

Evolution in Groups: A Deeper Look at Synaptic Cluster Driven Evolution of Deep Neural Networks

Mohammad Javad Shafiee[†]
Department of
Systems Design Engineering,
University of Waterloo
Waterloo, Ontario, Canada
Email: mjshafiee@uwaterloo.ca

Elnaz Barshan
Department of
Systems Design Engineering,
University of Waterloo
Waterloo, Ontario, Canada
Email: ebarshan@uwaterloo.ca

Alexander Wong[†]
Department of
Systems Design Engineering,
University of Waterloo
Waterloo, Ontario, Canada
Email: a28wong@uwaterloo.ca

Abstract—A promising paradigm for achieving highly efficient deep neural networks is the idea of *evolutionary deep intelligence*, which mimics biological evolution processes to progressively synthesize more efficient networks. A crucial design factor in evolutionary deep intelligence is the genetic encoding scheme used to simulate heredity and determine the architectures of offspring networks. In this study, we take a deeper look at the notion of synaptic cluster-driven evolution of deep neural networks which guides the evolution process towards the formation of a highly sparse set of synaptic clusters in offspring networks. Utilizing a synaptic cluster-driven genetic encoding, the probabilistic encoding of synaptic traits considers not only individual synaptic properties but also inter-synaptic relationships within a deep neural network. This process results in highly sparse offspring networks which are particularly tailored for parallel computational devices such as GPUs and deep neural network accelerator chips. Comprehensive experimental results using four well-known deep neural network architectures (LeNet-5, AlexNet, ResNet-56, and DetectNet) on two different tasks (object categorization and object detection) demonstrate the efficiency of the proposed method. Cluster-driven genetic encoding scheme synthesizes networks that can achieve state-of-the-art performance with significantly smaller number of synapses than that of the original ancestor network. (~ 125 -fold decrease in synapses for MNIST). Furthermore, the improved cluster efficiency in the generated offspring networks (~ 9.71 -fold decrease in clusters for MNIST and a ~ 8.16 -fold decrease in clusters for KITTI) is particularly useful for accelerated performance on parallel computing hardware architectures such as those in GPUs and deep neural network accelerator chips.

Keywords—EvoNet; deep learning; evolution; deep neural network; embedded systems

I. INTRODUCTION

One of the key factors in revitalizing deep neural networks [1]–[4] and their tremendous success is the significant growth in computational power. The proliferation of massively parallel computing devices such as graphics processing units (GPUs) and distributed computing has revolutionized the training and inference of deep neural networks due to their highly parallelizable nature. This incredible rise in the computational power enables researchers to design and build increasingly larger and deeper neural networks to boost modeling accuracy. This on-going growth in architectural complexity, however, has become a bottleneck in the widespread adoption of such

deep neural networks in many operational scenarios. For example, in many applications such as self-driving cars, smart-phone applications, and surveillance cameras, the computational resources are limited to low-power embedded GPUs, CPUs and deep learning accelerator chips with strong constraints on the memory usage. Furthermore, there are many situations where the use of cloud computing is intractable due to transmission cost, bandwidth issues, as well as privacy concerns.

Considering the obstacles associated with the architectural complexity of deep neural networks, we have witnessed a growing attention towards learning highly efficient deep neural network architectures that are able to provide strong modeling power for operational scenarios where limited memory and computational resources are available. One of the first approaches in this area is the optimal brain damage method [5] in which synapses are pruned based on their strength. Gong *et al.* [6] proposed a network compression framework where vector quantization is used to shrink the storage requirements of deep neural networks. Han *et al.* [7] utilized Huffman coding in addition to pruning and vector quantization to further reduce the memory requirements. Hashing is another related trick employed by Chen *et al.* [8] for network compression. Low-rank approximation [9]–[11] and sparsity learning [12]–[14] are other strategies used to build smaller and more efficient deep neural networks.

Recently, Shafiee *et al.* [15] introduced an evolutionary synthesis framework to progressively learn more efficient deep neural networks along successive generations¹. The proposed *evolutionary deep intelligence* approach mimics biological evolution mechanisms such as random mutation, natural selection, and heredity within a probabilistic graphical modeling paradigm to successively synthesize more efficient offspring network architectures. A crucial design factor in evolutionary deep intelligence is the genetic encoding scheme used to simulate heredity and affects the architectural traits that are passed to the offspring networks in a significant way. Therefore, a more effective genetic encoding scheme can enable better transfer of genetic information from the ancestor network to its offspring networks to build a more efficient yet powerful future generation.

