

Evolving Transferable Neural Pruning Functions

Yuchen Liu

Princeton University

yl16@princeton.edu

S.Y. Kung

Princeton University

kung@princeton.edu

David Wentzlaff

Princeton University

wentzlaf@princeton.edu

ABSTRACT

Structural design of neural networks is crucial for the success of deep learning. While most prior works in evolutionary learning aim at directly searching the structure of a network, few attempts have been made on another promising track, channel pruning, which recently has made major headway in designing efficient deep learning models. In fact, prior pruning methods adopt human-made pruning functions to score a channel’s importance for channel pruning, which requires domain knowledge and could be sub-optimal. To this end, we pioneer the use of genetic programming (GP) to discover strong pruning metrics automatically. Specifically, we craft a novel design space to express high-quality and transferable pruning functions, which ensures an end-to-end evolution process where no manual modification is needed on the evolved functions for their transferability after evolution. Unlike prior methods, our approach can provide both compact pruned networks for efficient inference and novel closed-form pruning metrics which are mathematically explainable and thus generalizable to different pruning tasks. While the evolution is conducted on small datasets, our functions shows promising results when applied to more challenging datasets, different from those used in the evolution process. For example, on ILSVRC-2012, an evolved function achieves state-of-the-art pruning results.

KEYWORDS

Convolutional Neural Network, Channel Pruning, Efficient Deep Learning, Genetic Programming

ACM Reference Format:

Yuchen Liu, S.Y. Kung, and David Wentzlaff. 2022. Evolving Transferable Neural Pruning Functions. In *Genetic and Evolutionary Computation Conference (GECCO '22), July 9–13, 2022, Boston, MA, USA*. ACM, Boston, MA, USA, 12 pages. <https://doi.org/10.1145/3512290.3528694>

1 INTRODUCTION

Convolutional neural networks (CNNs) have demonstrated superior performance on various computer vision tasks [12, 13, 23, 64]. However, CNNs require huge storage space, high computational budget, and large memory utilization, which could far exceed the resource limit of edge devices like mobile phones and embedded gadgets. As a result, many methods have been proposed to reduce their cost, such as weight quantization [7, 10, 25], tensor factorization [34, 39], weight pruning [26, 81], and channel pruning [31, 48].

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

GECCO '22, July 9–13, 2022, Boston, MA, USA

© 2022 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-9237-2/22/07.

<https://doi.org/10.1145/3512290.3528694>

Among them all, channel pruning is the preferred approach to learn dense compact models, which has been receiving increased focus from the research community.

Channel pruning is usually achieved in three steps: (1) score channels’ importance with a hand-crafted pruning function; (2) remove redundant channels based on the scores; (3) retrain the network. The performance of channel pruning largely depends on the pruning function used in step (1). Current scoring metrics are mostly handcrafted to extract crucial statistics from channels’ feature maps [32, 80] or kernel parameters [31, 42] in a labelless [29, 51] or label-aware [38, 82] manner. However, the design space of pruning functions is so large that hand-crafted metrics are usually sub-optimal, and enumerating all functions with human labor under the space is impossible. While prior evolutionary learning works aim to automate the design for the structure of the network directly [67, 72], no attempts, to the best of our knowledge, have been made to evolve the pruning metrics. These raise the question: can we leverage evolutionary strategies to automatically develop strong pruning functions to advance channel pruning?

To this end, we take the first step to adopt genetic programming (GP) to learn transferable pruning functions, as shown in Fig. 1. In particular, a population of functions is evolved by applying them to pruning tasks of small image classification datasets, and the evolved functions can later be transferred to larger and more challenging datasets. Our closed-form, explainable, learned functions are transferable and generalizable: (1) They are applicable to pruning tasks of different image datasets and networks, and can also be used for other machine learning tasks, e.g., feature selection; (2) they demonstrate competitive pruning performance on datasets and networks that are different from those used in the evolution process. Such transferability and generalizability provides a unique advantage to our method, where prior meta-pruning methods like MetaPruning [52] and LFPC [28] are learned and evolved on the same tasks with no transferability and perform inferior to our approach.

Specifically, we encode pruning functions using expression trees where we carefully design our search space to allow transferability of the evolved functions. For example, we propose a uni-tree search space for label-aware pruning metrics, which makes them applicable to different datasets. Such a design space ensures an end-to-end evolution process, where the learned functions are transferable to other datasets without any manual modification after evolution. Moreover, under our encoding space, we are able to build a group of competitive hand-crafted pruning functions, which we name as SOAP (state-of-the-art population), and we find the use of SOAP considerably improves the evolution performance. The populations of the functions are evolved with two different pruning tasks, LeNet on MNIST and VGGNet on CIFAR-10. We observe that evolving on two tasks produces better functions than only evolving on one of them, and more surprisingly, our scheme can even produce more effective pruning functions than direct evolution on a large dataset,

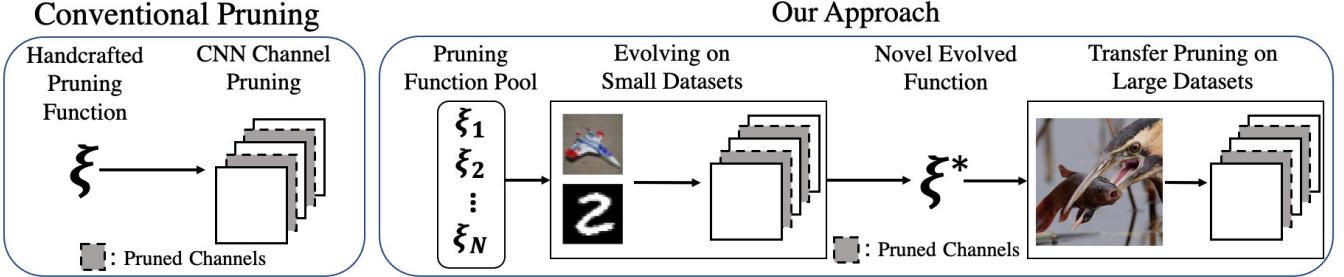


Figure 1: Illustration of our approach. Compared to conventional methods which mainly use handcrafted pruning functions, we aim to learn the pruning functions automatically via an evolution strategy, genetic programming. The evolved functions are transferable and generalizable, further enhancing the pruning performance.

e.g., ILSVRC-2012, under the same computational budget. We analyze the merits of an evolved function both mathematically and visually and transfer it to three larger datasets, CIFAR-100, SVHN, and ILSVRC-2012, where it exceeds the state-of-the-art pruning results on all of them.

Our main contributions are summarized as follows:

- We propose a novel paradigm where we leverage genetic programming to learn transferable channel pruning functions which advance pruning efficacy.
- We develop a novel design space to allow an end-to-end co-evolution process for searching transferable pruning functions. Such a space also enables us to express SOAP, which helps improve the effectiveness of the evolution.
- We provide an analysis on our closed-form evolved functions, which could further streamline the design of pruning metrics. The evolved functions also show generalizability to other machine learning tasks, e.g., feature selection.
- When transferred to datasets unseen by the evolution, our evolved functions achieve state-of-the-art pruning results. For example, with 26.9% and 53.4% FLOPs¹ reduction from MobileNet-V2, we achieve top-1 accuracies of 71.90% and 69.16% on ILSVRC-2012, outperforming the state of the art.

2 RELATED WORK

Hand-Crafted Channel Pruning. Channel pruning [31, 38, 80, 82] is generally realized by using a handcrafted pruning function to score channels' saliency and remove redundant ones. Based on the scoring procedure, it can be categorized into labelless pruning and label-aware pruning.

Labelless channel pruning typically adopts the norm-based property of the channel's feature maps or associated filters as pruning criterion [29, 31, 32, 42, 44, 51, 53, 54, 78]. For example, Liu et al. [51] and Ye et al. [78] use the absolute value of scaling factors in the batch-norm, while ℓ_1 -norm and ℓ_2 -norm of channels' associated filters are computed in [29, 42, 44] as channels' importance. On the other hand, researchers have designed metrics to evaluate class discrepancy of channels' feature maps for label-aware pruning [38, 49, 82]. Zhuang et al. [82] insert discriminant losses in the network and remove channels that are the least correlated to the losses after iterative optimization. Kung et al. [38] and Liu et

al. [49, 50] adopt closed-form discriminant functions to accelerate the scoring process.

While these works use handcrafted scoring metrics, we learn transferable and generalizable pruning functions automatically.

Meta-Learning. Our work falls into the category of meta-learning, where prior works have attempted to optimize machine learning components, including hyper-parameters [5, 17, 68], optimizers [4, 8, 76], and neural network structures [46, 47, 62, 63, 74, 77, 83, 84].

