

Promoting Reproductive Isolation Through Diversity in On-line Collective Robotics

Amine Boumaza

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France

amine.boumaza@loria.fr

ABSTRACT

We present a behavioral diversity selection scheme that favors reproductive isolation to promote the learning of multiple task in on-line embodied evolutionary robotics (EER). The scheme estimates the behavior of the controllers without the need to access the agent experience, respecting thus the online, distributed properties EER. Reproductive isolation is assessed through coalescence trees and task specialization is tested on a concurrent foraging setting.

CCS CONCEPTS

• Computing methodologies → Evolutionary robotics;

KEYWORDS

Swarm Robotics, On-line Embodied Evolution, Diversity.

ACM Reference Format:

Amine Boumaza. 2021. Promoting Reproductive Isolation Through Diversity in On-line Collective Robotics. In *2021 Genetic and Evolutionary Computation Conference Companion (GECCO '21 Companion)*, July 10–14, 2021, Lille, France. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3449726.3459557>

1 MOTIVATIONS

It has been shown that evolving specialized behaviors in Embodied Evolutionary Robotics (EER) requires reproductive isolation, and tailored selection operators [2, 6]. The fact that agents spread their genetic material without any barriers does not help the conservation of specialized skills. One way to limit mating encounters and the preservation of learn skills is through geographical isolation (*Allopatric speciation*) where contact and genetic exchange between agent is limited. If reproductive isolation is a requirement to behavior specialization, in this study we ask the following question: can this isolation be enforced in other ways than geographically? Indeed, there exists in nature some evidence that behavioral separation can be promoted through *sympatric speciation*; situations where multiple species share a single environment and exchange genes but reproductive barriers maintain species distinct [5].

In this study, we propose to explore this idea for promoting isolation. Agents share the same geographical region, and the addition of tailored selection methods ensures the isolation needed for specialized skills to be preserved. Our goal is to verify the following hypothesis: *promoting diversity favors reproduction isolation, and*

thus favors the emergence of specialization. To that end, we propose a selection scheme that promotes diversity by selecting solutions based on their originality in the behavioral space. The selection scheme respects the distributed nature of EER, it operates locally on the agents and uses generic behavioral descriptors measure.

2 METHODS

We use a modified version of minimally Environment-driven Distributed EA (mEDEA) [1] to which we add a selection operator. We considered a concurrent foraging task in which agents have to collect items randomly placed in the environment. There exist two types of items (red and black) and agents are rewarded one unit of fitness every time they collect one item regardless of its type. When picked, an item of the same type is randomly placed in the environment. This task has been extensively studied in the context of task specialization and of diversity preservation [2, 3]. Collecting one type of items or the other is considered a distinct task, because agent perceive the types of items through different sensors.

The selection scheme we propose selects solutions at the agents level by choosing the genome that has the most different behavior from other genomes in the agent's list. It can be considered as an instance of Novelty Search [4], without the archive. Since agents live specific experiences depending on their encounters (obstacles, other agents, etc.), it is difficult to generalize a behavior measurement outside from the situation that it was measured in. We propose to estimate the behavior of the genomes regardless of the experience of the agents the originated from. Since the goal is to identify the most "different" behavior in the selection pool, we propose to simulate behaviors and measure their difference on random sensory data. The difference in the outputs will be the basis for the selection criterion.

To be more specific, let us consider a typical agent of the swarm, let it be a . We note $L^a = \{x_1 \dots x_\mu\}$ the set of genomes collected by agent a during its lifetime. The selection scheme aims at selecting one genome $\bar{x} \in L^a$ that has the most different behavior to replace the current active genome of a . At each selection step, agent a generates ι_k with $k = 1, \dots, K$ "fake" inputs s.a. $\iota_k = \text{Uniform}(0,1)^{|s|}$, where $|s|$ is number of sensors and $[0, 1]$ is the range of values they can take. All neuro-controllers in L^a are presented the same "fake" sensory data, as if they experienced the same situation. For each input ι_k , we compute the "would be" output $o_{i,k} = C_i(x_i, \iota_k)$ of each controller in L^a , where $C(x, \iota)$ computes the output of the neuro-controller with weights x on input ι . The behavior of genome x_i is then defined as the concatenation of its controller outputs from all inputs or: $B(x_i) = \{o_{i,1}, o_{i,2}, \dots, o_{i,K}\}$, which represent a vector of $K \times |e|$ elements, $|e|$ being the number of effectors each agent has. The agent then selects the genome with the most different behavior defined as the behavior that has the largest average distance to all

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 '21 Companion, July 10–14, 2021, Lille, France

© 2021 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-8351-6/21/07...\$15.00

<https://doi.org/10.1145/3449726.3459557>

others or: $\bar{x} = \arg \max_{x \in L^a} \left(\frac{1}{|L^a|} \sum_{y \in L^a} \|B(x) - B(y)\| \right)$. To assess if the diversity selection scheme promotes isolation, we compare a purely objective selection method (number of items) with the above diversity selection method (novelty only).

