Combination of Novelty Search and Fitness-Based Search Applied to Robot Body-Brain Co-Evolution

Peter Krčah*

Computer Center
Charles University
Prague, Czech Republic
peter.krcah@ruk.cuni.cz

Daniel Toropila*

Computer Center
Charles University
Prague, Czech Republic
daniel.toropila@ruk.cuni.cz

Abstract: Evolutionary algorithms are a frequently used technique for designing morphology and controller of a robot. However, a significant challenge for evolutionary algorithms is premature convergence to local optima. Recently proposed Novelty Search algorithm introduces a radical idea that premature convergence can be avoided by ignoring the original objective and searching for any novel behaviors instead. In this paper, we apply novelty search to the problem of body-brain co-evolution. We show that novelty search significantly outperforms fitness-based search in a deceiving barrier avoidance task. Furthermore, we demonstrate an unexpected result that switching from novelty search to fitness-based search after the deceptive barrier is overcome does not significantly improve overall search performance.

Keywords: evolutionary algorithms, novelty search, body-brain co-evolution, artificial neural networks

1 Introduction

Evolutionary algorithms are a frequently used technique for designing morphology and controller of a robot [1, 4, 9]. The advantage of using evolutionary algorithms compared to other optimization methods is that only a high level fitness function and a set of genetic operators (mutation, crossover) need to be specified. While such fitness-driven search can be very successful, a common problem in evolutionary algorithms is premature converge to local optima. This is often caused by too much focus on exploitation of already discovered areas of the search space as opposed to exploration of yet unknown solutions. The trade-off between exploration and exploitation is well understood and several methods for addressing this issue have been proposed [3]. Several of the diversity maintaining methods have also been applied to evolutionary robotics [1, 5, 8].

*The research is supported by the Czech Science Foundation under the project no. 201/09/H057 and by the Charles University Grant Agency under contract no. 9710/2010.

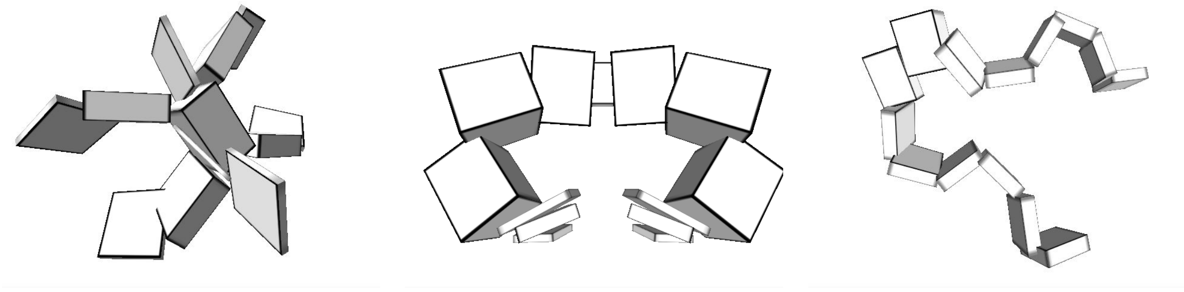

(a) (b) (c)

Figure 1: **Examples of Evolved Robot Morphology.** Several evolved robots exhibit symmetry (1b) and segmentation (1c).

However, recent works suggest that the underlying problem lies not just in the balance of exploration and exploitation, but in the deceitfulness of the fitness function itself [2,6,7]. Novelty search, a recently proposed approach [6], introduces a radical idea that for some problems convergence to local optima can be avoided by simply ignoring the original objective altogether and instead only searching for *any* novel behaviors regardless of their quality with respect to the fitness function. While seemingly counter-intuitive, this approach has already been successfully applied in multiple domains [2, 6, 7].

In this paper, we use search for behavioral novelty as a method for avoiding premature convergence to local optima in the evolution of body and controller of a simulated robot. We also investigate the effects of switching from novelty search to fitness based search after a varying number of generations. The goal of this experiment is to examine if combined search improves overall performance by accelerating the convergence after the robot overcomes the barrier.

2 Methods

2.1 The Robot

The robot is composed of boxes of varying sizes connected by joints (see figure 1 for examples of evolved robots). Genetic representation of a robot is a directed graph of nodes representing body parts. Robot is created from this genome by first adding the body part represented by the root node of the graph and then recursively traversing connections in depth-first order, adding encountered nodes and connections to the robot. To prevent infinite recursion, each node has a *recursive limit* which limits the number of passes through the given node during transcription. Each node also specifies the *size* of the resulting body part and a type of joint connecting body part to its parent. A connection between two body parts specifies the position, rotation and scale of the target body part relative to the source body part. Optional reflection flag is also specified which (if enabled) causes a mirror copy of the body part to be added to the robot along with the original body part.