¹This framework is fundamentally different from the past attempts for utilizing evolutionary methods in training neural networks where genetic algorithm has been used to create neural networks [16]–[21] with high modeling capabilities in a direct but highly computationally expensive manner.

[†]DarwinAI, Waterloo, Ontario, Canada.




Fig. 1. Overview of the evolutionary deep intelligence framework. The architectural traits of ancestor networks are encoded via probabilistic ‘DNA’ sequences. A new offspring network is synthesized in each generation based on the probabilistic ‘DNA’ sequences (heredity), environmental factors (natural selection), and random mutation. The synthesized offspring deep neural network is then trained using data. The evolutionary process is repeated over successive generations to synthesize progressively more efficient deep neural networks.

The introduced genetic encoding scheme in [15] merely considers the individual synaptic properties in the sense that the probability of synthesizing each synapse within the network is independent of the rest of the synapses and thus it ignores the dependence between different synapses. However, there are neurobiological evidences that support the increasing probability of co-activation for *correlated* synapses which encode similar information and locate close to each other on the same dendrite—synaptic clustering [22]–[26]. Inspired by this observation, incorporating synaptic clustering in the genetic encoding scheme of evolutionary deep models is potentially a fruitful direction to investigate. Moreover, synthesizing the offspring networks based on synaptic clusters (instead of basing it purely on individual synapses) can increase the efficiency of the offspring deep neural networks running on parallel computing devices such as GPUs and deep neural network accelerator chips.

In this study, we take a deeper look at the notion of synaptic cluster-driven evolution of deep neural networks which guides the evolution process towards the formation of a highly sparse set of synaptic clusters in the offspring networks. This process results in highly sparse offspring networks which are particularly tailored for parallel computing devices such as GPUs and deep neural network accelerator chips. We introduce a multi-factor genetic encoding scheme in which the synaptic probability considers both the probability of synthesis for the cluster of synapses that includes a particular synapse and the probability of synthesis for that particular synapse within the cluster. This genetic encoding scheme effectively promotes the formation of synaptic clusters over successive generations during the evolution process while supporting the formation of highly efficient deep neural networks.

II. METHODOLOGY

Inspired by the neurobiological evidences, we propose a synaptic cluster-driven evolution framework for deep neural networks in which the ancestor network is guided towards the formation of a highly sparse set of synaptic *clusters* in the offspring networks. The key idea here is to design a genetic encoding scheme that considers the *inter-synaptic* relationships as well as the individual properties of each synapse. More formally, the synthesis probability distribution for each synapse is a product of the formation likelihood of the corresponding synaptic cluster and the synthesis probability of that particular synapse within that cluster.

A. Evolutionary Deep Intelligence

Prior to describing the notion of synaptic cluster-driven evolutionary synthesis and the correspondingly proposed genetic encoding scheme, let us first provide an overview of the evolutionary deep intelligence framework introduced by Shafiee *et al.* [15] for synthesizing progressively more efficient deep neural networks over successive generations within a probabilistic graphical modeling paradigm. In the evolutionary deep intelligence framework as shown in Fig. 1, the architectural traits of ancestor deep neural networks are encoded via probabilistic ‘DNA’ sequences. New offspring networks possessing diverse network architectures are then synthesized stochastically based on the ‘DNA’ from the ancestor networks and probabilistic computational environmental factor models, thus mimicking random mutation, heredity, and natural selection. These offspring networks are then trained, much like one would train a newborn, and have more efficient, more diverse network architectures while achieving powerful modeling capabilities. A crucial design factor in evolutionary deep intelligence is the genetic encoding scheme used to mimic heredity, which can have a significant impact on the architecture of the evolved offspring networks. In this study, we aim at exploring more effective genetic encoding schemes to guide the synthesis process so that the modeling capabilities of the ancestor network are faithfully captured by the more efficient offspring networks along successive generations.

