAS) [Cai *et al.*, 2019; Liu *et al.*, 2019; Tan *et al.*, 2018] has attracted attention by achieving state-of-the-art performance on several tasks. Search space selection serves as a crucial predecessor to NAS by defining the space of architectures being considered during search. Large and complex search spaces have an increased likelihood of containing better-performing candidates. Additionally, technical advancements in operation choices, architecture complexity, and network depth/width have resulted in an increased need for such spaces. However, these large and complex search spaces come with a slew of optimization and efficiency challenges [Hu *et al.*, 2020].

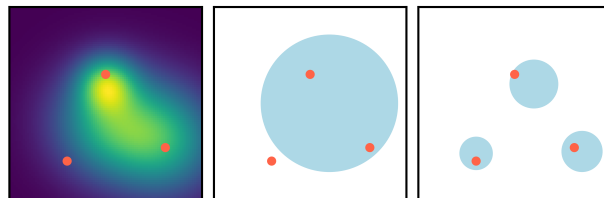


Figure 1: The graphs represent a 2-D embedding of a high-dimensional search space. The red dots represent the 3 best-performing networks in the space. *Left*: A heatmap with lighter colors representing regions with higher accuracy. *Middle*: The blue region represents the shrunk search space produced by other methods. *Right*: We produce a smaller search space (indicated by smaller blue area) that is a combination of the 3 regions.

Past efforts have focused on improving NAS search techniques to accommodate increasing search space size and complexity. Nevertheless, search space improvements itself have been relatively underutilized for driving improvement. More recently, research has shown that search space shrinkage algorithms can significantly improve the performance of a variety of NAS methods [Hu *et al.*, 2020]. As a result, shrinkage methods were proposed to shrink the search space, while retaining some top-performing networks. Most of these approaches shrink the original space to a single region consisting of a relatively higher proportion of well-performing networks. To do so, these algorithms utilize *operational pruning* - either explicitly or implicitly pruning lower-ranked operational choices from the space. Here, ranking and pruning are holistic - all architectures within the space containing the operation contribute to its perceived value. From this standpoint, if an operation is frequently found in poorly-performing architectures then the rank of the operation is likely lower. This results in all networks containing lower-ranked operations being removed from the space. However, as illustrated in Figure 1, operational pruning is too aggressive. Inevitably, this technique removes well-performing networks from other regions of the space that typically contain a lower proportion of optimal networks. Hence, we aim to better differentiate and retain optimal architectures throughout the space to improve search performance and alleviate the challenges brought by large search spaces.

We propose a new paradigm, LISSNAS: Locality-based Iterative Search space Shrinkage for Neural Architecture Search. In contrast to existing approaches, our method comprehensively extracts a smaller search space with an extremely high proportion of diverse, well-performing networks from the original space. This directly results in improved NAS performance. LISSNAS aims to incorporate more well-performing networks, prune poor-performing networks, and maintain diversity across the FLOP range in our search space.

To accomplish these goals, our algorithm hinges on *architectural pruning* by leveraging locality, the property that structural similarity between architectures correlates to performance similarity. We begin by sampling networks from the original space. Then, using a predictor, we evaluate the samples and leverage locality to construct more clusters of networks that are likely to be well-performing. Eventually, we extract and selectively query pockets of likely well-performing networks in contrast to querying all architectures in the space. We note that even with a predictor, querying all architectures in the space is time-consuming and computationally expensive as spaces can contain upwards of 10^{18} architectures. After several automated iterations of refinement, our output, final, search space consists of a combination of well-performing pockets of networks.

We showcase our method on a vast array of search spaces with various sizes across datasets. We accentuate the effectiveness of our shrunk spaces when used in one-shot search by achieving the best Top-1 accuracy in two different search spaces. Our method achieves a SOTA Top-1 accuracy of 77.6% in ImageNet under mobile constraints, best-in-class Kendall-Tau, architectural diversity, and search space size. Overall, our contributions are as follows:

1. We successively establish locality preservation quantitatively within 3 popular search spaces (ShuffleNetV2, NASBench101, TransNASBench) and 6 datasets (ImageNet, CIFAR10, 4 datasets in TransNASBench - scene, object detection, jigsaw, and semantic segmentation).
2. We introduce a new metric for quantitative comparison across different shrinkage algorithms, the *shrink index*. This index captures the improved probability of sampling a well performing network in our shrunk search space.
3. We demonstrate that LISSNAS improves overall NAS performance (Top-1) in addition to several search space quality metrics. We efficiently achieve SOTA Top-1 accuracy of 77.6% on ImageNet under mobile constraints.

2 Related Work

2.1 General NAS

Designing a CNN poses a challenge to researchers looking to optimize for many constraints such as latency, model size, and accuracy. Recent work in NAS attempts to automate the process of designing and discovering CNN architectures that meet these constraints. In 2017, NAS with Reinforcement Learning [Zoph and Le, 2017] introduced the concept of NAS utilizing an RL-based controller that iteratively proposed and trained candidates. Training the controller and candidates

was computationally exorbitant. Since then, several utilized genetic algorithms [Real *et al.*, 2019; Real *et al.*, 2017; Xue *et al.*, 2021], Bayesian optimization [Bergstra *et al.*, 2013; Mendoza *et al.*, 2016], and predictors [Wen *et al.*, 2020; White *et al.*, 2021; Wei *et al.*, 2022].

