Benefits of Proportionate Selection in Embodied Evolution: a Case Study with Behavioural Specialization

Nicolas Bredeche* Sorbonne Universités UPMC Univ Paris 06, CNRS, ISIR Paris, France F-75005 nicolas.bredeche@upmc.fr Jean-Marc Montanier Softbank Robotics Europe Paris, France jmmontanier@softbankrobotics.com Simon Carrignon Barcelona Supercomputing Center (BSC) Barcelona, Spain simon.carrignon@bsc.es

ABSTRACT

We, as well as others, have already shown in previous works that reproductive isolation and a large population size are critical to achieve behavioral specialization in embodied evolutionary robotics. Here, we extend our previous work from [3] by experimentally demonstrating that fitness-proportionate and rank-based selection operators largely outperform other selection operators when it comes to evolving behavioral specialization.

CCS CONCEPTS

•Computer systems organization \rightarrow Robotics; •Computing methodologies \rightarrow Artificial intelligence;

KEYWORDS

evolutionary robotics, embodied evolution, swarm robotics

ACM Reference format:

Nicolas Bredeche, Jean-Marc Montanier, and Simon Carrignon. 2017. Benefits of Proportionate Selection in Embodied Evolution: a Case Study with Behavioural Specialization. In *Proceedings of GECCO '17 Companion, Berlin, Germany, July 15-19, 2017, 2 pages.*

DOI: http://dx.doi.org/10.1145/3067695.3082551

1 INTRODUCTION

Embodied evolution (EE), as a subfield of evolutionary robotics, is dedicated to the design of distributed on-line learning algorithms for swarm/collective robotics [2, 4]. One major challenge is posed by the acquisition of behavioral specialization, i.e. where two (or more) sub-parts of a group of robots can evolve different decisionmaking strategies when required. As shown in previous works, this is far from trivial and may only be achieved under very specific conditions.

In our own previous work [3], we showed that *reproductive isolation* plays a major role to help bootstrap and maintain more than one genotypic signature in the population.

GECCO '17 Companion, Berlin, Germany

@ 2017 Copyright held by the owner/author (s). Publication rights licensed to ACM. 978-1-4503-4939-0/17/07... \$15.00

DOI: http://dx.doi.org/10.1145/3067695.3082551

In this paper, we extend our previous work to consider proportionate selection operators, either based on absolute fitness values (i.e. fitness proportionate selection) or relative fitness values (i.e. rank-based selection). While other selection methods are now generally preferred by the community, proportionate selection retains an interesting feature when it comes to embodied evolution: by design, these selection methods are able to maintain polymorphism (i.e. population with more than one family of genotypes) by sampling from the whole population rather than from a truncated part.

We claim, and show in the following, that proportionate selection methods offer *very* efficient off-the-shelf solutions to achieve behavioral specialization in embodied evolutionary robotics.

2 METHOD

We use the *exact* same method we previously described in [3]¹. We provide a short summary here.Khepera-like robots are placed in a round environment where two resources can be found. Each robot may forage energy from one particular resource, but both resources must be foraged to allow the whole population to survive (i.e. each resource can only feed half the population). Foraging energy from resources serves two purposes: (1) to compute the fitness value of the robot and (2) to enable the robot to move (i.e. a robot is removed from the experiment if its energy level drops to zero).

Each robot is driven by a simple perceptron (hyperbolic tangent activation function, no hidden layer) that maps sensory inputs to two motor output neurons (left and right motor speed). Perceptron's weights are evolved, as well as a particular gene, termed g_{skill} , that defines the capacity to synthesize energy from each resource (i.e. $g_{skill} \in [-1, +1]$, with each extreme value defining a particular ability to synthesize from one of the two resources). The rationale is that to forage energy from one resource, a robot must both stand in the area where the resource is available and to be able to synthesize energy from it (which depends on the value of g_{skill}).

The evolutionary algorithms used are simplistic versions of embodied evolution, with exchange of genetic material occuring in an on-line, distributed fashion among robots which are close enough to one another. A comparaison between several selection operators was performed in our original work, including random selection (as implemented in mEDEA [1]) and performance-based selection (tournament and elitist selection operators).

We devise four experimental simulation setups, each defined by the population size considered (either 200 or 500 robots) and the position of the resources with respect to one another (the *SeperateEnv* environment, where resources are found in different areas, and the

^{*}Corresponding author

Permission to make digital or hard copies of all or part 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 components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

 $^{^{1}}http://journal.frontiersin.org/article/10.3389/frobt.2016.00038/full$

more challenging *ColocateEnv* environment, where resources are co-located). In all setups, evolutionary adaptation is expected to converge towards an equilibrium between two families of genotypes, each specialized on one resource, whether reproductive isolation is possible or not.