B. Cluster-driven Genetic Encoding

Let the network architecture of a deep neural network be expressed by $\mathcal{H}(N, S)$, with N denoting the set of possible neurons and S is the set of possible synapses in the network. Each neuron $n_i \in N$ is connected via a set of synapses $\bar{s} \subset S$ to neuron $n_j \in N$ such that the synaptic connectivity $s_i \in S$ is associated with a $w_i \in W$ denoting its strength. Given the network architecture at the previous generation, i.e., \mathcal{H}_{g-1} , the architectural traits of a deep neural network in generation g , is encoded by the conditional probability $P(\mathcal{H}_g | \mathcal{H}_{g-1})$, which can be treated as the probabilistic ‘DNA’ sequence of a deep neural network.

We assume that synaptic connectivity characteristics in an ancestor network are desirable traits to be inherited by the descendant networks. Therefore, the genetic information of a deep neural network is encoded in synaptic probability $P(S_g | W_{g-1})$, where $w_{k,g-1} \in W_{g-1}$ is the synaptic strength of each synapse

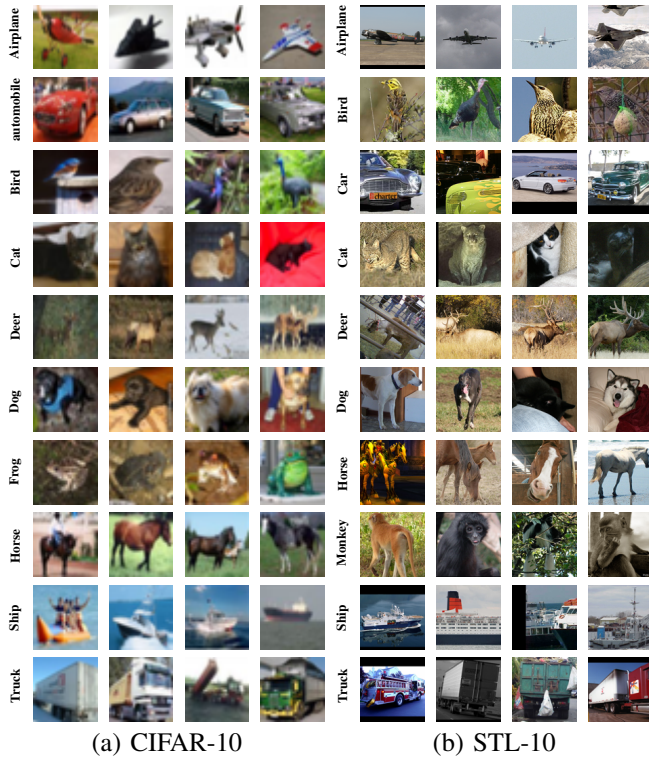


Fig. 2. Example images from two of the tested object categorization benchmark datasets: a) CIFAR-10, and b) STL-10.

$s_{k,g} \in S_g$. In the proposed genetic encoding scheme for synaptic cluster-driven evolution, we aim at incorporating the neurobiological phenomenon of synaptic clustering [22]–[26], where the probability of synaptic co-activation increases for correlated synapses which encode similar information and are close together on the same dendrite.

To promote the formation of synaptic clusters over successive generations of more efficient offspring networks, we introduce a multi-factor synaptic probability model defined as follows:

$$P(S_g|W_{g-1}) = \prod_{c \in C} \left[P(\bar{s}_{g,c}|W_{g-1}) \cdot \prod_{i \in c} P(s_{g,i}|w_{g-1,i}) \right] \quad (1)$$

where the first factor (first conditional probability) models the synthesis probability of a particular cluster of synapses, $\bar{s}_{g,c}$, while the second factor models the synthesis probability of a particular synapse, $s_{g,i}$, within synaptic cluster c . More specifically, the probability $P(\bar{s}_{g,c}|W_{g-1})$ represents the likelihood that a particular synaptic cluster, $\bar{s}_{g,c}$, be synthesized as a part of the network architecture in generation g given the synaptic strength in generation $g-1$. For example, in a deep convolutional neural network, the synaptic cluster c can be any subset of synapses such as a kernel or a set of kernels within the deep neural network. The probability $P(s_{g,i}|w_{g-1,i})$ represents the likelihood of existence of synapse i within the cluster c in generation g given its synaptic strength in generation $g-1$. As such, the proposed synaptic probability model not only promotes the persistence of strong synaptic connectivity in offspring deep neural networks over successive generations,

but also promotes the persistence of strong synaptic clusters in offspring deep neural networks over successive generations.