The next significant milestone was the introduction of weight sharing through one-shot networks, also referred to as supernetworks. SMASH [Brock *et al.*, 2018] employs a One-Shot Network that uses a set of shared weights to train many candidates simultaneously. At the time, this approach significantly reduced training time but limited architecture exploration, impacting performance. To alleviate this limitation, OFA [Cai *et al.*, 2019], Attentive NAS [Wang *et al.*, 2020], and several other works train these networks through novel path sampling methods and a host of optimization techniques [Wang *et al.*, 2021; Luo *et al.*, 2019a; Luo *et al.*, 2019b; Pham *et al.*, 2018; Xie *et al.*, 2018; Zhang *et al.*, 2020]. Overall, one-shot networks proved to be efficient and resulted in SOTA NAS performances across a variety of tasks.

2.2 Search Space

Search Space Shrinkage Motivation. Angle-Based Search (ABS) [Hu *et al.*, 2020] investigated behaviors of NAS methods on various shrunk search spaces. The results show that shrunk search spaces have the potential to enhance existing NAS algorithms across many different search methods. Neural Search Space Evolution (NSE) [Ci *et al.*, 2020] presents the drawbacks of large search spaces and acknowledges the importance of the size and design of the search space in supernet training. Additionally, several others [Sciuto *et al.*, 2019; Zhang *et al.*, 2018b] have shown that enlarging a search space typically results in significant accuracy loss for the final searched architecture. Lastly, SPOS [Guo *et al.*, 2019] determined that search spaces of reduced cardinality could achieve higher correlations.

Search Space Shrinkage. Several methods have attempted to shrink the search space while preserving the top-performing networks [Hu *et al.*, 2020; Xia and Ding, 2020; You *et al.*, 2020; Huang *et al.*, 2021; Radosavovic *et al.*, 2020; Tyagi *et al.*, 2022]. GreedyNAS [You *et al.*, 2020] proposes an automated NAS optimization strategy that greedily samples well-performing networks during training. NSE [Ci *et al.*, 2020] simply segments the search space and trains in a guided way similar to GreedyNAS. These works fundamentally use a controller to efficiently sample networks in training and evaluation without truly constructing a new search space. ABS drops the lowest-performing operation choices until the search space is shrunk to the desired size. Similar to ABS, PADNAS [Xia and Ding, 2020] also prunes operations but does so on a layer-wise basis.

In 2020, several manual (human-dependent) techniques were proposed to build optimal search spaces. Designing Network Design Spaces (DDS) [Radosavovic *et al.*, 2019; Radosavovic *et al.*, 2020] utilizes statistical analysis to assist researchers with iterative search space shrinkage. Additionally, Visual Steering for One-Shot [Tyagi *et al.*, 2022], uses promising structures proposed by specialists to build a search space and eventually converges at the optimal network.

Both of these approaches are time-consuming and require iterations of domain expert input for success.

3 Definitions and Theoretical Motivation

3.1 Definitions, Instantiations, and Assumptions

Our work utilizes specialized terminology such as locality, search space, model distribution, edit distance, isomorphism, and neighbors. For clarity on the definition of these terms please refer to Appendix B.

This work makes the following assumptions:

1. All potential changes that can be made to a network in a search space are of equal value (Appendix G).
2. To determine the statistics of our final shrunk search spaces, we assume that a distribution of networks is more indicative than point sampling networks [Radosavovic *et al.*, 2019].

3.2 Benchmark

We investigate locality and our algorithm’s performance on NASBench101, TransNASBench Macro, NASBench301, and ShuffleNetV2. For a more detailed description of these search spaces and respective datasets, please reference the Appendix C.

3.3 Theoretical Motivation

Generally speaking, NAS is more optimal if a comparable (or better) top-1 accuracy is obtained from search in less time. We now explain how a smaller search space with a higher density of well-performing networks increases NAS search performance and efficiency (optimality). To do so, we use the example of training a supernet.

Most large search spaces contain some proportion of weak-performing architectures. In a standard supernet training method, training involves sampling architectures from the search space. Let us assume that we define our search space as A , and that A can be cleanly partitioned into A_{good} and A_{bad} [Huang *et al.*, 2021]. These partitions refer to potentially optimal and suboptimal networks respectively. Definitionally:

$$A = A_{good} \cup A_{bad}, A_{good} \cap A_{bad} = \emptyset \quad (1)$$

Assume that we are given an algorithm F to remove relatively poor-performing architectures from a search space. Since $F(A) = A_{shrunk}$, the shrunk search space has reduced cardinality as well as contains a higher proportion of optimal networks. During training, the supernet will now sample more optimal architectures at an increased frequency. Firstly, since the cardinality of the search space decreases, the training efficiency improves. Secondly, an increased proportion of optimal networks in A_{shrunk} enables the supernet to focus training and searching on superior regions of the search space. This results in better performance ranking.