Prior works on neural architecture search (NAS) have leveraged reinforcement learning (RL) to discover high-performing network structures [2, 6, 73, 74, 83, 84]. Recently, NAS has also been adopted to find efficient network structures [73, 74]. Another line of works adopts evolution strategies (ES) to explore the space of network structures [11, 16, 47, 56, 60, 62, 63, 67, 69, 72, 75, 77], which demonstrates competitive performance to RL methods. This notion is pioneered by neuro-evolution [18, 70, 71] which evolves the topology of small neural networks. In the era of deep learning, Suganuma et al. [72] leverage Cartesian genetic programming to find competitive network structures. Real et al. [62] evolve networks that improve over the ones found by RL-based NAS [84]. Dai et al. [11] apply ES to design efficient and deployable networks for mobile platforms. Templier et al. [75] propose a geometric encoding scheme for more efficient parameter search.

Compared to prior works, we employ evolutionary learning from a new angle for efficient network design, where we learn transferable pruning functions that produce state-of-the-art pruning results. Our work is orthogonal to prior works, for example, our evolved functions can be potentially applied on evolutionary NAS-learned networks to further enhance their efficiency.

Meta-Pruning. Prior works [9, 28, 30, 33, 52] have also adopted a similar notion of learning to prune a CNN. We note that an evolution strategy is used in LeGR [9] and MetaPruning [52] to search for a pair of pruning parameters and network encoding vectors, respectively. However, our approach is drastically different from them in terms of search space and search candidates, where we search for effective combinations of operands and operators to build transferable and powerful pruning functions. He et al. propose LFPC [28] to learn network pruning criteria (functions) across layers by training a differentiable criteria sampler. However, rather than learning new pruning functions, their goal is to search within a pool of existing pruning criteria and find candidates that are good for a

¹Number of floating points operations for an image inference.

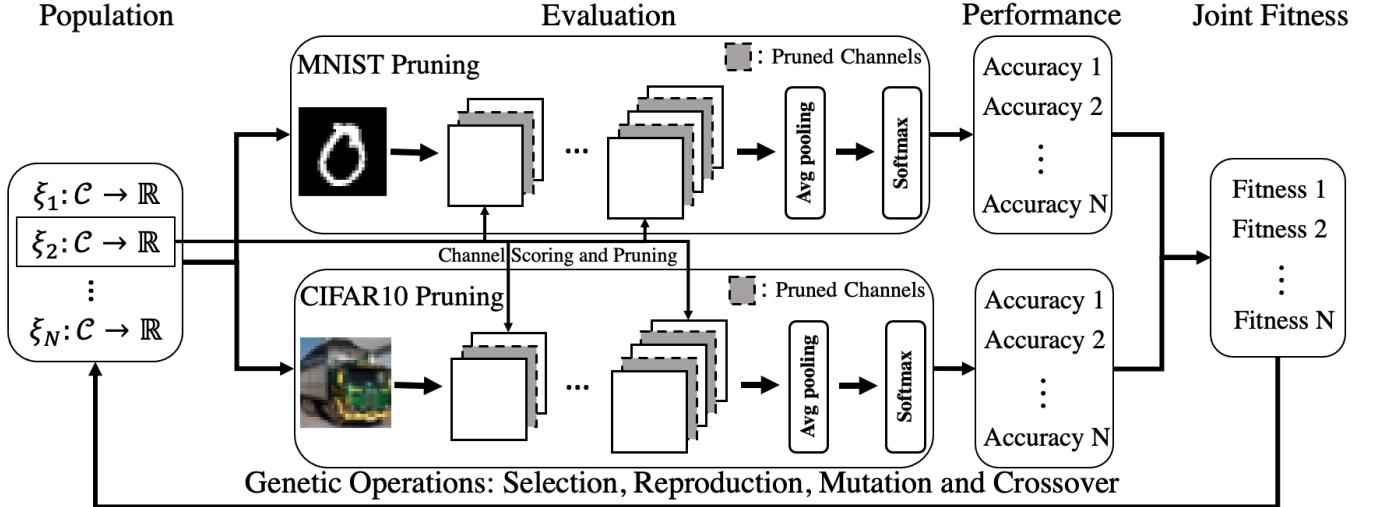


Figure 2: Illustration of our approach to evolve channel pruning functions. A population of functions is applied to conduct pruning tasks on two datasets, MNIST and CIFAR-10. Each function receives a fitness value by combining its pruned networks' accuracies. The population will then go through a natural selection process to improve the functions' effectiveness.

Filter-based operands	whole layer's filter $\mathcal{W} \in \mathbb{R}^{c_{out} \times c_{in} \times h \times w}$, channel's incoming filter $\mathcal{W}_I \in \mathbb{R}^{c_{in} \times h \times w}$, channel's batch-normed parameter $\mathcal{B} \in \mathbb{R}^4$
Map-based operands	Feature maps collection $\mathcal{F} = \{(f_i, y_i) f_i \in \mathbb{R}^{H \times W}, y_i \in [1 : C], i \in [1 : N]\}$, two partitions of feature maps collections $\mathcal{F}^+ = \{f_i y_i = k, k \in [1 : C]\}$ and $\mathcal{F}^- = \{f_i y_i \neq k, k \in [1 : C]\}$

Table 1: Operand Space

Elementwise operators	addition, subtraction, multiplication, division, absolute value, square, square root, adding ridge factor
Matrix operators	matrix trace, matrix multiplication, matrix inversion, inner product, outer product, matrix/vector transpose
Statistics operators	summation, product, mean, standard deviation, variance, counting measure
Specialized operators	rbf kernel matrix getter, geometric median getter, tensor slicer

Table 2: Operator Space

certain layer's pruning. On the contrary, our evolution recombines the operands and operators and produces novel pruning metrics, which are generally good for all layers.

We also notice that MetaPruning [52], LFPC [28], and other methods [9, 30, 33] are all learned on one task (dataset and network) and applied only on the same task with no transferability. In contrast, we only need one evolution learning process, which outputs evolved functions that are transferable across multiple tasks and demonstrate competitive performance on all of them.

3 METHODOLOGY

In Fig. 2, we present our evolution framework, which leverages genetic programming [36] to learn high-quality channel pruning functions. We first describe the design space to encode channel scoring functions. Next, we discuss the pruning tasks to evaluate the functions' effectiveness. Lastly, genetic operators are defined to traverse the function space for competitive solutions.

3.1 Function Design Space

Expression Tree. In channel pruning, a pruning function $\xi : \mathcal{C} \mapsto \mathbb{R}$ scores the channels to determine their importance/redundancy, where \mathcal{C} denotes feature maps, filters, and their statistics associated with the channels. This scoring process can be viewed as a series of operations with operators (addition, matrix multiplication, etc.) and operands (feature maps, filters, etc.). We thus adopt an expression tree encoding scheme to represent a pruning function where inner nodes are operators, and leaves are operands.

As shown in Tab. 1 and 2, our function design space includes two types of operands (6 operands in total) and four types of operators (23 operators in total), via which a vast number of pruning functions can be expressed. The statistics operators can compute the statistics of an operand in two dimensions, namely, global dimension (subscript with 'g') and sample dimension (subscript with 's'). The global dimension operators flatten operands into a 1D sequence and extract corresponding statistics, while the sample dimension operators compute statistics on the axis of samples. For example,

$\text{sum}_g(\mathcal{W})$ returns the summation of all entries of a kernel tensor, while $\text{mean}_s(\mathcal{F})$ returns $\bar{f} \in \mathbb{R}^{H \times W}$, which is the sample average of all feature maps. We also include specialized operators which allow us to build complicated but competitive metrics like maximum mean discrepancy (MMD) [24] and filter’s geometric median [31]. **Function Encoding.** The channel scoring functions can be categorized into two types: labelless metrics and label-aware metrics. For labelless functions like filter’s ℓ_1 -norm, we adopt a direct encoding scheme as $\text{sum}_g(\text{abs}(\mathcal{W}_I))$ with the expression tree shown in Fig. 3.

For label-aware metrics such as the one in [38] and MMD [24], which measure class discrepancy of the feature maps, we observe a common computation graph among them, as shown in Fig. 3: (1) partition the feature maps in a labelwise manner; (2) apply the same operations on each label partition and all feature maps; (3) average/sum the scores of all partitions to obtain a single scalar. These metrics can be naively encoded as C -branch trees (C : number of class labels in the dataset). However, directly using the naive encoding scheme will result in data-dependent non-transferable metrics because: (1) C varies from dataset to dataset (e.g., metrics for CIFAR-10 is not transferable to CIFAR-100); (2) mutating the subtrees differently could make the metric overfit to a specific label numbering scheme (e.g., for a metric with different subtrees on class-1 and class-2, renumbering the labels would mean the metric would compute something different, which is undesirable).

To combat the above issues, we express a label-aware function by a uni-tree which encodes the common operations that are applied to each label partition, as explained in Fig. 3. Instead of directly encoding the operands from a specific label partition, like \mathcal{F}^{1+} (feature maps with labels equal to 1) and \mathcal{F}^{1-} (feature maps with labels not equal to 1), we use a symbolic representation of \mathcal{F}^+ and \mathcal{F}^- to generically encode the partition concept. In the actual scoring process, the uni-tree is compiled back to a C -branch computation graph, with \mathcal{F}^+ and \mathcal{F}^- converted to the specific map partitions. Such uni-tree encoding allows us to evolve label-aware metrics independent of C and label numbering schemes, which ensures their transferability to datasets unseen by the evolution process.

