

Is Novelty Search Good for Evolving Morphologically Robust Robot Controllers?

Extended Abstract

Ruben Putter, Geoff Nitschke
 Department of Computer Science
 University of Cape Town, Cape Town, South Africa
 pttand010@myuct.ac.za, gnitschke@cs.uct.ac.za

ABSTRACT

This study evaluates comparative behavioral search methods for evolutionary controller design in robot teams, where the goal is to evaluate the *morphological robustness* of evolved controllers. That is, where controllers are evolved for specific robot sensory-motor configurations (morphologies) but must continue to function as these morphologies degrade. Robots use neural controllers where behavior evolution is directed by developmental *Neuro-Evolution* (HyperNEAT). Guiding evolutionary controller design we use objective (*fitness function*) versus non-objective (*novelty*) search. The former optimizes for behavioral fitness and the latter for behavioral novelty. These search methods are evaluated across varying robot morphologies and increasing task complexity. Results indicate that both novelty and objective search evolve team controllers (behaviors) that are morphologically robust given degrading robot morphologies and increasing task complexity. Results thus suggest that novelty search is not necessarily suitable for generating robot team behaviors that are robust to changes in robot morphologies (for example, due to damaged or disabled sensors and actuators).

KEYWORDS

Morphological Robustness, Robot Teams, Objective and Non-Objective Evolutionary Controller Design

ACM Reference Format:

Ruben Putter, Geoff Nitschke. 2018. Is Novelty Search Good for Evolving Morphologically Robust Robot Controllers?. In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), Stockholm, Sweden, July 10-15, 2018*, IFAAMAS, 3 pages.

MORPHOLOGICALLY ROBUST CONTROLLERS

Autonomous robots are increasingly being applied to remote and hazardous environments [1, 3], environments damage to sensory-actuator systems (*morphologies* [17]) cannot be easily repaired if damaged. An unsolved problem in the controller design for such autonomous robots is having controllers continue to effectively function given unexpected changes, such as damage, to robot morphology. Currently, robotic systems recover from damage via self-diagnosis and selection from pre-designed contingency plans in order to continue functioning [2, 10, 21]. Though robots using such self-diagnosis and recovery are problematic systems as they are expensive, requiring sophisticated monitoring sensors, and difficult

to design as *a priori* knowledge of all necessary contingency plans is assumed [5]. Addressing this, recent work in *Evolutionary Robotics* [8] elucidated the efficacy of population based stochastic *trial and error* methods for *online* damage recovery in autonomous robots operating in physical environments [5]. This was demonstrated as being akin to self-adaptation and injury recovery of animals observed in nature.

This study further contributes to this research area, focusing on *evolutionary controller design* [11] within the broader context of *collective* [15] and *swarm* [1] robotics. That is, evolutionary controller design for robot groups that must continue to accomplish tasks given that damage is sustained to the morphologies of some or all of the robots in the group. Given previous non-objective controller evolution work in collective robotics [12-14], and previous research demonstrating the efficacy of *novelty search* [16] for evolving behaviors that operate across various robotic morphologies [4], the following research objectives formed the focus of this study.

- (1) Novelty search behaviors out-perform those of objective search in terms of average team task performance in increasingly complex collective construction tasks.
- (2) Novelty search is suitable for evolving morphologically robust controllers in collective construction tasks.

To test this objective, experiments evaluated various robot morphologies in increasing complex collective construction tasks. That is, evolved controller morphological robustness was evaluated in terms of the given controller's task performance when coupled with alternate robot morphologies. This study's contribution was thus to elucidate the impact of objective versus non-objective search on the evolution of controllers that exhibit *morphologically robust* behaviors in collective robotic systems [8]. Specifically, collective robotic systems that must effectively adapt to unforeseen morphological change, such as the loss or damage of sensors without significant task performance degradation [2, 5].

METHODS AND EXPERIMENTS

HyperNEAT [19] was applied to evolve robot team (collective) behaviors, where teams were behaviorally and morphologically homogenous teams meaning all robots in a given team used the same controller and sensory configuration. HyperNEAT extends NEAT (*Neuro-Evolution of Augmented Topologies*) [20], where *Artificial Neural Network* (ANN) controllers were indirectly encoded using a CPPN (*Compositional Pattern Producing Network*) [18].

