



Exploring Affordances in Open-Ended Online Collective Robotics

Amine Boumaza

► To cite this version:

Amine Boumaza. Exploring Affordances in Open-Ended Online Collective Robotics. Proceedings of the 2024 Artificial Life Conference, International Society for Artificial Life, Jul 2024, Copenhagen, Denmark. pp.317-325, 10.1162/isal_a_00834 . hal-04674518

HAL Id: hal-04674518

<https://inria.hal.science/hal-04674518v1>

Submitted on 30 Sep 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Exploring Affordances in Open-Ended Online Collective Robotics

Amine Boumaza

Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France
`amine.boumaza@loria.fr`

Abstract

Online Embodied Evolution is a distributed learning method for collective heterogeneous robotic swarms, in which evolution is carried out in a decentralized manner. These algorithms are well suited for open-ended evolution where the goal is to evolve efficient survival strategies in a priori unknown and changing environments. In this work, we are interested in analyzing the features of the environment that favors evolution of certain behaviors. We hypothesize that certain types of perceivable artifacts in the environment provide affordances that exert a selection pressure toward the fixation of certain behaviors. We test this hypothesis in simulation in two different settings: an open-ended environment without any selection pressure and compare it to one in which originality of behaviors is enforced. The experiments show that in the open-ended environment agents evolved a certain type of behavior that exploited affordances, hinting at the existence of a selection pressure from the affordances as it is claimed from ecological psychology in the case of natural systems. We present different objective and subjective measures to support this claim.

1 Introduction

Open-ended evolution has been an important topic of research in recent decades. One of its goals, beyond modeling and understanding the conditions to create artificial life, is to design algorithms that allow swarms of robots to self-adapt to their environments with little or no human intervention.

In this context, embodied evolutionary robotics (EER), a subfield of evolutionary robotics (ER) (Nolfi and Floreano, 2000), aims to design collective behaviors for swarms of agents evolving online (Ficci et al., 1999). In a nutshell, these are algorithms in which evolution is decentralized, where each agent, typically a mobile robot, runs an EA on board and exchanges genetic material with other agents when they meet, using local communication channels. Selection and variation are performed locally on the agent, in contrast to traditional ER, where usually a central process governs evolution. These algorithms can best be described as a metaphor of the “selfish gene” of Richard Dawkins; genes compete for survival using agents as vehicles to spread. They have been studied in different contexts and are considered good candidates to achieve open-ended evolution (Bredeche et al., 2018).

EER algorithms share many properties with their ER siblings. They however have some differences. Since the number of agents is limited, solutions must create as many offspring as possible to survive. This pushes evolution to create behaviors that favor the spread of genetic material; behaviors that should maximize mating encounters in a physically distributed swarm and therefore overcome any pressure from the environment that could hinder this ability. In a sense EER can be viewed as a constant search for an equilibrium between evolving fit behaviors that are robust to changes in the environment.

Past experiments (Bredeche et al., 2012) in open-ended evolution with EER, showed that it is possible to evolve interesting behaviors without the need for a fitness. For example, learning to navigate while avoiding obstacles, learning to share altruistically resources, or to specialize and coordinate. In these situations, the human designer did not specify a task to be optimized but rather, in a quest to understand minimal conditions, designed specific environments that exerted the needed pressures to make behaviors evolve. The behaviors emerge from complex dynamics and interaction between the agents and their environment, and between different selection pressure forces that drive evolution.

Analyzing and understanding these dynamics is challenging, and many questions remain open. The challenges are inherent to the distributed and online nature of the algorithms and to the intricate interactions between the agents and the environment. In this paper, we are interested in this latter issue by studying the features of the environment that drive evolution or at least shape its result in a specific way. We revisit a specific setting in open-ended environments that presented interesting results, where a swarm of robots evolved unexpected behaviors (Bredeche, 2014), in an attempt to understand the elements that help the evolution of these behaviors. We argue that behaviors are favored by the existence of certain features in the environment, otherwise known as *affordances* (Gibson, 2015). And we rely on a key hypothesis from ecological psychology, one which states that behaviors are regulated by affordances. According to Edward Reed (Reed, 1996):

Affordances and only the relative availability (or nonavailability) of affordances create selection pressure on the behavior of individual organisms; hence, behavior is regulated with respect to the affordances of the environment for a given animal. [...] Affordances in the animal's niche are not relations; they are instead resources—in this case, resources for obtaining value from the environment through regulated action. (Reed, 1996, 18)

According to Reed, the availability of affordances creates selection pressure on the behavior of the agents. Affordances do not cause behavior but make it possible; they are opportunities for behavior. He further adds: “*The ability to encounter an affordance requires a perceptual system attuned to the use of information enabling that affordance to regulate action.*”

We present few results with a swarm of robots on a simple open-ended survival scenario. The simulation environment contains specific perceivable objects (Bredeche, 2014). We conduct two sets of experiments: 1) one in which selection at the level of the agent is random without any selection pressure towards performance and 2) a second one where selection favors originality with a selection pressure towards novelty (Lehman and Stanley, 2011). Our goal is to test if in the absence of any selection pressure, except from the one exerted by the environment, affordances create the needed selection pressure for certain types of behaviors to fixate in the population. The second setting, considered as a control experiment, should break the fixation of any behavior, as it always seeks for novelty.

