
Mitigating Architectural Mismatch During the Evolutionary Synthesis of Deep Neural Networks

Audrey G. Chung*, Paul Fieguth, Alexander Wong

Vision and Image Processing Research Group, University of Waterloo, Waterloo, ON, Canada
Waterloo Artificial Intelligence Institute, University of Waterloo, Waterloo, ON, Canada

* agchung@uwaterloo.ca

Abstract

Evolutionary deep intelligence has recently shown great promise for producing small, powerful deep neural network models via the organic synthesis of increasingly efficient architectures over successive generations. Existing evolutionary synthesis processes, however, have allowed the mating of parent networks independent of architectural alignment, resulting in a mismatch of network structures. We present a preliminary study into the effects of architectural alignment during evolutionary synthesis using a gene tagging system. Surprisingly, the network architectures synthesized using the gene tagging approach resulted in slower decreases in performance accuracy and storage size; however, the resultant networks were comparable in size and performance accuracy to the non-gene tagging networks. Furthermore, we speculate that there is a noticeable decrease in network variability for networks synthesized with gene tagging, indicating that enforcing a like-with-like mating policy potentially restricts the exploration of the search space of possible network architectures.

1 Introduction

Deep neural networks (DNNs) [3, 1, 2, 4] have recently garnered widespread interest due to their demonstrated ability to improve state-of-the-art performance in many challenging areas of research. This boost in performance, however, has largely been attributed to increasingly large and complex network architectures, resulting in storage and memory requirements beyond the resources available in practical scenarios. As such, methods for producing highly efficient DNNs have been developed to reduce the memory and computational needs while maintaining performance accuracy.

Inspired by nature, Shafiee *et al.* [5] introduced *evolutionary deep intelligence* as a biologically-motivated alternative to compressing existing DNNs directly by allowing networks to organically synthesize new and increasingly efficient network architectures over successive generations. While existing evolutionary deep intelligence methods have leveraged asexual evolutionary synthesis [5–9] and m -parent sexual evolutionary synthesis [10–12], the evolutionary synthesis process combines architectural structures sequentially regardless of their relative positions in each parent network. As a result, current evolutionary deep intelligence methods have employed evolutionary synthesis processes that combine parent networks independent of architectural alignment, resulting in a mismatch of network structures.

In this paper, we present a preliminary study into the effects of architectural alignment during evolutionary synthesis via the introduction of a gene tagging system. Gene tagging is explored within the context of m -parent sexual evolutionary synthesis and evaluated over a range of environmental resource models. The gene tagging system allows for the proper alignment of architectural structures that originated from the same location in the ancestor network, and enforces a like-with-like mating policy during evolutionary synthesis.

2 Methods

In this study, we modify the existing m -parent evolutionary synthesis approach [11] to investigate the efficacy of structural alignment and its effect on the synthesis of new network architectures across a range of simulated environmental resource models.

2.1 m -Parent Evolutionary Synthesis

Let the network architecture be formulated as $\mathcal{H}(N, S)$, where N is the set of possible neurons and S denotes the set of possible synapses in the network. Each neuron $n_j \in N$ is connected to neuron $n_k \in N$ via a set of synapses $\bar{s} \subset S$, such that the synaptic connectivity $s_j \in S$ has an associated $w_j \in W$ to denote the connection’s strength. In the seminal evolutionary deep intelligence paper [5], the synthesis probability $P(\mathcal{H}_g | \mathcal{H}_{g-1}, \mathcal{R}_g)$ of a new network at generation g is approximated by the synaptic probability $P(S_g | W_{g-1}, R_g)$ to emulate heredity through the generations of networks. $P(\mathcal{H}_g | \mathcal{H}_{g-1}, \mathcal{R}_g)$ is also conditional on an environmental factor model \mathcal{R}_g to imitate natural selection via simulated environmental resources.

Generalizing to m -parent evolutionary synthesis [11], a newly synthesized network $\mathcal{H}_{g(i)}$ can be dependent on a subset of all previously synthesized networks \mathcal{H}_{G_i} , where the set of network indices G_i corresponds to the set of previous networks on which $\mathcal{H}_{g(i)}$ is dependent, and $g(i)$ represents the generation number corresponding to the i^{th} network. Note that in the general case, the number of networks in subset \mathcal{H}_{G_i} and the range of generational dependency $g(G_i)$ is only constrained by the number and generational range of already synthesized networks.

