GeneCAI: Genetic Evolution for Acquiring Compact AI

Mojan Javaheripi
ECE Dept., UCSD, USA
mojan@ucsd.edu

Mohammad Samragh
ECE Dept., UCSD, USA
msamragh@ucsd.edu

Tara Javidi
ECE Dept., UCSD, USA
tjavidi@ucsd.edu

Farinaz Koushanfar
ECE Dept., UCSD, USA
farinaz@ucsd.edu

ABSTRACT

In the contemporary big data realm, Deep Neural Networks (DNNs) are evolving towards more complex architectures to achieve higher inference accuracy. Model compression techniques can be leveraged to efficiently deploy these compute-intensive architectures on resource-limited mobile devices. Such methods comprise various hyperparameters that require per-layer customization to ensure high accuracy. Choosing the hyperparameters is cumbersome as the pertinent search space grows exponentially with model layers. This paper introduces GeneCAI, a novel optimization method that automatically *learns* how to tune per-layer compression hyperparameters. We devise a bijective translation scheme that encodes compressed DNNs to the *genotype* space. Each genotype's optimality is measured using a multi-objective score based on the accuracy and number of floating-point operations. We develop customized genetic operations to iteratively evolve the non-dominated solutions towards the optimal Pareto front, thus, capturing the optimal trade-off between model accuracy and complexity. GeneCAI optimization method is highly scalable and can achieve a near-linear performance boost on distributed multi-GPU platforms. Our extensive evaluations demonstrate that GeneCAI outperforms existing rule-based and reinforcement learning methods in DNN compression by finding models that lie on a better accuracy/complexity Pareto curve.

CCS CONCEPTS

• Computing methodologies → Neural networks; Genetic algorithms; Search methodologies; • Hardware → Hardware-software codesign; • Theory of computation → Bio-inspired optimization;

KEYWORDS

Deep Learning, Genetic Algorithms, Multi-objective Optimization, Computer aided/automated design, Parallel Optimization

ACM Reference Format:

Mojan Javaheripi, Mohammad Samragh, Tara Javidi, and Farinaz Koushanfar. 2020. GeneCAI: Genetic Evolution for Acquiring Compact AI. In *Genetic and Evolutionary Computation Conference (GECCO '20)*, July 8–12, 2020, Cancún, Mexico. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3377930.3390226>

1 INTRODUCTION

With the growing range of applications for Deep Neural Networks (DNNs), the demand for higher accuracy has led to a continuous increase in the complexity of state-of-the-art models. Such high

execution cost hinders the deployment of DNNs in real-time applications on commodity hardware. Fortunately, modern neural networks have been shown to incur high redundancies that can be eliminated without compromising inference accuracy [15]. Effective identification and removal of such redundancies have fueled research in two interlinked domains: (i) Developing model compression techniques, e.g., pruning [9, 11, 17, 21, 24, 40], quantization [47], and coding [7]. (ii) Devising automated policies that learn how to configure compression techniques to jointly achieve accuracy and compactness [10, 16, 33, 34, 39, 43]. In this paper, we focus on the latter.

The effectiveness of contemporary compression techniques relies on the careful tuning of several hyperparameters across DNN layers. These hyperparameters directly control the trade-off between accuracy and execution cost on a constrained device. The question to be answered is how to find an optimal hyperparameter configuration that results in a high compression rate while minimally affecting inference accuracy. Figure 1 shows how an intelligent hyperparameter selection policy can better estimate the geometry of the optimal Pareto front for the same compression technique (Pruning).

Figure 1: Accuracy/FLOPs Pareto curves for pruning a pre-trained VGG network on CIFAR-10 benchmark.

Finding the optimal set of hyperparameters is quite challenging as the space of possibilities grows exponentially with the number of DNN layers. Such a large search-space renders manual or computerized greedy algorithms sub-optimal or infeasible for hyperparameter tuning. Existing work in automated compression suggests the use of heuristic methods [9, 11, 17] or Reinforcement Learning (RL) [10, 42, 43]. To tackle the high-dimensionality of the search space, heuristics and RL algorithms specify the hyperparameters one layer at a time, thus, overlooking the layer inter-dependencies.

In this paper, we take an alternative approach, called GeneCAI, that *simultaneously* tunes the compression hyperparameters for all DNN layers using genetic algorithms (GA). We devise a customized translator that encapsulates the hyperparameter configuration of all DNN layers as *one* fixed-length vector $\vec{x} \in \mathbb{R}^d$, called an *individual*. Using this translation scheme, each hyperparameter vector translates to a unique compressed DNN. In this setting, the search for compression hyperparameters directly translates to optimizing an objective function $f(\vec{x}) : \mathbb{R}^d \rightarrow \mathbb{R}$ over $\vec{x}$ where $f(\cdot)$ is an arbitrary measure of quality. We suggest a unified formulation for $f(\cdot)$, called the *score*, which assesses the quality of individuals by combining inference accuracy and a desired execution cost, e.g., runtime.

GeneCAI iteratively *evolves* the initial *population* towards superior individuals. To promote a high-quality initialization, we propose a context-aware boundary characterization that significantly contributes to convergence and search results. After initialization, each evolution iteration comprises three consecutive genetic operations,

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

GECCO '20, July 8–12, 2020, Cancún, Mexico

© 2020 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.

ACM ISBN 978-1-4503-7128-5/20/07...\$15.00

<https://doi.org/10.1145/3377930.3390226>

namely, *selection*, *crossover*, and *mutation*. We balance GeneCAI exploration and exploitation by devising a diversity-guided GA. In particular, we define and measure diversity at the genotype level and adapt our crossover and mutation operations accordingly. By maintaining a diversified population, GeneCAI ensures a fast and stable search and prevents premature convergence.

In high-level, GeneCAI offers several benefits: **(I)** it is highly cost-effective as it takes a post-processing approach where a pre-trained DNN is compressed while no training is required amid GA iterations. **(II)** GeneCAI exploits parallelism to reduce optimization time by concurrent sample evaluations; We develop an accompanying API that manages the infrastructure for multi-GPU execution of GeneCAI search. **(III)** GeneCAI makes no assumptions about the objective function; by adjusting the cost, one can apply GeneCAI to various DNN compression tasks and hardware constraints. **(IV)** GeneCAI encoding and genetic operations accommodate optimization of both discrete and continuous-valued hyperparameters.

To demonstrate GeneCAI generalizability, we evaluate our methodology on four compression techniques and several architectures. Our evaluations unveil the full potential of GA in learning to effectively combine multiple methods and push the limits of network compression. For VGG-16 on ImageNet, GeneCAI pushes the state-of-the-art FLOPs reduction from $5\times$ to $7.1\times$ with higher accuracy. For expert-designed efficient MobileNets, GeneCAI obtains on average 1.2% higher top-1 accuracy than the MobileNet Pareto curve [12]. We further show that GeneCAI computational flow is highly scalable and enjoys a linear reduction in search time with an increased number of distributed computing resources.

2 BACKGROUND AND RELATED WORK

Neural Network Compression. Existing model compression techniques can be roughly categorized in two classes: post-training compression and compression-aware training. The post-training approach is applied to pre-trained DNNs, without modifying the standard training loss or the back-propagation routines [9, 11, 17, 18, 21, 23, 24, 27, 32, 38]. Alternatively, in compression-aware training [1, 5], the training loss or back-propagation algorithm is altered to create models that are inherently *compressible*.

