

Augmenting neuro-evolutionary adaptation with representations does not incur a speed accuracy trade-off

Douglas Kirkpatrick

Department of Computer Science and Engineering
BEACON Center for the Study of Evolution in Action
Michigan State University
East Lansing, Michigan
kirkpa48@msu.edu

Arend Hintze

Department of Integrative Biology
Department of Computer Science and Engineering
BEACON Center for the Study of Evolution in Action
Michigan State University
East Lansing, Michigan
hintze@msu.edu

ABSTRACT

Representations, or sensor-independent internal models of the environment, are important for any type of intelligent agent to process and act in an environment. Imbuing an artificially intelligent system with such a model of the world it functions in remains a difficult problem. However, using neuro-evolution as the means to optimize such a system allows the artificial intelligence to evolve proper models of the environment. Previous work has found an information-theoretic measure, R , which measures how much information a neural computational architecture (henceforth loosely referred to as a brain) has about its environment, and can additionally be used speed up the neuro-evolutionary process. However, it is possible that this improved evolutionary adaptation comes at a cost to the brain's ability to generalize or the brain's robustness to noise. In this paper, we show that this is not the case; to the contrary, we find an improved ability of the to evolve in noisy environments when the neuro-correlate R is used to augment evolutionary adaptation.

CCS CONCEPTS

• **Computing methodologies** → **Knowledge representation and reasoning**; **Genetic algorithms**; *Artificial life*;

KEYWORDS

Genetic Algorithms, Artificial Life, Representations, Artificial Intelligence, Robustness of Solutions

ACM Reference Format:

Douglas Kirkpatrick and Arend Hintze. 2019. Augmenting neuro-evolutionary adaptation with representations does not incur a speed accuracy trade-off. In *Genetic and Evolutionary Computation Conference Companion (GECCO '19 Companion)*, July 13–17, 2019, Prague, Czech Republic. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3319619.3322047>

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
GECCO '19 Companion, July 13–17, 2019, Prague, Czech Republic
© 2019 Copyright held by the owner/author(s). Publication rights licensed to the Association for Computing Machinery.
ACM ISBN 978-1-4503-6748-6/19/07...\$15.00
<https://doi.org/10.1145/3319619.3322047>

1 INTRODUCTION

Neuro-evolution has shown promising results in creating more capable and adaptable solutions for problems in artificial intelligence [5, 11]. One of the downsides to this process is that the search itself is inherently random and as such, an optimal agent may not evolve quickly. Similarly, it is possible that the evolutionary search gets stuck, and results in a less than optimal solution. To combat this defect, many methods have been proposed to accelerate neuro-evolution. The method to reward solutions that are more modular [4] is technically the closest to the approach used here.

One issue that has already been identified to be specifically troublesome in deep learning is representations. It seems as if neural networks are easily fooled [9] either because they are over fitted, or because they simply can not generalize well. They lack a robust internal model, instead relying too heavily on training irregularities. However, it has been shown that neuro-evolution can very well overcome this limitation, and allows systems to create cognitive agents that have meaningful and robust representations about their environment [8]. Furthermore, it has been shown that one can augment evolutionary adaptation [10] and the use of representations by first measuring representations and then adding said measurement to the fitness function. This is similar to a multiple objective optimization [12], but where both parameters are mostly in line with the same goal.

We investigate two aspects of intelligence that might suffer from trade-offs and are relevant for artificial intelligence applications: the brain's ability to generalize and to be robust against noise. Often in machine learning the distinction between robustness and generalization is hard to make, and the terms are often used interchangeably. Here generalization is defined as the ability to deal with a new task that the system was not explicitly trained on. Robustness on the other hand is the ability to perform the same task under noise.

In this paper, we demonstrate that neither the ability of Markov Brains to generalize nor the brain's robustness to noise is reduced when the evolutionary process is augmented by using R . Moreover, we found co-evolving with R resulted in better performance when evolving R in noisy environments.

2 MATERIALS AND METHODS

All experiments were carried out using evolvable Markov Brains [6] in the MABE framework [3]. For these brains R (representations) and the smearedness (structure) of these representations can be




Figure 1: Mean performance of agents evolved under different levels of noise (x axis) with (solid line) and without (dashed line) augmenting the GA with R . The performance ($\bar{W}$) is from agents from 350 replicate experiments for each level of noise, after 40,000 generations of evolution. Error bars indicate the standard error. The results were significantly different when tested with a Wilcoxon Rank-Sum test ($p = 5.95e - 5$). These agents were evolved to catch blocks of size 2 and avoid blocks of size 4 and had a gap width of 2.

easily quantified [7, 8, 10]. Markov Brains were optimized using a genetic algorithm in 350 replicates over 40,000 generations to perform an active categorical perception task [1, 2]. In this task, blocks of different sizes fall towards the agent who has to catch small and avoid large blocks. In order to augment evolution the performance of the agents was calculated using the following fitness function:

$$W = 1.10^{(C-I)} \left(1 + \frac{R}{R_{max}}\right) \quad (1)$$

where C is the number of correct actions, and I the number of incorrect ones. R is the measure of the agent’s representations, and R_{max} is the calculated maximum value of R .