C. Cluster-driven Evolutionary Synthesis

In the seminal paper on evolutionary deep intelligence by Shafiee *et al.* [15], the synthesis probability $P(\mathcal{H}_g)$ is composed of the synaptic probability $P(S_g|W_{g-1})$, which mimic heredity, and environmental factor model $\mathcal{F}(\mathcal{E})$ which mimic natural selection by introducing quantitative environmental conditions that offspring networks must adapt to:

$$P(\mathcal{H}_g) = \mathcal{F}(\mathcal{E}) \cdot P(S_g|W_{g-1}) \quad (2)$$

In this study, (2) is reformulated in a more general way to enable the incorporation of different quantitative environmental factors over both the synthesis of synaptic clusters as well as each individual synapse, thus facilitating for synaptic cluster-driven evolution of deep neural networks:

$$P(\mathcal{H}_g) = \prod_{c \in C} \left[\mathcal{F}_c(\mathcal{E}) P(\bar{s}_{g,c}|W_{g-1}) \cdot \prod_{i \in c} \mathcal{F}_s(\mathcal{E}) P(s_{g,i}|w_{g-1,i}) \right] \quad (3)$$

where $\mathcal{F}_c(\cdot)$ and $\mathcal{F}_s(\cdot)$ represents environmental factors enforced at the cluster and synapse levels, respectively.

D. Realization of Cluster-driven Genetic Encoding

To demonstrate the benefits of the proposed cluster-driven genetic encoding scheme, a simple realization of this scheme is presented in this study. Here, since we wish to promote the persistence of strong synaptic clusters in offspring deep neural networks over successive generations, the synthesis probability of a particular cluster of synapses $\bar{s}_{g,c}$ is modeled as

$$P(\bar{s}_{g,c} = 1|W_{g-1}) = \exp \left(\frac{\sum_{i \in c} \lfloor \omega_{g-1,i} \rfloor}{Z} - 1 \right) \quad (4)$$

where $\lfloor \cdot \rfloor$ encodes the truncation of a synaptic weight and Z is a normalization factor to make (4) a probability distribution, $P(\bar{s}_{g,c}|W_{g-1}) \in [0, 1]$. The truncation of synaptic weights in the model reduces the influence of very weak synapses within a synaptic cluster on the genetic encoding process. The probability of a particular synapse, $s_{g,i}$, within synaptic cluster c , denoted by $P(s_{g,i} = 1|w_{g-1,i})$ can be expressed as:

$$P(s_{g,i} = 1|w_{g-1,i}) = \exp \left(\frac{\omega_{g-1,i}}{z} - 1 \right) \quad (5)$$

where z is a layer-wise normalization constant. By incorporating both of the aforementioned probabilities in the proposed scheme, the relationships amongst synapses as well as their individual synaptic strengths are taken into consideration in the genetic encoding process.

III. EXPERIMENTAL RESULTS

To study the efficacy of synaptic cluster-driven evolution for synthesizing highly efficient deep neural networks, evolutionary synthesis of deep neural networks across several generations was performed using the proposed genetic encoding scheme across four well-known deep neural network architectures for two different tasks:




Fig. 3. Example images from the MNIST character recognition dataset.

- Object categorization using three benchmark datasets: MNIST [27], STL-10 [28] and CIFAR10 [29].
- Object detection using one benchmark dataset: KITTI [30].

A. Datasets

The CIFAR-10 image dataset [29] comprises of 50,000 training and 10,000 test natural images with 10 different classes which equally distributed. The image size is 32×32 . The MNIST image dataset [27] comprises of 60,000 training images and 10,000 test images of 28×28 handwritten digits. The STL-10 image dataset [28] has 5,000 labeled training and 8,000 labeled test natural images categorized into 10 different classes. Although the original image size is 96×96 the images were shrunk in to 32×32 for consistency purposes. It is also should be noted that the 100,000 unlabeled images associated to this dataset were not used in training process. Fig. 2 demonstrates some examples of two natural images categorization datasets (i.e., CIFAR-10 and STL-10). The example images of MNIST dataset are shown in Fig. 3.