Experiments tested 15 robots in a bounded two dimensional continuous simulation environment (20 x 20 units) with randomly

Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2018), M. Dastani, G. Sukthankar, E. André, S. Koenig (eds.), July 10-15, 2018, Stockholm, Sweden. © 2018 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

Table 1: Sensory configuration (number and type of sensor) for each robot morphology.

Morphology ID	Proximity Sensors	Ultrasonic Sensors	Color Ranged Sensors	Low-Resolution Camera	Construction Zone Sensors
1	5	3	1	1	1
2	3	2	1	1	1
3	1	0	1	1	1

distributed type *A* and *B* blocks. Robots and blocks were initialized with random orientations and positions throughout the environment. The experimental objective was to evaluate the *morphological robustness* of HyperNEAT evolved controllers for robot teams given collective construction tasks. We measured average task performance of controllers evolved for three team morphologies and two levels of task complexity. Each experiment comprised a team *controller evolution* and *re-evaluation* stage, where the latter was the morphological robustness test.

In the controller evolution stage, each experiment applied HyperNEAT, where evolutionary search (to evolve team behavior) was directed by either objective-based [9] or novelty search [16], running for 100 generations. Each generation comprised three team *lifetimes* (1000 simulation iterations), where each team lifetime tested different robot starting positions, orientations and building-block locations.

Teams that achieved an average task performance that was not significantly lower across all *re-evaluated* morphologies were considered to be *morphologically robust*. Specifically, the fittest controller evolved for a given morphology and level of task complexity was re-evaluated in the other morphologies for the same level of task complexity. For example, the fittest controller evolved for morphology 1 was re-evaluated in morphologies 2 and 3 and the average task performance calculated across all re-evaluation runs. Re-evaluation runs were *non-evolutionary*, meaning controllers were not further evolved, and each re-evaluation was equivalent to one team lifetime. There were 20 re-evaluation runs for each morphology to account for random variations in robot and block starting positions and orientations.

Collective Construction Task

The task required robots to search the environment for building blocks and cooperatively push the blocks together into a structure. Task complexity was the level of cooperation required to optimally solve the task, that is, connect all the blocks in construction zones. For complexity levels 1 and 2, there were 15 type *A* and *B* blocks, respectively. In the case of level 1, a single robot could push each block, but in the case of level 2, three robots were required to cooperatively push and connect blocks. A construction zone was formed via at least two blocks being pushed together thus forming a structure. Once a construction zone was created, all blocks attached to it were fixed in position and could not be disconnected. The task used a maximum of three construction zones and unconnected blocks had to be pushed and connected to one of these construction zones. Team task performance was calculated as the number of blocks connected in construction zones during a team's lifetime,

where average task performance was the highest performance at the end of each run (100 generations), averaged over 20 runs.

RESULTS AND DISCUSSION

Results indicate that, for both novelty and objective search, given increasing task complexity, there was no significant difference in average task performance between controllers evolved in any morphology and re-evaluated in other morphologies. However, there were two exceptions to this result. First, novelty search applied in task complexity 1 to evolve controllers in morphology 1, where the fittest controller re-evaluated in morphology 3 yielded a significantly lower average task performance. Second, objective search applied in task complexity 2 to evolve controllers in morphology 1, where similarly, the fittest controller re-evaluated in morphology 3 yielded a significantly lower task performance.

In these cases, one may observe that there is a large difference in number of sensors used by morphology 1 versus morphology 3 (table 1), indicating that controllers evolved for the high sensor complement (and thus functionality) of morphology 1, are not readily transferable to a simpler sensory configuration (with much less functionality). However, this was not the case for novelty search applied in task level 2 or objective-search applied in task level 1.

The key results were thus two-fold. First (supporting the first hypothesis), team behaviors evolved by novelty-search, for all morphologies and task complexity levels, significantly out-performed team behaviors evolved with objective search for corresponding morphologies and task complexity levels. This result is supported by previous work [12], [14], [13], [6], [7]. Second (refuting the second hypothesis), results indicated that, for all team morphologies and levels of task complexity evaluated, both novelty and objective search were effective in evolving morphologically robust controllers.

That is, for any given morphology, there was no significant difference in average task performance between the fittest novelty search evolved team controllers and re-evaluation of these controllers in the other morphologies. The same result was observed for objective search. This second result indicates that while novelty search yields advantages over objective search in terms of evolving high task performance team behaviors, it yields no benefits over objective search for evolving morphologically robust controllers in the collective construction task for the given team morphologies. However, the suitability of novelty search for evolving morphologically robust controllers in complex collective behavior tasks is the topic of ongoing research.

