Impact of Speciation Heuristic on Crossover and Search in NEAT

Phillip Verbancsics
Space and Naval Warfare Systems Center Pacific
San Diego, CA 92152-5001, USA
phillip.verbancsics@navy.mil

ABSTRACT

Crossover is an important genetic operator that can combine beneficial genes together. Unfortunately, neuro-evolution (NE) has not experienced the benefits of crossover, despite significant efforts that enabled crossover for neural networks. Orthogonally, speciation has become an important feature in NE for diversity maintenance; however, speciation research has focused on *what* measure is driving speciation versus *how* the measure determines species. This research posits that an appropriate speciation heuristic can enable effective crossover in NE by determining potential mating partners. This paper investigates these concepts and presents empirical evidence that demonstrates; (1) the impact of the speciation heuristic, (2) crossover's negative effect, and (3) a speciation heuristic that enables effective crossover in NE.

Categories and Subject Descriptors

I.2.6 [**Artificial Intelligence**]: Learning – *Connectionism and neural nets*

Keywords

Speciation; Crossover; Artificial Neural Networks; NEAT

1. INTRODUCTION

Neuro-evolutionary (NE) approaches that rely on crossover have shown diminished performance versus NE that emphasizes mutation [4] and thus some researchers dismiss crossover as an operator in NE. However, crossover's effectiveness is influenced by the pool of mating partners. Such pools in nature are limited by reproductive isolation, which is reinforced by speciation [1]. Speciation has proven to be popular in NE to encourage "niches" to preserve diversity, such as genotypic, phenotypic, and behavioral [2]. Thus research has focused on "niches" rather than mating pools.

This paper explores how the heuristic that is applied to the speciation metric can influence performance. Four different heuristics for selecting species are investigated: First Compatible, Most Compatible, Parental, and Uncanny Valley (described in Section 3). Results from experiments demonstrate the negative effect that crossover has on performance, with the majority of speciation heuristics performing worse under such conditions. However, changing the heuristic results in different performance profiles. The results indicate the importance of the speciation heuristic in performance of crossover and reveal the potential to unlock crossover's power through improved mating pool selection.

GECCO '15 July 11-15, 2015, Madrid, Spain

ACM ISBN 978-1-4503-3488-4/15/07.

DOI: http://dx.doi.org/10.1145/2739482.2764667

2. NEAT

The NeuroEvolution of Augmenting Topologies (NEAT) evolutionary algorithm [3] is a popular method that evolves neural networks. NEAT evolves connection weights as well as adds new nodes and connections over generations, thereby increasing solution complexity. NEAT has proven to be effective in challenging control and decision making tasks. The NEAT implementation for this paper differs from the canonical NEAT in a few details. First, crossover is applied on a neuron basis rather than connection basis. Next, crossover and mutation are mutually exclusive. Finally, only one new species can be created each generation.

3. SPECIATION HEURISTICS

This paper investigates four different genotypic-based speciation heuristics. First Compatible (original NEAT heuristic) wherein genomes are placed with the first species with compatibility below a threshold. Most Compatible places genomes with the species with the closest compatibility; Parental looks *only* at the genome's parents for compatibility, if the all the parent species are above threshold, either a new species is created or the genome is placed with the most compatible of the parent and new species. Uncanny Valley is the novel speciation approach that that is described in further detail in the next paragraph.

Uncanny Valley first looks at the species of the genome's parents. If a parent species is below threshold, then the remaining species are searched for a non-parent species below the threshold. If there exists such a species, the genome is placed in that species, otherwise the genome is placed with the parent species. If the genome is not compatible with its parent species, either a new species is created for the genome or it is placed with the most compatible of the parent and new species.

4. EXPERIMENTAL APPROACH

Each of the speciation heuristics are tested in the XOR and double-pole balancing (DPB) domains under 0.1 and

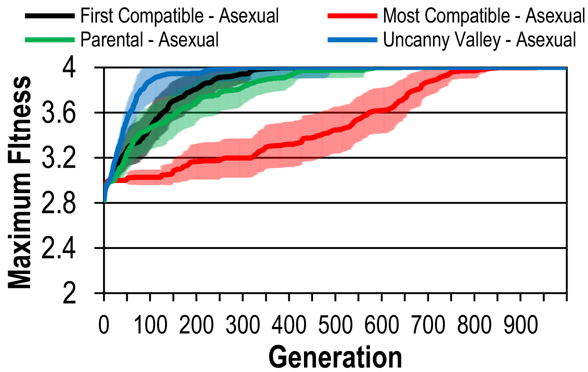

Figure 1: XOR with Low Crossover.