The Kitti dataset [30] is utilized for the purpose of object detection. We specifically apply the examined network structure for car detection in the scene. The utilized dataset comprises of 6373 training and 1108 images with 1242×375 in size as testing set. As seen in Fig. 4, the goal is to find a bounding box around every car in the scene.

B. Experimental Setup

For the MNIST and STL-10 experiments for object categorization, the LeNet-5 architecture [27] is selected as the network architecture of the original, first generation ancestor network. For the CIFAR10 experiment for object categorization, two different network architectures were explored as the first generation ancestor network. First, the AlexNet architecture [31] is utilized as one of the ancestor networks for CIFAR10, with the first layer modified to utilize $5 \times 5 \times 3$ kernels instead of $11 \times 11 \times 3$ kernels given the smaller image size in CIFAR10. Second, the ResNet-56 architecture [32] is utilized as another ancestor network for CIFAR10, which allows us to study the behaviour of the proposed scheme for two very different deep neural network architectures. For the KITTI experiment for object detection, the DetectNet architecture which is derived from GoogleNet [33] that is tuned for object

detection is selected as the network architecture of the original, first generation ancestor network.

The environmental factor model being imposed at different generations in this study is designed to form deep neural networks with progressively more efficient network architectures than its ancestor networks while maintaining modeling accuracy. More specifically, $\mathcal{F}_c(\cdot)$ and $\mathcal{F}_s(\cdot)$ is formulated in this study such that an offspring deep neural network should not have more than 80% of the total number of synapses in its direct ancestor network. Furthermore, in this study, each kernel in the deep neural network is considered as a synaptic cluster in the synapse probability model. In other words, the probability of the synthesis of a particular synaptic cluster (i.e, $P(\bar{s}_{g,c}|W_{g-1})$) is modeled as the truncated summation of the weights within a kernel.

C. Architectural Efficiency Over Generations

In this study, offspring deep neural networks were synthesized in successive generations until the accuracy of the offspring network exceeded 4%, so that we can better study the changes in architectural efficiency in the descendant networks over multiple generations. Table I and Table II show the architectural efficiency (defined in this study as the total number of synapses of the original, first-generation ancestor network divided by that of the current synthesized network) versus the modeling accuracy at several generations for three datasets and based on three different network architectures (LeNet-5, AlexNet, and ResNet-56). As observed in Table I, the descendant network at the 13th generation for MNIST based on the LeNet-5 architecture was a staggering ~ 125 -fold more efficient than the original, first-generation ancestor network without exhibiting a significant drop in the test accuracy ($\sim 1.7\%$ drop). This trend was consistent with that observed with the STL-10 results based on the LeNet-5 architecture, where the descendant network at the 10th generation was ~ 56 -fold more efficient than the original, first-generation ancestor network without a significant drop in test accuracy ($\sim 1.2\%$ drop). It also worth noting that since the training dataset of the STL-10 dataset is relatively small, the descendant networks at generations 2 to 8 actually achieved higher test accuracies when compared to the original, first-generation ancestor network, which illustrates the generalizability of the descendant networks compared to the original ancestor network as the descendant networks had fewer parameters to train.



Fig. 4. Example images from the KITTI object detection dataset. Ground truth object bounding boxes are overlaid onto the images for context.

TABLE I. ARCHITECTURAL EFFICIENCY VS. TEST ACCURACY FOR DIFFERENT GENERATIONS OF SYNTHESIZED NETWORKS FOR MNIST AND STL-10 USING LeNet-5. “GEN.,” “A-E” AND “ACC.” DENOTE GENERATION, ARCHITECTURAL EFFICIENCY, AND ACCURACY, RESPECTIVELY

MNIST - LeNet-5			STL-10 - LeNet-5		
Gen.	A-E	ACC.	Gen.	A-E	ACC.
1	1.00X	99.47	1	1.00X	57.74
5	5.20X	99.41	3	2.37X	59.33
7	12.09X	99.28	7	14.99X	60.51
11	62.74X	98.49	9	38.22X	57.44
13	125.09X	97.75	10	56.27X	56.58