It follows that we aim to shrink and improve our search space by pruning networks from A_{bad} . Unfortunately, we are typically unaware of the exact partition in A and are hence unable to precisely determine which networks to prune, without ground truth accuracies.

4 Methods

4.1 Assumptions

To reiterate, the goal of the proposed pipeline is to shrink a large search space by identifying and extracting well-performing search space subsets. We begin by stating the assumptions of our method and empirically verify locality preservation. We then utilize our observations to develop our method and obtain results. Lastly, to bolster our approach, we also perform some ablation experiments.

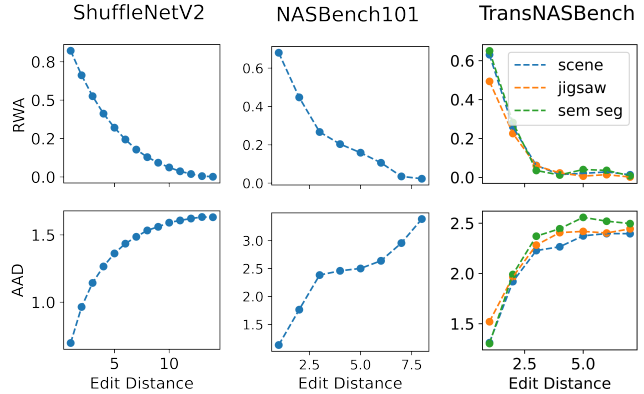


Figure 2: The top row of graphs demonstrate RWA. The correlation drop-off is steep and the correlations are weak ($< .3$) at one-third of the total edit distance available. The bottom row of graphs show the AAD. The trends are consistent across all search spaces and tasks.

4.2 Locality Preservation

Before locality exploitation, we empirically verify that locality is preserved across our search space types, sizes, and tasks. To do so, we illustrate Random Walk Autocorrelation (RWA) and compute the averaged Absolute Accuracy Difference (AAD) between generated neighbors and original networks.

To compute the AAD and RWA, we define our metric of similarity between networks as the *edit distance* between networks - the minimum number of changes required to change one network into another. Traditionally, computing the edit distance between two networks is computationally expensive. To alleviate the burden of calculating pairwise edit distances, we formulate our problem in terms of neighbor generation as opposed to edit distance computation, i.e: instead of computing the edit distance between two sampled networks, our method randomly samples a network and makes changes (up to a certain edit distance) to obtain another network. The generated network is referred to as a *neighbor*.

Firstly, we demonstrate the RWA between the accuracies of the sampled network and the generated neighbor [Ying *et al.*, 2019]. RWA is defined as the autocorrelation of the accuracies of networks as a random walk of changes is made. As the number of changes increases the correlation of accuracies with the original network should decrease. Our observations match our hypothesis - the RWA decreases as the number of changes made to a network increases as shown in Figure 2.

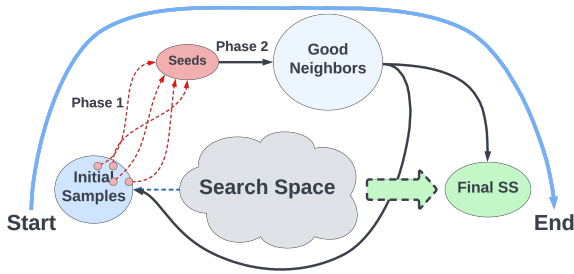


Figure 3: LISSNAS begins by sampling networks from the space. Next in Phase 1, the best samples are taken as the seed networks. These seeds are used in Phase 2 to generate good neighbors through locality exploitation. Finally, we determine whether to continue sampling and refining the space or terminate with our current shrunk search space.

Our RWA results are corroborated by determining the AAD between networks at increasing edit distances. We observe that as the number of changes made to a network increases, the accuracy gap (between the original and the generated neighbor) increases. This directly demonstrates locality by accentuating that networks that have more structural similarity perform similarly. Notably, all of our search spaces preserve locality and we summarize our observations in Figure 2. Our observations give rise to the premise of our algorithm. If we find a relatively well-performing network within a search space, we generate neighbor networks to add to our new shrunk space. The next section outlines the proposed locality-guided shrinkage algorithm.

4.3 Locality-based Iterative Search Space Shrinkage

Figure 3 appropriately highlights the pipeline of our shrinkage approach. We simplify our approach to 2 main stages: seed search and neighbor generation.