In this paper, we implement fitness-proportionate and rankbased selection operators. To some extent, proportionate selection operators share with random selection the propensity to sample from the whole population. They also share with other performancebased selection operators the incentive to bias selection towards better performance w.r.t. the objective function.

3 RESULTS

Table 1 (resp. Table 2) gives results with 200 (resp. 500) robots per population, for both environments. Each run falls in one of three categories, whether evolutionary adaptation completely failed (*extinct*), succeeded partly (*one group*, which implies only half the population was able to survive) or succeeded completely (*two groups*, meaning the whole population survived, with each group foraging energy from a particular resource).

		# runs		
selection	environment	2 groups	1 group	extinct
random	separateEnv	18	32	0
tournament-5 tournament-20	separateEnv separateEnv	8 3	37 35	5 12
elitist	separateEnv	1	31	18
fitness-prop rank-prop	separateEnv separateEnv	28 13	22 37	0 0
random	collocateEnv	2	42	6
tournament-5 tournament-20	collocateEnv collocateEnv	1 0	33 16	16 34
elitist	collocateEnv	0	18	32
fitness-prop rank-prop	collocateEnv collocateEnv	1 4	44 46	5 0

Table 1: Classification of the outcome of runs with a population of 200 robots. Classes are determined using the value of the skill's gene. 50 runs per experiment. Results in bold are new, other results are taken from [3].

Results clearly show that proportionate selection methods retain the benefits of pure environment-driven selection (cf. *random* selection) *and* benefits from optimizing a user-defined objective function in even the most difficult *collocateEnv* environment.

This is particularly prominent when the population size is large as selection pressure (resp. genetic drift) is weaker (resp. stronger) in smaller populations. With a large population, proportionate selection methods outperform other methods by a large margin in the more difficult *ColocateEnv* environment, i.e. when there is no reproductive isolation to help evolve separate species in separate locations.

		# runs		
selection	environment	2 groups	1 group	extinct
random	separateEnv	48	2	0
tournament-5	separateEnv	38	12	0
tournament-20	separateEnv	27	19	4
elitist	separateEnv	11	35	4
fitness-prop	separateEnv	49	1	0
rank-prop	separateEnv	44	6	0
random	collocateEnv	1	43	6
tournament-5	collocateEnv	4	35	11
tournament-20	collocateEnv	0	26	24
elitist	collocateEnv	0	30	20
fitness-prop	collocateEnv	9	37	4
rank-prop	collocateEnv	19	31	0

Table 2: Classification of the outcome of runs with a population of 500 robots. Classes are determined using the value of the skill's gene. 50 runs per experiment. All results are new.

The **take-home message** of this paper is thus the following: proportionate selection operators are very relevant when it comes to evolve behavioural specialization, thanks to a population-wide sampling method and performance-biased selection. This message, of course, has to be considered in conjunction with previous results [3]: reproductive isolation and a large population size still play important roles in the outcome of evolutionary adaptation.

An open issue is that there is no clear distinction as to what is best between fitness-proportionate and rank-based selection methods as each excels in specific setups. Given the only difference concerns the weight given to each individual using either absolute or relative fitness values, there may be an optimal way to distribute the selection probability over the population, e.g. using a particular of non-linear rank-based fitness assignement method. But the shape of this particular function may well depend on the task and the environment at hand.

ACKNOWLEDGMENT

This work is supported by the European Union's Horizon 2020 research and innovation programme under grant agreement No.640891.

REFERENCES

- Nicolas Bredeche and Jean-Marc Montanier. Environment-driven embodied evolution in a population of autonomous agents. *Parallel Problem Solving from Nature*, *PPSN XI*, pages 290–299, 2010.
- [2] Stephane Doncieux, Nicolas Bredeche, Jean-Baptiste Mouret, and Agoston E. (Gusz) Eiben. Evolutionary robotics: What, why, and where to. Frontiers in Robotics and AI, 2:4, 2015.
- [3] Jean-Marc Montanier, Simon Carrignon, and Nicolas Bredeche. Behavioral specialization in embodied evolutionary robotics: Why so difficult? Frontiers in Robotics and AI, 3:38, 2016.
- [4] Richard A Watson, Sevan G Ficici, and Jordan B Pollack. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1–18, 2002.