Dynamic Behavioral Diversity

Stephane Doncieux
ISIR, Université Pierre et Marie Curie-Paris 6,
CNRS UMR 7222
4 place Jussieu, F-75252, Paris Cedex 05,
France
doncieux@isir.upmc.fr

Jean-Baptiste Mouret
ISIR, Université Pierre et Marie Curie-Paris 6,
CNRS UMR 7222
4 place Jussieu, F-75252, Paris Cedex 05,
France
mouret@isir.upmc.fr

ABSTRACT

Behavioral Diversity (BD) aims at improving the evolution of robots by fostering exploration on the basis of their behavior, whereas evolutionary algorithms typically consider the diversity on a genotypic level. Several Behavioral Similarity Measures (BSM), the key component to improve behavioral diversity, have been investigated in the litterature. Current benchmarks show that (1) most tested BSM improve the final performance, (2) they do not lead to the same improvements and, (3) it is hard to predict a priori which BSM will work the best. Instead of trying to find the best BSM, a different approach is proposed here: assuming that several BSM are available, we propose to randomly switch between then each K generations (e.g. K=20). This new approach is tested on a ball collecting task. Results show that better fitness values are obtained with the random switch approach than with any single measure. In effect, the present contribution show that it is possible use behavioral diversity while avoiding to choose between behavioral similarity measures.

Categories and Subject Descriptors: I.2.6 [Artificial intelligence]: Learning, Robotics

General Terms: Algorithms.

Keywords: Evolutionary Algorithm, Evolutionary Robotics, Neuro-Evolution, Behavior, Exploration, Diversity, Novelty.

In Evolutionary Robotics (ER), it has been widely observed that many genotypes and phenotypes lead to similar behaviors, whereas a very small difference in a genotype can substantially modify the behavior of the robot [8]. Since ER is ultimately seeking behaviors, explicitly encouraging diversity in the behavioral space revealed to drastically improve convergence of ER experiments in multiple contexts [7, 6, 8]. Behavioral diversity methods foster diversity by rewarding individuals whose behavior is different from the rest of the population. The similarity is computed using behavior similarity measures (BSM) that can be problem-specific [6, 5, 8, 9] or designed to be valid for a large class of robotic tasks [2, 10, 9, 3, 8]. Empirical comparisons between BSM have been made in several papers [3, 8], but, although different BSM yielded different results, no measure has proved to be "the one" for every contexts. The present work investigates another approach: if the best BSM is not known, *is it possible to efficiently exploit all the different BSM an experimenter can think of without having to choose among them?*

GECCO'12 Companion, July 7–11, 2012, Philadelphia, PA, USA.
ACM 978-1-4503-1178-6/12/07.

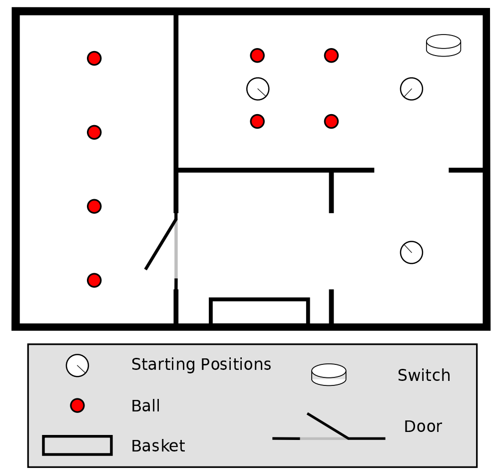

Figure 1: Overview of the arena and of the robot. The goal of the experiment is to place as many balls as possible into the basket.

The general framework of this approach is to add an *helper* objective besides objectives rewarding the performance on the task [4]. The helper objective measures the mean distance of the evaluated solution to the rest of the population on the basis of *behavior* comparisons [7, 6, 8]. Instead of using a single BSM during the whole run, the proposed method consists in changing the BSM used to compute the behavioral diversity along the generations. For each generation, a single BSM is used, but at a given generation period K, a BSM is randomly selected from a set of M BSM and kept for the next K generations.

The selection algorithm is NSGA-II [1] with a population of 200. The considered task is the collect ball task described in [9] (figure 1). The robot is controlled by a neural network with twelve inputs and three outputs. Both structure and parameters of the neural network are evolved with DNN [8]. DNN is a simple direct encoding with no crossover. Mutations can add or remove neurons and connections and change connection weights.