1. **Seed Search:** search for well-performing networks within a given search space
 - (a) Our initial networks are obtained by randomly sampling the original space. Note that we sample at least 1000 networks from each of our search spaces based on our understanding of network distributions.
 - (b) Since we are unaware of the ground truth accuracies of our initial samples, an accuracy predictor determines performance. Similar to NASBench301, we utilize XGBoost as our accuracy predictor. NASBench301 outlines the effectiveness of this predictor in modeling performances while providing low-cost inferences.
 - (c) Based on our predicted accuracies, we select the top-performing networks and refer to these as *seeds*.
2. **Neighbor Generation:** generate neighbors for our seeds
 - (a) For each seed, we generate several neighbor networks at small edit distances. To do so, we de-

termine a threshold edit distance. This is the edit distance up to which we can generate neighbors for each seed. Through our analysis of RWA, we found a good threshold is 1/3 of the total available edit distance as shown in Figure 2. Note: Evaluating all architectures is prohibitively expensive. We exploit locality to generate neighbors in order to guide search. Details on this are in the ablation study section (Appendix H).

- (b) The seed networks along with their neighbors constitute the shrunk search space with a higher proportion of optimal nets. We compute the mean accuracy of the constructed search space and perform more iterations if the average accuracy increases. This entails going back to the seed selection step and determining new seeds. Since locality is preserved within the space, a neighbor is likely to perform similarly to its seed.
- (c) If the average accuracy decreases or plateaus, we output our previous iteration’s search space. From this perspective, our algorithm shrinks the search space to the maximum extent possible.

4.4 Optimality Verification

Once we obtain the output shrunk search space, we verify its performance. Primarily, our algorithm hinges on reduced cardinality and an increased proportion of well-performing networks. We verify these metrics in live time (as the algorithm runs) by ensuring that the size of the search space reduces while the average search space accuracy increases. However, for easier comprehension and generalizability, we would like to encapsulate these baseline metrics through quantitative parameters. Given the lack of standardized metrics (or approaches) for search space quality evaluation, how do we verify that a search space has shrunk effectively?

We first clarify our theoretical definition of *effective shrinkage*. Understanding this term is crucial for a preliminary indication of the improvement and effectiveness of our algorithm across iterations. We consider a shrinkage algorithm to be both *optimal* and *effective* if these conditions are met: 1) the cardinality of the search space decreases with each iteration; and 2) the shrunk search space has a higher probability of containing well-performing networks than the original. The second condition implies that the probability of the search space containing a well-performing network increases as the iterations increase until algorithm termination.

Secondly, to verify that our search space has shrunk effectively, we determine the quality of our output shrunk search spaces using a myriad of other metrics and classical statistics. However, in terms of our baseline conditions, our first condition is easily quantified as the difference in size between the output shrunk search space and the original. For the second condition, we define the *shrink index*:

The Shrink Index. If we uniformly sample a network n from A , then let’s say that the probability that $n \in A_{good}$ is $p = |A_{good}|/|A|$. Likewise, the probability that $n \in A_{bad}$ is then $1 - p$. If we proceed to sample multiple networks independently and based on the results of a binomial distribution, we have the theorem:

Theorem 1. *If n paths are sampled uniformly i.i.d. from A , then it holds that at least k ($k \leq n$) paths are from A_{good} with probability*

$$\sum_{i=k}^n \binom{i}{n} p^i (1-p)^{(n-i)},$$

where $p = |A_{good}|/|A|$.

Conservatively, we determined that we would like our search space to have at least 20% of networks from A_{good} . In addition, we can estimate p for a search space by querying a small subset networks. When using the shrink index for comparison, the threshold accuracy used to generate p should be kept constant.

By using this threshold and p , we compute the P_{good} (probability of sampling a good network) initially for the entire initial search space, the final output search space, and at all intermediate iterations. An increase in this probability across iterations implies that the proportion of networks from A_{good} is increasing in the current search space and, as a result, increasing our odds of sampling a well-performing network. We term the difference between the initial probability and the probability in our shrunk search space as the *shrink index* (s.i) as seen in equation 2:

$$s.i = P(good | init_ss) - P(good | shrunk_ss) \quad (2)$$

Intuitively, the shrink index is dependent on the initial search space as well as the dataset of consideration. Shrink indices indicate the increase in the probability of sampling a well-performing network and serve as a good comparison metric across different shrinkage approaches. Generally, higher indices indicate more optimal shrinkage. We note that the shrink index can only be used for comparing algorithms on the same dataset and search space. To illustrate this, consider that we have search space A which is reduced to A^* , and search space B which is reduced to B^* . Assume the shrink index of search space A^* is 0.5 and the shrink index of search space B^* is 0.2. However, the actual probability of sampling a good network from A goes from 20% to 70% in A^* in contrast to B which goes from 70% to 90% in B^* . In this scenario, while the shrink index of A^* is higher, search space B^* is likely to be a better search space of choice. This example highlights the importance of using shrink index to compare algorithms on the same starting search space since the initial concentration of networks is important. As shown in Table 1, our shrink index is significantly higher than other works when compared on ShuffleNetV2. Other search spaces also follow this trend.