SOAP. Using the above described function encoding, we can implement a broad range of competitive pruning functions: filter’s ℓ_1 -norm [42], filter’s ℓ_2 -norm [29], batch norm’s scaling factor [51], filter’s geometric median [31], Discriminant Information [38], Maximum Mean Discrepancy [24], Absolute SNR [22], Student’s T-Test [41], Fisher Discriminant Ratio [61], and Symmetric Divergence [55]. For the last four metrics, we adopt the scheme in [49] for channel scoring. We name this group of functions the state-of-the-art population (SOAP), which helps our evolution in many aspects. For instance, in Sec. 6, we find that initializing the population with SOAP evolves better pruning functions than random initialization. Detailed implementation of SOAP is included in Supplementary.

3.2 Function Effectiveness Evaluation

The encoded functions are applied to empirical pruning tasks to evaluate their effectiveness. To avoid overfitting on certain data patterns and increase the generality of the evolved functions, we evolve the population of functions on two different pruning tasks, LeNet-5 [40] on MNIST [40] and VGG-16 [66] on CIFAR-10 [37]. In

both pruning tasks, we adopt a one-shot pruning scheme and report the retrained accuracies on validation sets. For each pruning task, we keep the pruning settings (layers’ pruning ratios, target pruning layers, etc.) and the retraining hyper-parameters (learning rate, optimizer, weight decay factor, etc.) the same for all evaluations throughout the evolution process. This guarantees a fair effectiveness comparison over different functions in all generations and ensures we are evolving better functions rather than better hyper-parameters. In this way, we can meta-learn powerful functions that perform well on both MNIST and CIFAR-10 and are generalizable to other datasets. Not surprisingly, evolving with both tasks produce stronger pruning functions than evolving on only one of them, shown in Sec. 3.3. Moreover, in Sec. 6, we find our strategy enjoys better cost-effectiveness compared to direct evolution on a large dataset, e.g., ILSVRC-2012.

3.3 Function Fitness

After evaluation, each encoded function receives two accuracies, $\text{Acc}_{\text{MNIST}}$ and $\text{Acc}_{\text{CIFAR}}$, from the pruning tasks. We investigate two accuracy combination schemes, weighted arithmetic mean (Eqn. 1) and weighted geometric mean (Eqn. 2), to obtain the joint fitness of a function. A free parameter $\alpha \in [0, 1]$ is introduced to control the weights of different tasks.

$$\text{Fitness} = \alpha \times \text{Acc}_{\text{MNIST}} + (1 - \alpha) \times \text{Acc}_{\text{CIFAR}} \quad (1)$$

$$\text{Fitness} = (\text{Acc}_{\text{MNIST}})^\alpha \times (\text{Acc}_{\text{CIFAR}})^{1-\alpha} \quad (2)$$

Ablation Study. To decide the fitness combination scheme for the main experiments, we conduct 10 small preliminary evolution tests using a grid of $\alpha \in \{0, 0.3, 0.5, 0.7, 1\}$ with both combination schemes. Note that when $\alpha \in \{0, 1\}$, the process degenerates to single dataset evolution. We empirically evaluate the generalizability of the best evolved functions from each test by applying them to prune a ResNet-38 on CIFAR-100. Note CIFAR-100 is not used in the evolution process, and thus the performance on it speaks well for evolved functions’ generalizability. In Fig. 4, we find that solely evolving on MNIST ($\alpha = 1$) would be the least effective option for CIFAR-100 transfer pruning. In addition, we find that functions evolved on two datasets ($\alpha \in [0.3, 0.5, 0.7]$) generally perform better than the ones that just evolve on a single dataset ($\alpha \in [0, 1]$). We observe that setting $\alpha = 0.5$ with weighted geometric mean leads to the best result, which we adopt in the main experiments.

3.4 Genetic Operations

Selection. After evaluation, the population will undergo a selection process, where we adopt tournament selection [21] to choose a subset of competitive functions.

Diversity Maintenance. This subset of functions is then used to reproduce individuals for the next generation. However, we observe shrinkage of the population’s genetic diversity when all children are reproduced from parents, as the selected parents only represent a small pool of genomes. Such diversity shrinkage would result in premature convergence of the evolution process. To combat this issue, we reserve a slot in the next generation and produce individuals in the slots by randomly cloning functions from SOAP or building random trees. We find this adjustment empirically useful to help the evolution proceed longer.

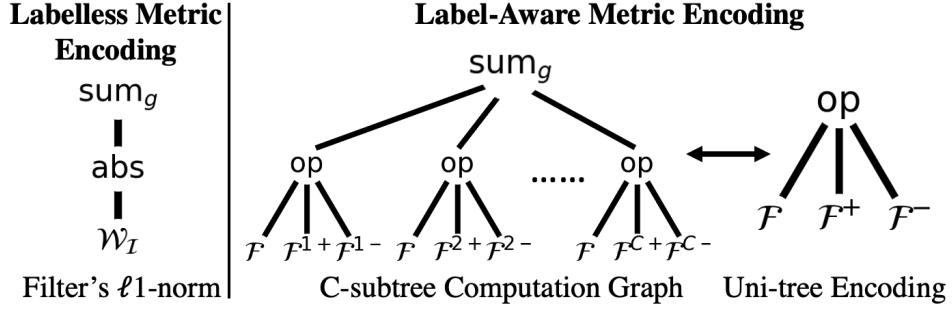


Figure 3: Illustration of the pruning function encoding. Left: For labelless scoring metrics like filter's ℓ_1 -norm, we adopt a direct tree encoding scheme. Right: For label-aware scoring metrics, we encode the C -subtree computation graph by a uni-tree (C : number of class labels). The uni-tree encodes the common operations (op) on each label partition ($\mathcal{F}^+, \mathcal{F}^-$) and all feature maps (\mathcal{F}). This scheme allows transferable function evolution.

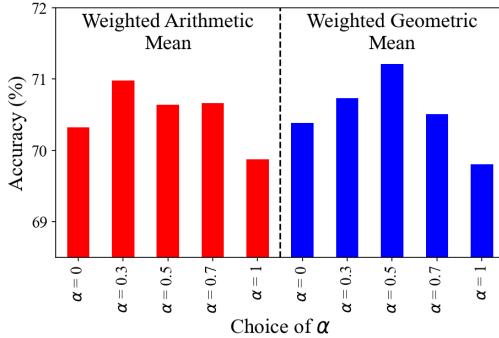


Figure 4: Preliminary evolution tests on the choice of fitness combination scheme. The best evolved function from each scheme is applied to conduct a pruning test on CIFAR-100 with ResNet-38, and their accuracies are plotted.

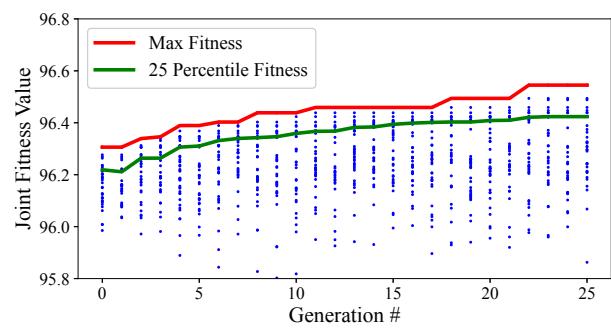


Figure 5: Progress of the evolution experiment. Each dot represents an individual function evaluation. The red curve shows functions with the best fitness over generations, while the green curve shows the individuals at the 25 percentile fitness. The effectiveness of the best function and the population's overall quality are both monotonically increasing.

Mutation and Crossover. We conduct mutation and crossover on the reproduced population to traverse the function design space for new expressions. We adopt the conventional scheme of random tree mutation and one point crossover [3]. After mutation and crossover, the population will go through the next evolution iteration.

Function Validity. The function expressions generated from mutation and crossover can be invalid (non-invertible matrix, dimension inconsistency, etc.) due to the random selections of operators, operands, and nodes in the expression trees. To combat this issue and enlarge our valid function space, some operators are deliberately modified from their standard definition. For instance, whenever we need to invert a positive semi-definite scatter matrix S , we automatically add a ridge factor ρI , and invert the matrix $S + \rho I$. For dimension inconsistency in elementwise operations, we have two options to pad the operand with a smaller dimension: (1) with 0 for $+$ and $-$, and 1 for \times , and \div , (2) with its own value if it is a scalar. Moreover, we conduct a validity test on the mutated/crossovered functions every time after the mutation/crossover process. The invalid expressions are discarded, and the mutation/crossover operations are repeated until we recover the population size with

all valid functions. These methods ensure we generate valid function expressions under our vast design space during the evolution process.