TABLE II. ARCHITECTURAL EFFICIENCY VS. TEST ACCURACY FOR DIFFERENT GENERATIONS OF SYNTHESIZED NETWORKS FOR CIFAR-10 USING ALEXNET AND RESNET-56. “GEN.,” “A-E” AND “ACC.” DENOTE GENERATION, ARCHITECTURAL EFFICIENCY, AND ACCURACY, RESPECTIVELY

CIFAR10 - AlexNet			CIFAR10 - ResNet-56		
Gen.	A-E	ACC.	Gen.	A-E	ACC.
1	1.00X	86.69	1	1.00X	93.00
2	1.64X	88.14	3	1.66X	91.95
3	2.82X	87.66	5	3.04X	90.27
5	8.39X	85.88	6	3.69X	89.76
6	14.39X	84.59	7	4.02X	89.79

TABLE III. ARCHITECTURAL EFFICIENCY VS. PRECISION AND RECALL FOR DIFFERENT GENERATIONS OF SYNTHESIZED NETWORKS FOR KITTI USING DETECTNET. “GEN.,” “A-E” DENOTE GENERATION AND ARCHITECTURAL EFFICIENCY, RESPECTIVELY

KITTI - DetectNet			
Gen.	A-E	Precision	Recall
1	1.00X	78.03	62.03
7	2.48X	81.93	69.31
10	5.24X	81.36	67.57
11	9.97X	78.84	63.73

For the case of CIFAR10 based on the AlexNet architecture, the descendant network at the 6th generation network was ~ 14.4 -fold more efficient than the original ancestor network with $\sim 2\%$ drop in test accuracy. For the case of CIFAR10 based on the ResNet-56 architecture, the descendant network at the 7th generation network was ~ 4 -fold more efficient than the original ancestor network with $\sim 3\%$ drop in test accuracy.

Finally, Table III shows the architectural efficiency versus modeling accuracy at several generations for the KITTI dataset and based on the DetectNet architecture. For this case, the descendant network at the 11th generation was ~ 10 -fold more efficient compared to the original, first-generation ancestor network while achieving a $\sim 0.8\%$ increase in precision and a $\sim 1.7\%$ increase in recall. These results not only demonstrate the efficacy of the proposed scheme, but also illustrate the applicability of the proposed scheme for a variety of different network architectures.

D. Cluster Efficiency

Table IV and V shows the cluster efficiency of the synthesized deep neural networks in the last generations, where cluster efficiency is defined in this study as the total number of kernels in a layer of the original, first-generation ancestor network divided by that of the current synthesized network. It can be observed that for MNIST, the cluster efficiency of

TABLE IV. CLUSTER EFFICIENCY VS. TEST ACCURACY FOR THE FIRST AND LAST REPORTED GENERATIONS OF SYNTHESIZED NETWORKS FOR MNIST AND STL-10 USING LeNet-5. "GEN.", "C-E" AND "ACC." DENOTE GENERATION, CLUSTER EFFICIENCY, AND ACCURACY, RESPECTIVELY

MNIST - LeNet-5			STL-10 - LeNet-5		
Gen.	C-E	ACC.	Gen.	C-E	ACC.
1	1.00X	99.47	1	1.00X	57.74
13	9.71X	97.75	10	5.96X	56.58

TABLE V. CLUSTER EFFICIENCY VS. TEST ACCURACY FOR THE FIRST AND LAST REPORTED GENERATIONS OF SYNTHESIZED NETWORKS FOR CIFAR-10 USING ALEXNET AND RESNET-56. "GEN.", "C-E" AND "ACC." DENOTE GENERATION, CLUSTER EFFICIENCY, AND ACCURACY, RESPECTIVELY

CIFAR10 - AlexNet			CIFAR10 - ResNet-56		
Gen.	C-E	ACC.	Gen.	C-E	ACC.
1	1.00X	86.69	1	1.00X	93.00
6	2.82X	84.59	6	2.25X	89.79

TABLE VI. CLUSTER EFFICIENCY VS. PRECISION AND RECALL FOR THE FIRST AND LAST REPORTED GENERATIONS OF SYNTHESIZED NETWORKS FOR KITTI USING DETECTNET. "GEN." AND "C-E" DENOTE GENERATION AND CLUSTER EFFICIENCY, RESPECTIVELY

KITTI - DetectNet			
Gen.	C-E	Precision	Recall
1	1.00X	78.03	62.03
11	8.16X	78.84	63.73

the last-generation descendant network is $\sim 9.7X$, which may result in a near 9.7-fold potential speed-up in running time on embedded GPUs and deep neural network accelerator chips by reducing the number of arithmetic operations by ~ 9.7 -fold compared to the first-generation ancestor network, though computational overhead in other layers such as ReLU may lead to a reduction in actual speed-up. The potential speed-up from the last-generation descendant network for STL-10 is lower compared to MNIST dataset, with the reported cluster efficiency in last-generation descendant network $\sim 6X$.