This representation permits very compact encoding of complex body structures allowing for features such as symmetry (using reflection flags) and repetitive segmentation (using recursive transcription). Examples of evolved robots are shown in figure 1. In our experiments, robots are placed in a simulated water environment. Water is simulated by applying a drag force opposing the movement of each body part and disabling gravity.

Robot's controller is distributed in the body of the robot. Each body part contains a local neuro-controller (an artificial neural network), as well as a local sensor and effector. A sensor in each body part is measuring current angle of each degree of freedom of a joint and an effector allows the robot to control the amount of torque applied to each degree of freedom of a joint. The amount of torque is limited to prevent instabilities in the physical simulation. Apart from standard sigmoidal transfer function, oscillatory transfer function was used to enable faster discovery of efficient swimming strategies.

2.2 Novelty Search

Novelty Search is an algorithm proposed recently as an alternative to standard fitness function-based approach [7]. Instead of following the gradient of the fitness function, novelty search directs the search towards any yet unexplored parts of the behavior space. This is achieved by modifying the search algorithm to use the measure of individual's behavioral novelty instead of the fitness function. No other modifications to the underlying search algorithm are required.

Measure of individual's novelty is computed using an archive of previously discovered behaviors [7] as an average distance of the individual's behavior from k closest behaviors recorded in the archive:

$$\rho(x) = \frac{1}{k}\sum_{i=0}^{k} dist(x, \mu_i),$$

where μ_i is the ith-nearest neighbor of x and with respect to distance measure $dist$. This measure provides an estimate of local sparseness in the vicinity of the behavior being measured. Novelty search thus promotes individuals which are further away from already discovered behaviors. If novelty of an individual exceeds some threshold value, individual is added to the archive. In this work the size of an archive was unlimited.

Hierarchical NEAT (NeuroEvolution of Augmenting Topologies) algorithm was used as an underlying search algorithm in all experiments (complete description is presented in [5]).

3 Experiments

By searching *only* for novelty in behavior, novelty search is less prone to falling into traps set up by the objective-based fitness function in the form of local optima. We test this hypothesis in a barrier avoidance experiment using a highly deceptive fitness function with a constrained behavioral space.

In the barrier avoidance experiment, robot is placed at one end of a 20m x 20m x 25m container while the target is placed at the other end of the container (see figure 3a). Robot's goal is to reach the target within the time limit of 60 seconds. Task is made deceptive by putting the barrier directly in front of the robot, so it obstructs the direct path to the target. Moreover, the shape of the barrier only makes it possible to reach the target if the robot first moves *away* from it, which makes this task highly deceptive. Fitness function is defined as

$$f = \max(0, d(p_{start}, p_{target}) - d(p_{final}, p_{target})) \tag{1}$$

where $d(p_{start}, p_{target})$ is the constant distance from the starting position of the robot to the target (in this case 20m) and $d(p_{final}, p_{target})$ is the distance from the final position of the robot to the target. Behavior of the robot is defined as robot's final position and behavior metrics is defined as the Euclidean distance between final positions of two robots. Target position is fixed for all experiments and robot has no information about the position of the target. Barrier avoidance experiment is performed in a simulated water environment.

The performance of novelty search was compared to the performance of standard fitness-based search. Random search was used as a baseline when comparing performance (random search was performed by assigning each individual a random fitness value). To further analyze both search methods, additional experiments were performed using a combination of the novelty

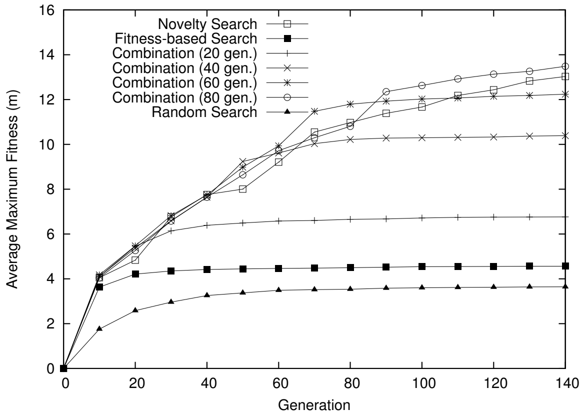

(a) Inverted Distance From Target (see equation (1))