The synthesis probability combining the probabilities of m parent networks \mathcal{H}_{G_i} is represented by some cluster-level mating function $\mathcal{M}_c(\cdot)$ and some synapse-level mating function $\mathcal{M}_s(\cdot)$:

$$P(\mathcal{H}_{g(i)} | \mathcal{H}_{G_i}, \mathcal{R}_{g(i)}) = \prod_{C \in \mathcal{C}} \left[P(s_{g(i), C} | \mathcal{M}_c(W_{\mathcal{H}_{G_i}}), \mathcal{R}_{g(i)}^c) \cdot \prod_{j \in C} P(s_{g(i), j} | \mathcal{M}_s(w_{\mathcal{H}_{G_i}, j}), \mathcal{R}_{g(i)}^s) \right]. \quad (1)$$

These mating functions are formulated using an intersection-based mating policy [11], i.e.,:

$$\mathcal{M}_c(W_{\mathcal{H}_{G_i}}) = \prod_{k=1}^m \alpha_{c,k} W_{\mathcal{H}_k} \quad \mathcal{M}_s(w_{\mathcal{H}_{G_i}, j}) = \prod_{k=1}^m \alpha_{s,k} w_{\mathcal{H}_k, j} \quad (2)$$

where $W_{\mathcal{H}_k}$ represents the cluster’s synaptic strength for the k^{th} parent network $\mathcal{H}_k \in \mathcal{H}_{G_i}$. Similarly, $w_{\mathcal{H}_k, j}$ represents the synaptic strength of a synapse j within a given cluster for the k^{th} parent network $\mathcal{H}_k \in \mathcal{H}_{G_i}$.

2.2 Mitigating Architecture Mismatch via Gene Tagging

To encourage like-with-like mating during evolutionary synthesis, this study introduces a gene tagging system to enforce structural alignment, i.e., only mating architectural clusters originating from the same location in the ancestor network. As such, the cluster-level and synapse-level mating functions proposed in [11] are reformulated as follows:

$$\mathcal{M}_c(\bar{W}_{\mathcal{H}_{G_i}}) = \prod_{k \in \mathcal{K}_c} \alpha_{c,k} \bar{W}_{\mathcal{H}_k} \quad \mathcal{M}_s(\bar{w}_{\mathcal{H}_{G_i}, j}) = \prod_{k \in \mathcal{K}_c} \alpha_{s,k} \bar{w}_{\mathcal{H}_k, j} \quad (3)$$

where \mathcal{K}_c is the subset of parent networks with existing architectural clusters corresponding to a single gene tagged cluster $c \in C$, C is the set of clusters that exists in $\mathcal{H}_{g(i)}$, and \bar{W} and \bar{w} are the gene tagged and structurally-aligned synaptic strengths.

This study also introduces a parameter-based lower bound on the subset of parent networks with existing architectural clusters \mathcal{K}_c . Prior studies [10–12] employed an intersection-based mating policy where a cluster was synthesized only if the cluster existed within all m parent networks (i.e., $\mathcal{K}_c = m$). To allow for more flexibility within the m -parent evolutionary synthesis, we propose a percent of population parameter χ to control the proportion of parent network architectures that must contain any given cluster during evolutionary synthesis. As such, only clusters that exist within the specified proportion of parent networks are synthesized (i.e., $m\chi \leq \mathcal{K}_c \leq m$), and the intersection-based policy can be achieved using $\chi = 1$.