Each compression method has its benefits and tradeoffs; In either case, the compression hyperparameters still need to be specified. Most of the existing works in DNN compression literature rely on hand-crafted or heuristic methods for per-layer configuration of the compression hyperparameters, which can be cumbersome and result in sub-optimal solutions. Different from heuristic methods that compress one layer at a time, GeneCAI performs whole-network compression where all layers are simultaneously optimized to account for inter-layer correlations. We target post-training compression and apply fine-tuning only once. This allows us to customize DNNs for various hardware platforms with little search/training overhead.

Automated Architecture Search. Modifying the DNN architecture graph in search of higher training/inference efficiency [15] has gained much research traction recently. Several works focus on designing automated methodologies for achieving such compact and accurate neural networks. GA has been applied to Neural Architecture Search (NAS) [14, 22, 26, 29, 37, 41] where the goal is to build a neural network architecture from scratch. Different from NAS, this paper focuses on learning hyperparameters for DNN customization, which acts on pre-trained DNNs and does not require

training-from-scratch. This property allows our algorithm to enjoy a fast search-space exploration using parallel computing resources.

Automated DNN Compression. Authors of [13] develop a pruning scheme that selects the pruned filters using GA rather than magnitude-based or gradient-based approaches [21, 28]. Nevertheless, the amount of pruning applied per layer is still unknown. GeneCAI addresses this remaining challenge by learning the compression hyperparameters. As a result, GeneCAI can be applied to generalized compression techniques and not just network pruning. Reinforcement Learning (RL) [10, 39] is another promising tool for DNN compression. Although effective in finding near-optimal solutions, RL applies compression one layer at a time to overcome the curse of dimensionality. Thus, the RL agent typically needs many learning episodes to identify the inherent inter-layer correlations.

GA is shown to produce competitive results with RL on the challenging tasks of neural architecture search and robotic control while enjoying higher scalability and significantly lower training overhead [29, 31, 41]. Inspired by this, we develop an optimization method based on GAs that achieves better/similar final solutions as prior art with much lower runtime and computational overhead. GeneCAI optimizes all hyperparameters at once with high scalability by exploiting parallelism, especially in distributed settings.

3 PROBLEM FORMULATION

DNN compression, in high-level, is a transformation $\mathcal{T}(M, \vec{x})$ that converts a pre-trained model M to a compressed model $\hat{M}_{\vec{x}}$ with lower computational complexity. In this process, the adjustable hyperparameter vector $\vec{x} \in \mathbb{R}^d$ controls the complexity and accuracy of the output model. A desirable compressed network $\hat{M}_{\vec{x}}$ should satisfy two properties: (i) the generalization capability of the compressed model should resemble the original network and (ii) the execution cost of $\hat{M}_{\vec{x}}$ on the target hardware platform should be as low as possible. These properties outline a multi-objective black-box optimization problem over the hyperparameter vector $\vec{x}$, dubbed an *individual*. We assume we have access to a scoring oracle, $f(\cdot)$, that assesses each individual based on the classification accuracy on a validation set and computational complexity of its corresponding compressed model. The scoring scheme is designed to capture the trade-off between model accuracy $A(\hat{M}_{\vec{x}})$ and hardware cost $C(\hat{M}_{\vec{x}})$. GeneCAI's customized score translates the original bi-objective problem of DNN compression to a single-objective optimization:

$$\max_{\vec{x} \in \mathbb{R}^d} f(A(\hat{M}_{\vec{x}}), C(\hat{M}_{\vec{x}})), \quad (1)$$

For simplicity, we show $f(A(\hat{M}_{\vec{x}}), C(\hat{M}_{\vec{x}}))$ as $f(\vec{x})$ in the rest of the paper. Since full knowledge about $f(\cdot)$ and/or its first derivatives cannot be assumed, often empirical evaluations and optimization is the only viable strategy. Brute-force evaluation, in general, is infeasible as the search-space grows exponentially with d . Instead, we propose an empirical zeroth order optimization algorithm based on GA, dubbed GeneCAI, for finding the maximum value in Eq. (1), i.e., f^* . We denote the population at the t^{th} step of the GA by $\mathbb{P}^t = \{\vec{x}_n^t\}_{n=1}^N$ where N is the population size. GeneCAI aims at finding a near-optimal population after T iterations, $\mathbb{P}^T = \{\vec{x}_n^T\}_{n=1}^N$, where the individuals are in vicinity of the optimal solution:

$$f(\vec{x}_n^T) \approx f^* \quad \forall n \in \{1, \dots, N\} \quad (2)$$

In particular, by adopting a guided search through genetic operations, GeneCAI iteratively samples better DNNs with higher $f(\vec{x})$.

Figure 2: Overview of GeneCAI genetic solution for hyperparameter customization. GeneCAI iteratively improves the non-dominated front by lowering the FLOPs while maintaining a desired accuracy threshold.

4 GENECAI APPROACH

We provide a generic solution to effectively compress a pre-trained DNN while maximally preserving model accuracy. GeneCAI automation policy acts on a pool of hyperparameters and explores the corresponding search-space using customized genetic operations. An overview of GeneCAI genetic solution is shown in Figure 2 and Algorithm 1 presents the corresponding pseudo-code. Below, we summarize the key steps in our optimization framework:

I. First, a pre-processing step characterizes the search-space boundaries within which the optimal solution can reside. These boundaries are specified based upon task-enforced constraints on inference accuracy. Using the acquired boundaries, an initial population of hyperparameter vectors $\mathbb{P}^0 = \{\vec{x}_1^0, \dots, \vec{x}_N^0\}$ are sampled (Sec. 4.4).

II. At each iteration, all individuals $\vec{x}$ are translated to their compressed DNNs $\hat{M}_{\vec{x}}$ (Sec. 4.1). The optimization scores $f(\vec{x})$ are then evaluated in parallel (Sec. 4.2).

III. Based on the new evaluations, GeneCAI genetic operations, i.e., selection, crossover, and mutation (Sec. 4.3), are performed to update the population towards a new, more competent, generation.

By iteratively applying steps II and III above, the non-dominated Pareto front for compressing a desired pre-trained DNN is obtained. In the following, we elaborate on each design component in detail.

Algorithm 1 GeneCAI Search Algorithm

Inputs: fitness oracle $f(\cdot)$, population size (N), iterations (T), crossover parameters (P_{cross}, P_{swap}), mutation parameter P_{tweak} .