The cluster efficiency for the last generation descendant networks for CIFAR10 using AlexNet and ResNet-56 are $\sim 2.8X$ and $\sim 2.25X$, respectively. Finally, Table VI demonstrates the cluster efficiency of the synthesized deep neural networks in the last generations for the KITTI dataset using DetectNet. As seen, the cluster efficiency of the last generation descendent network is $\sim 8.16X$, while achieving a $\sim 0.8\%$ increase in precision and a $\sim 1.7\%$ increase in recall compared to the original, first-generation ancestor network. These results demonstrate that not only can the proposed genetic encoding scheme promotes the synthesis of deep neural networks that are highly efficient yet maintains modeling accuracy, but also promotes the formation of highly sparse synaptic clusters that make them highly tailored for devices designed for highly parallel computations such as GPUs and deep neural network accelerator chips.

IV. CONCLUSION

In this study, we took a deeper look at a synaptic cluster-driven strategy for the evolution of deep neural networks, where synaptic probability within a deep neural network is driven towards the formation of highly sparse synaptic clusters. Experimental results for the tasks of object categorization and object detection demonstrated that the synthesized 'evolved' offspring networks using this evolution strategy via a synaptic cluster-driven genetic encoding scheme can achieve state-of-the-art performance while having network architectures

that are not only significantly more efficient compared to the original ancestor network, but also highly tailored for operational machine learning applications and scenarios using devices designed for highly parallel computations such as GPUs and deep neural network accelerator chips. Future work involves investigating alternative realizations of this new genetic encoding scheme beyond the simple realization presented in this study, and study the network architecture evolutions of deep neural networks synthesized by these realizations over successive generations to better understand their effectiveness and efficiency.

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning." in *Nature*, pp. 436-444, 2015.
- [2] A. Graves, A. Abdel-rahman, and G. Hinton, "Speech recognition with deep recurrent neural networks." in *IEEE international conference on acoustics, speech and signal processing*, pp. 6645-6649, 2013.
- [3] Y. Bengio, "Learning deep architectures for AI." in *Foundations and trends® in Machine Learning*, vol. 2, no. 1, pp. 1-127, 2009.
- [4] J. Tompson, A. Jain, Y. LeCun, and C. Bregler, "Joint training of a convolutional network and a graphical model for human pose estimation." in *Advances in Neural Information Processing Systems (NIPS)*, pp. 1799-1807, 2014.
- [5] Y. LeCun, J. S. Denker, S. A. Solla, R. E. Howard, and L. D. Jackel, "Optimal brain damage." in *Advances in Neural Information Processing Systems (NIPS)*, 1989.
- [6] Y. Gong, L. Liu, M. Yang, and L. Bourdev, "Compressing deep convolutional networks using vector quantization," *CoRR, abs/1412.6115*, 2014.
- [7] S. Han, J. Pool, J. Tran, and W. Dally, "Learning both weights and connections for efficient neural network," in *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [8] W. Chen, J. T. Wilson, S. Tyree, K. Q. Weinberger, and Y. Chen, "Compressing neural networks with the hashing trick," *CoRR, abs/1504.04788*, 2015.
- [9] Y. Ioannou, D. Robertson, J. Shotton, R. Cipolla, and A. Criminisi, "Training cnns with low-rank filters for efficient image classification," *arXiv preprint arXiv:1511.06744*, 2015.
- [10] M. Jaderberg, A. Vedaldi, and A. Zisserman, "Speeding up convolutional neural networks with low rank expansions," *arXiv preprint arXiv:1405.3866*, 2014.
- [11] E. Denton, W. Zaremba, J. Bruna, Y. LeCun and R. Fergus, "Exploiting Linear Structure within Convolutional Networks for Efficient Evaluation," in *Advances in Neural Information Processing Systems*, 2014.
- [12] J. Feng and T. Darrell, "Learning the structure of deep convolutional networks," in *Proceedings of the IEEE International Conference on Computer Vision*, pp. 2749-2757, 2015.