GreedyNAS	0.3
ABS	0.1
NSE	0.2
PADNAS	0.1
Ours	0.4

Table 1: Shrink Index in ShuffleNetV2 Search Space on ImageNet dataset

4.5 NAS Methodology

Now that we have a shrunk search space with a verifiable proportion of well-performing networks, we proceed to determine the effectiveness of our shrunk spaces on search performance. Starting with the ImageNet under mobile constraints space, we shrank the FairNAS space using LISSNAS in combination with an XGBoost predictor. From the shrunk search space, we proceeded to run one-shot NAS following the same scheme as K-shot [Su *et al.*, 2021] with several minor optimizations for hardware. For comparisons against other shrinkage methods, we again utilize a predictor in combination with one-shot NAS to report both Pearson Correlation and Top-1 accuracy. Special care is taken to ensure that all works that are compared receive fair treatment. We use the same predictor for shrinking the search spaces using different methods and then train the one-shot network using the same hyperparameter optimization scheme.

5 Results

We now present, evaluate, and discuss extensive experimental results to show the effectiveness of LISSNAS. We begin by evaluating NAS performance on ImageNet under mobile constraints before directly comparing with other search space shrinkage works. Finally, we compare the search space statistics, size, and architectural diversity. Like many works in NAS, shrinkage methods are implemented in different search spaces making direct comparison challenging. In the following sections, we implement, optimize, and tune competing methods in order to compare our experimental results.

5.1 ImageNet Classification

ImageNet under mobile constraints serves as a canonical benchmark for all NAS techniques. As such, we are able to compare not only search space shrinkage techniques but also modern general NAS works. For a fair comparison, we elect to only include standard methods on CNN architectures without pretraining, knowledge distillation, or attention-based mechanisms. We note that Top-1 performance comparisons are feasible despite methods using other search spaces. In Table 2, we see that LISSNAS achieves SOTA accuracy. To obtain these results, we ran LISSNAS on the FairNAS search space containing 6^{19} architectures. With this size, querying every architecture becomes unfeasible even with fast inferences. Through leveraging locality, we are able to search this space in only 4 GPU days.

Comparison Against Shrinkage Works. We now directly compare the performance of our method to other shrinkage approaches on the same search spaces - NASBench101 and ShuffleNetV2 search space. To do so, we optimized and implemented these approaches before running one-shot NAS for search. NASBench101 and ShuffleNetV2 search spaces vary greatly in size and complexity - enabling diverse and fair comparison.

In Table 3, we display the Pearson correlation coefficients and Top-1 performances for each of the methods. Our search space has a significantly higher Pearson correlation in both search spaces demonstrating the value of training a supernet on our shrunk search space compared to other techniques. We

Model	Flops/Params	Top-1
MobileNetV3*	219M/5.4M	74.0
NASNet-A	564M/5.3M	74.0
ShuffleNetV2	299M/3.5M	72.6
SPOS*	319M/3.3M	74.3
OFA	230M/-	76.9
FBNet-C	375M/5.5M	74.9
FairNAS-C	325M/5.6M	76.7
K-shot-NAS-B	332M/6.2M	77.2
GreedyNASv2-S	324M/5.7M	77.5
DDS	400M/4.3M	74.1
ABS	472M/-	75.9
NSENet	330M/4.6M	75.5
PAD-NAS-S-B	315M/4.1M	74.9
LISSNAS	329M/6.1M	77.6

Table 2: ImageNet under mobile constraints. * Verified reproducible

Search Space	NASBench101		ShuffleNetV2	
	Pearson	Top-1	Pearson	Top-1
Full Search Space	0.467	91.0	0.334	70.0
ABS	0.58	93.0	0.397	72.8
NSE	0.56	92.1	0.406	73.0
PADNAS	0.59	92.4	0.355	69.9
LISSNAS	0.659	93.4	0.429	74.6

Table 3: Search Space Shrinkage NAS Performance

repeated our experiments and computed standard deviations across runs and found all results presented to be statistically significant.

Lastly, in both search spaces, LISSNAS also produces the highest Top-1 accuracies. For NASBench, LISSNAS is able to find the best network in the space. We note that this network is not present in the shrunk search space produced by ABS, NSE, or PADNAS.

5.2 Search Space Statistics

Aside from search performance metrics, statistical measures provide meaningful insight into the quality of our search space. We compare our output search space to the output search spaces produced by other shrinkage algorithms using these metrics. In Table 4, we accentuate that both the average accuracy and maximum accuracy of our output spaces are the highest, indicating an improvement in the number of well-performing networks. This also indicates a higher proportion of well-performing networks, mirroring the results from our shrink index.

