

Exploring Divergence in Soft Robot Evolution

Daniele Gravina
Institute of Digital Games
University of Malta
daniele.gravina@um.edu.mt

Antonios Liapis
Institute of Digital Games
University of Malta
antonios.liapis@um.edu.mt

Georgios N. Yannakakis
Institute of Digital Games
University of Malta
georgios.yannakakis@um.edu.mt

ABSTRACT

Divergent search is a recent trend in evolutionary computation that does not reward proximity to the objective of the problem it tries to solve. Traditional evolutionary algorithms tend to converge to a single good solution, using a fitness proportional to the quality of the problem's solution, while divergent algorithms aim to counter convergence by avoiding selection pressure towards the ultimate objective. This paper explores how a recent divergent algorithm, surprise search, can affect the evolution of soft robot morphologies, comparing the performance and the structure of the evolved robots.

CCS CONCEPTS

•Computing methodologies → Search methodologies; Genetic algorithms; Evolutionary robotics;

KEYWORDS

Surprise search; novelty search; divergent search; deception; fitness-based evolution; soft robots; CPPN; artificial life

ACM Reference format:

Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2017. Exploring Divergence in Soft Robot Evolution. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017*, 2 pages. DOI: <http://dx.doi.org/10.1145/3067695.3076072>

1 INTRODUCTION

Evolving virtual creatures is a popular testbed used to assess the potential creativity of evolutionary algorithms (EAs). The robots' many degrees of freedom and the difficulty of the task makes evolutionary computation ideal for tackling this problem. In the last years several environments have been proposed; while it has been proven that evolving interesting and efficient artificial creatures is feasible, these works still fall short when compared to the complexity found in nature. In [1], Cheney et al. propose to achieve this ambitious objective by evolving the robots' morphologies by means of different materials, creating "soft" robots.

While the environment and encoding directly affect the design space and the expressivity of evolved robots, a poorly designed reward system can limit the potential for creative discovery in the design space [6]. For example, local optima in the fitness function can strongly bias the search towards less interesting morphologies. By explicitly ignoring the objective, open-ended evolution

can instead overcome the limit of traditional fitness-based search. Divergent search is a recent paradigm that pushes for the intrinsic properties of the search by avoiding selection pressure towards the ultimate objective. In this paper, we show how a recent divergent algorithm, surprise search, is able to discover diverse and well-performing solutions in the search space of virtual creatures.

2 SURPRISE SEARCH

Built on the principles of open-ended evolution and inspired by cognitive science and computational creativity literature, the *surprise search* algorithm favours *unexpected* solutions for selection [2, 10], mimicking a self-surprise process. Surprise search uses a prediction model to identify behavioural patterns in previous generations and predict behaviours in the current generation; following that, it rewards behaviours that deviate from the predicted ones [10].

Surprise search consists of two key phases: the prediction phase and the deviation phase. In the first phase, the algorithm tries to predict the behaviours of the next generation, based on a number of past generations (h) and the locality of the behavioural information (k). These two pieces of information are used to build the prediction model m . Once the predicted behaviours p are found, the algorithm rewards the unexpectedness of the individuals based on the distance of the n closest predicted behaviours of the current generation. More details can be found in [2, 10].

3 SOFT ROBOT EVOLUTION

In [1], Cheney et al. propose to create soft robot morphologies via Compositional Pattern Producing Networks (CPPNs) [8], evolved using neuroevolution of augmenting topologies [9]. The evaluation is based on data collected through simulations run on VoxCad [4]. As in [1], soft robots in this testbed consist of four materials, two active (red and green) and two passive (cyan and blue). Green voxels expand and contract following a signal at a predefined frequency, while a counter-phase signal actuates the red voxels. Passive materials are not actuated but are deformed by nearby materials: cyan voxels are soft and blue voxels are stiffer. These voxels are placed on a 3D-lattice with a predefined resolution; the evolved morphologies are simulated via VoxCad and the resulting behaviour is used to compute the fitness of the evolved robot. As in [1], a CPPN is used to specify the material (if any) of each voxel. Each x, y, z coordinate of the cubic lattice is provided as input to the CPPN: its first output determines whether the voxel is empty, while the highest score of the last three outputs decides the material of that voxel.

4 EXPERIMENTS

Previous work has explored the effectiveness of divergent algorithms in soft robot locomotion tasks [7]. However, in this work we want to assess how divergence can affect the outcome of soft robot

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).

GECCO '17 Companion, Berlin, Germany

© 2017 Copyright held by the owner/author(s). 978-1-4503-4939-0/17/07...\$15.00

DOI: <http://dx.doi.org/10.1145/3067695.3076072>

evolution in respect to the performance and variety of morphologies evolved. Therefore we propose to analyse the differences between surprise search and two baselines: objective (i.e. fitness-based) search and novelty search.

In *objective search*, the objective is to evolve robots capable of moving as far as possible: the chosen performance metric is the Euclidean distance of the robot's centre of mass at the start and at the end of simulation time in body lengths [1, 7]. Objective search attempts to maximise this distance.

Novelty search [5] pushes for divergent behaviours by selecting individuals based on how different their behaviours are compared to the current population and an archive of past novel individuals. In every generation, individuals may be added to the archive if their novelty score is above a dynamic threshold. The novelty score of an individual is evaluated based on the average distance (d_n) of the n nearest neighbours in the current population and the novel archive. Novelty search uses the same parameters as in [7], with $n = 10$.