4 EVOLUTION ON MNIST AND CIFAR-10

Experiment Settings. We conduct the experiment with a population size of 40 individuals over 25 generations. The population is initialized with 20 individuals chosen by randomly cloning functions from SOAP and 20 random expression trees. The size of the selection tournament is 4 and we select 10 functions in each generation. 24 individuals are reproduced from the selected functions, while 6 individuals are from SOAP or randomly built. The mutation and crossover probability are both set to be 0.75. We prune 92.4% of FLOPs from a LeNet-5 (baseline acc: 99.26%) and 63.0% of FLOPs from a VGG-16 (baseline acc: 93.7%), respectively. Such aggressive pruning schemes help us better identify functions' effectiveness. We use the weighted geometric mean in Eqn. 2 to combine two validation accuracies with $\alpha = 0.5$. Our codes are implemented with DEAP [19] and TensorFlow [1] for the genetic operations and the neural network pruning. The experiments are carried out on a cluster with SLURM job scheduler [79] for workload parallelization.

Network	Method	Test Acc (%)	Acc ↓ (%)	FLOPs	Pruned (%)	Parameters	Pruned (%)
ResNet 164	SLIM [51]	98.22 → 98.15	0.07	172M	31.1	1.46M	14.5
	Ours	98.22 → 98.26	-0.04	92M	63.2	0.64M	63.0

Table 3: SVHN Transfer Pruning Results

Network	Method	Test Acc (%)	Acc ↓ (%)	FLOPs	Pruned (%)	Parameters	Pruned (%)
VGG19	SLIM [51]	73.26 → 73.48	-0.22	256M	37.1	5.0M	75.1
	G-SD [49]	73.40 → 73.67	-0.27	161M	59.5	3.2M	84.0
	Ours	73.40 → 74.02	-0.62	155M	61.0	2.9M	85.5
ResNet 56	SFP [29]	71.33 → 68.37	2.96	76M	39.3	-	-
	FPGM [31]	71.40 → 68.79	2.61	59M	52.6	-	-
	LFPC [28]	71.33 → 70.83	0.58	61M	51.6	-	-
	LeGR [9]	72.41 → 71.04	1.37	61M	51.4	-	-
	Ours	72.05 → 71.70	0.35	55M	56.2	0.38M	54.9
ResNet 110	LCCL [14]	72.79 → 70.78	2.01	173M	31.3	1.75M	0.0
	SFP [29]	74.14 → 71.28	2.86	121M	52.3	-	-
	FPGM [31]	74.14 → 72.55	1.59	121M	52.3	-	-
	TAS [15]	75.06 → 73.16	1.90	120M	52.6	-	-
	Ours	74.40 → 73.85	0.55	111M	56.2	0.77M	55.8
ResNet 164	LCCL [14]	75.67 → 75.26	0.41	195M	21.3	1.73M	0.0
	SLIM [51]	76.63 → 76.09	0.54	124M	50.6	1.21M	29.7
	DI [38]	76.63 → 76.11	0.52	105M	58.0	0.95M	45.1
	Ours	77.15 → 77.77	-0.62	92M	63.2	0.66M	61.8

Table 4: CIFAR-100 Transfer Pruning Results

$$\xi^*(C) = \frac{\text{var}_g(\mathcal{F}^-)}{\text{var}_g(\mathcal{F}^+)} + \frac{\text{var}_g(\mathcal{F}^+)}{\text{var}_g(\mathcal{F}^-)} + \frac{\|\text{std}_g(\bar{f}) \times \text{var}_g(\mathcal{F}^-) \times \bar{f} + (\text{var}_g(\mathcal{F}^+) - \text{mean}_g(\mathcal{F}^-))\mathbf{1}\|_2^2}{\text{var}_g(\mathcal{F}^+) + \text{var}_g(\mathcal{F}^-)} \quad (3)$$

Experiment Result. Our evolution progress is shown in Fig. 5, where the red curve denotes the functions with the maximum fitness while the green curve plots the ones with the top 25 percentile fitness. Both curves increase monotonically over generations, indicating that the quality of both the best function and the entire population improves over time. This demonstrates the effectiveness of our scheme. Specifically, the best pruned LeNet-5/VGG-16 in the first generation have accuracies of 99.15%/93.55% while the best accuracies in the last generation are 99.25%/94.0%. As the first generation is initialized with SOAP functions, such results suggest that the algorithm derives metrics that outperform handcrafted functions in SOAP. The whole evolution takes 98 GPU-days on P100, which is a reasonable amount of computation for modern evolution learning. While this is a pioneering work², we envision that future work could further reduce the evolution computation.

Evolved Function. We present the winning function in Eqn. 3, where $\bar{f} = \text{mean}_s(\mathcal{F})$ denotes sample average of the feature maps and $\mathbf{1}$ is a vector with all entries set to 1. The first two terms of the function award a high score to channels with class-diverged feature maps whose $\text{var}_g(\mathcal{F}^+)$ or $\text{var}_g(\mathcal{F}^-)$ is significantly smaller than the other. Channels with these feature maps contain rich class information as it generates distinguishable responses to different classes. The third term’s denominator computes the sum of the feature maps variances while its numerator draws statistics from the average feature maps and the distance between \mathcal{F}^+ and \mathcal{F}^- ,

²Compared to initial works on NAS, which take 2000 GPU-days [84] and 3000 GPU-days [63], we are 20/30x faster.

which resembles the concept of signal-to-noise ratio. Two points worth mentioning for this function: (1) it identifies important statistical concepts from human-designed metrics, where it learns from Symmetric Divergence [55] to measure the divergence of class feature maps. (2) it contains unique math concepts that are empirically good for channel importance measurement, which is shown in the novel statistics combination of the feature maps in the third term’s numerator. Our visual result in Sec. 6 shows ξ^* can be further applied to feature selection, which represents another machine learning task.

5 TRANSFER PRUNING

Benchmarks. To show the generalizability of our evolved pruning function, we apply ξ^* in Eqn. 3 to more challenging datasets that are not used in the evolution process: CIFAR-100 [37], SVHN [58], and ILSVRC-2012 [12]. We compare our method with metrics from SOAP, e.g., L1 [42], FPGM [31], G-SD [49], and DI [38], where ξ^* outperforms all these handcrafted metrics. We also include other “learn to prune” methods like Meta [52] and LFPC [28] and other state-of-the-art methods like DSA [59] and CC [43] for comparison. The results are summarized in Tab. 3, 4, and 5, where the accuracies are shown as “baseline acc. → pruned acc.” and the numbers for all other methods are copied from their papers. On ILSVRC-2012, we report our pruned models at different FLOPs reduction levels and add a suffix specifying their FLOPs pruning ratios (e.g., Ours 60%-pruned). This is because different prior arts report their compressed models at different rates, and we want to make a fair comparison

Network	Method	Top-1 Acc. (%)	Top-1 ↓ (%)	Top-5 Acc. (%)	Top-5 ↓ (%)	FLOPs (B) Pruned (%)	Params (M) Pruned (%)
VGG16	L1 [42]	-	-	89.90 → 89.10	0.80	7.74 (50.0)	-
	CP [32]	-	-	89.90 → 89.90	0.00	7.74 (50.0)	-
	G-SD [49]	71.30 → 71.88	-0.58	90.10 → 90.66	-0.56	6.62 (57.2)	133.6 (3.4)
	Ours 59%-pruned	71.30 → 72.37	-1.07	90.10 → 91.05	-0.95	6.34 (59.0)	133.5 (3.5)
	RNP [45]	-	-	89.90 → 86.67	3.23	5.16 (66.7)	138.3 (0.0)
	SLIM [51]	-	-	89.90 → 88.53	1.37	5.16 (66.7)	-
	FBS [20]	-	-	89.90 → 89.86	0.04	5.16 (66.7)	138.3 (0.0)
ResNet 18	Ours 17%-pruned	70.05 → 70.08	-0.03	89.40 → 89.24	0.16	1.50 (16.8)	11.2 (3.9)
	SLIM [51]	68.98 → 67.21	1.77	88.68 → 87.39	1.29	1.31 (28.0)	-
	LCCL [14]	69.98 → 66.33	3.65	89.24 → 86.94	2.30	1.18 (34.6)	11.7 (0.0)
	Ours 37%-pruned	70.05 → 69.09	0.96	89.40 → 88.59	0.81	1.14 (36.7)	9.3 (20.1)
	SFP [29]	70.28 → 67.10	3.18	89.63 → 87.78	1.85	1.06 (41.8)	-
	DCP [82]	69.64 → 67.35	2.29	88.98 → 87.60	1.38	0.98 (46.0)	6.2 (47.0)
	FPGM [31]	70.28 → 68.41	1.87	89.63 → 88.48	1.15	1.06 (41.8)	-
MobileNet V2	Meta [52]	69.72 → 68.61	1.11	89.07 → 88.35	0.72	1.09 (40.0)	-
	Ours 41%-pruned	70.05 → 68.85	1.20	89.40 → 88.45	0.95	1.07 (41.0)	8.8 (24.5)
	Uniform [65]	71.80 → 69.80	2.00	-	-	0.22 (26.9)	-
	AMC [30]	71.80 → 70.80	1.00	-	-	0.22 (26.9)	-
	CC [43]	71.88 → 70.91	0.89	-	-	0.22 (28.3)	-
	LeGR [9]	72.70 → 71.20	1.50	-	-	0.22 (27.9)	-
	Ours 27%-pruned	72.18 → 71.90	0.28	90.49 → 90.38	0.11	0.22 (26.9)	2.8 (20.4)
	DCP [82]	70.11 → 64.22	5.89	-	3.77	0.17 (44.7)	2.6 (25.9)
	Meta [52]	72.70 → 68.20	4.50	-	-	0.14 (53.4)	-
	Ours 53%-pruned	72.18 → 69.16	3.02	90.49 → 88.66	1.83	0.14 (53.4)	2.1 (39.3)

Table 5: ILSVRC-2012 Transfer Pruning Results. We report our pruned models at different FLOPs levels to ensure a fair comparison with different prior arts. We add a suffix specifying FLOPs pruning percentage for each of our pruned model.

to all of them. We find that our evolved function achieves state-of-the-art results on all datasets.