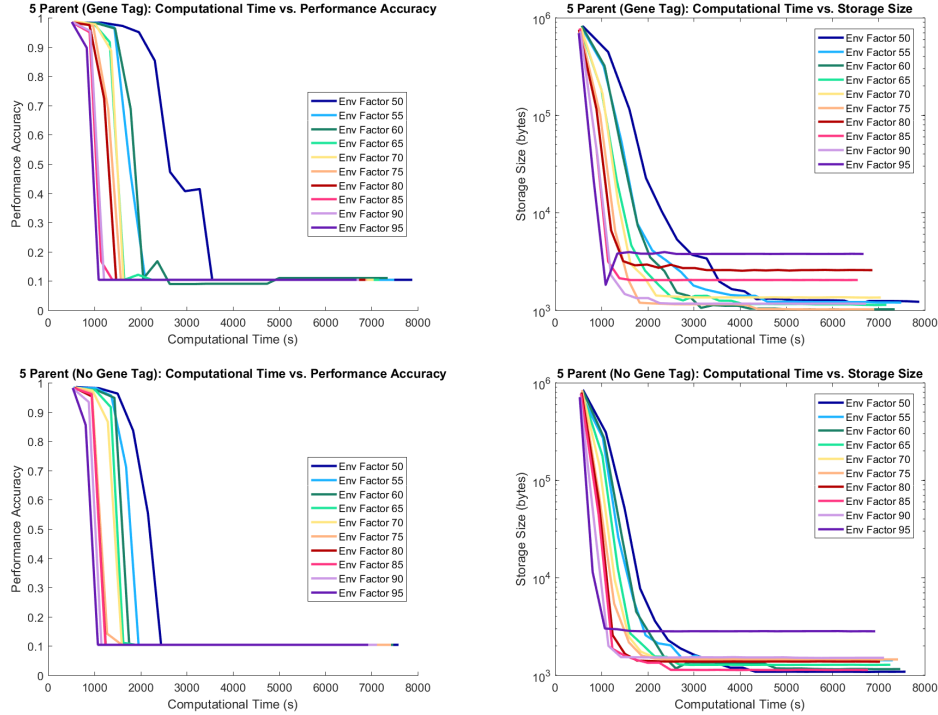


Figure 1: Performance accuracy (left) and storage size (right) for 5-parent evolutionary synthesis with gene tagging (top row) and no gene tagging (bottom row) using various environmental factor models. Plots best viewed in colour.

3 Results

3.1 Experimental Setup

The effect of structural alignment during the m -parent evolutionary synthesis was evaluated using 10% of the MNIST [13] hand-written digits dataset to increase the training speed of the synthesized network architectures and increase the inherently low intra-class variation within the MNIST dataset. For this preliminary study, 5-parent evolutionary synthesis and a χ of 60% was used, and \mathcal{H}_{G_i} was restricted to networks in the immediately preceding generation, i.e., for a newly synthesized network $\mathcal{H}_{g(i)}$, the $m = 5$ parent networks in \mathcal{H}_{G_i} are from generation $g(i) - 1$.

The first generation ancestor network was trained using the LeNet-5 architecture [14], and the synthesized offspring networks were assessed using performance accuracy on the MNIST dataset and storage size (representative of the architectural efficiency of a network) of the networks with respect to the computational time required. Similar to [12], each filter was considered as a synaptic cluster in the multi-factor synapse probability model, and the cluster-level environmental factor model $\mathcal{R}_{g(i)}^c$ and the synapse-level environmental factor model $\mathcal{R}_{g(i)}^s$ were varied together from 50% to 95% at 5% increments, i.e.,:

$$\mathcal{R}_{g(i)}^c, \mathcal{R}_{g(i)}^s = \{50, 55, 60, 65, 70, 75, 80, 85, 90, 95\}\%, \quad (4)$$

to simulate the varying availability of environmental resources during m -parent evolutionary synthesis. The synthesized networks were assessed using performance accuracy on the MNIST dataset and storage size (representative of the architectural efficiency of a network) of the networks with respect to the computational time required.

3.2 Experimental Results

Figure 1 shows the performance accuracy and storage size for 5-parent evolutionary synthesis with gene tagging (top row) and no gene tagging (bottom row) given various cluster-level and synapse-level




Figure 2: Performance accuracy as a function of storage size for 5-parent sexual evolutionary synthesis using various environmental factor models. Networks synthesized using gene tagging (diamond) show minimal to no noticeable difference relative to networks synthesized without gene tagging (round) in terms of maintaining performance accuracy while decreasing storage size. Plots best viewed in colour.

environmental factor models. The incorporation of gene tagging resulted in a more gradual decrease in both performance accuracy and storage size relative to computational time, while the omission of gene tagging produced network architectures with more rapidly decreasing performance accuracy and storage size. Note that the bottom plateau in performance accuracy is at 10% akin to random guessing, as the MNIST dataset consists of 10 classes of handwritten digits. While there is no inherent limit in network reduction when using the evolutionary deep intelligence approach, Figure 1 shows that there are natural plateaus in network storage size, particularly with the higher environmental factor models, i.e., environmental factor models of 80%, 85%, and 95%. In addition, note that these plateaus in network storage size tend to occur more when synthesizing networks using gene tagging.