Additionally, DDS [Radosavovic *et al.*, 2019], demonstrated the effectiveness of error Empirical Distribution Functions (EDF) in evaluating search spaces. Given this prior, we generated error EDFs based on the distribution of networks in our output space, the baseline search space, and the output search spaces of other shrinkage algorithms. These results are presented in Table 4. By performing the Kolmogorov-Smirnov (KS) test [Radosavovic *et al.*, 2020], we established

Search Space	Size	Avg Acc	Max Acc
Full SS - NASBench101	423K	90.4	94.6
ABS	100K	80.2	93.2
NSE	25K	90.1	93.6
PADNAS	15K	90.4	94.1
LISSNAS	13K	92.1	94.6
Full SS - ShuffleNetV2	1T	64.7	75.6
ABS	100K	65.1	74.9
NSE	1M	64.7	75.2
PADNAS	1M	66.2	73.1
LISSNAS	1M	67.5	75.6

Table 4: Shrunk Search Space Statistics for NASBench101 and ShuffleNetV2

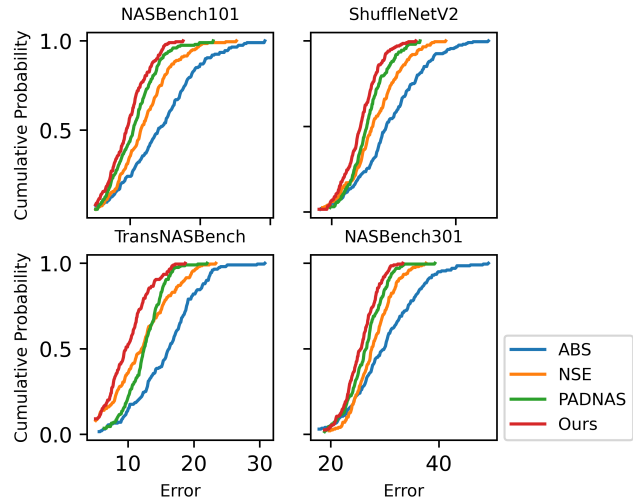


Figure 4: Error EDF captures the quality of a search space. Our algorithm maintains a concentration of better networks across the board.

statistically significant improvement with our approach. Notably, this is the case even between our approach and PADNAS which have more visually similar curves. Furthermore, search spaces with a greater area under the curve (AUC) typically have greater concentrations of good networks. Based on this, we can infer that our search space likely contains a larger proportion of good models than the other spaces as our curve covers a larger area.

5.3 Cardinality and Extent of Shrinkage

We begin by comparing our method to ABS and DDS. These methods were implemented on the MobileNet and AnyNetX space respectively. In Table 5, we compare the cardinality/size and factor of reduction across algorithms. We reiterate that smaller search spaces are feasible solutions for optimization and efficiency issues in NAS. The illustrated reduction factors describe the extent of shrinkage achieved rel-

		SS Size	Reduction Factor
MobileNet	Original	10^{12}	N/A
	ABS	10^5	$10^7 \times$
	LISSNAS	10^6	$10^6 \times$
AnyNetX	Original	10^{18}	N/A
	DDS	10^8	$10^{10} \times$
	LISSNAS	10^9	$10^{11} \times$

Table 5: Size Reduction Comparison

	Full SS Size	Our Size	Reduction
TransNASBench	7.3K	1K	$7 \times$
NASBench101	423K	13K	$33 \times$
ShuffleNetV2	4^{20}	10^6	$10^6 \times$
NASBench301	10^{18}	10^7	$10^{11} \times$

Table 6: LISSNAS Size Reduction

ative to the initial size of the search space. Note that DDS is manual design shrinkage method performed by experts and proved difficult to beat. However, our search space offers fairly close cardinality while being generated without expert input.

Table 6 shows the effectiveness of our algorithm on a range of search spaces of varying initial sizes. From this table, we observe that larger search spaces are pruned more as indicated by larger reduction factors. We attribute this to the disparity between the cardinality of well-performing networks in the space and cardinality of the entire space. In more detail, the number of networks that lie in the region next to the Pareto frontier (region of well-performing networks) is much smaller than the overall number of nets in the space. Large and small search spaces have a similar number of networks that lie in the Pareto frontier. However, due to increased cardinality, larger spaces can be pruned to a greater degree, leading to higher reduction factors.

5.4 Architectural Diversity

Lastly, while size provides a good baseline for evaluating our algorithm, exploring diversity is equally important. Multi-objective NAS relies on maintaining the same FLOP/Param range of the original space. As an example, deploying in a mobile setting requires smaller models whereas cloud GPUs with powerful processors utilize large and complex architectures to improve performance. Unexpectedly, we observed that naive search space shrinkage methods tend to drop the lower-performing architectures without considering the need for architectural diversity. From this perspective, we ensured that our algorithm covers the full range of architectures spanned by the original space.

From Figure 5, we can see that our search space does a better job of preserving networks with fewer FLOPs. Without special tuning, other shrinkage algorithms are incapable of preserving the smaller networks. Our method inherently preserves architectural diversity in contrast to being tuned to

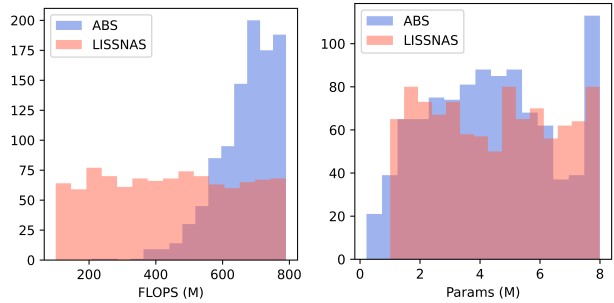


Figure 5: Histograms in the ShuffleNetV2 search space show bias for larger networks in other algorithms for larger models. More results are in the appendix.

run in specific FLOP ranges.