Settings. We adopt a one-shot pruning scheme with a uniform pruning ratio across layers for our transfer pruning and use the SGD optimizer with Nesterov Momentum [57] for retraining. The weight decay factor and the momentum are set to be 0.0001 and 0.9, respectively. On SVHN/CIFAR-100, we use a batch size of 32/128 to fine-tune the network with 20/200 epochs. The learning rate is initialized at 0.05 and multiplied by 0.14 at 40% and 80% of the total number of epochs. On ILSVRC-2012, we use a batch size of 128 to fine-tune VGG-16/ResNet-18/MobileNet-V2 for 30/100/100 epochs. For VGG-16/ResNet-18, the learning rate is started at 0.0006 and multiplied by 0.4 at 40% and 80% of the total number of epochs. We use a cosine decay learning rate schedule for MobileNet-V2 [65] with an initial rate of 0.03.

SVHN. We first evaluate ξ^* on SVHN with ResNet-164. Ours outperforms SLIM [51] by 0.1% in accuracy with significant hardware resource savings: 32.1% more FLOPs saving and 48.5% more parameters saving, which well demonstrates the effectiveness of ξ^* .

CIFAR-100. On VGG-19, our pruned model achieves an accuracy gain of 0.35% with respect to G-SD [49]. Compared to LFPC [28] and LeGR [9], our pruned ResNet-56 achieves an accuracy gain of 0.87% and 0.66%, respectively, while having 5% less FLOPs. On ResNet-110, our method outperforms FPGM [31] and TAS [15] by 1.30% and 0.69% in terms of accuracy with 4% less FLOPs. In comparison with LCCL [14], SLIM [51], and DI [38], our pruned ResNet-164 achieves

an accuracy of 77.77% with 63.2% FLOPs reduction which advances all prior methods.

ILSVRC-2012. On VGG-16, our approach improves over the baseline by nearly 1.1% in top-1 accuracy with 2.4 \times acceleration. Our 3.3 \times -accelerated model advances the state of the art by achieving top-1/top-5 accuracies of 71.64%/90.60%. On ResNet-18, our approach reduces 16.8% of the FLOPs without top-1 accuracy loss. Compared to LCCL [14], our method achieves a 2.72% top-1 accuracy gain with a higher FLOPs reduction ratio. We demonstrate top-1 accuracy gains of 1.75% and 1.50% with respect to SFP [29] and DCP [82] with over 40% FLOPs reduction. We finally show our performance on a much more compact network, MobileNet-V2, which is specifically designed for mobile deployment. When 26.9% of FLOPs are pruned, our approach outperforms AMC [30], Meta [52], and LeGR [9] with a top-1 accuracy of 71.90%. At a higher pruning ratio, our method advances DCP [82] and Meta [52] by top-1 accuracies of 4.94% and 0.96%, with 53.4% FLOPs reduction.

6 ABLATION STUDY

Random Initial Population. In Fig. 6, we conduct a control experiment which initializes all individuals as random expression trees to study the effectiveness of initializing our population with SOAP. We also turn off the SOAP function insertion in the reproduction process for the control experiment. All other parameters (number of generations, size of population, α , etc.) are kept to be

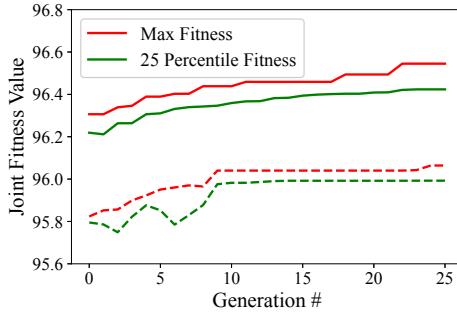


Figure 6: Comparing random initial population evolution (dashed line) with the evolution in Sec. 4 (solid line). Thanks to the expressiveness of our function space, the evolution with randomly-initialized functions also achieve good pruning fitness. However, we observe that it converges very early around the 8th generation and stalls at the plateau for a long period. Moreover, its final fitness has a clear performance gap with respect to the one in Sec. 4.

the same as in Sec. 4 for a fair comparison. We find that evolving with random population also achieves a good pruning fitness, which indicates that our design space is of powerful expressiveness. However, we observe early convergence and a final performance gap in the control experiment compared to the main experiment in Sec. 4, demonstrating the advantage of using SOAP for evolution. **Evolution on ILSVRC-2012.** In contrast to our evolution strategy with a joint fitness function on MNIST and CIFAR-10, we conduct an evolution on only ILSVRC-2012 as a control experiment. We restrict the total computation budget to be the same as Sec. 4, i.e. 98 GPU-days, and evolve on ResNet-18 with a population size of 40 over 25 generations. Due to the constrained budget, each pruned net is only retrained for 4 epochs. We include detailed evolution settings and results in Supplementary. Two major drawbacks are found with this evolution strategy: (1) **Imprecise evaluation.** Due to the lack of training epochs, the function’s actual effectiveness is not precisely revealed. We take two functions with fitness 63.24 and 63.46 from the last generation, and use them again to prune ResNet-18 but fully retrain for 100 epochs. We find that the one with lower fitness in evolution achieves an accuracy of 68.27% in the full training, while the higher one only has an accuracy of 68.02%. Such result indicates that the evaluation in this evolution procedure could be inaccurate, while our strategy ensures a full retraining for precise effectiveness assessment. (2) **Inferior performance.** The best evolved function with this method, $\xi_{ImageNet}$ (in Supplementary), performs inferior to ξ^* shown in Eqn. 3 when transferred to a different dataset. In particular, when applied to pruning 56% FLOPs from ResNet-110 on CIFAR-100, $\xi_{ImageNet}$ only achieves an accuracy of 72.51% while ξ^* reaches 73.85%. These two issues suggest that evolving on two small datasets would have better cost-effectiveness than using a single large scale dataset like ILSVRC-2012.

Feature Selection. We further apply ξ^* to another machine learning task, feature selection, to visually understand our evolved function. In particular, we compare ξ^* (right) vs. DI [38] (middle) on MNIST feature selection in Fig. 7. The red pixels indicate the important features evaluated by the metrics, while the blue ones are

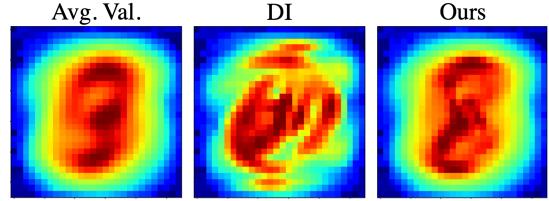


Figure 7: Feature selection by DI [38] (middle) and ξ^* (right) for MNIST, where ξ^* tends to preserve features with higher means and more robust pattern in reference of the average feature values map (left).

redundant. Taking the average feature values map (left) for reference, we find that our evolved function tends to select features with higher means, where the MNIST pattern is more robust.

7 CONCLUSION

In this work, we propose a novel paradigm integrating evolutionary learning with channel pruning, which first learns novel channel pruning functions from small datasets, and then transfers them to larger and more challenging datasets. We develop an end-to-end genetic programming framework to automatically search for transferable pruning functions over our novel function design space without any manual modification after evolution. We present and analyze a closed-form evolved function which can offer strong pruning performance and further streamline the design of our pruning strategy. The learned pruning function exhibits remarkable generalizability to datasets different from those in the evolution process. Specifically, on SVHN, CIFAR-100, and ILSVRC-2012, we achieve state-of-the-art pruning results.

8 ACKNOWLEDGEMENT

We thank a former colleague from Princeton Parallel Group, Yanqi Zhou, for her help with parallel implementation of the evolution. This material is based upon work supported by the National Science Foundation under Grant No. CCF-1822949. Any opinions, findings, and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

REFERENCES

- [1] Martin Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. 2016. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 16)*. 265–283.
- [2] Bowen Baker, Otkrist Gupta, Nikhil Naik, and Ramesh Raskar. 2016. Designing neural network architectures using reinforcement learning. *arXiv preprint arXiv:1611.02167* (2016).
- [3] Wolfgang Banzhaf, Peter Nordin, Robert E Keller, and Frank D Francone. 1998. *Genetic programming*. Springer.
- [4] Irwan Bello, Barret Zoph, Vijay Vasudevan, and Quoc V Le. 2017. Neural optimizer search with reinforcement learning. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 459–468.
- [5] James Bergstra, Daniel Yamins, and David Daniel Cox. 2013. Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures. (2013).
- [6] Han Cai, Ligeng Zhu, and Song Han. 2018. Proxylessnas: Direct neural architecture search on target task and hardware. *arXiv preprint arXiv:1812.00332* (2018).