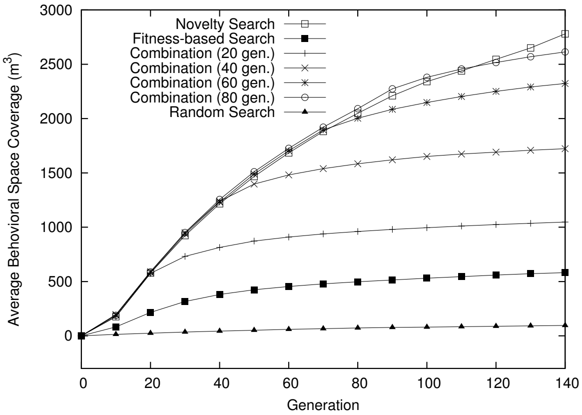

(b) Behavior Space Coverage

Figure 2: **Comparison of Different Search Methods.**

search and the fitness-based search. The combined search starts with novelty search and switches to fitness-based search after 20, 40, 60 or 80 generations. The motivation for switching the methods in the middle of the search is to more efficiently exploit strengths of each method. Once novelty search overcomes the barrier (by focusing on the exploration part of the search), fitness based search may be able to quickly converge to a solution (by focusing on the exploitation).

3.1 Parameter Settings

The number of individuals in a population is set to 300, with 150 generations per run. All parameters of the underlying search algorithm (hierarchical NEAT [5]) were set to the same values in both the standard fitness-based configuration and in the novelty-based configuration. Parameter k for computing the novelty of an individual was set to 15 and the novelty threshold for adding an individual to the archive was set to 0.1. Each configuration was tested independently at least 25 times. All significance levels were computed using Student's t-test.

4 Results

In barrier avoidance experiment, fitness-based search never successfully found a way out of the barrier, reaching an average maximum fitness of 4.56m (see figure 2a). Performance of the fitness-based search was in this case only marginally better than random search which achieved maximum average fitness of 3.68m (although the difference was statistically significant; $p < 10^{-8}$). All experiments involving the novelty search (either combined with fitness-based search or standalone), consistently escaped the trap. For combined experiments, the later the switch was made, the higher the final achieved fitness value was. The best average fitness values were reached by the novelty search (13.16m), and the combined method switching at 60 (12.26m) and 80 generations (13.65m). The differences between these best methods were not statistically significant ($p > 0.25$). Two other two combined search methods achieved average fitness value of 10.39m (switch after 40 generations) and 6.81m (switch after 20 generations).

4.1 Analysis of Behavioral Diversity

To provide better insight into how individual search methods explore the space of behaviors, the coverage of behavior space was computed as a total number of cells from a 1m x 1m x 1m grid that contained the final position of at least one robot. This statistic was computed after each

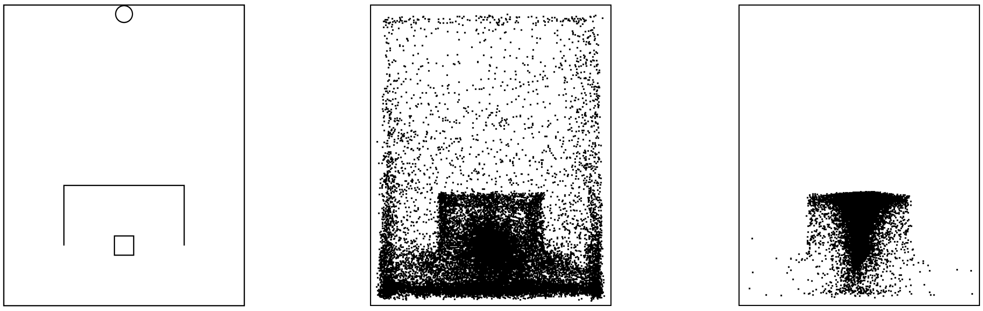

(a) Barrier Avoidance Experiment Layout

(b) Barrier Avoidance (Novelty Search)

(c) Barrier Avoidance (Fitness Search)

Figure 3: **Final Positions Of The Robot Visited In Typical Runs.** Final positions (3b, 3c) are three-dimensional and are shown here using a two-dimensional orthogonal projection. Dimensions of the environment are 20m x 20m x 25m. The barrier is located between the robot and the target and is formed by a hollow 10m x 10m x 5m box with the bottom face left out. The target (indicated by the circle in figure 3a) is located at the top of the container at the initial distance of 20m from robot's initial position (indicated by the square).

generation using all behaviors seen since the start of the run. Resulting values were averaged over 25 runs (see figure 2b).

