Evolving Diverse Collective Behaviors Independent of Swarm Density

Payam Zahadat
Artificial Life Lab
Department of Zoology
Karl-Franzens University Graz
Graz, Austria
payam.zahadat@uni-graz.at

Heiko Hamann
Department of Computer Science
University of Paderborn
Paderborn, Germany
heiko.hamann@uni-paderborn.de

Thomas Schmickl
Artificial Life Lab
Department of Zoology
Karl-Franzens University Graz
Graz, Austria
thomas.schmickl@uni-graz.at

There are multiple different ways of implementing artificial evolution of collective behaviors. Besides a classical offline evolution approach, there is, for example, the option of environment-driven distributed evolutionary adaptation in the form of an artificial ecology [2] and more generally there is the approach of embodied evolution [1, 3, 6]. Another recently reported approach is the application of novelty search to swarm robotics [5]. In the following, we report an extension of the approach of [7]. The underlying concept is an information-theoretic analogon to thermodynamic (Helmholtz) free energy [8]. The assumption is that the brain is permanently trying to predict future perceptions and that minimizing the prediction error is basically inherent to brains. This is defined by the 'free-energy principle' of [4]. The struggle for prediction success requires a complementary force that represents curiosity and exploration. In this abstract we present an extended method called *diverse-prediction* that rewards not only for correct predictions but also for each visited sensory state. This proves to be a better approach compared to the method *prediction* that was reported before [7].

1. SETUP

As in the preliminary work [7], here we also evolve pairs of artificial neural networks (ANN). The prediction network, see Fig. 1(d), predicts future sensor input and the action network, see Fig. 1(c), outputs the agent's next action. We simulate a homogeneous swarm (all agents share the same genome) of $N = 20$ agents that move on a ring of circumference L, see Fig. 1(a). Note that there are two different concepts of populations, the population of simulated agents and the population of genomes. The agents have four, discrete sensors covering four regions of the agent's vicinity and output 1 for 'there is at least one neighbor' or 0 otherwise, see Fig. 1(b). The available actions are: move forward or invert the heading. We evolve pairs of ANN with a population of size 50 (2nd concept of population) for 75 generations in 200 independent runs for each tested setting. The *prediction* method rewards for good prediction of each sensor value using fitness function

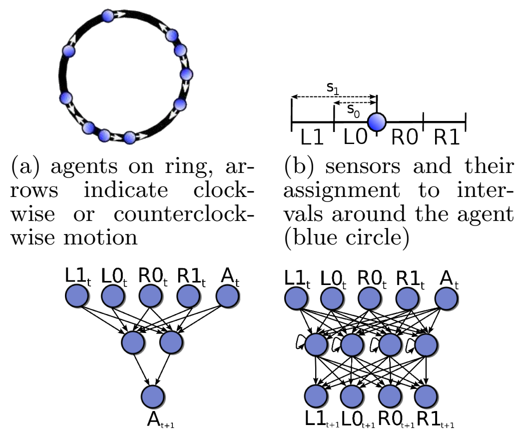

(a) agents on ring, arrows indicate clockwise or counterclockwise motion

(b) sensors and their assignment to intervals around the agent (blue circle)

(c) action network

(d) prediction network

Figure 1: Setup of the collective system, sensor setup, action network, and prediction network [7].

$$F_{\text{pred}} = \frac{1}{4NT} \sum_{t=0}^{T-1} \sum_{i=0}^{N-1} \sum_{j \in \{L1, L0, R0, R1\}} c_{i,j}(t), \qquad (1)$$

whereas N is the swarm size, T is the length of the evaluation in time steps, and $c_{i,j}(t) = 1$ if the prediction of the previous time step for sensor j of agent i matches the current value of sensor j, otherwise $c_{i,j}(t) = 0$. The fitness is averaged over 10 independent simulation runs for every evaluated genome.

The fitness function for *diverse-prediction* rewards for visiting more sensor states (combinations of sensor values) and making good predictions in those states. For that, we sum up the ratios between the number of correct predictions for each visited sensor state and the number of visits of that state over the whole swarm over time:

$$F_{\text{divPred}} = \frac{1}{2^4} \sum_{j \in V} \left(\sum_{i=0}^{N-1} \sum_{t=0}^{T-1} c_{i,j}(t) / \sum_{i=0}^{N-1} \sum_{t=0}^{T-1} n_{i,j}(t) \right), \quad (2)$$

whereas V is the set of sensor states which were visited at least once by an agent.

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

DOI: http://dx.doi.org/10.1145/2739482.2768492

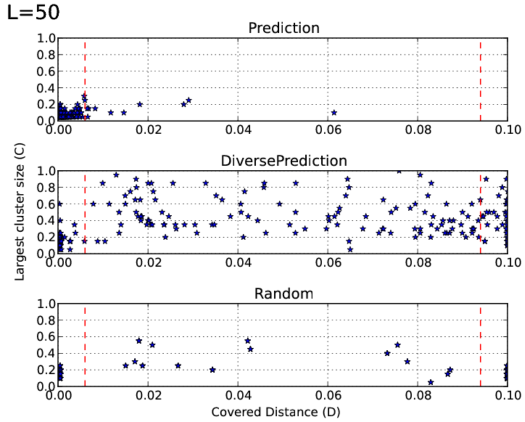

(a) 2-d projection of behavior space, evolved behaviors and random networks as control for ring circumference $L = 50$.