We will show that in the case of the random setting, robots learn to exploit the features of the environment and that evolution designed behaviors that tend to be similar. These behaviors are sufficient to overcome environmental pressures and survive. In the second setting, agents evolve other behaviors and manage to also survive, they however, did not fully exploit the features of the environment. We start by describing the evolutionary algorithm and our experimental procedure along with the tested scenarios, and finally conclude with a discussion of the results.

2 Methods

We consider a variant of the Minimally Environment-driven Distributed EA (mEDEA) (Bredeche and Montanier, 2010). This algorithm, as its name indicates, is an open-ended evolutionary algorithm where the goal is survival. Genomes compete at the level of the population, and those that do not maximize the number of mating opportunities become extinct, and only those that do so, survive. In contrast, we confront this open-ended version of the algorithm to one that favors novelty. Both these instances are described below.

2.1 The Embodied Evolutionary Algorithm

Algorithm 1 describes the variant of mEDEA we use. It considers a swarm of λ mobile agents a^j with $j = 1, \dots, \lambda$ each of which executes a neuro-controller whose parameters are \mathbf{x}^j (the active genome). Each agent keeps a list L^j , which is initially empty, in which it stores other genomes received from other agents.

Algorithm 1: mEDEA

```
1 for  $1 \leq j \leq \lambda$  in parallel do
2    $\mathbf{x}^j \leftarrow \text{random}()$ 
3    $a^j$  is active
4    $\sigma^j \leftarrow \sigma_0$ 
5 repeat
6   for  $1 \leq j \leq \lambda$  in parallel do
7      $t \leftarrow 0, L^j \leftarrow \emptyset$ 
8     while  $t < t_{\max}$  do
9        $t \leftarrow t + 1$ 
10      if  $a^j$  is active then
11         $\text{execute}(\mathbf{x}^j)$ 
12         $\text{broadcast}(\mathbf{x}^j, \sigma^j)$ 
13       $L^j \leftarrow L^j \cup \text{listen}()$ 
14      if  $L^j \neq \emptyset$  then
15         $\bar{\mathbf{x}} \leftarrow \text{select}(L^j)$ 
16         $\sigma \leftarrow \text{adapt}(\sigma^{\bar{\mathbf{x}}})$ 
17         $\mathbf{x}^j \leftarrow \text{mutate}(\bar{\mathbf{x}}, \sigma)$ 
18      else  $a^j$  is not active
19 until termination condition met
```

At each time step $t < t_{\max}$, an agent executes its active controller and broadcasts its genome within a limited range. At the same time, it listens for genomes from other agents, and when a genome is received (a mating event), it is saved in the agent's list L^j (its local population). This procedure is executed in parallel on all agents during t_{\max} steps, the evaluation period or one generation.

At the end of a generation, agents select a genome $\bar{\mathbf{x}} \in L^j$, and replace their active genome with a mutated copy of the selected one. In the case of open-ended environments as considered here, selection is random. In other settings, handcrafted fitness functions can be used. After this selection step, the list is emptied, and a new generation begins. Genomes that do not survive this selection step become extinct. In this work, we consider genomes as vectors $\mathbf{x} \in \mathbb{R}^N$, which encode the parameters of the neuro-controller. Only the weights undergo evolution (fixed topology).

Each genome has its own mutation step-size σ whose initial values is σ_0 , and when broadcasting, this value is sent along with the genome. On the receiving end, when a genome is selected from the local list, it is mutated using its σ value. Step-sizes are adapted before the mutation takes place, where each agent chooses with probability (0.5) to either increase or decrease the step-size. This rule increases the likelihood that the most adapted value between the two will survive and spread in the swarm. The update rule is defined as:

$$\sigma = \begin{cases} \min(\sigma(1 + \gamma), \sigma^+) & \text{if increase} \\ \max(\sigma(1 - \gamma), \sigma^-) & \text{if decrease} \end{cases}$$

where σ^+ and σ^- are the allowable upper and lower bounds, and γ is a positive constant (Table 1). Mutation is Gaussian using the updated step-size:

$$\mathbf{x}^j := \bar{\mathbf{x}} + \sigma^2 \times \mathcal{N}(1, 0).$$

The number of genomes the agent collects $\mu^j = |L^j|$ ($0 \leq \mu^j \leq \lambda$) is conditioned by its mating encounters. In the event that an agent does not have any mating opportunities and finishes its evaluation period with an empty list $L^j = \emptyset$, it becomes inactive, a state during which the agent is motionless and which can last several generations. During this period, the inactive agent continues to listen for incoming genomes from other agents passing by, and once $L^j \neq \emptyset$ the agent becomes active again at the start of the next generation. The algorithm fails if all agents become inactive, as there is no way to gather any new genomes.




Figure 1: Comparing the behaviors of two neuro-controller on the same input (10^3 randomly generated sensor values). Scatter plots (left) of the two motors outputs values of both controllers and their corresponding behavioral maps (right). The red lines on the left show the convex hull of the output values to emphasize their spread.