Output: Population after T iterations, $\mathbb{P}^T = \{\vec{x}_n^T\}_{n=1}^N$.

```

1:  $\mathbb{P}^0 = \{\vec{x}_n^0\}_{n=1}^N$  ▷ Directed Initialization
2:  $div_{thr} = \frac{div(\mathbb{P}^0)}{2}$ 
3: for  $t = 0 \dots T - 1$  do
4:   Calculate scores:  $\{f_n = f(\vec{x}_n^t)\}_{n=1}^N$  ▷ Evaluation
5:    $\mathbb{S}^t = \text{sample}(\mathbb{P}^t | \{f_n\}_{n=1}^N)$  ▷ Selection
6:   Derive list of parents:  $\mathbb{S}^t \rightarrow \hat{\mathbb{S}}^t = \{(p_1^n, p_2^n)\}_{n=1}^{\frac{N}{2}}$ 
7:   for  $(p_1, p_2)$  in  $\hat{\mathbb{S}}^t$  do ▷ Crossover
8:      $\mathbb{C}^t.append(\text{crossover}(p_1, p_2 | P_{cross}, P_{swap}))$ 
9:   if  $div(\mathbb{C}^t) \leq div_{thr}$  then ▷ Mutate
10:    Adjust  $P_{mutate}$ 
11:     $\mathbb{P}^{t+1} = \text{Mutate}(\mathbb{C}^t | P_{mutate}, P_{tweak})$ 

```

4.1 Genetic Encoding

An initial step for the application of GeneCAI is defining the pertinent search-space for black-box optimization. To this end, we need a global encoding scheme $\mathcal{T}(M, \vec{x})$ that translates a vector of compression hyperparameters $\vec{x}$ to its corresponding compressed architecture

$\hat{M}_{\vec{x}}$. To ensure an effective search, our proposed translation possesses the following characteristics:

- **Distinctness.** Our encoding is bijective: given an individual $\vec{x}$ and model M , the compressed $\hat{M}_{\vec{x}}$ is uniquely determined.
- **Continuity.** A small change in an individual translates to a similarly small alternation in the corresponding compressed DNN. This is particularly important to ensure search convergence.
- **Scalability.** GeneCAI encoding offers a compact, low dimensional representation, enabling effective search. The search-space dimensionality scales linearly with the number of DNN layers and is invariant to the number of DNN parameters/operations.

We direct our focus on *four* compression tasks, namely, structured [11] and non-structured [6] Pruning, Singular Value Decomposition [46], and Tucker-2 approximation [18]. Figure 3 shows a high-level view of our encoding scheme for a 4-layer neural network. For pruning, we allocate 1 hyperparameter per layer: $\vec{x} \in \mathbb{R}^4$. For low-rank approximation, we allocate 2 hyperparameters for Tucker-2 and 1 hyperparameter for SVD, per layer, resulting in $\vec{x} \in \mathbb{R}^6$. In the following, we explain each method’s hyperparameters and their interpretation as a compressed model in detail.

Pruning. Pruning reduces the model size by setting a percentage of low priority parameters/channels to zero [6, 11, 21]. We allocate one continuous value $p \in [0, 1]$, per layer, to represent the ratio of non-zero values. For the example of ResNet-50, each sample will thus be a vector $\vec{x} \in [0, 1]^{50}$. We consider two contemporary DNN pruning methods, namely, structured and non-structured pruning. Structured pruning aims at removing a portion of feature-map channels while non-structured pruning removes a subset of DNN layer weights.

Having defined the vector of hyperparameters $\vec{x}$ for pruning, we implement the transformation $\mathcal{T}(M, \vec{x})$ following common practice in prior art. For structured pruning, we use the sum of absolute gradients of model loss with respect to ReLU feature-map channels for pruning priority [28]. For a ReLU layer with c feature-map channels and pruning rate p , the $\lfloor p \times c \rfloor$ channels with lowest priorities (lower absolute gradients) are removed. For non-structured pruning, we use the absolute value of weight parameters to prioritize them [6]. Given a weight $W \in \mathbb{R}^{k \times k \times c \times f}$ and pruning rate p , the $\lfloor p \times k \times k \times c \times f \rfloor$ elements with lowest absolute values are pruned.

Figure 3: Vectorized representation of an example 4-layer DNN for Pruning and Decomposition. Here, CONV and FC denote convolutional and fully-connected layers, respectively.

SVD. We apply SVD on weight parameters (W) of fully-connected layers ($W \in \mathbb{R}^{c \times f}$) and point-wise convolutions ($W \in \mathbb{R}^{1 \times 1 \times c \times f}$). For SVD, the compression parameter is the decomposition rank which takes an integer value in $\{1, \dots, R\}$ where $R = \min(c, f)$. For efficiency, we discretize the space of approximation ranks in each layer into 64 values and encode them as follows:

$$\text{rank} \in \{1 \leftarrow \frac{R}{64}, \dots, 64 \leftarrow R\}, R = \min\{c, f\} \quad (3)$$

Tucker-2. Tucker decomposition is a generalized Higher Order SVD (HOSVD) for arbitrary-shaped tensors. We apply this method on 4-way weight tensors in convolutional layers, $W \in \mathbb{R}^{k \times k \times c \times f}$. We focus on Tucker-2 which only decomposes the tensor along c and f , i.e., output and input channels. For Tucker-2 decomposition, the compression hyperparameter for each convolution layer is a tuple of integer-valued approximation ranks (r_1, r_2) where $r_1 \in \{1, \dots, c\}$ and $r_2 \in \{1, \dots, f\}$. To increase search efficiency, we quantize the space of decomposition ranks to 8 bins per-way as follows:

$$r_1 \in \{1 \leftarrow \frac{c}{8}, \dots, 8 \leftarrow c\} \quad r_2 \in \{1 \leftarrow \frac{f}{8}, \dots, 8 \leftarrow f\} \quad (4)$$

4.2 Objective Evaluation

We develop a customized scoring mechanism to assess the quality of individuals in each iteration of GeneCAI algorithm. The multi-objective score of GeneCAI simultaneously reflects the compressed DNN's accuracy and computational complexity. The objective of DNN compression then translates to maximizing this score. We formalize DNN compression as a constrained optimization as follows:

$$\begin{aligned} \max_{\vec{x} \in \mathbb{R}^d} \Delta C(M, \vec{x}) \\ \text{subject to } A(\hat{M}_{\vec{x}}) > \text{acc}_{thr} \end{aligned} \quad (5)$$

where $\Delta C(M, \vec{x})$ represents the normalized difference in hardware cost, e.g., number of floating-point operations (FLOPs), between the uncompressed network, M , and the compressed model, $\hat{M}_{\vec{x}}$. The function $A(\cdot)$ denotes the inference accuracy and acc_{thr} is a task-enforced threshold on the post-compression accuracy. Having an accuracy constraint is crucial since the optimization algorithm will converge to a model size of zero otherwise. To solve the constrained optimization problem in Eq. (5), we formulate it as the following primal unconstrained optimization using penalty methods [2]:

$$\max_{\vec{x} \in \mathbb{R}^d} \Delta C(M, \vec{x}) - \log(\text{PEN}_{acc}(\vec{x})) \quad (6)$$

where the term $\log(\text{PEN}_{acc}(\vec{x}))$ is the exterior penalty function [4] that enforces a constraint on the accuracy, i.e., $A(\hat{M}_{\vec{x}}) > \text{acc}_{thr}$. The function $\text{PEN}_{acc}(\vec{x})$ measures the accuracy degradation as follows:

$$\text{PEN}_{acc}(\vec{x}) = \begin{cases} A(M) - A(\hat{M}_{\vec{x}}) & A(\hat{M}_{\vec{x}}) \geq \text{acc}_{thr} \\ A(M) - A(\hat{M}_{\vec{x}}) + e^{\text{acc}_{thr} - A(\hat{M}_{\vec{x}})} & A(\hat{M}_{\vec{x}}) < \text{acc}_{thr} \end{cases} \quad (7)$$