3 RESULTS AND DISCUSSIONS

One way to measure if reproductive isolation occurs is to inspect the population dynamics. We take a gene perspective where at each generation, we record the dissemination of the genes between agents and track their descendants. We then extract the corresponding *coalescence tree* which synthesize the genealogy of the last surviving genomes. This tree is constructed from the last generation up to the first where, at each step lineages are merged whenever two or more genomes share the same parent (a coalescence event).

The most recent common ancestor (MRCA) is the genome that all the surviving genomes (those that form the last generation) descend from and its height (Time to MRCA) in the tree is a good indicator of genetic closeness, since it is related to the number of generations and thus to the number of mutations between the MRCA and a current population. The larger it is, the farther apart are the genomes in the population, the smaller it is, the more its descendants are related. We measure also, the length of the branches of the coalescence tree as it is related to of how much history genomes share. Genome would share the least history if they come from a common ancestor far back in time and then evolved along distinct lineages *i.e.* higher reproductive isolation.




Figure 1: TMRCA (left) and lineage branch length (right).

To assess if reproductive isolation is favored by diversity, we compare both instances on the basis of the coalescence measures we described above. If we inspect the left of Figure 1, it is clear that the objective based instance has the lowest TMRCA. On the other hand, in the diversity selection instance, there was no MRCA! This is indicated by the flat box-plot at the value 500 (the number of generation). This reflects the fact that in this case there were distinct lineages that coexisted since the start of the simulation that never coalesced. In parallel, if we look at the branch length, we notice that the higher values are those for diversity selection indicating that the end-survivors of those instances share the least genetic history. The difference is statistically significant (p -value $< 10^{-4}$, 64 runs). To illustrate the above results Figure 2 shows coalescence trees from typical runs from both instances.

To measure if both behaviors (specializing on one item type) are present in the swarm, we measure the histogram of the ratio of collected black and red items by each agent. When all agents collect both items indistinguishably, this histogram is centered around 0.5. On the other hand, when agents favor one color over the other




Figure 2: Coalescence trees of typical runs.

the histogram is spread toward the edges. Figure 3, clearly shows that the diversity selection scheme create diverse behaviors on all the spectrum. We notice also that there is a number of agents that collect only items of one type (on the left or the right edge). On the other hand objective based selection create the least diversity most agents collect equally red and black items. The difference is statistically significant (p -value $< 10^{-4}$, 64 runs).




Figure 3: Histogram of red/black across the population.

We proposed the idea of favoring reproductive isolation without geographical constraints through the use of a diversity selection scheme and tested the hypothesis on a concurrent foraging task. The results suggest that it is the case. Furthermore, we introduced a procedure that allows estimating behavioral distance on board the agent, and we argue that it provides sound measurements of the behaviors independent of the agents' experiences, which in the case of EER is not accessible.

REFERENCES

- [1] Nicolas Bredeche and Jean-Marc Montanier. 2010. Environment-driven Embodied Evolution in a Population of Autonomous Agents. In *PPSN 2010*. Krakow, 290–299.
- [2] Evert Haasdijk, Nicolas Bredeche, and A. E. Eiben. 2014. Combining Environment-Driven Adaptation and Task-Driven Optimisation in Evolutionary Robotics. *PLoS ONE* 9, 6 (June 2014), e98466. <https://doi.org/10.1371/journal.pone.0098466>
- [3] Emma Hart, Andreas S. W. Steyven, and Ben Paechter. 2018. Evolution of a functionally diverse swarm via a novel decentralised quality-diversity algorithm. In *Proceedings of the Genetic and Evolutionary Computation Conference, GECCO 2018, Kyoto, Japan, July 15-19, 2018*, Hernán E. Aguirre and Keiki Takadama (Eds.). ACM, 101–108. <https://doi.org/10.1145/3205455.3205481>
- [4] Joel Lehman and Kenneth O. Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation* 19, 2 (June 2011), 189–223. https://doi.org/10.1162/EVCO.a_00025
- [5] John Maynard Smith. 1966. Sympatric Speciation. *The American Naturalist* 100, 916 (1966).
- [6] Jean-Marc Montanier, Simon Carrignon, and Nicolas Bredeche. 2016. Behavioral Specialization in Embodied Evolutionary Robotics: Why So Difficult? *Frontiers in Robotics and AI* 3 (2016), 38.