REFERENCES

- [1] G. Beni. 2004. From Swarm Intelligence to Swarm Robotics. In *Proceedings of the First International Workshop on Swarm Robotics*. Springer, Santa Monica, USA.
- [2] J. Bongard, V. Zykov, and H. Lipson. 2006. Resilient machines through continuous self-modeling. *Science* 314(5802) (2006), 1118–1121.
- [3] R. Brooks and A. Flynn. 1989. Fast, Cheap and Out of Control: Robot Invasion of the Solar System. *Journal of the British Interplanetary Society* 1, 1 (1989), 478–485.
- [4] F. Corucci, M. Calisti, H. Hauser, and C. Laschi. 2015. Novelty-Based Evolutionary Design of Morphing Underwater Robots. In *Proceedings of the Genetic Evolutionary Computation Conference*. ACM, Madrid, Spain, 145–152.
- [5] A. Cully, J. Clune, D. Tarapore, and J. Mouret. 2015. Robots that can adapt like animals. *Nature* 521(1) (2015), 503–507.
- [6] S. Didi and G. Nitschke. 2016. Hybridizing Novelty Search for Transfer Learning. In *Proceedings of the IEEE Symposium Series on Computational Intelligence*. IEEE Press, Athens, Greece, 10–18.
- [7] S. Didi and G. Nitschke. 2016. Multi-Agent Behavior-Based Policy Transfer. In *Proceedings of the European Conference on the Applications of Evolutionary Computation*. Springer, Porto, Portugal, 181–197.
- [8] S. Doncieux, N. Bredeche, J-B Mouret, and A. Eiben. 2015. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI | Evolutionary Robotics* 2(4) (2015), 1–18.
- [9] A. Eiben and J. Smith. 2003. *Introduction to Evolutionary Computing*. Springer, Berlin, Germany.
- [10] W. Fenton, T. McGinnity, and L. Maguire. 2001. Fault diagnosis of electronic systems using intelligent techniques: A Review. *IEEE Transactions on Systems, Man, and Cybernetics, Part C: Applications and Reviews* 31(3) (2001), 269–281.
- [11] D. Floreano, P. Dürri, and C. Mattiussi. 2008. Neuroevolution: from architectures to learning. *Evolutionary Intelligence* 1, 1 (2008), 47–62.
- [12] J. Gomes and A. Christensen. 2013. Generic Behaviour Similarity Measures for Evolutionary Swarm Robotics. In *Proceedings of the 15th Annual Conference on Genetic and Evolutionary Computation*. ACM, Amsterdam, Netherlands, 199–206.
- [13] J. Gomes, P. Mariano, and A. Christensen. 2015. Devising Effective Novelty Search Algorithms: A Comprehensive Empirical Study. In *Proceedings of the Genetic Evolutionary Computation Conference*. ACM, Madrid, Spain, 943–950.
- [14] J. Gomes, P. Urbano, and A. Christensen. 2013. Evolution of swarm robotics systems with novelty search. *Swarm Intelligence* 7 (2013), 115–144.
- [15] R. Kube and H. Zhang. 1994. Collective robotics: from social insects to robots. *Adaptive Behaviour* 2, 2 (1994), 189–218.
- [16] J. Lehman and K. Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19(2) (2011), 189–223.
- [17] R. Pfeifer and J. Bongard. 2006. *How the body shapes the way we think*. MIT Press, Cambridge, USA.
- [18] K. Stanley. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic Programming and Evolvable Machines: Special Issue on Developmental Systems* 8, 2 (2007), 131–162.
- [19] K. Stanley, D'Ambrosio, and J. Gauci. 2009. Hypercube-based indirect encoding for evolving large-scale neural networks. *Artificial Life* 15, 1 (2009), 185–212.
- [20] K. Stanley and R. Miikkulainen. 2002. Evolving Neural Networks Through Augmenting Topologies. *Evolutionary Computation* 10, 2 (2002), 99–127.
- [21] G. Gordon V. Verma, R. Simmons and S. Thrun. 2004. Real-time fault diagnosis. *IEEE Robotics and Automation Magazine* 11(2) (2004), 56–66.