Figure 4 visualizes the accuracy penalty. To prevent undesirable drop of accuracy, we greatly diminish the score of individuals that cause lower accuracies than the set constraint, acc_{thr} . The $\log$ penalty term can thus be estimated as follows:

$$\log(\text{PEN}_{acc}(\vec{x})) = \begin{cases} \log(A(M) - A(\hat{M}_{\vec{x}})) & A(\hat{M}_{\vec{x}}) \geq \text{acc}_{thr} \\ \text{acc}_{thr} - A(\hat{M}_{\vec{x}}) & A(\hat{M}_{\vec{x}}) < \text{acc}_{thr} \end{cases} \quad (8)$$

For accuracy values satisfying the threshold, this term enforces the accuracy maximization objective and applying the logarithm smoothens the accuracy variations by damping sudden changes. For

Figure 4: GeneCAI accuracy penalty with $\text{acc}_{thr} = 80\%$ when original model accuracy is $A(M) = 93.5\%$.

accuracy values below acc_{thr} , a linear penalty is applied to prevent further accuracy loss. To prevent numerical instability, we define our score function as the exponential of the primal optimization of Eq. (6) which is formalized in Eq. (9):

$$f(\vec{x}) = \frac{e^{\Delta C(M, \vec{x})}}{\text{PEN}_{acc}(\vec{x})} \quad (9)$$

Note that maximizing the score function of Eq. (9) is equivalent to maximizing its logarithm value in Eq. (6). The developed score function encourages a reduction in hardware cost (in the numerator) while penalizing the decrease in model accuracy caused by compression (in the denominator). Our proposed scoring mechanism is globally applicable to various compression tasks and can be modified to reflect a variety of hardware costs, e.g., power, memory footprint, and runtime. To ensure efficiency, inference accuracies are measured on a small held-out portion of the training samples. GeneCAI scoring function successfully models the ultimate goal of high compression with minimal accuracy degradation.

4.3 Optimization with Genetic Operations

GeneCAI genetic optimization is a metaheuristic approach inspired by natural evolution and the notion of “survival of the fittest” and can be leveraged to explore large search-spaces. Our GA works on a *population of individuals*. The core idea is to encourage the creation of superior individuals and elimination of the inferior ones. To this end, an iterative process evolves the previous generation into a new, more competent population by performing a set of bio-inspired actions, i.e., *selection*, *cross-over*, and *mutation*.

Figure 5 illustrates the process of evolving a new population from the previous one. After the current population is evaluated (see Eq. (9)), individuals are assigned scores representing their quality, i.e., *fitness*. The selection step then performs a sampling (with replacement) based on individuals' fitness scores. Crossover and mutation aim to explore the proximity of the selected individuals by injecting small random patterns. To promote a diverse non-dominated front of solutions, we tune the hyperparameters associated with mutation based on a diversity measure. This, in turn, maintains the balance between *exploration* and *exploitation*.

Figure 5: Genetic operations for creating a new population.

GeneCAI genetic operations are devised to be compatible with the translation method. In particular, our designed operations satisfy the “closure” property: the output vector of mutation/cross-over translates to a valid compressed DNN. Note that selection and cross-over are unified across different compression techniques while the mutation operation needs to be customized per compression method. **Selection.** GeneCAI selection stage chooses high-quality individuals based on the ordering of Pareto dominance to generate the next population. GeneCAI adopts a fitness-proportionate selection mechanism. Let us denote the population at the beginning of t -th iteration by $\mathbb{P}^t = \{\vec{x}_n^t\}_{n=1}^N$ with fitness scores $\{f_n\}_{n=1}^N$ obtained from Eq. (9). To select superior samples, we construct the following probability distribution over the population using the fitness scores:

$$p_n = \frac{f_n - f_{\min}}{\sum_{n=1}^N (f_n - f_{\min})}, \quad (10)$$

where $f_{\min}$ is the minimum score. Subtraction of the minimum score ensures that the probability of selecting the weakest individual is zero and it is always eliminated. The superior individuals, $\mathbb{S}^t$, are selected by performing a non-uniform random sampling of the old population where the probability of selecting an individual is derived from Eq. (10). The total number of individuals in the population remains unchanged across iterations. Superior individuals can be selected multiple times for the new population (sampling with replacement) while ones with low scores are rarely chosen. The combination of GeneCAI evaluation and selection stages preserves high-quality individuals and eliminates weak ones.

Crossover. We design a crossover operation that creates new individuals, i.e., new compression configurations, $\mathbb{C}^t$ by inheriting and combining compression rates from a pair of parent individuals. The crossover is performed by randomly swapping corresponding elements of parent genotypes. Given the selected population $\mathbb{S}_t$, we sort the individuals in descending order based on their fitness score. To form the cross-over pairs, we pick the fittest individual available as the first parent. We then choose the individual which has the highest *distance* with the first parent as the second parent. Such a distance-based selection of pairs is motivated by increasing diversity among the newly generated offspring and promoting exploration. For two individuals $\vec{x}_1, \vec{x}_2 \in \mathbb{R}^d$, the distance is calculated as follows:

$$\text{dist}(\vec{x}_1, \vec{x}_2) = \frac{1}{d \times K} \sqrt{\sum_{i=1}^d (\vec{x}_1[i] - \vec{x}_2[i])^2} \quad (11)$$

where K is a constant equal to the maximum allowed value for the corresponding compression hyperparameter, e.g., $K = 64$ for SVD decomposition. We use two parameters to control the crossover operation: p_{cross} determines the probability of applying crossover between the parents, and p_{swap} is the per-element swapping probability. The proposed crossover allows superior individuals to exchange learned patterns and transfer knowledge across the population.

Mutation. Mutation randomly tweaks each individual in the crossoverd population $\mathbb{C}^t$ to create a new population $\mathbb{P}^{t+1}$ for the next iteration of the GA. We devise two variations of mutation for continuous and discrete-valued individuals.

- **Continuous-valued Individuals.** Each element of a continuous-valued individual is mutated by adding a random value drawn from a zero-mean Normal distribution $\mathcal{N}(0, 0.2)$. We then clip the values to ensure they remain in the allowed interval, i.e., $[0, 1]$.

- **Discrete-valued Individuals.** Discrete individuals are mutated by randomly incrementing/decrementing vector elements such that the value remains in the valid discrete range in Sec. 4.1.

Mutation allows exploration of the neighborhood of candidate points in the search-space. Similar to cross-over, we define two control parameters: P_{mutate} is the probability that the individual gets mutated and P_{tweak} determines the per-element tweaking probability. To maintain the balance between exploration and exploitation and avoid premature convergence, we tune the mutation parameters based on the diversity of individuals in $\mathbb{C}^t$. We use the dispersion of individuals in the population and define diversity as follows:

$$\text{div}(\mathbb{C}^t) = \frac{1}{N} \sum_{n=1}^N \left[\sum_{i=1}^d (\vec{x}_n^t[i] - \bar{\mu}_x^t[i])^2 \right] \quad (12)$$

where $\bar{\mu}_x^t \in \mathbb{R}^d$ is the mean of all individuals in $\mathbb{C}^t$. As can be seen, the diversity function is closely tied to the variance. The diversity can, therefore, be adjusted by controlling the per-element variance of the individuals in the crossoverd population $\text{Var}[X_i]$. Let us denote an arbitrary individual after crossover by $\vec{x} \in \mathbb{C}^t$, which transforms to $\vec{y} \in \mathbb{P}^{t+1}$ after mutation. The i^{th} element of $\vec{y}$ is thus sampled from a random variable Y_i with the following probability distribution:

$$P(Y_i = \vec{y}[i]) = \begin{cases} p_M & \vec{y}[i] = \vec{x}[i] + \eta \\ 1 - p_M & \vec{y}[i] = \vec{x}[i] \end{cases} \quad (13)$$

where $p_M = P_{\text{mutate}} \times P_{\text{tweak}}$ and η is the random perturbation applied during mutation. Note that $\vec{x}$ is also a random variable. The per-element variance of $\vec{y}$ is thus:

$$\text{Var}[Y_i] = \text{Var}[X_i] + p_M \sigma_\eta^2 \quad (14)$$

Here, σ_η^2 is the variance of the added perturbation. Summing up the vector values in Eq. (14) provides the new population diversity:

$$\text{div}(\mathbb{P}^{t+1}) = \text{div}(\mathbb{C}^t) + d \times p_M \sigma_\eta^2 \quad (15)$$

