

Assessing Architectural Similarity in Populations of Deep Neural Networks

Audrey G. Chung
Paul Fieguth
Alexander Wong

University of Waterloo, ON, Canada
University of Waterloo, ON, Canada
University of Waterloo, ON, Canada

Abstract

Evolutionary deep intelligence has recently shown great promise for producing small, powerful deep neural network models via the synthesis of increasingly efficient architectures over successive generations. However, little has been done to directly assess architectural similarity between networks during the synthesis process. We present a preliminary study into quantifying architectural similarity via the percentage overlap of architectural clusters. Results show that networks synthesized using architectural alignment (via gene tagging) maintain higher architectural similarities within each generation, potentially restricting the search space of highly efficient network architectures.

1 Introduction

The use of deep neural networks (DNNs) [1, 2] has become ubiquitous over the last few years due to their demonstrated efficacy in many challenging application areas; however, the modelling accuracy of high-performance DNNs is a result of increased model size and complexity, rendering them impractical for real-world scenarios with limited computational and memory resources and motivating the need for methods that reduce the computational requirements of DNNs while maintaining performance accuracy. Inspired by nature, *evolutionary deep intelligence* [3] is a one such method, synthesizing increasingly efficient and compact network architectures over successive generations from existing high-performance DNNs.

While the seminal papers in evolutionary deep intelligence [3, 4] formulated the synthesis process as asexual evolutionary synthesis, recent work [5, 6] has investigated the use of sexual evolutionary synthesis to produce populations of increasingly compact DNNs at each generation. Most recently, Chung *et al.* [7] conducted an initial study into mitigating architectural mismatch during sexual evolutionary synthesis via a gene tagging system. While results showed no notable difference in performance accuracy, it raises an interesting question: how can we assess the architectural similarity of DNNs in a meaningful and useful way?

2 Methods

We investigate the quantification of architectural similarity using generations of networks synthesized via multi-parent evolutionary synthesis with and without gene tagging [7]. Using the idea of architectural clusters from [4], we propose the use of percentage overlap of architectural clusters as an intuitive representation of network architecture similarity made viable in the context of multi-parent evolutionary synthesis via the gene tagging system [7]. As such, gene tagging (which allows for architectural alignment during evolutionary synthesis) can similarly be used to calculate percentage overlap of existing architectural clusters originating from the same location in the ancestor network. Percentage overlap is indicative of network population diversity within a generation, e.g., relatively low average percentage overlap would indicate a generation of synthesized networks with comparatively higher architectural variability.

3 Results

In this study, we used the network architectures synthesized in [7] with the least aggressive environmental factor model ($R_{g(i)}^c, R_{g(i)}^s = 50$) and trained on the MNIST dataset [8]. Architectural similarity was assessed on the first seven generations of networks (after which the performance accuracy degraded to random guessing) synthesized with and without gene tagging.

Synthesizing networks with and without gene tagging (Figure 1) both produced architectures that increase in variability over successive generations; however, networks synthesized with gene tagging diversify more slowly than those without gene tagging (as shown in Table 1). Figure 1 and Table 1 also suggest that generations of networks approaching an optimal tradeoff between performance accuracy and storage size tend to also have the highest architectural variability, e.g., in generations 3 and 4. Lastly, note that the increasing percentage overlap in generations 6 and 7 of networks synthesized without gene tagging is a result of sparse, low-variability architectures that can no longer represent the problem space.

Table 1: Average percentage overlap of architectural clusters in network models for the first seven generations of 5-parent sexual evolutionary synthesis.

Gen No.	Gene Tagging	No Gene Tagging
1	93.75%	93.71%
2	87.59%	78.11%
3	83.49%	68.84%
4	71.81%	66.64%
5	73.17%	68.44%
6	69.09%	82.74%
7	73.48%	91.05%

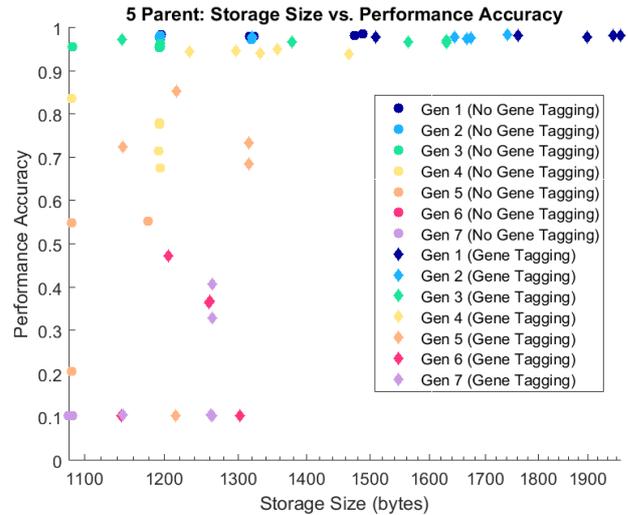


Fig. 1: Performance accuracy as a function of storage size for the first seven generations of 5-parent sexual evolutionary synthesis for networks synthesized with gene tagging (diamond) and without gene tagging (round). Plots best viewed in colour.

4 Discussion

We presented a preliminary study in assessing architectural similarity between DNNs to improve the sexual evolutionary synthesis process. Results show that networks synthesized using gene tagging have less architectural variability than networks synthesized without gene tagging, as quantified by relatively higher overlap percentages of architectural clusters. This indicates that the use of gene tagging is potentially restricting the exploration of highly efficient network architectures in the search space. Future work includes further investigation into quantities of information, e.g., mutual information, as well as the development of a custom similarity metric for optimal architectural similarity during sexual evolutionary synthesis.

Acknowledgments

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 via the Nvidia Hardware Grant Program.

References

- [1] Y. Bengio, "Learning deep architectures for AI," 2009.
- [2] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," 2015.
- [3] M. J. Shafiee *et al.*, "Deep learning with Darwin: Evolutionary synthesis of deep neural networks," 2018.
- [4] M. J. Shafiee and A. Wong, "Evolutionary synthesis of deep neural networks via synaptic cluster-driven genetic encoding," 2016.
- [5] A. G. Chung *et al.*, "The mating rituals of deep neural networks: Learning compact feature representations through sexual evolutionary synthesis," 2017.
- [6] A. G. Chung *et al.*, "Polyploidism in deep neural networks: m-parent evolutionary synthesis of deep neural networks in varying population sizes," 2017.
- [7] A. G. Chung *et al.*, "Mitigating architectural mismatch during the evolutionary synthesis of deep neural networks," 2018.
- [8] Y. LeCun *et al.*, "The MNIST database of handwritten digits," 1998.