Figure 2 shows performance accuracy as a function of storage size for 5-parent sexual evolutionary synthesis using various cluster-level and synapse-level environmental factor models, where the best synthesized networks are closest to the top left corner, i.e., high performance accuracy and low storage size. Networks synthesized using gene tagging (diamond points) appear to be minimally worse relative to networks synthesized without gene tagging (round points) in terms of maintaining performance accuracy while decreasing storage size.

4 Discussion

In this work, we presented a preliminary study into the effects of architectural alignment during evolutionary synthesis via the introduction of a gene tagging system to enforce a like-with-like mating policy. We speculate that the use of gene tagging inherently decreases overall network variability across the synthesized network architectures, indicating that enforcing a like-with-like mating policy via the use of gene tagging potentially restricts the exploration of the search space of possible network architectures. As a result, synthesized networks may fail to achieve the same increases in architectural efficiency as architectures synthesized with more random approaches. Future work includes further investigation into the effects of architectural alignment via the use of different χ s and varying the number of parents networks m , and the exploration of how gene tagging potentially affects the architectural variability in synthesized networks.

Acknowledgements

This research has been supported by the Canada Research Chairs Program and Natural Sciences and Engineering Research Council of Canada (NSERC). The authors also thank Nvidia for the GPU hardware used in this study through the Nvidia Hardware Grant Program.

References

- [1] Y. Bengio, “Learning deep architectures for AI,” *Foundations and trends in Machine Learning*, vol. 2, no. 1, pp. 1–127, 2009.
- [2] A. Graves, A.-r. Mohamed, and G. Hinton, “Speech recognition with deep recurrent neural networks,” in *2013 IEEE international conference on acoustics, speech and signal processing*. IEEE, 2013, pp. 6645–6649.
- [3] Y. LeCun, Y. Bengio, and G. Hinton, “Deep learning,” *Nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [4] J. J. Tompson, A. Jain, Y. LeCun, and C. Bregler, “Joint training of a convolutional network and a graphical model for human pose estimation,” in *Adv ances in neural information processing systems*, 2014, pp. 1799–1807.
- [5] M. J. Shafiee, A. Mishra, and A. Wong, “Deep learning with Darwin: Evolutionary synthesis of deep neural networks,” *Neural Processing Letters*, 2017.
- [6] M. J. Shafiee and A. Wong, “Evolutionary synthesis of deep neural networks via synaptic cluster-driven genetic encoding,” in *Advances in Neural Information Processing Systems*, 2016.
- [7] M. J. Shafiee, E. Barshan, and A. Wong, “Evolution in groups: a deeper look at synaptic cluster driven evolution of deep neural networks,” *arXiv preprint arXiv:1704.02081*, 2017.
- [8] M. J. Shafiee, F. Li, and A. Wong, “Exploring the imposition of synaptic precision restrictions for evolutionary synthesis of deep neural networks,” *arXiv preprint arXiv:1707.00095*, 2017.
- [9] M. J. Shafiee, E. Barshan, F. Li, B. Chwyl, M. Karg, C. Scharfenberger, and A. Wong, “Learning efficient deep feature representations via transgenerational genetic transmission of environmental information during evolutionary synthesis of deep neural networks,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 979–986.
- [10] A. G. Chung, M. Javad Shafiee, P. Fieguth, and A. Wong, “The mating rituals of deep neural networks: Learning compact feature representations through sexual evolutionary synthesis,” in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1220–1227.
- [11] A. G. Chung, P. Fieguth, and A. Wong, “Polyploidism in deep neural networks: m-parent evolutionary synthesis of deep neural networks in varying population sizes,” *Journal of Computational Vision and Imaging Systems*, vol. 3, no. 1, 2017.
- [12] A. G. Chung, P. Fieguth, and A. Wong, “Nature vs. nurture: The role of environmental resources in evolutionary deep intelligence,” *Computer and Robot Vision (CRV), 2018 15th Conference on*, 2018.
- [13] Y. LeCun, C. Cortes, and C. J. Burges, “The MNIST database of handwritten digits,” 1998.
- [14] 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.