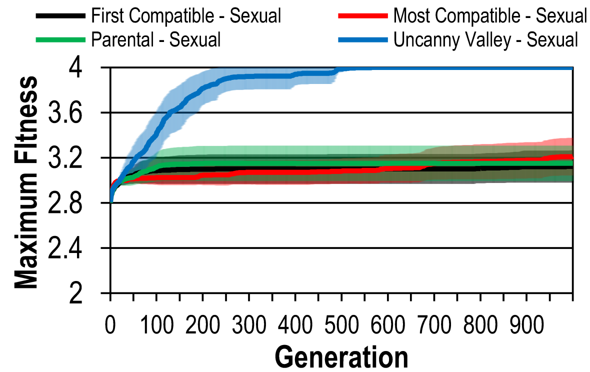

Figure 2: XOR with High Crossover.

0.9 crossover rates. XOR is an important test domain to ensure NE can correctly solve non-linear functions. XOR experiments have evolution limited to 1000 generations and does not terminate at the first solution. The DPB domain is a well-known benchmark in reinforcement learning domains. In double-pole balancing, two poles are attached at a hinge to a movable cart. The learning agent must learn how to keep both poles elevated by instructing the cart to move at particular velocity. Fitness is the number of time-steps the poles are kept elevated. Runs are limited to 500 generations and stop once a solutions reaches 100000 time-steps.

5. RESULTS

Results are averaged over 40 runs and only vary speciation heuristic and crossover rate. In XOR with high asexual rates, all speciation heuristics find a perfect solution (figure 1), but demonstrate different learning trajectories. Most Compatible and Parental are substantially similar, the Uncanny Valley converges on a solution the fastest, and Most Compatible is slowest to the optimal. The Uncanny Valley finds a solution to XOR significantly ($p < 0.01$) faster than the other speciation heuristics, First Compatible and Parental are not significantly different from each other, and Most Compatible is significantly ($p < 0.01$) worse.

In contrast, high crossover results in difficulties for First Compatible, Most Compatible, and Parental with statistically identical performances (figure 2). However, Uncanny Valley significantly ($p < 0.01$) outperforms the other approaches. Indeed, the time to a solution for Uncanny Valley with high crossover is not significantly different from approaches under with low crossover. Results for DPB are similar for low and high crossover rates (figures 3, 4).

6. DISCUSSION & CONCLUSION

Neuro-evolutionary approaches have embraced speciation, but have not been able to exploit the power of crossover. Results demonstrated that crossover can be detrimental to performance in NE. However, different speciation heuristics can allow NE to harness crossover's power. Note, this paper does not claim one speciation heuristic is better than another, rather the implication is that implementation details (e.g. speciation heuristic) significantly impact performance.

7. ACKNOWLEDGMENTS

This work was supported and funded by the SSC Pacific Naval Innovative Science and Engineering (NISE) Program.

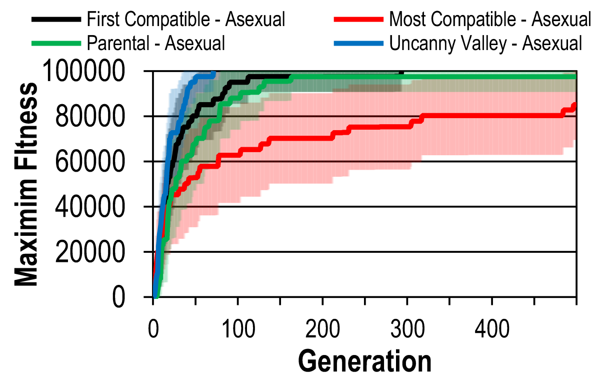

Figure 3: DPB with Low Crossover.

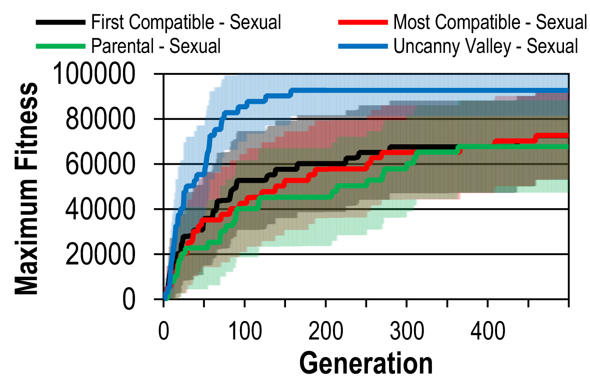

Figure 4: DPB with High Crossover.

8. REFERENCES

[1] E. Mayr. *The growth of biological thought: diversity, evolution, and inheritance.* Harvard Univ. Press, 1982.

[2] J.-B. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary computation*, 20(1):91–133, 2012.

[3] K. O. Stanley and R. Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary Computation*, 10:99–127, 2002.

[4] X. Yao. Evolving artificial neural networks. *Proceedings of the IEEE*, 87(9):1423–1447, 1999.