The ability of the evolved brains to generalize was tested by quantifying their performance using blocks of sizes agents did not encounter during evolution. Robustness was quantified by measuring the performance of the agents while their hidden states were infused with noise in the form of random zeros and ones.

3 RESULTS

We asked if augmenting the search of a genetic algorithm (GA) by using R has negative side effects, specifically on the ability of agents to be robust and to generalize. We first affirmed that agents indeed evolve faster and have better final performance when the GA was augmented by using R . We found no significant loss of the agents ability to generalize evolved by a GA augmented with R versus agents evolved without. In rare cases we found a slight improvement in the agents ability to generalize when evolved using R to augment the GA, however without any identifiable pattern.

Similarly, agents evolved by a GA augmented by R showed no significant loss or improvement towards their ability to deal with noise. Their robustness was unaffected. However, when evolving agents not only by augmenting R but by also infusing noise into their sensor states during evolution, we found an improved performance under noise (see Figure 1). This suggests that augmenting a GA with R in conjunction with noise allows agents to become

more robust during evolution compared to agents evolved by a GA not augmented with R .

4 CONCLUSIONS

The work described here reaffirms the benefits of R as a neuro-correlate capable of accelerating rates of adaptation, and prompts several new avenues of research to follow. First of all, the use of R does not compromise the ability of an evolved agent to generalize or to be robust against sensor noise. Unfortunately, performance optimization by a GA is very different from gradient descent or other optimization methods in deep learning. As such, R can not simply be applied in that context, even though these systems seem to struggle with creating internal models and representations.

The most interesting aspect we observed pertains to the effect that augmenting the GA with R has on evolution under noisy conditions. While using R made no difference to sensor noise robustness when agents were evolved without noise, evolution under noise resulted in significantly better adaptation when the GA was augmented with R . This implies two things: evolving representations under noise might be a harder challenge, and the augmentation therefore is advantageous. However, more interesting is the implication that one might be able to take advantage of the observation that augmenting the GA with R allows better performance under noise. Sensor noise introduces error into the fitness function; reducing the number of trials performed reduces the sample size and similarly introduces error into the fitness function. Perhaps instead of testing an agents performance thoroughly on all possible test cases, one can use a subset of the possible cases while augmenting the GA with R . We seek to investigate this possibility further.

REFERENCES

- [1] Randall D Beer. 2003. The dynamics of active categorical perception in an evolved model agent. *Adaptive Behavior* 11, 4 (2003), 209–243.
- [2] Randall D Beer et al. 1996. Toward the evolution of dynamical neural networks for minimally cognitive behavior. *From animals to animats 4* (1996), 421–429.
- [3] Clifford Bohm, Nitash CG, and Arend Hintze. 2017. MABE (Modular Agent Based Evolver): A Framework for Digital Evolution Research. In *Proceedings of the 14th European Conference on Artificial Life*, C. Knibbe et al. (Ed.). *Proceedings of the 14th European Conference of Artificial Life*, 76–83.
- [4] Jeff Clune, Jean-Baptiste Mouret, and Hod Lipson. 2013. The evolutionary origins of modularity. *Proc. R. Soc. B* 280, 1755 (2013), 20122863.
- [5] Dario Floreano, Peter Dürri, and Claudio Mattiussi. 2008. Neuroevolution: from architectures to learning. *Evolutionary Intelligence* 1, 1 (2008), 47–62.
- [6] Arend Hintze, Jeffrey A Edlund, Randal S Olson, David B Knoester, Jory Schossau, Larissa Albantakis, Ali Tehrani-Saleh, Peter Kvam, Leigh Sheneman, Heather Goldsby, et al. 2017. Markov brains: A technical introduction. *arXiv preprint arXiv:1709.05601* (2017).
- [7] Arend Hintze, Douglas Kirkpatrick, and Christoph Adami. 2018. The structure of evolved representations across different substrates for artificial intelligence. In *Artificial Life Conference Proceedings*. MIT Press, 388–395.
- [8] Lars Marstaller, Arend Hintze, and Christoph Adami. 2013. The evolution of representation in simple cognitive networks. *Neural computation* 25 (2013), 2079–2107.
- [9] Anh Nguyen, Jason Yosinski, and Jeff Clune. 2015. Deep neural networks are easily fooled: High confidence predictions for unrecognizable images. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 427–436.
- [10] Jory Schossau, Christoph Adami, and Arend Hintze. 2015. Information-theoretic neuro-correlates boost evolution of cognitive systems. *Entropy* 18, 1 (2015), 6.
- [11] Xin Yao. 1999. Evolving artificial neural networks. *Proc. IEEE* 87, 9 (1999), 1423–1447.
- [12] Aimin Zhou, Bo-Yang Qu, Hui Li, Shi-Zheng Zhao, Ponnuthurai Nagarathnam Suganthan, and Qingfu Zhang. 2011. Multiobjective evolutionary algorithms: A survey of the state of the art. *Swarm and Evolutionary Computation* 1, 1 (2011), 32–49.