For a desired threshold on population diversity, we can, therefore, determine the mutation parameters using Eq. (15). Since tweaking multiple elements of the individual vector can result in drastic changes in the corresponding compressed DNN’s architecture and accuracy, we restrict P_{tweak} to a small value ($P_{\text{tweak}} = 0.05$) and merely adjust P_{mutate} for diversity control. In our experiments, we set the diversity threshold to be half the diversity for the randomly initialized population. Such adaptive tuning of mutation allows for a diversity-guided search and ensures a fast and stable convergence.

4.4 Directed Initialization

A naïve random initialization of samples in the first iteration can result in a slow and sub-optimal convergence. To address this, we utilize boundary characterization as a pre-processing step to enable a targeted sample initialization. The general behavior of inference accuracy for a compressed DNN $\hat{M}_{\vec{x}}$ is monotonic: as the compression rates increase, the accuracy drops. As such, for any task-enforced threshold on the accuracy, we can characterize the boundaries of $\vec{x}[i]$ on a per-layer basis. Boundary Characterization allows for a directed search that eliminates unnecessary exploration of outlier subspaces, i.e., regions that are unlikely to contain the optimization solution. Figure 6 visualizes the hyperparameter search-space and the outlier regions for pruning a 2-layer neural network. The horizontal plane corresponds to different configurations of per-layer pruning rates and the vertical axis represents their quality (score $f(\cdot)$).

Figure 6: Search space for pruning a 2-layer network.

The optimal solution to the search problem is a configuration with the highest score. The outlier regions in Figure 6 are therefore the flat sectors of the space. By filtering out the non-optimal outlier regions, GeneCAI sampling can find the solution within the region bounded by the initialization margins. We find a threshold vector $\vec{\theta}$ where each element constrains a single hyperparameter that corresponds to a certain layer in the compressed DNN. In the example of Figure 6, $\vec{\theta}$ has two elements, each presented by a dashed line ($\vec{\theta}[1]$ and $\vec{\theta}[2]$). Below we describe how the boundaries $\vec{\theta}[i]$ are obtained given an accuracy threshold acc_{thr} for each compression task.

Pruning. We obtain a threshold vector $\vec{\theta} \in \mathbb{R}^d$ with the i^{th} element $\vec{\theta}[i]$ specifying the maximum pruning rate of the i^{th} layer such that the inference accuracy of the compressed model $\hat{M}_{\vec{x}}$ does not violate the accuracy threshold acc_{thr} :

$$\vec{\theta}[i] = \max\{p\} \text{ s.t. } \vec{x}[j] = \begin{cases} p & j = i \\ 0 & j \neq i \end{cases} \& A(\hat{M}_{\vec{x}}) > acc_{thr} \quad (16)$$

Decomposition. For SVD and Tucker-2 decompositions, the threshold vector $\vec{\theta}$ represents per-layer minimum ranks satisfying acc_{thr} :

$$\vec{\theta}[i] = \min\{r\} \text{ s.t. } \vec{x}[j] = \begin{cases} r & j = i \\ K_j & j \neq i \end{cases} \& A(\hat{M}_{\vec{x}}) > acc_{thr} \quad (17)$$

where $K_j \in \{64, 8\}$ is the maximum allowed encoded rank for the j^{th} hyperparameter (see equations 3 & 4).

Note that real-world neural networks have many more layers and the pertinent search-space is of much higher dimensionality than the 2-dimensional plane shown Figure 6. For a d -dimensional search-space, the proposed boundary characterization scheme reduces the effective (continuous) search volume from 1 to $\prod_{i=1}^d \vec{\theta}[i]$ for pruning. For decomposition, the (discrete) search-space size is reduced from $\prod_{i=1}^d K_i$ to $\prod_{i=1}^d (K_i - \vec{\theta}[i])$. This approach significantly improves search convergence time and solution quality.

After performing boundary characterization, we randomly sample the initial population from the space enclosed by the threshold vector $\vec{\theta}$. For pruning, the i^{th} element $\vec{x}[i]$ in each sample vector is drawn from a Normal distribution $\mathcal{N}(\vec{\theta}[i]/2, \vec{\theta}[i]/2)$. For decomposition, the i^{th} element $\vec{x}[i]$ is randomly selected from $\{\vec{\theta}[i], \vec{\theta}[i]+1, \dots, K_i\}$.

5 EXPERIMENTS

We provide extensive evaluations on CIFAR-10 [20] and ImageNet [30] benchmarks. The evaluated network architectures include AlexNet, VGG, ResNet family, and MobileNets. All models are implemented using Pytorch library. The networks are trained from scratch following the parameter setup and training schedule adopted by the original papers [8, 12, 19, 35, 36]. For CIFAR-10, we use a VGG-variant as used in [17]. We compress the models using GeneCAI and fine-tune them for 60 and 20 epochs on CIFAR-10 and ImageNet, respectively. The genetic parameters are set to

$P_{tweak} = 0.05$, $P_{cross} = 0.8$, $P_{swap} = 0.2$ following [41]. P_{mutate} is adaptively determined based on the population diversity, as discussed in Sec. 4.3. All results are from a single run of the GA.

5.1 Quantitative Results on CIFAR-10

We apply GeneCAI to pre-trained CIFAR-10 architectures and compare our results with prior art in Table 1. Specifically, we conduct experiments with non-structured pruning (P_n), structured pruning (P_s), SVD and Tucker decomposition (D), and combination of multiple compression methods ($D + P_s$). We set the population size to 100 for ResNet-56 and ResNet-50, 200 for ResNet-110, and 50 for VGG. We randomly select 1,000 images from the training data to use as validation set for score computation. The threshold (acc_{thr}) is set to 90% for ResNet-X and 60% for VGG. For all networks, we let GeneCAI run for 50 iterations.

Non-structured Pruning (P_n in Table 1). We perform non-structured pruning on ResNet-50 and report the ratio of non-zero model parameters. acc_{thr} is set to 93% and we do not perform any fine-tuning on the compressed model. As shown, GeneCAI achieves $1.34\times$ lower parameters compared to state-of-the-art RL method [10]. Note that the lower FLOPs and comparable accuracy of [25] are due to training the model from scratch whereas GeneCAI and [10] do not include any fine-tuning.

Structured Pruning (P_s in Table 1). We implement structured pruning by adding masks after *ReLU* activation layers. We base our comparisons on FLOPs per inference relative to the uncompressed baseline. With similar or higher accuracy, GeneCAI compressed models achieve $1.35\times$ lower FLOPs than prior art (on average).