Since vector embeddings are common for architecture representations, we utilize embeddings to compute the maximum cosine distance between architectures in a search space. We utilize the same embedding scheme across all architectures for fair comparison. Our results in Table 7, substantiate our previous diversity observations. From this standpoint, dissimilar networks have a higher cosine distance and indicate diversity within the space. Interestingly, we also found that our shrunk search space is a negligible distance away from covering the complete original space.

NASBench101	PADNAS	135.2
	ABS	167.3
	NSE	201.5
	LISSNAS	255.6
	Original	260
ShuffleNetV2	PADNAS	17.9
	ABS	21.2
	NSE	21.6
	LISSNAS	28.0
	Original	30

Table 7: Maximum Cosine Distance

6 Conclusion

In summary, we present a novel search space shrinkage method that improves NAS performance. Through exploiting locality, LISSNAS is able to maintain architectural diversity while improving quality as measured through a variety of statistical approaches. The *shrink index*, can be used to compare search space shrinkage works in the future. We believe that search space optimization holds much promise for driving future improvement in NAS.

Acknowledgments

This work was supported in part by the National Science Foundation through NSF grants CNS-2112562 and IIS-

2140247. Special thanks to Shakthi Visagan for the feedback and editing.

Contribution Statement

Both first authors, Arjun Sridhar and Bhavna Gopal, contributed equally to the production of this work.

References

- [Bergstra *et al.*, 2013] James Bergstra, Daniel Yamins, and David Cox. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 115–123, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR.
- [Brock *et al.*, 2018] Andrew Brock, Theo Lim, J.M. Ritchie, and Nick Weston. SMASH: One-shot model architecture search through hypernetworks. In *International Conference on Learning Representations*, 2018.
- [Cai *et al.*, 2019] Han Cai, Chuang Gan, and Song Han. Once for all: Train one network and specialize it for efficient deployment. *CoRR*, abs/1908.09791, 2019.
- [Chu *et al.*, 2019] Xiangxiang Chu, Bo Zhang, Ruijun Xu, and Jixiang Li. Fairnas: Rethinking evaluation fairness of weight sharing neural architecture search. *CoRR*, abs/1907.01845, 2019.
- [Ci *et al.*, 2020] Yuanzheng Ci, Chen Lin, Ming Sun, Boyu Chen, Hongwen Zhang, and Wanli Ouyang. Evolving search space for neural architecture search. *CoRR*, abs/2011.10904, 2020.
- [Guo *et al.*, 2019] 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. *CoRR*, abs/1904.00420, 2019.
- [Howard *et al.*, 2019] Andrew Howard, Mark Sandler, Grace Chu, Liang-Chieh Chen, Bo Chen, Mingxing Tan, Weijun Wang, Yukun Zhu, Ruoming Pang, Vijay Vasudevan, Quoc V. Le, and Hartwig Adam. Searching for mobilenetv3. *CoRR*, abs/1905.02244, 2019.
- [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. *CoRR*, abs/2004.13431, 2020.
- [Huang *et al.*, 2021] Tao Huang, Shan You, Fei Wang, Chen Qian, Changshui Zhang, Xiaogang Wang, and Chang Xu. Greedynasv2: Greedier search with a greedy path filter. *CoRR*, abs/2111.12609, 2021.
- [Liu *et al.*, 2019] Hanxiao Liu, Karen Simonyan, and Yiming Yang. DARTS: Differentiable architecture search. In *International Conference on Learning Representations*, 2019.
- [Luo *et al.*, 2019a] Renqian Luo, Tao Qin, and Enhong Chen. Balanced one-shot neural architecture optimization. *arXiv preprint arXiv:1909.10815*, 2019.
- [Luo *et al.*, 2019b] Renqian Luo, Tao Qin, and Enhong Chen. Understanding and improving one-shot neural architecture optimization. *CoRR*, abs/1909.10815, 2019.
- [Mendoza *et al.*, 2016] Hector Mendoza, Aaron Klein, Matthias Feurer, Jost Tobias Springenberg, and Frank Hutter. Towards automatically-tuned neural networks. In *Workshop on automatic machine learning*, pages 58–65. PMLR, 2016.
- [Pham *et al.*, 2018] Hieu Pham, Melody Guan, Barret Zoph, Quoc Le, and Jeff Dean. Efficient neural architecture search via parameters sharing. In *International conference on machine learning*, pages 4095–4104. PMLR, 2018.
- [Radosavovic *et al.*, 2019] Ilija Radosavovic, Justin Johnson, Saining Xie, Wan-Yen Lo, and Piotr Dollar. On network design spaces for visual recognition. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- [Radosavovic *et al.*, 2020] 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.
- [Real *et al.*, 2017] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Suematsu, Jie Tan, Quoc V Le, and Alexey Kurakin. Large-scale evolution of image classifiers. In *International Conference on Machine Learning*, pages 2902–2911. PMLR, 2017.
- [Real *et al.*, 2019] 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.
- [Sciuto *et al.*, 2019] Christian Sciuto, Kaicheng Yu, Martin Jaggi, Claudiu Musat, and Mathieu Salzmann. Evaluating the search phase of neural architecture search. *CoRR*, abs/1902.08142, 2019.
- [Su *et al.*, 2021] Xiu Su, Shan You, Mingkai Zheng, Fei Wang, Chen Qian, Changshui Zhang, and Chang Xu. K-shot NAS: learnable weight-sharing for NAS with k-shot supernets. *CoRR*, abs/2106.06442, 2021.
- [Tan *et al.*, 2018] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, and Quoc V. Le. Mnasnet: Platform-aware neural architecture search for mobile. *CoRR*, abs/1807.11626, 2018.
- [Tyagi *et al.*, 2022] Anjul Tyagi, Cong Xie, and Klaus Mueller. NAS-navigator: Visual steering for explainable one-shot deep neural network synthesis. *IEEE Transactions on Visualization and Computer Graphics*, pages 1–11, 2022.
- [Wang *et al.*, 2020] Dilin Wang, Meng Li, Chengyue Gong, and Vikas Chandra. Attentiveness: Improving neural architecture search via attentive sampling. *arXiv preprint arXiv:2011.09011*, 2020.
- [Wang *et al.*, 2021] Dilin Wang, Chengyue Gong, Meng Li, Qiang Liu, and Vikas Chandra. Alphanet: Improved