Finally, *surprise search* relies on a prediction model and a distance function (d_s). The surprise score is computed as the Euclidean distance between the individual's trajectory and the four closest predicted trajectories ($n = 4$). The predicted trajectories are computed by using linear regression of the sampled points of the previous two generations' trajectories. The local behaviours are computed via the k -means clustering algorithm, where $k = 15$, found empirically.

Divergent algorithms require a different behaviour characterization in order to compute the distance between individuals. As in [7], both algorithms use the two-dimensional trajectory of the soft robots while the distance characterization is the average of the Euclidean distance of the sampled points of two trajectories.

Reported results are collected from the 30 fittest individuals in 30 independent evolutionary runs. The simulation in VoxCad and the CPPN algorithm use the same parameters as in [7], while the lattice resolution is fixed to $5 \times 5 \times 5$.

4.1 Results

Observing the performance of the solutions in terms of the objective function, both surprise search and novelty search outperform objective search significantly: respectively they obtain a median of 10.77 ($MAD = 1.62$), 10.61 ($MAD = 1.21$) and 6.26 ($MAD = 0.77$) body lengths. On the other hand, to assess the diversity of the fittest individuals, Fig. 1 shows the relationship between the minimum (nearest neighbour) distance and corresponding fitness of each fittest individual per run. The distance used is the structural distance, i.e. the number of voxels that have different materials at the same position between two soft robots. Surprise search is able to evolve well-performing structures and at the same time keep a higher minimum distance than both baselines. Novelty, instead, tends to evolve very efficient and diverse outliers, but overall the fittest robots are less diverse: it obtains a median minimum distance of 22 ($MAD = 7.41$), while surprise and objective reach respectively have a median minimum distance of 35.5 ($MAD = 12.36$) and 33 ($MAD = 14.82$). Finally, Fig. 2 shows the most common structure evolved by each algorithm, obtained by running the k -medoids algorithm ($k = 5$) and selecting the medoids from the largest clusters. The most common structure evolved by objective search is a pyramid-shaped structure: based on the simulations, this structure

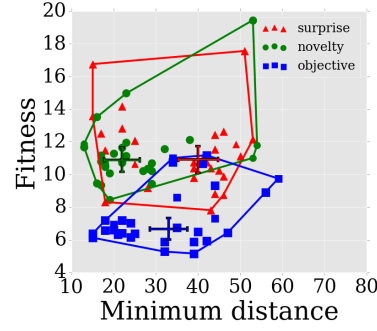


Figure 1: Scatter plots of fitness and structural minimum distance.




Figure 2: Most common structures.

shows a really stable behaviour, but at the same time is particularly slow, as it can only slither on the floor. On the other hand, the most common structure evolved by novelty and surprise is a structure composed of active materials only, which shows an unexpected “tumbleweed” behaviour that demonstrates to be really efficient in this setting due to their “hopping” locomotion strategy.

5 FUTURE WORK

We foresee several directions for future work, primarily exploring how different behaviour characterizations and different soft robot lattice sizes affect the performance of surprise search. The diversity of individuals within the same population should also be assessed, to compare convergence of the different methods. Finally it is interesting to compare novelty and surprise search, in terms of structural and behavioural diversity, with a hybrid of the two [3].

6 ACKNOWLEDGMENTS

This work has received funding from the FP7 Marie Curie CIG project AutoGameDesign (project no: 630665) and the Horizon 2020 project CrossCult (project no: 693150).

REFERENCES

- [1] Nick Cheney, Robert MacCurdy, Jeff Clune, and Hod Lipson. 2013. Unshackling evolution: evolving soft robots with multiple materials and a powerful generative encoding. In *Proceedings of the Genetic and Evolutionary Computation Conference*.
- [2] Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2016. Surprise Search: Beyond Objectives and Novelty. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM.
- [3] Daniele Gravina, Antonios Liapis, and Georgios N. Yannakakis. 2017. Coupling Novelty and Surprise for Evolutionary Divergence. In *Proceedings of the Genetic and Evolutionary Computation Conference*.
- [4] Jonathan Hiller and Hod Lipson. 2012. Dynamic simulation of soft heterogeneous objects. *arXiv preprint arXiv:1212.2845* (2012).
- [5] Joel Lehman and Kenneth O Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19, 2 (2011).
- [6] Joel Lehman and Kenneth O Stanley. 2011. Evolving a diversity of virtual creatures through novelty search and local competition. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 211–218.
- [7] Georgios Methenitis, Daniel Hennes, Dario Izzo, and Arnoud Visser. 2015. Novelty search for soft robotic space exploration. In *Proceedings of the Genetic and Evolutionary Computation Conference*. ACM, 193–200.
- [8] Kenneth O Stanley. 2007. Compositional pattern producing networks: A novel abstraction of development. *Genetic programming and evolvable machines* 8, 2 (2007), 131–162.
- [9] Kenneth O Stanley and Risto Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary computation* 10, 2 (2002), 99–127.
- [10] Georgios N. Yannakakis and Antonios Liapis. 2016. Searching for Surprise. In *Proceedings of the International Conference on Computational Creativity*.