Decomposition and Pruning ($D + P_s$ in Table 1). To unveil the full optimization potential of our method, we allow GeneCAI to learn and combine multiple compression techniques, namely, pruning, SVD, and Tucker. The $D + P_s$ experiments are conducted by first decomposing the network and then applying structured pruning. As shown in Table 1, GeneCAI pushes the limits of compression by $2.45\times$ on average with less than 1% drop in accuracy compared to state-of-the-art works. We also report FLOPs reduction by only applying decomposition (SVD and Tucker), shown by D in Table 1.

5.2 Quantitative Results on ImageNet

Table 1 summarizes GeneCAI results for ImageNet dataset. Population size is 100 for ResNet-50, 50 for VGG-16, and 20 for AlexNet. We let GeneCAI run for 50 iterations with $acc_{thr} = 10\%$ for all models. This target accuracy is compensated by fine-tuning.

Non-structured Pruning (P_n in Table 1). We perform non-structured pruning on AlexNet and report the ratio of non-zero model parameters. GeneCAI achieves higher accuracy with 3% higher parameter size compared to a Bayesian Optimization approach [3].

Structured Pruning (P_s in Table 1). On ResNet-50, GeneCAI compresses the model to $1.51\times$ less FLOPs on average while achieving similar/higher test accuracy compared to best prior works. On VGG-16, GeneCAI outperforms all heuristic methods and gives competing results with [25] and the state-of-the-art RL method [10]. Note that [25] does not propose a hyperparameter optimization algorithm and merely focuses on the training of already-compressed DNNs. As such, their approach is orthogonal to GeneCAI and can be combined with our method to further improve final accuracy.

Decomposition and Pruning ($P_s + D$ in Table 1). Using a combination of decomposition and pruning, GeneCAI achieves $2.04\times$ lower

FLOPs than related work (on average) with slightly higher accuracy on ResNet-50. On VGG-16, GeneCAI pushes the state-of-the-art RL-based FLOPs reduction from $5\times$ to $7.04\times$ with higher accuracy.

Table 1: Comparison with contemporary compression methods based on the non-zero parameter ratio/FLOPs.

CIFAR-10			Top1	Cost	ImageNet			Top1	Top5	Cost
		(%)	(%)			(%)	(%)	(%)	(%)	(%)
ResNet-50	Baseline	93.7	100	AlexNet	Baseline	60.7	80.0	100		
	AMC [10]	93.5	40.0		CAC [3]	54.8	-	5.0		
	RP [25]	93.4	20.0		GeneCAI (P_n)	56.1	78.2	8.5		
	GeneCAI (P_n)	93.6	30.0		Baseline	75.1	93.0	100		
ResNet-56	Baseline	93.6	100	ResNet-50	SFP [9]	62.1	84.6	58.2		
	RP [25] (Res56B)	93.1	72.4		SSS [14]	71.8	90.8	57.0		
	CP [11]	91.8	50.0		RP [25] (Res50-2x)	75.0	-	50.0		
	AMC [10]	91.9	50.0		CP [11]	-	90.8	50.0		
	SFP [9]	93.3	47.4		GDP [24]	71.9	90.7	48.7		
	PF [40]	92.8	40.0		ThiNet [27]	71.0	90.0	44.0		
	GeneCAI (P_s)	93.2	44.0		RP [25] (Res50-30)	71.6	-	30.0		
	GeneCAI (D)	93.5	59.1		GeneCAI (P_s)	72.7	91.1	31.9		
ResNet-110	GeneCAI ($D + P_s$)	93.2	36.9	GeneCAI (D)	74.3	92.1	44.5			
	Baseline	94.0	100	GeneCAI ($D + P_s$)	72.1	90.9	23.7			
	RP [25] (Res110B)	93.6	61.4	VGG-16	Baseline	71.1	90.0	100		
	FP [21]	93.3	61.4		GDP [24]	67.5	87.9	24.5		
AMC [10]	93.8	59.2	RNP [23]		-	86.3	20.0			
GeneCAI (P_s)	93.6	41.2	SPP [38]		-	87.6	20.0			
GeneCAI (D)	93.9	55.3	AMC [10]		-	88.2	20.0			
GeneCAI ($D + P_s$)	92.6	21.9	RP [25] (VGG-6-5x)		71.0	-	20.0			
VGG	Baseline	93.6	100	GeneCAI (P_s)	68.1	88.1	20.0			
	RP [25] (VGG16A)	93.7	65.8	GeneCAI (D)	70.6	90.1	31.1			
	ThiNet [27]	93.4	39.4	GeneCAI ($D + P_s$)	68.4	88.5	14.2			
	NRE [17]	93.4	32.4							
	GeneCAI (P_s)	93.3	32.8							
	GeneCAI (D)	93.5	25.5							
	GeneCAI ($D + P_s$)	93.1	14.5							

5.3 Compressing Compact Networks

To further demonstrate the effectiveness of GeneCAI optimization, we apply compression to MobileNet architectures trained on ImageNet dataset. These networks are specifically designed for embedded applications with a strict efficiency constraint. As such, MobileNets inherently have very low complexity/redundancy which renders their compression quite challenging. We apply pruning to MobileNetV1 and MobileNetV2 with a population size of 50, and let GeneCAI run for 100 iterations. We compare the compression rate and accuracy achieved by GeneCAI with the FLOPs-accuracy Pareto curve of the original MobileNet architectures [12, 35], the state-of-the-art AutoML approach [10, 42], and compression-aware training methodologies, i.e., US-Nets [44, 45].

Table 2: GeneCAI pruning for MobileNetV1&V2 on ImageNet.

	Policy	Top1 (%)	Top5 (%)	FLOPs
MobileNetV1	Baseline (1x)	70.6	89.5	569 M
	MobileNetV1 (0.75x) [12]	68.4	88.2	325 M
	US-Nets [44]	69.5	-	325 M
	GeneCAI	70.5	89.3	323 M
	US-Nets [44]	68.8	-	287 M
	AMC [10]	70.5	89.1	285 M
	NetAdapt [42]	69.1	-	284 M
	GeneCAI	70.4	89.1	283 M
	US-Nets [44]	66.8	-	217 M
	GeneCAI	67.9	88.1	210 M
	MobileNetV1 (0.5x) [12]	63.7	-	149 M
	US-Nets [44]	63.5	-	136 M
MobileNetV2	GeneCAI	64.1	85.4	136 M
	Baseline (1x)	71.6	90.3	313 M
	MobileNetV2 (0.75x) [35]	69.8	88.3	220 M
	US-Nets [44]	70.0	-	222 M
	AMC [10]	-	89.3	220 M
	GeneCAI	70.1	89.5	220 M

Table 2 encloses the results of applying structured pruning to MobileNetV1 and MobileNetV2. We benchmark several target FLOPs and compare with prior work with similar computational complexities. On average, GeneCAI achieves 1.2% better accuracy than the MobileNetV1 Pareto curve. Compared to US-Nets, GeneCAI achieves an average of 1.0% higher accuracy. Under $\sim 50\%$ FLOPs, GeneCAI achieves 1.3% higher accuracy than NetAdapt. Compared to AMC, GeneCAI achieves lower FLOPs with comparable accuracy (-0.1%). On MobileNetV2, for a 30% FLOPs reduction, GeneCAI achieves lower FLOPs and higher accuracy than US-Nets and higher accuracy with same FLOPs compared to AMC and MobileNetV2.