- training of supernet with alpha-divergence. In Marina Meila and Tong Zhang, editors, *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pages 10760–10771. PMLR, 18–24 Jul 2021.
- [Wei *et al.*, 2022] Chen Wei, Chuang Niu, Yiping Tang, Yue Wang, Haihong Hu, and Jimin Liang. Npenas: Neural predictor guided evolution for neural architecture search. *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [Wen *et al.*, 2020] Wei Wen, Hanxiao Liu, Yiran Chen, Hai Li, Gabriel Bender, and Pieter-Jan Kindermans. Neural predictor for neural architecture search. In *European Conference on Computer Vision*, pages 660–676. Springer, 2020.
- [White *et al.*, 2021] Colin White, Arber Zela, Robin Ru, Yang Liu, and Frank Hutter. How powerful are performance predictors in neural architecture search? *Advances in Neural Information Processing Systems*, 34:28454–28469, 2021.
- [Wu *et al.*, 2018] Bichen Wu, Xiaoliang Dai, Peizhao Zhang, Yanghan Wang, Fei Sun, Yiming Wu, Yuandong Tian, Peter Vajda, Yangqing Jia, and Kurt Keutzer. Fbnet: Hardware-aware efficient convnet design via differentiable neural architecture search. *CoRR*, abs/1812.03443, 2018.
- [Xia and Ding, 2020] Xin Xia and Wenrui Ding. HNAS: hierarchical neural architecture search on mobile devices. *CoRR*, abs/2005.07564, 2020.
- [Xie *et al.*, 2018] Sirui Xie, Hehui Zheng, Chunxiao Liu, and Liang Lin. Snas: stochastic neural architecture search. *arXiv preprint arXiv:1812.09926*, 2018.
- [Xue *et al.*, 2021] Yu Xue, Pengcheng Jiang, Ferrante Neri, and Jiayu Liang. A multi-objective evolutionary approach based on graph-in-graph for neural architecture search of convolutional neural networks. *International Journal of Neural Systems*, 31(09):2150035, 2021.
- [Ying *et al.*, 2019] Chris Ying, Aaron Klein, Esteban Real, Eric Christiansen, Kevin Murphy, and Frank Hutter. Nas-bench-101: Towards reproducible neural architecture search. *CoRR*, abs/1902.09635, 2019.
- [You *et al.*, 2020] Shan You, Tao Huang, Mingmin Yang, Fei Wang, Chen Qian, and Changshui Zhang. Greedynas: Towards fast one-shot nas with greedy supernet. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 1996–2005, 2020.
- [Zhang *et al.*, 2018a] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018.
- [Zhang *et al.*, 2018b] Xinbang Zhang, Zehao Huang, and Naiyan Wang. You only search once: Single shot neural architecture search via direct sparse optimization. *CoRR*, abs/1811.01567, 2018.
- [Zhang *et al.*, 2020] Xinbang Zhang, Zehao Huang, Naiyan Wang, Shiming Xiang, and Chunhong Pan. You only search once: Single shot neural architecture search via direct sparse optimization. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 43(9):2891–2904, 2020.
- [Zoph and Le, 2017] Barret Zoph and Quoc Le. Neural architecture search with reinforcement learning. In *International Conference on Learning Representations*, 2017.
- [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.