Comparison of the amount of behavior space covered by individual methods shows that novelty search was exploring the behavior space faster than the fitness based search ($p < 10^{-10}$). This was confirmed by combined experiments, where later switch to fitness-based search consistently resulted in higher coverage of behavior space. The difference in how novelty search and fitness-based search explore the behavior space can best be seen in the set of typical runs shown in figures 3b and 3c. While novelty search thoroughly and evenly explores the search space, fitness-based search tends to exhaustively search a small number of promising directions (3c). Fitness based search reached average behavior space coverage of 5.89% of the container, which closely matches the volume of the interior space of the barrier (5% of the container). Combined search methods switching at generation 20, 40, 60 and 80 reached 10.57%, 17.33%, 23.48% and 26.49% of the container, respectively. Standalone novelty search reached the highest behavior space coverage of all methods: 28.87%. Example of a typical run of novelty search is shown in figure 3b. All differences in behavior coverage are statistically significant ($p < 0.05$) except the difference between novelty search and combined search switching after 80 generations ($p < 0.06$).

5 Discussion

The deceptiveness of the task lies in the fact that fitness-based search leads the robot directly to the target until it reaches the barrier. At that point, in order to find better solutions, the search needs to explore behaviors that are temporarily further away from target. In the standard objective-based approach, getting further away from the target decreases the fitness value; fitness function thus has a local optimum inside the barrier where fitness-based search is likely to be trapped. This was confirmed by the results of the fitness-based search, where individuals never fully escaped the barrier. In the same experiment, the standalone novelty search was able to explore the entire container and was capable of consistently finding the target.

The hypothesis that switching to fitness-based search after the barrier is overcome will improve the performance of the search was not confirmed by the experiments. Experiments with combined search switching at generations 60 and 80 indicate that switching to fitness search may provide a speedup. However, the speedup was only temporary and the difference from novelty-search was not sufficiently significant ($p > 0.1$). Moreover, behavior space analysis shows that after switching to fitness-based search, the exploration rate slowed significantly.

For future work, it would be interesting to investigate long-term trends in behavior space exploration by novelty search, which in our experiments did not reach a stable plateau in the first 150 generations (longer runs were not feasible due to large demands on processing power).

6 Conclusion

In this paper, we applied recently proposed novelty search algorithm to the evolution of body and brain of a simulated robot. We demonstrated advantages and disadvantages of novelty search in a barrier avoidance experiment. We have shown that if the range of possible behaviors is limited (by putting the robot inside a container) and fitness function is deceptive then the novelty search algorithm significantly outperforms fitness-based search. We also examined the combinations of novelty search and fitness-based search. Results have shown that switching to the fitness-based search doesn't significantly improve final fitness or increase exploration of the behavior space. In summary, this paper confirms that the advantages of novelty search demonstrated previously in other domains can also be leveraged in the co-evolution of body and brain of robots.

References

[1] J. C. Bongard and G. S. Hornby. Guarding against premature convergence while accelerating evolutionary search. In *GECCO '10: Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 111–118, New York, NY, USA, 2010. ACM.

[2] J. Doucette and M. I. Heywood. Novelty-based fitness: An evaluation under the santa fe trail. In *Genetic Programming, 13th European Conference, EuroGP 2010*, Lecture Notes in Computer Science, pages 50–61. Springer, 2010.

[3] D. E. Goldberg. *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison-Wesley Professional, January 1989.

[4] F. Gruau. Cellular encoding for interactive evolutionary robotics. In *Fourth European Conference on Artificial Life*, pages 368–377, Cambridge, MA, USA, 1997. MIT Press.

[5] P. Krcah. Towards efficient evolutionary design of autonomous robots. In *ICES '08: Proceedings of the 8th international conference on Evolvable Systems: From Biology to Hardware*, pages 153–164, Berlin, Heidelberg, 2008. Springer-Verlag.

[6] J. Lehman and K. O. Stanley. Exploiting open endedness to solve problems through the search for novelty. In *In Proceedings of the Eleventh International Conference on Artificial Life*, Cambridge, MA, 2008. MIT Press.

[7] J. Lehman and K. O. Stanley. Efficiently evolving programs through the search for novelty. In *GECCO '10: Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 837–844, New York, NY, USA, 2010. ACM.

[8] J.-B. Mouret and S. Doncieux. Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity. In *CEC'09: Proceedings of the Eleventh conference on Congress on Evolutionary Computation*, pages 1161–1168, Piscataway, NJ, USA, 2009. IEEE Press.

[9] K. Sims. Evolving virtual creatures. In *SIGGRAPH '94: Proceedings of the 21st annual conference on Computer graphics and interactive techniques*, pages 15–22, New York, NY, USA, 1994. ACM Press.