[7] Wenlin Chen, James Wilson, Stephen Tyree, Kilian Weinberger, and Yixin Chen. 2015. Compressing neural networks with the hashing trick. In *International Conference on Machine Learning*. 2285–2294.

[8] Yutian Chen, Matthew W Hoffman, Sergio Gómez Colmenarejo, Misha Denil, Timothy P Lillicrap, Matt Botvinick, and Nando De Freitas. 2017. Learning to learn without gradient descent by gradient descent. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org, 748–756.

[9] Ting-Wu Chin, Ruizhou Ding, Cha Zhang, and Diana Marculescu. 2020. Towards Efficient Model Compression via Learned Global Ranking. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 1518–1528.

[10] Matthieu Courbariaux, Itay Hubara, Daniel Soudry, Ran El-Yaniv, and Yoshua Bengio. 2016. Binarized neural networks: Training deep neural networks with weights and activations constrained to +1 or -1. *arXiv preprint arXiv:1602.02830* (2016).

[11] Xiaoliang Dai, Peizhao Zhang, Bichen Wu, Hongxu Yin, Fei Sun, Yanghan Wang, Marat Dukhan, Yunqing Hu, Yiming Wu, Yangqing Jia, et al. 2019. Chamnet: Towards efficient network design through platform-aware model adaptation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 11398–11407.

[12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*. Ieee, 248–255.

[13] Chao Dong, Chen Change Loy, Kaiming He, and Xiaou Tang. 2015. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence* 38, 2 (2015), 295–307.

[14] Xuanyi Dong, Junshi Huang, Yi Yang, and Shuicheng Yan. 2017. More is less: A more complicated network with less inference complexity. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 5840–5848.

[15] Xuanyi Dong and Yi Yang. 2019. Network pruning via transformable architecture search. In *Advances in Neural Information Processing Systems*. 759–770.

[16] Chrisantha Fernando, Dylan Banarse, Malcolm Reynolds, Frederic Besse, David Pfau, Max Jaderberg, Marc Lanctot, and Daan Wierstra. 2016. Convolution by evolution: Differentiable pattern producing networks. In *Proceedings of the Genetic and Evolutionary Computation Conference 2016*. ACM, 109–116.

[17] Matthias Feurer, Aaron Klein, Katharina Eggensperger, Jost Springenberg, Manuel Blum, and Frank Hutter. 2015. Efficient and robust automated machine learning. In *Advances in neural information processing systems*. 2962–2970.

[18] Dario Floreano, Peter Dür, and Claudia Mattiussi. 2008. Neuroevolution: from architectures to learning. *Evolutionary Intelligence* 1, 1 (2008), 47–62.

[19] Félix-Antoine Fortin, François-Michel De Rainville, Marc-André Gardner, Marc Parizeau, and Christian Gagné. 2012. DEAP: Evolutionary Algorithms Made Easy. *Journal of Machine Learning Research* 13 (jul 2012), 2171–2175.

[20] Xitong Gao, Yiren Zhao, Łukasz Dudziak, Robert Mullins, and Cheng-zhong Xu. 2018. Dynamic channel pruning: Feature boosting and suppression. *arXiv preprint arXiv:1810.05331* (2018).

[21] David E Goldberg and Kalyanmoy Deb. 1991. A comparative analysis of selection schemes used in genetic algorithms. In *Foundations of genetic algorithms*. Vol. 1. Elsevier, 69–93.

[22] Todd R Golub, Donna K Slonim, Pablo Tamayo, Christine Huard, Michelle Gaasenbeek, Jill P Mesirov, Hilary Coller, Mignon L Loh, James R Downing, Mark A Caligiuri, et al. 1999. Molecular classification of cancer: class discovery and class prediction by gene expression monitoring. *science* 286, 5439 (1999), 531–537.

[23] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. *Advances in neural information processing systems* 27 (2014).

[24] Arthur Gretton, Karsten M Borgwardt, Malte J Rasch, Bernhard Schölkopf, and Alexander Smola. 2012. A kernel two-sample test. *Journal of Machine Learning Research* 13, Mar (2012), 723–773.

[25] Song Han, Huizi Mao, and William J Dally. 2015. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149* (2015).

[26] Song Han, Jeff Pool, John Tran, and William Dally. 2015. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*.

[27] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.

[28] Yang He, Yuhang Ding, Ping Liu, Linchao Zhu, Hanwang Zhang, and Yi Yang. 2020. Learning Filter Pruning Criteria for Deep Convolutional Neural Networks Acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2009–2018.

[29] Yang He, Guoliang Kang, Xuanyi Dong, Yanwei Fu, and Yi Yang. 2018. Soft filter pruning for accelerating deep convolutional neural networks. *arXiv preprint arXiv:1808.06866* (2018).

[30] Yihui He, Ji Lin, Zhijian Liu, Hanrui Wang, Li-Jia Li, and Song Han. 2018. Amc: Automl for model compression and acceleration on mobile devices. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 784–800.

[31] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. 2019. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4340–4349.

[32] Yihui He, Xiangyu Zhang, and Jian Sun. 2017. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 1389–1397.

[33] Qiangui Huang, Kevin Zhou, Suya You, and Ulrich Neumann. 2018. Learning to prune filters in convolutional neural networks. In *2018 IEEE Winter Conference on Applications of Computer Vision (WACV)*. IEEE, 709–718.

[34] Max Jaderberg, Andrea Vedaldi, and Andrew Zisserman. 2014. Speeding up convolutional neural networks with low rank expansions. *arXiv preprint arXiv:1405.3866* (2014).

[35] Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).

[36] John R Koza and John R Koza. 1992. *Genetic programming: on the programming of computers by means of natural selection*. Vol. 1. MIT press.

[37] Alex Krizhevsky and Geoffrey Hinton. 2009. *Learning multiple layers of features from tiny images*. Technical Report. Citeseer.

[38] S.Y. Kung, Zeyjiang Hou, and Yuchen Liu. 2019. Methodical Design and Trimming of Deep Learning Networks: Enhancing External BP Learning with Internal Omnipresent-supervision Training Paradigm. In *ICASSP 2019-2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 8058–8062.

[39] Vadim Lebedev, Yaroslav Ganin, Maksim Rakhuba, Ivan Osledets, and Victor Lempitsky. 2014. Speeding-up convolutional neural networks using fine-tuned cp-decomposition. *arXiv preprint arXiv:1412.6553* (2014).

[40] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. 1998. Gradient-based learning applied to document recognition. *Proc. IEEE* 86, 11 (1998).

[41] Erich L Lehmann and Joseph P Romano. 2006. *Testing statistical hypotheses*. Springer Science & Business Media.

[42] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. 2016. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710* (2016).

[43] Yuchao Li, Shaohui Lin, Jianzhuang Liu, Qixiang Ye, Mengdi Wang, Fei Chao, Fan Yang, Jincheng Ma, Qi Tian, and Rongrong Ji. 2021. Towards Compact CNNs via Collaborative Compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 6438–6447.

[44] Yuchao Li, Shaohui Lin, Baochang Zhang, Jianzhuang Liu, David Doermann, Yongjian Wu, Feiyue Huang, and Rongrong Ji. 2019. Exploiting kernel sparsity and entropy for interpretable CNN compression. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2800–2809.

[45] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. 2017. Runtime neural pruning. In *Advances in Neural Information Processing Systems*. 2181–2191.

[46] Chenxi Liu, Barret Zoph, Maxim Neumann, Jonathon Shlens, Wei Hua, Li-Jia Li, Li Fei-Fei, Alan Yuille, Jonathan Huang, and Kevin Murphy. 2018. Progressive neural architecture search. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 19–34.

[47] Hanxiao Liu, Karen Simonyan, Oriol Vinyals, Chrisantha Fernando, and Koray Kavukcuoglu. 2017. Hierarchical representations for efficient architecture search. *arXiv preprint arXiv:1711.00436* (2017).

[48] Yuchen Liu, Zhixin Shu, Yijun Li, Zhe Lin, Federico Perazzi, and Sun-Yuan Kung. 2021. Content-aware gan compression. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 12156–12166.

[49] Yuchen Liu, David Wentzlaff, and SY Kung. 2020. Rethinking Class-Discrimination Based CNN Channel Pruning. *arXiv preprint arXiv:2004.14920* (2020).

[50] Yuchen Liu, David Wentzlaff, and SY Kung. 2021. Class-Discriminative CNN Compression. *arXiv preprint arXiv:2110.10864* (2021).

[51] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. 2017. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE International Conference on Computer Vision*. 2736–2744.