Measured Speedup. We present measured hardware speedups of GeneCAI compressed MobileNets on an embedded GPU (NVIDIA Pascal) and an embedded CPU (ARM Cortex-A57) in Table 3. Measurements are averaged on 100 runs using a batch size of 32. GeneCAI successfully models hardware cost to achieve real speedups on par with theory.

Table 3: GeneCAI compressed MobileNets Speedup on embedded CPU and GPU for structured pruning on ImageNet.

Model	Theoretical Speedup	Real Speedup	
		Cortex-A57 (CPU)	Pascal (GPU)
MobileNetV1	1.7x	1.6x	1.3x
	2.0x	1.7x	1.4x
	2.7x	2.5x	1.7x
	4.0x	3.4x	1.9x
MobileNetV2	1.4x	1.4x	1.4x

5.4 Search Overhead and Scalability

The core computational load in GeneCAI algorithm corresponds to evaluating the genetic population which contains N individuals. For each individual $\vec{x}_i$, the evaluation phase comprises transforming $\vec{x}_i$ to its corresponding compressed DNN $\hat{M}_{\vec{x}_i}$, measuring inference accuracy on the validation data, and emulating the execution cost.

Since individuals in the population are independent, evaluation can be well-parallelized on multiple GPU devices to achieve faster search convergence. Table 4 summarizes the runtime of GeneCAI algorithm for several benchmarks and datasets. Runtimes are measured on a machine with an Intel Xeon E5 CPU and four NVIDIA Titan Xp GPUs. The results show high scalability: runtime drops almost linearly with the number of GPUs. The state-of-the-art RL algorithm reports ~ 1 hour to compress CIFAR-10 architectures [10]. For their most complex benchmark, i.e., ResNet-56, GeneCAI achieves a search time of only ~ 12 minutes on a single GPU and ~ 3 minutes on four GPUs.

Table 4: Search runtime of GeneCAI for pruning on various benchmarks. Here, N denotes the number of individuals in the population and N_{iters} is the number of search iterations.

Dataset	Arch.	N	N_{iters}	Search Time (minutes)			
				1 GPU	2 GPU	3 GPU	4 GPU
ImageNet	AlexNet	50	50	10	5	3	3
	VGG-16	50	50	112	57	38	28
	ResNet-50	100	50	145	73	49	36
	MobileNetV1	50	100	97	48	32	24
	MobileNetV2	50	100	116	57	38	25
	VGG	50	50	3	2	1	1
CIFAR-10	ResNet-50	100	50	35	19	13	11
	ResNet-56	100	50	12	7	5	3
	ResNet-110	200	50	55	30	22	16

5.5 Analysis and Discussion

To illustrate GeneCAI methodology, we consider VGG architecture trained on CIFAR-10 compressed with structured (filter) pruning for $acc_{thr} = 60\%$. The initial samples obtained from our directed initialization scheme are shown in Figure 7-a, where each column corresponds to an individual and each row represents a model layer. After applying GeneCAI for 50 iterations, the population pattern of Figure 7-b is learned. Upon convergence, individuals have similarly high scores. GeneCAI successfully learns expert-designed rules: first and last layers of the network (first and last rows in Figure 7-b) are given high densities to preserve accuracy.

Figure 7: (a) Initialized population for structured pruning. (b) Population upon convergence.

GeneCAI performs *whole-network* compression by encoding all layers’ hyperparameters in one genotype. As such, GeneCAI effectively learns which hyperparameter configuration least affects model accuracy and most reduces the overall FLOPs. To show this capability, we present the per-layer FLOPs for VGG-16 network trained on ImageNet in Figure 8. The bars show the per-layer FLOPs percentage in the original model; the curve shows the per-layer percentage of pruned FLOPs in the compressed DNN. Different from prior art [17], GeneCAI prunes the first convolutions more and relaxes pruning for later layers as they have a minor role in FLOPs.

5.6 Ablation Study

In this section, we study the effect of various GeneCAI parameters on algorithm convergence and final FLOPs/accuracy. For brevity, we only focus on the task of structured pruning for VGG on CIFAR-10. We show the trend lines as well as a fraction of individuals (black dots) across GeneCAI iterations.

Effect of Initialization. Figure 9-a shows the evolution of FLOPs ratio for two initialization policies, one with uniformly random samples and one with our proposed initialization scheme discussed in Sec. 4.4. As seen, naive initialization greatly harms the convergence rate and final FLOPs.

Effect of Population Size. Figure 9-b presents the effect of population size on GeneCAI convergence. A higher number of individuals

Figure 8: Per-layer FLOPs (bar charts) and percentage of pruned FLOPs (curve) for VGG16 trained on ImageNet. GeneCAI prunes more from layers with higher FLOPs.

results in a smoother convergence and lower final FLOPs. This effect saturates for a large enough population. In our experiments, we observed that the population size should be larger for searching vector spaces with higher dimensionality.

Effect of Mutation and Cross-over Parameters. We conduct two experiments, one with $(P_{mutate}, P_{cross}) = (0.8, 0.2)$ and the other with $(0.4, 0.1)$ and compare the convergence in Figure 9-c. Higher (P_{mutate}, P_{cross}) allows the GA to explore more possibilities, leading to faster convergence while smaller probabilities result in a more stable evolution. Thus, GeneCAI dynamically tunes the GA parameters to ensure balanced exploration and exploitation (see Sec. 4.3).

6 CONCLUSION

This paper introduces GeneCAI, a method to automate DNN compression using genetic algorithms. Our algorithm learns how the compression hyperparameters should be set across layers to achieve a better performance than models designed by human experts. The core idea behind GeneCAI is to translate compression hyperparameters into a vector of genes and explore the corresponding search-space using genetic operations. This approach allows GeneCAI to be generic and applicable to any combination of post-processing DNN compression methods. We showcase GeneCAI effectiveness on four DNN compression methods, namely, structured and non-structured pruning, SVD, and Tucker decomposition. Experiments show that GeneCAI achieves superior results compared to state-of-the-art heuristic and reinforcement-learning based algorithms on a variety of benchmarks. The proposed method is highly efficient as it does not require gradient computation and is entirely parallelizable.

7 ACKNOWLEDGMENT

This work was supported in part by Qualcomm Innovation Fellowship (QIF2019-US), NSF Grants CCF-1719133 and CCF-1513883, ARO (W911NF1910317), and SRC-Auto (2019-AU-2899).

Figure 9: Ablation studies for VGG on CIFAR-10: (a) Effect of initialization method. (b) Effect of population size. (c) Effect of mutation and cross-over probabilities. Here, individuals shown in grey are those with lower accuracy than $acc_{thr} = 60\%$ while individuals shown in black are those that meet the accuracy threshold. Red dots correspond to the highest-quality individual in that iteration.

