

General Intelligence through Prolonged Evolution of Densely Connected Neural Networks

Padmini Rajagopalan
Dept. of Computer Science
University of Texas at Austin
Austin, TX 78712
padmini@cs.utexas.edu

Aditya Rawal
Dept. of Computer Science
University of Texas at Austin
Austin, TX 78712
aditya@cs.utexas.edu

Kay E. Holekamp
Dept. of Zoology
Michigan State University
East Lansing, MI 48824
holekamp@msu.edu

Risto Miikkulainen
Dept. of Computer Science
University of Texas at Austin
Austin, TX 78712
risto@cs.utexas.edu

ABSTRACT

Different species of animals have vast differences in how general their learning abilities and behaviors are. This paper analyzes the effect of network connection density and prolonged evolution on general intelligence. Using the NEAT algorithm for neuroevolution, network structures with different connectivities were evaluated in recognizing digits and their mirror images. These experiments show that general intelligence, i.e. recognition of previously unseen examples, increases with increase in connectivity. General intelligence also increases with the number of generations in prolonged evolution, even when performance no longer improves in the known examples. This outcome suggests that general intelligence depends on specific anatomical and environmental factors.

Categories and Subject Descriptors

I.2.11 [Artificial Intelligence]: Distributed Artificial Intelligence—*Intelligent agents*

General Terms

Experimentation

Keywords

Neuroevolution, Intelligence

1. INTRODUCTION

All species have special abilities that can be termed domain-specific intelligence. Some species are very good at solving new tasks that they have never seen before. This ability is called domain-general intelligence.

A prevailing theory is that general intelligence emerges from the interaction of multiple processes or modules in the brain [5]. In this paper, the NEAT algorithm [4] was used to evolve simple neural networks. These networks were evolved for one task, but tested on a different one to study their ability to adapt. Network connection densities were varied to resemble variations in interconnectivities of brain modules. The length of prolonged evolution after the old tasks had been solved was also varied. It was discovered that the performance on the new task increased with increasing connection density, up to a point. That performance also improved with prolonged evolution even though fitness in the old tasks no longer improved.

2. EXPERIMENTS

Simple two-layer neural networks were evolved using the NEAT algorithm. In the testing phase, new input was presented to the neural network. The test performance was taken as the "general intelligence" of that particular network. The networks were initialized with different connection densities. After initialization, only the connection weights were evolved, keeping the structure fixed. Each population consisted of either 30 or 100 neural networks with identical connection densities. Each experiment was run 20 times for each population and the results were averaged.

The handwritten digit recognition task was adapted for a test for general intelligence. Inputs were not only images of the digits (taken from the NIST database), but also their mirror images. The mirror image of a single digit, 7, was left out of the training set and only used as a test input. If the neural network had evolved to recognize that some of its inputs were mirror images of others, it should be able to recognize a previously unseen set of mirror images.

The influence of prolonged evolution on test performance could be verified by varying the number of generations for which the neural network was evolved after its performance had plateaued on the inputs shown during evolution.

3. RESULTS AND CONCLUSIONS

The test performance increased with increase in the connection density of the neural network. But after 65% connection density, it started to decrease again (Figure 1). At the

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

GECCO'14, July 12–16, 2014, Vancouver, BC, Canada.

ACM 978-1-4503-2881-4/14/07.

<http://dx.doi.org/10.1145/2598394.2598434>.




Figure 1: The average performance on the test phase after evolving the network for 100 generations on the other inputs. Connection density fraction is on the x -axis and percentage of test input images the neural network identified correctly on the y -axis.




Figure 2: The average test performance of neural networks with connection density 0.65 after evolution for 50, 100, 200, and 300 generations on the other inputs. Number of generations is on the X-axis and percentage of test input images the neural network identified correctly on the Y-axis.

same time, the fitness on the training inputs also decreased, indicating that the decrease was not due to overfitting, but from the difficulty of evolving large numbers of connection weights. Excessive complexity would also explain why recurrent neural networks generally have a lower test performance than feed-forward neural networks. Not only is the number of connections larger, but the number of timesteps required for the network to settle is also higher.

The test performance also increased with increasing numbers of generations for which the neural network was evolved. Experiments were conducted on networks with the optimal connection density, 0.65 (Figure 2). It should be noted that there was no significant increase in the fitness over the evolution examples during the prolonged training.

These results confirm that both denser connectivity and prolonged evolution established network processes that implemented more general solutions and therefore, more general intelligence. Hence, it can be concluded that general intelligence is determined by specific anatomical and environmental factors that affect the evolution of an animal species.

The simulations from this paper can be extended in the future to more realistic tasks in environments resembling those in which real-life animal species live, gaining insight into differences observed in biology. Neuroevolution has already been used to evolve complex animal behaviors such as group hunting, communication with conspecifics and evading predators [2, 1, 3]. Eventually, the same approach could be useful in creating intelligent behaviors for artificial agents in video games or robotics as well.

4. ACKNOWLEDGMENTS

This research was supported in part by NSF grants DBI-0939454 and IIS-0915038, and in part by NIH grant R01-GM105042.

5. REFERENCES

- [1] P. Rajagopalan, A. Rawal, R. Miikkulainen, M. A. Wiseman, and K. E. Holekamp. The role of reward structure, coordination mechanism and net return in the evolution of cooperation. In *Proceedings of the IEEE Conference on Computational Intelligence and Games (CIG 2011)*, pages 258–265, 2011.
- [2] A. Rawal, P. Rajagopalan, and R. Miikkulainen. Constructing competitive and cooperative agent behavior using coevolution. In *Proceedings of the IEEE Conference on Computational Intelligence and Games (CIG 2010)*, pages 107–114, August 2010.
- [3] A. Rawal, P. Rajagopalan, R. Miikkulainen, and K. Holekamp. Evolution of a communication code in cooperative tasks. In *Artificial Life*, volume 13, pages 243–250, 2012.
- [4] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10(2):99–127, 2002.
- [5] H. L. J. van der Maas, C. V. Dolan, R. P. P. P. Grasman, J. M. Wicherts, H. M. Huizenga, and M. E. J. Raijmakers. A dynamical model of general intelligence: The positive manifold of intelligence by mutualism. *Psychological Review*, 113(4):842–861, 2006.