[52] Zechun Liu, Haoyuan Mu, Xiangyu Zhang, Zichao Guo, Xin Yang, Kwang-Ting Cheng, and Jian Sun. 2019. Metapruning: Meta learning for automatic neural network channel pruning. In *Proceedings of the IEEE International Conference on Computer Vision*. 3296–3305.

[53] Christos Louizos, Max Welling, and Diederik P Kingma. 2017. Learning Sparse Neural Networks through L_0 Regularization. *arXiv preprint arXiv:1712.01312* (2017).

[54] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. 2017. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*. 5058–5066.

[55] Man-Wai Mak and Sun-Yuan Kung. 2006. A solution to the curse of dimensionality problem in pairwise scoring techniques. In *International Conference on Neural Information Processing*. Springer, 314–323.

[56] Risto Miikkulainen, Jason Liang, Elliot Meyerson, Aditya Rawal, Daniel Fink, Olivier Francon, Bala Raju, Hormoz Shahrazad, Arshak Navruzyan, Nigel Duffy, et al. 2019. Evolving deep neural networks. In *Artificial Intelligence in the Age of*

Neural Networks and Brain Computing. Elsevier, 293–312.

[57] Yurii E Nesterov. 1983. A method for solving the convex programming problem with convergence rate $O(1/k^2)$. In *Dokl. akad. nauk Ssr*, Vol. 269. 543–547.

[58] Yuval Netzer, Tao Wang, Adam Coates, Alessandro Bissacco, Bo Wu, and Andrew Y Ng. 2011. Reading digits in natural images with unsupervised feature learning. (2011).

[59] Xuefei Ning, Tianchen Zhao, Wenshuo Li, Peng Lei, Yu Wang, and Huazhong Yang. 2020. Dsa: More efficient budgeted pruning via differentiable sparsity allocation. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part III* 16. Springer, 592–607.

[60] Damien O’Neill, Bing Xue, and Mengjie Zhang. 2020. Neural architecture search for sparse DenseNets with dynamic compression. In *Proceedings of the 2020 Genetic and Evolutionary Computation Conference*. 386–394.

[61] Paul Pavlidis, Jason Weston, Jinsong Cai, and William Noble Grundy. 2001. Gene functional classification from heterogeneous data. In *Proceedings of the fifth annual international conference on Computational biology*. 249–255.

[62] Esteban Real, Alok Aggarwal, Yanping Huang, and Quoc V Le. 2019. Regularized evolution for image classifier architecture search. In *Proceedings of the aaai conference on artificial intelligence*, Vol. 33. 4780–4789.

[63] Esteban Real, Sherry Moore, Andrew Selle, Saurabh Saxena, Yutaka Leon Sue-matsu, Jie Tan, Quoc V Le, and Alexey Kurakin. 2017. Large-scale evolution of image classifiers. In *Proceedings of the 34th International Conference on Machine Learning–Volume 70*. JMLR.org, 2902–2911.

[64] Shaogang Ren, Kaiming He, Ross Girshick, and Jian Sun. 2015. Faster r-cnn: Towards real-time object detection with region proposal networks. In *Advances in neural information processing systems*. 91–99.

[65] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 4510–4520.

[66] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

[67] Nilopal Sinha and Kuan-Wen Chen. 2021. Evolving neural architecture using one shot model. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 910–918.

[68] Jasper Snoek, Oren Rippel, Kevin Swersky, Ryan Kiros, Nadathur Satish, Narayanan Sundaram, Mostafa Patwary, Mr Prabhat, and Ryan Adams. 2015. Scalable bayesian optimization using deep neural networks. In *International conference on machine learning*. 2171–2180.

[69] Kenneth O Stanley, Jeff Clune, Joel Lehman, and Risto Miikkulainen. 2019. Designing neural networks through neuroevolution. *Nature Machine Intelligence* 1, 1 (2019), 24–35.

[70] Kenneth O Stanley, David B D’Ambrosio, and Jason Gauci. 2009. A hypercube-based encoding for evolving large-scale neural networks. *Artificial life* 15, 2 (2009), 185–212.

[71] Kenneth O Stanley and Risto Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* 10, 2 (2002), 99–127.

[72] Masanori Suganuma, Shinichi Shirakawa, and Tomoharu Nagao. 2017. A genetic programming approach to designing convolutional neural network architectures. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 497–504.

[73] Mingxing Tan, Bo Chen, Ruoming Pang, Vijay Vasudevan, Mark Sandler, Andrew Howard, and Quoc V Le. 2019. Mnasnet: Platform-aware neural architecture search for mobile. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2820–2828.

[74] Mingxing Tan and Quoc V Le. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. *arXiv preprint arXiv:1905.11946* (2019).

[75] Paul Templier, Emmanuel Rachelson, and Dennis G Wilson. 2021. A geometric encoding for neural network evolution. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 919–927.

[76] Olga Wichrowska, Niru Maheswaranathan, Matthew W Hoffman, Sergio Gomez Colmenarejo, Misha Denil, Nando de Freitas, and Jascha Sohl-Dickstein. 2017. Learned optimizers that scale and generalize. In *Proceedings of the 34th International Conference on Machine Learning–Volume 70*. JMLR.org, 3751–3760.

[77] Lingxi Xie and Alan Yuille. 2017. Genetic cnn. In *Proceedings of the IEEE International Conference on Computer Vision*. 1379–1388.

[78] Jianbo Ye, Xin Lu, Zhe Lin, and James Z Wang. 2018. Rethinking the smaller-norm-less-informative assumption in channel pruning of convolution layers. *arXiv preprint arXiv:1802.00124* (2018).

[79] Andy B Yoo, Morris A Jette, and Mark Grondona. 2003. Slurm: Simple linux utility for resource management. In *Workshop on Job Scheduling Strategies for Parallel Processing*. Springer, 44–60.

[80] Ruichi Yu, Ang Li, Chun-Fu Chen, Jui-Hsin Lai, Vlad I Morariu, Xintong Han, Mingfei Gao, Ching-Yung Lin, and Larry S Davis. 2018. Nisp: Pruning networks using neuron importance score propagation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 9194–9203.

[81] Tianyun Zhang, Shaokai Ye, Kaili Zhang, Jian Tang, Wujie Wen, Makan Fardad, and Yanzhi Wang. 2018. A systematic dnn weight pruning framework using alternating direction method of multipliers. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 184–199.

[82] Zhuangwei Zhuang, Mingkui Tan, Bohan Zhuang, Jing Liu, Yong Guo, Qingyao Wu, Junzhou Huang, and Jinhui Zhu. 2018. Discrimination-aware channel pruning for deep neural networks. In *Advances in Neural Information Processing Systems*. 875–886.

[83] Barret Zoph and Quoc V Le. 2016. Neural architecture search with reinforcement learning. *arXiv preprint arXiv:1611.01578* (2016).

[84] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. 2018. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 8697–8710.

	addition	add(+)
	subtraction	sub(-)
	multiplication	mul(x)
	division	div(÷)
	absolute value	abs
	square	sq
	square root	sqrt
	adding ridge factor	ridge
Elementwise operators	matrix trace	tr
	matrix multiplication	matmul
	matrix inversion	inv
	inner product	dot
	outer product	outprod
	matrix/vector transpose	tran
Statistics operators	summation	sum _{s,g}
	product	prod _{s,g}
	mean	mean _{s,g}
	standard deviation	std _{s,g}
	variance	var _{s,g}
	counting measure	count _{s,g}
Specialized operators	rbf kernel matrix getter	rbf
	geometric median getter	geo
	tensor slicer	slice

Table 6: Detailed Operator Space

We organize our supplementary material as follows. In Sec. 9, we present a more detailed table for the operator space and our implementation of the state-of-the-art population (SOAP). In Sec. 10, we include more experimental details of our evolution and pruning study. We discuss more detailed settings and results of evolution on ILSVRC-2012 in Sec. 11. Lastly, we present extra evolved functions in Sec. 12.

9 SOAP IMPLEMENTATION

9.1 Operator Space

In Tab. 6, we present the detailed operator space with operators and their abbreviations.