REFERENCES

- [1] Jose M Alvarez and Mathieu Salzmann. 2017. Compression-aware training of deep networks. In *NIPS*. 856–867.
- [2] Thomas Bäck, David B Fogel, and Zbigniew Michalewicz. 1997. *Handbook of evolutionary computation*. CRC Press.
- [3] Changan Chen et al. 2018. Constraint-aware deep neural network compression. In *ECCV*. 400–415.
- [4] Edwin KP Chong and Stanislaw H Zak. 2013. *An introduction to optimization*. Vol. 76. John Wiley & Sons.
- [5] Mohammad Ghasemzadeh, Mohammad Samragh, and Farinaz Koushanfar. 2018. Rebnnet: Residual binarized neural network. In *2018 IEEE 26th Annual International Symposium on Field-Programmable Custom Computing Machines (FCCM)*. IEEE, 57–64.
- [6] Song Han et al. 2015. Learning both weights and connections for efficient neural network. In *NIPS*. 1135–1143.
- [7] 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).
- [8] Kaiming He et al. 2016. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 770–778.
- [9] 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).
- [10] 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.
- [11] Yihui He, Xiangyu Zhang, and Jian Sun. 2017. Channel pruning for accelerating very deep neural networks. In *International Conference on Computer Vision (ICCV)*, Vol. 2.
- [12] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861* (2017).
- [13] Yiming Hu, Siyang Sun, Jianquan Li, Xingang Wang, and Qingyi Gu. 2018. A novel channel pruning method for deep neural network compression. *arXiv preprint arXiv:1805.11394* (2018).
- [14] Zehao Huang and Naiyan Wang. 2018. Data-driven sparse structure selection for deep neural networks. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 304–320.
- [15] Mojan Javaheripi, Bitar Darvish Rouhani, and Farinaz Koushanfar. 2019. SWNet: Small-World Neural Networks and Rapid Convergence. *arXiv preprint arXiv:1904.04862* (2019).
- [16] Mojan Javaheripi, Mohammad Samragh, and Farinaz Koushanfar. 2019. Peeking Into the Black Box: A Tutorial on Automated Design Optimization and Parameter Search. *IEEE Solid-State Circuits Magazine* 11, 4 (2019), 23–28.
- [17] Chunhui Jiang, Guiying Li, Chao Qian, and Ke Tang. 2018. Efficient DNN Neuron Pruning by Minimizing Layer-wise Nonlinear Reconstruction Error. In *IJCAI*. 2–2.
- [18] Yong-Deok Kim, Eunhyeok Park, Sungjoo Yoo, Taelim Choi, Lu Yang, and Dongjun Shin. 2015. Compression of deep convolutional neural networks for fast and low power mobile applications. *arXiv preprint arXiv:1511.06530* (2015).
- [19] Alex Krizhevsky et al. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [20] Alex Krizhevsky and Geoffrey Hinton. 2009. *Learning multiple layers of features from tiny images*. Technical Report. Citeseer.
- [21] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. 2016. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710* (2016).
- [22] Jason Liang, Elliot Meyerson, Babak Hodjat, Dan Fink, Karl Mutch, and Risto Miikkulainen. 2019. Evolutionary neural automl for deep learning. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 401–409.
- [23] Ji Lin, Yongming Rao, Jiwen Lu, and Jie Zhou. 2017. Runtime neural pruning. In *Advances in Neural Information Processing Systems*. 2181–2191.
- [24] Shaohui Lin, Rongrong Ji, Yuchao Li, Yongjian Wu, Feiyue Huang, and Baochang Zhang. 2018. Accelerating Convolutional Networks via Global & Dynamic Filter Pruning. In *IJCAI*. 2425–2432.
- [25] Zhuang Liu et al. 2018. Rethinking the value of network pruning. *arXiv:1810.05270* (2018).
- [26] Zhichao Lu, Ian Whalen, Vishnu Boddeti, Yashesh Dhebar, Kalyanmoy Deb, Erik Goodman, and Wolfgang Banzhaf. 2019. NSGA-Net: neural architecture search using multi-objective genetic algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 419–427.
- [27] 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.
- [28] Pavlo Molchanov, Stephen Tyree, Tero Karras, Timo Aila, and Jan Kautz. 2016. Pruning convolutional neural networks for resource efficient transfer learning. *arXiv preprint arXiv:1611.06440* 3 (2016).
- [29] 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.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. 2015. Imagenet large scale visual recognition challenge. *International journal of computer vision* 115, 3 (2015), 211–252.
- [31] Tim Salimans, Jonathan Ho, Xi Chen, Szymon Sidor, and Ilya Sutskever. 2017. Evolution strategies as a scalable alternative to reinforcement learning. *arXiv preprint arXiv:1703.03864* (2017).
- [32] Mohammad Samragh, Mohammad Ghasemzadeh, and Farinaz Koushanfar. 2017. Customizing neural networks for efficient fpga implementation. In *Field-Programmable Custom Computing Machines (FCCM)*, 2017 IEEE 25th Annual International Symposium on. IEEE, 85–92.
- [33] Mohammad Samragh, Mojan Javaheripi, and Farinaz Koushanfar. 2019. AutoRank: Automated rank selection for effective neural network customization. In *Proceedings of the ML-for-Systems Workshop at the 46th International Symposium on Computer Architecture (ISCA'19)*.
- [34] Mohammad Samragh, Mojan Javaheripi, and Farinaz Koushanfar. 2019. CodeX: Bit-Flexible Encoding for Streaming-based FPGA Acceleration of DNNs. *arXiv preprint arXiv:1901.05582* (2019).
- [35] 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.
- [36] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [37] Bin Wang, Yanan Sun, Bing Xue, and Mengjie Zhang. 2019. Evolving deep neural networks by multi-objective particle swarm optimization for image classification. In *Proceedings of the Genetic and Evolutionary Computation Conference*. 490–498.
- [38] Huan Wang, Qiming Zhang, Yuehai Wang, and Haoji Hu. 2017. Structured probabilistic pruning for convolutional neural network acceleration. *arXiv preprint arXiv:1709.06994* (2017).
- [39] Kuan Wang et al. 2018. HAQ: hardware-aware automated quantization. *arXiv preprint arXiv:1811.08886* (2018).
- [40] Jiaxiang Wu et al. 2018. PocketFlow: An Automated Framework for Compressing and Accelerating Deep Neural Networks. (2018).
- [41] Lingxi Xie and Alan Yuille. 2017. Genetic cnn. *arXiv preprint arXiv:1703.01513* (2017).
- [42] Tien-Ju Yang, Andrew Howard, Bo Chen, Xiao Zhang, Alec Go, Mark Sandler, Vivienne Sze, and Hartwig Adam. 2018. Netadapt: Platform-aware neural network adaptation for mobile applications. In *Proceedings of the European Conference on Computer Vision (ECCV)*. 285–300.
- [43] Amir Yazdanbakhsh et al. 2018. ReLeQ: An Automatic Reinforcement Learning Approach for Deep Quantization of Neural Networks. *arXiv preprint arXiv:1811.01704* (2018).
- [44] Jiahui Yu and Thomas S. Huang. 2019. Universally Slimmable Networks and Improved Training Techniques. In *The IEEE International Conference on Computer Vision (ICCV)*.
- [45] Jiahui Yu, Linjie Yang, Ning Xu, Jianchao Yang, and Thomas S Huang. 2019. Slimmable neural networks. 7th International Conference on Learning Representations, ICLR 2019 ; Conference date: 06-05-2019 Through 09-05-2019.
- [46] Hao Zhou, Jose M Alvarez, and Fatih Porikli. 2016. Less is more: Towards compact cnns. In *European Conference on Computer Vision*. Springer, 662–677.
- [47] Shuchang Zhou, Yuxin Wu, Zekun Ni, Xinyu Zhou, He Wen, and Yuheng Zou. 2016. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160* (2016).