- [13] B. Liu, M. Wang, H. Foroosh, M. Tappen, and M. Pensky, "Sparse convolutional neural networks," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2015, pp. 806–814.
- [14] W. Wen, C. Wu, Y. Wang, Y. Chen, and H. Li, "Learning structured sparsity in deep neural networks," *arXiv preprint arXiv:1608.03665*, 2016.
- [15] M. J. Shafiee, A. Mishra, and A. Wong, "Deep learning with Darwin: Evolutionary synthesis of deep neural networks," *arXiv preprint arXiv:1606.04393*, 2016.
- [16] D. White and P. Ligomenides, "Gannet: A genetic algorithm for optimizing topology and weights in neural network design," in *International Workshop on Artificial Neural Networks*. Springer, pp. 322–327, 1993.
- [17] K. O. Stanley and R. Miikkulainen, "Evolving neural networks through augmenting topologies," *Evolutionary computation*, vol. 10, no. 2, pp. 99–127, 2002.
- [18] P. J. Angeline, G. M. Saunders, and J. B. Pollack, "An Evolutionary Algorithm that Constructs Recurrent Neural Networks," *IEEE Transactions on Neural Networks*, 1994.
- [19] K. O. Stanley, B. D. Bryant, and R. Miikkulainen, "Real-time neuroevolution in the NERO video game," *IEEE Transactions on Evolutionary Computation*, 2005.
- [20] J. Stanley, and K. O. Stanley, "Generating Large-Scale Neural Networks Through Discovering Geometric Regularities," *Genetic and Evolutionary Computation Conference*, 2007.
- [21] S. S. Tirumala, S. Ali, and C. P. Ramesh, "Evolving deep neural networks: A new prospect," *Genetic and Evolutionary Computation Conference*, 2016.
- [22] O. Welzel, C. H. Tischbirek, J. Jung, E. M. Kohler, A. Svetlitchny, A. W. Henkel, J. Kornhuber, and T. W. Groemer, "Synapse clusters are preferentially formed by synapses with large recycling pool sizes," *PloS one*, vol. 5, no. 10, p. e13514, 2010.
- [23] G. Kastellakis, D. J. Cai, S. C. Mednick, A. J. Silva, and P. Poirazi, "Synaptic clustering within dendrites: An emerging theory of memory formation," *Progress in neurobiology*, vol. 126, pp. 19–35, 2015.
- [24] M. E. Larkum and T. Nevian, "Synaptic clustering by dendritic signalling mechanisms," *Current opinion in neurobiology*, vol. 18, no. 3, pp. 321–331, 2008.
- [25] N. Takahashi, K. Kitamura, N. Matsuo, M. Mayford, M. Kano, N. Matsuki, and Y. Ikegaya, "Locally synchronized synaptic inputs," *Science*, vol. 335, no. 6066, pp. 353–356, 2012.
- [26] J. Winnubst and C. Lohmann, "Synaptic clustering during development and learning: the why, when, and how," *Frontiers in molecular neuroscience*, vol. 5, 2015.
- [27] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [28] A. Coates, A. Y. Ng, and H. Lee, "An analysis of single-layer networks in unsupervised feature learning," in *International conference on artificial intelligence and statistics*, 2011, pp. 215–223.
- [29] A. Krizhevsky and G. Hinton, "Learning multiple layers of features from tiny images," 2009.
- [30] A. Geiger, P. Lenz, and R. Urtasun, "Are we ready for Autonomous Driving? The KITTI Vision Benchmark Suite," in *Conference on Computer Vision and Pattern Recognition*, 2012.
- [31] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems (NIPS)*, 2012.
- [32] K. He, X. Zhang, S. Ren, and J. Sun, "Deep Residual Learning for Image Recognition," in *Conference on Computer Vision and Pattern Recognition*, 2016.
- [33] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going Deeper with Convolutions," in *Conference on Computer Vision and Pattern Recognition*, 2015.