9.2 SOAP Functions

With the abbreviations of operators in Tab. 6 and the symbols of operands presented in Tab. 1 of the main paper, we can thus give the precise expressions of the functions in SOAP:

- Filter’s ℓ_1 -norm: $\text{sum}_g(\text{abs}(\mathcal{W}_I))$
- Filter’s ℓ_2 -norm: $\text{sqrt}(\text{sum}_g(\text{sq}(\mathcal{W}_I)))$
- Batch normalization’s scaling factor: $\text{abs}(\text{slice}(\mathcal{B}))$
- Filter’s geometric median: $\text{sqrt}(\text{sum}_g(\text{sq}(\mathcal{W}_I - \text{geo}(\mathcal{W}))))$
- Discriminant Information:
 $\text{count}_s(\mathcal{F}^+) \times \text{matmul}(\text{tran}(\text{mean}_s(\mathcal{F}^+) - \text{mean}_s(\mathcal{F})), \text{inv}(\text{ridge}(\text{matmul}(\text{tran}(\mathcal{F} - \text{mean}_s(\mathcal{F})), \mathcal{F} - \text{mean}_s(\mathcal{F})))), \text{mean}_s(\mathcal{F}^+) - \text{mean}_s(\mathcal{F}))$
- Maximum Mean Discrepancy:
 $\text{div}(\text{sum}_g(\text{rbf}(\mathcal{F}^+, \mathcal{F}^+)), \text{sq}(\text{count}_s(\mathcal{F}^+))) + \text{div}(\text{sum}_g(\text{rbf}(\mathcal{F}^-, \mathcal{F}^-)), \text{sq}(\text{count}_s(\mathcal{F}^-))) - \text{div}(\text{sum}_g(\text{rbf}(\mathcal{F}^+, \mathcal{F}^-)), \text{mul}(\text{count}_s(\mathcal{F}^+), \text{count}_s(\mathcal{F}^-)))$

- Generalized Absolute SNR:
 $\text{div}(\text{sum}_g(\text{rbf}(\mathcal{F}^+, \mathcal{F}^-)), \text{mul}(\text{count}_s(\mathcal{F}^+), \text{count}_s(\mathcal{F}^-)))$
- Generalized Student’s T-Test:
 $\text{div}(\text{abs}(\text{mean}_g(\mathcal{F}^+) - \text{mean}_g(\mathcal{F}^-)), \text{std}_g(\mathcal{F}^+) + \text{std}_g(\mathcal{F}^-))$
- Generalized Fisher Discriminat Ratio:
 $\text{div}(\text{sq}(\text{mean}_g(\mathcal{F}^+) - \text{mean}_g(\mathcal{F}^-)), \text{var}_g(\mathcal{F}^+) + \text{var}_g(\mathcal{F}^-))$
- Generalized Symmetric Divergence:
 $\text{div}(\text{var}_g(\mathcal{F}^+), \text{var}_g(\mathcal{F}^-)) + \text{div}(\text{var}_g(\mathcal{F}^-), \text{var}_g(\mathcal{F}^+)) + \text{div}(\text{sq}(\text{mean}_g(\mathcal{F}^+) - \text{mean}_g(\mathcal{F}^-)), \text{var}_g(\mathcal{F}^+) + \text{var}_g(\mathcal{F}^-))$

10 EXPERIMENTAL DETAILS

10.1 Study on Fitness Combination Scheme

Preliminary Evolution. We conduct 10 preliminary experiments, where the variables are: $\alpha \in \{0, 0.3, 0.5, 0.7, 1\}$ and combination scheme $\in \{\text{weighted geometric mean, weighted arithmetic mean}\}$. For each experiment, we have a population of 15 functions which are evolved for 10 generations. The population is initialized with 10 individuals randomly cloned from SOAP and 5 random expression trees. The tournament size is 3, and the number of the selected functions is 5. The next generation is reproduced only from the selected functions. Other settings are the same as the main evolution experiment.

CIFAR-100 Pruning. We apply the best evolved functions from each preliminary evolution test to prune a ResNet-38 [27] on CIFAR-100 [37]. The baseline ResNet-38 adopts the bottleneck block structure with an accuracy of 72.3%. We use each evolved function to prune 40% of channels in all layers uniformly, resulting in a 54.7%/52.4% FLOPs/parameter reduction. The network is then fine-tuned by the SGD optimizer with 200 epochs. We use the Nesterov Momentum [57] with a momentum of 0.9. The mini-batch size is set to be 128, and the weight decay is set to be 1e-3. The training data is transformed with a standard data augmentation scheme [27]. The learning rate is initialized at 0.1 and divided by 10 at epoch 80 and 160.

10.2 Main Evolution Experiment

MNIST Pruning. On MNIST [40] pruning task, we prune a LeNet-5 [40] with a baseline accuracy of 99.26% from shape of 20-50-800-500 to 5-12-160-40. Such pruning process reduces 92.4% of FLOPs and 98.0% of parameters. The pruned network is fine-tuned for 300 epochs with a batch size of 200 and a weight decay of 7e-5. We use the Adam optimizer [35] with a constant learning rate of 5e-4.

CIFAR-10 Pruning. For CIFAR-10 [37] pruning, we adopt the VGG-16 structure from [42] with a baseline accuracy of 93.7%. We uniformly prune 40% of the channels from all layers resulting in 63.0% FLOPs reduction and 63.7% parameters reduction. The fine-tuning process takes 200 epochs with a batch size of 128. We set the weight decay to be 1e-3 and the dropout ratio to be 0.3. We use the SGD optimizer with Nesterov momentum [57], where the momentum is set to be 0.9. We augment the training samples with a standard data augmentation scheme [27]. The initial learning rate

$$\xi_{ImageNet}(C) = \left(\frac{\text{var}_g(\text{mean}_s(\mathcal{F}^+))}{\text{std}_g(\text{tr}(\mathcal{F}^+)) \times \text{mean}_g(\mathcal{F}^-)} \right)^4 \div \text{var}_g(\text{sqrt}(\mathcal{F})) \quad (4)$$

is set to be 0.006 and multiplied by 0.28 at 40% and 80% of the total number of epochs.

10.3 Transfer Pruning

We implement the pruning experiments in TensorFlow [1] and carry them out with NVIDIA Tesla P100 GPUs. CIFAR-100 contains 50,000/10,000 training/test samples in 100 classes. SVHN is a 10-class dataset where we use 604,388 training images for network training with a test set of 26,032 images. ILSVRC-2012 contains 1.28 million training images and 50 thousand validation images in 1000 classes. We adopt the standard data augmentation scheme [27] for CIFAR-100 and ILSVRC-2012.

10.4 Channel Scoring

As many of our pruning functions require activation maps of the channels to determine channels' importance, we need to feed-forward the input images for channel scoring. Specifically, for pruning experiments on MNIST, CIFAR-10, and CIFAR-100, we use all their training images to compute the channel scores. On SVHN and ILSVRC-2012, we randomly sample 20 thousand and 10 thousand training images for channel scoring, respectively.

11 EVOLUTION ON ILSVRC-2012

Evolution. We use ResNet-18 as the target network for pruning function evolution on ILSVRC-2012. Since only one task is evaluated, we directly use the retrained accuracy of the pruned network as the function's fitness. Other evolution settings for population, selection, mutation, and crossover are kept the same as Sec. 4 of the main paper.

Evaluation. We uniformly prune 30% of channels in each layer from a pretrained ResNet-18, resulting in a FLOPs reduction of 36.4%. Due to the constrained computational budget, we only fine-tune it for 4 epochs using the SGD optimizer with Nesterov momentum [57]. We use a batch size of 128 and initialize our learning rate at 0.001. The learning rate is multiplied by 0.4 at epoch 1 and 2.

Result. We show the evolution progress in Fig. 8. Due to the lack of training budget, the pruned net is clearly not well retrained as they only achieve around 63.5% accuracy, much lower than the performance of methods shown in Tab. 5 of the main paper at the similar pruning level. Such inadequate training results in a imprecise function fitness evaluation evidenced in Sec. 6 of the main paper. Moreover, the best evolved function from this strategy, $\xi_{ImageNet}$ (Eqn. 4), performs inferior to the co-evolved function ξ^* when transferred for CIFAR-100 pruning. These results demonstrate the advantage of our small dataset co-evolution strategy in cost-effectiveness.

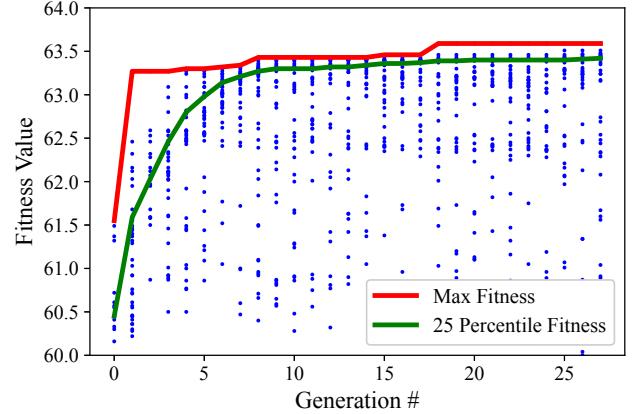


Figure 8: Function evolution on ImageNet.

12 EXTRA EVOLVED FUNCTIONS

We present additional evolved functions from our co-evolution strategy:

$$\xi_1(C) = \frac{\|\bar{f} - \text{var}_g(\mathcal{F}^-)\mathbf{1}\|_2^2}{\text{var}_g(\mathcal{F}^+) + \text{var}_g(\mathcal{F}^-)} + \text{var}_g(\mathcal{F}^+) \quad (5)$$

$$\xi_2(C) = \text{var}_g(\mathcal{F}^+) \quad (6)$$

$$\xi_3(C) = \text{var}_g(\mathcal{W}_I) \quad (7)$$

Eqn. 5 presents a metric with the concept of SNR for classification, while having a novel way of statistics combination. Moreover, our evolution experiments find that measuring the variance across all elements in \mathcal{F}^+ (Eqn. 6) and \mathcal{W}_I (Eqn. 7) would help us identify important channels empirically. These two functions are simple and effective yet remain undiscovered from the literature.