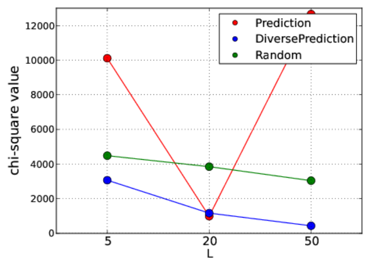

(b) Comparison 2-d projection to uniform distribution, chi-square test for $L \in \{5, 20, 50\}$ (low values are better).

Figure 2: Results of evolved behaviors for $L = 50$ and comparison between projection of behavior space and uniform distribution.

2. RESULTS

Generally a desired result is a well explored behavior space – similar, for example, to the approach of novelty search. We check the degree of exploration by investigating a 2-d projection of behavior space. One dimension is covered distance which is a measure of mobility of the agents in the swarm and is computed as the sum of the distances covered by the agents normalized by the considered time period and swarm size. Second dimension is the largest cluster size which is the maximum number of agents within sensor range normalized by swarm size over a period at the end of the evaluation.

We compare the distribution of the 2-d projection of the collective behaviors evolved by the *prediction* method, the *diverse-prediction* method, and a control population of random ANN. Fig. 2(a) shows results for a low agent density setup ($L = 50$). Behaviors evolved by the *prediction* method are biased towards low mobility of agents that stay apart from each other (dispersed). The control experiment with random networks mostly generates behaviors with very low mobility (moving back and forth) or in fewer cases agents move constantly in a fixed direction. In contrast, the *diverse-prediction* method evolves diverse behaviors as clearly seen in Fig. 2(a). The three methods were tested for 3 settings: high ($L = 5$), medium ($L = 20$), and low agent density ($L = 50$). For statistical comparison of the diversity of behaviors, the projected behavior space is divided into a 5×15 grid in each setup and Pearson's χ^2 test is used over the cells of the grid to compare the distributions of the behaviors with a uniform distribution (i.e., ideal case). We find that all the three distributions of *prediction*, *diverse-prediction*, and random are still far from a uniform distribution ($P > 0.99$). Fig. 2(b) shows the value of the χ^2 statistics for the three methods. The lower the value is, the closer is the tested distribution to a uniform distribution. We find that the *prediction* method performs well when the density is medium as reported before [7]. Otherwise it fails to generate diverse behaviors. The *diverse-prediction* method performs best and is independent of the agent density.

3. CONCLUSION

Our new method *diverse-prediction* implements a selection pressure to visit many different sensor states as well as to predict these states correctly. As a result, the evolved collective behaviors are more diverse and this diversity is less dependent on appropriate swarm densities. In future work, we plan to extend this approach and to investigate necessary conditions for the emergence of more complex behaviors also in environments that are closer to robotics.

4. ACKNOWLEDGMENTS

Acknowledgments. Supported by: EU-H2020 project 'florarobotica', no. 640959; EU-H2020 project 'subCULTron', no. 640967; EU-ICT project 'ASSISI_bf', no. 601074;

5. REFERENCES

[1] N. Bredeche, E. Haasdijk, and Á. E. Eiben. On-line, on-board evolution of robot controllers. In *9th International Conference on Artificial Evolution*, Oct. 2009.

[2] N. Bredeche, J.-M. Montanier, W. Liu, and A. F. Winfield. Environment-driven distributed evolutionary adaptation in a population of autonomous robotic agents. *Mathematical and Computer Modelling of Dynamical Systems*, 18(1):101–129, 2012.

[3] Á. E. Eiben, E. Haasdijk, and N. Bredeche. Embodied, on-line, on-board evolution for autonomous robotics. In P. Levi and S. Kernbach, editors, *Symbiotic Multi-Robot Organisms*, volume 7 of *Cognitive Systems Monographs*, pages 362–384. Springer, 2010.

[4] K. Friston. The free-energy principle: a unified brain theory? *Nature Reviews Neuroscience*, 11(2):127–138, 2010.

[5] J. Gomes, P. Urbano, and A. L. Christensen. Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*, 7(2-3):115–144, 2013.

[6] E. Haasdijk, N. Bredeche, and Á. E. Eiben. Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics. *PLoS ONE*, 9(6):e98466, 2014.

[7] H. Hamann. Evolution of collective behaviors by minimizing surprise. In H. Sayama, J. Rieffel, S. Risi, R. Doursat, and H. Lipson, editors, *14th Int. Conf. on the Synthesis and Simulation of Living Systems (ALIFE 2014)*, pages 344–351. MIT Press, 2014.

[8] H. von Helmholtz. *Handbuch der physiologischen Optik*. Ludwig Voss, Leipzig, Germany, 1867.