The fitness function is the number of collected balls divided by the maximum number of balls available. The performance is evaluated in three different contexts (with differ-

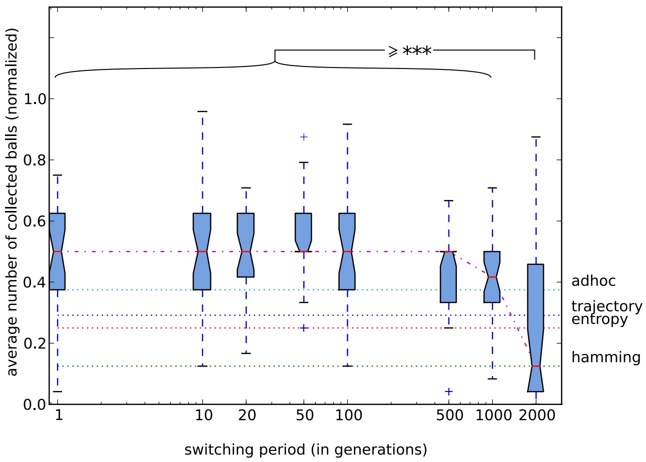

Figure 2: Influence of the switching period between BSM. Experiments with 4 BSM (adhoc, hamming, trajectory, entropy). Medians of 30 runs after 4000 generations. The medians of the control runs have been reported on the plot to ease comparisons. Only statistical difference between $K = 2000$ and the other treatments are shown. '*' means $p - value < 10^{-3}$**

ent starting positions, as shown on figure 1). The maximum number of balls is then 24.

Four different BSM have been considered: (1) adhoc [9]: the behavior of a robot is described by the final position of the balls;(2) hamming [8]: the value of each sensor and effector is discretized and memorized during 4000 time steps. It results in a binary vector of 60000 bits ($4000 \times (12$ sensors $+$ 3 effectors)). The BSM is then the Hamming distance that measures the number of bits that differ between the two vectors; (3) trajectory [9]: the discretized position of the robot is recorded each 50 time-steps. This vector is the behavior descriptor;(4) entropy: the behavior descriptor is the vector of the entropies of each sensors and effectors.

Every K generations, one of the four BSM is randomly selected and used for the next K generations. Treatments with $K \in \{1, 10, 20, 50, 100\}$ give better results than any of the single BSM treatments (statistically significant, Mann-Whitney U test, $p - value < 0.01$, see figure 2)[1]. The results of treatments with $K \in \{1, 10, 20, 50, 100\}$ are not statistically different between each other (Mann-Whitney U test, $p - value > 0.05$) and they are all different from those with $K \in \{1000, 2000\}$. As the runs last for 4000 generations, it is not surprising that above a period of 500, the performance decreases as the number of BSM switches will become negligible.

The proposed dynamic behavioral diversity method significantly improves behavioral diversity method. It requires to know several BSM, which is not a problem in practice, and the performance is robust to K, the only new parameter introduced, as long as enough BSM switches are performed during a run.

[1]Treatments with $K \in \{500, 1000\}$ are also significantly better than single BSM treatments, except for adhoc BSM.

Acknowledgments

This project was funded by the ANR EvoNeuro project, ANR-09-EMER-005-01.

1. REFERENCES

[1] K. Deb, S. Agrawal, A. Pratab, and T. Meyarivan. A fast elitist non-dominated sorting genetic algorithm for multi-objective optimization: Nsga-ii. In *Proceedings of the Parallel Problem Solving from Nature VI Conference*, pages 849–858, Paris, France, 2000. Springer. Lecture Notes in Computer Science No. 1917.

[2] S. Doncieux and J.-B. Mouret. Behavioral diversity measures for Evolutionary Robotics. In *IEEE Congress on Evolutionary Computation, 2010 (CEC 2010)*, pages 1303–1310, 2010.

[3] F. J. Gomez. Sustaining diversity using behavioral information distance. In *GECCO'09: Proceedings of the 11th annual conference on Genetic and Evolutionary Computation*, pages 113–120. ACM, 2009.

[4] J. D. Knowles, R. A. Watson, and D. W. Corne. Reducing Local Optima in Single-Objective Problems by Multi-objectivization. *First International Conference on Evolutionary Multi-Criterion Optimization*, 1993:268–282, 2001.

[5] J. Lehman and K. O. Stanley. Abandoning objectives: evolution through the search for novelty alone. *Evolutionary computation*, 19(2):189–223, Jan. 2011.

[6] J.-B. Mouret and S. Doncieux. Overcoming the bootstrap problem in evolutionary robotics using behavioral diversity. In *IEEE Congress on Evolutionary Computation, 2009 (CEC 2009)*, pages 1161–1168, 2009.

[7] J.-B. Mouret and S. Doncieux. Using Behavioral Exploration Objectives to Solve Deceptive Problems in Neuro-evolution. In *GECCO'09: Proceedings of the 11th annual conference on Genetic and evolutionary computation*, pages 627–634. ACM, 2009.

[8] J.-B. Mouret and S. Doncieux. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary Computation*, 20(1):91–133, 2012.

[9] C. Ollion and S. Doncieux. Why and How to Measure Exploration in Behavioral Space. In *GECCO '11: Proceedings of the 13 th annual conference on Genetic and Evolutionary Computation*, pages 267–294, 2011.

[10] L. Trujillo, G. Olague, E. Lutton, F. Fernández de Vega, L. Dozal, and E. Clemente. Speciation in Behavioral Space for Evolutionary Robotics. *Journal of Intelligent & Robotic Systems*, 64(3):323–351, Jan. 2011.