2.2 Selecting for Novelty

We confront the above algorithm to a version that favors behavioral novelty. This version differs only in the selection step (Line 15 in Algorithm 1). Introducing a selection pressure toward novelty to an online distributed EER can be challenging. Traditionally, novelty refers to originality throughout the course of evolution, whereas in our context, since we do not keep track of history, diversity should refer to originality at a given point in time. The goal is to identify genomes with the most different behavior in the agents' selection pools at the time of selection. We follow (Boumaza, 2023) to simulate behaviors and measure their differences on *fake* sensory data and then use these measurements as a selection criterion. The important step here is to present all genomes in the selection pool with the same *fake* sensory data, as if they experienced the same situation, and compare their *would be* response to that fake situation. We use the terms *fake* and *would be* to emphasize the fact that the sensory data is not measured by agents moving in the environment; it is a random sample from the sensory space, and the outputs are not to be executed by the agents.

To describe the selection procedure, let us focus on what happens on one agent. We note $L = \{\mathbf{x}_1 \dots \mathbf{x}_\mu\}$ the set of genomes it collected during its lifetime. The goal is to identify $\bar{\mathbf{x}} \in L$ that has the most different behavior. This genome is defined as the one that is the farthest away from all the other genomes in the list, according to some metric on the space of behaviors.

If we note \mathcal{S} (resp. \mathcal{M}) the space of sensor inputs (resp. motor outputs), we define $g : \mathbb{R}^N \times \mathcal{S} \rightarrow \mathcal{M}$ the function of the neuro-controller that computes the output $\mathbf{O} = g(\mathbf{x}, I)$ from input I of the neural network whose parameters are \mathbf{x} .

When selecting a new genome, each agent sample randomly K *fake* inputs from \mathcal{S} . For each genome $\mathbf{x}_i \in L$, it computes a behavioral descriptor that corresponds to the *fake outputs*, $\mathcal{B}_{\mathbf{x}_i} = g(\mathbf{x}_i, \mathbf{S})$. Finally, the agent select the new active genome $\bar{\mathbf{x}}$ as follows:

$$\bar{\mathbf{x}} = \arg \max_{\mathbf{x} \in L} \left(\frac{1}{|L|} \sum_{\mathbf{y} \in L} D(\mathcal{B}_{\mathbf{x}}, \mathcal{B}_{\mathbf{y}}) \right),$$

where D is the Jensen-Shannon distance (Endres and Schindelin, 2003), a symmetric bounded metric on distributions, which in our case, are histograms of $\mathcal{B}_{\mathbf{x}_i}$ (the behavioral maps).

Figure 1 illustrates this process. In there, two neuro-controllers (with the same architecture as in our experiments, a perceptron with 27 inputs and 2 outputs) with distinct sets of weights are given the same set of randomly generated input values. The output values are plotted on the left side of the figure, where we can see that both controllers have a different response to the same input i.e. robots with these controllers would behave differently in similar situations. The distance between both behaviors is estimated with D on the histograms in the right side of the figure.

To illustrate the sensitivity of this distance on genetically close genomes. We computed the distance as described above on genomes that are a few mutations away from each other. Figure 2 shows the




Figure 2: Evolution of the Jensen-Shannon distance (purple) and the cosine distance (light blue) for comparison wrt. successive mutations.

distance between a reference genome and mutated versions with different successive mutations (with a small step-size of $\sigma = 0.1$). Initially, both genomes are identical. The figure shows that small genetic differences can create sensibly different behaviors, and those differences are distinguishable using the above metric. For comparison, we present also the distances computed using the cosine metric. Both distances are upper-bounded by one.

3 Experiments

In the following experiments, we consider the same setting as in (Bredeche, 2014), where a swarm of agents evolves in an empty square arena. There are 10 randomly placed landmarks that agents can perceive; these however do not provide any advantage or disadvantage to agents and can be considered useless. We do not consider a specific task to solve but rather perform an open-ended evolution and observe the behaviors that emerge in different settings.

3.1 Simulation

The experiments¹ were performed on the Roborobo simulator (Bredeche et al., 2013), an environment that allows us to run experiments on large swarms of agents. In this simulator (Figure 3), agents are e-puck-like mobile robots with limited-range sensors and two differential drive wheels. Sensors are placed at 12 locations around the agent’s body to perceive obstacles and other agents. Two additional sensors provide the distance and orientation of the closest landmark with an unlimited range. Finally, agents can move using 2 motors. All sensor values are in the range $[0, 1]$ while actuator values are in the range $[-1, 1]$. The neuro-controller we consider here is a simple feed-forward perceptron with a hyperbolic tangent activation function. Its weights are encoded in the genome as a real vector $([1 \text{ bias neuron} + 12 \text{ sensors} \times (1 \text{ agent} + 1 \text{ obstacle}) + \text{land. distance} + \text{land. orientation}]) \times 2 \text{ outputs} = 54 \text{ weights}$). All the parameters of the experiments are summarized in Table 1.

¹The code can be downloaded at: <https://gitlab.inria.fr/boumaza/public-code>

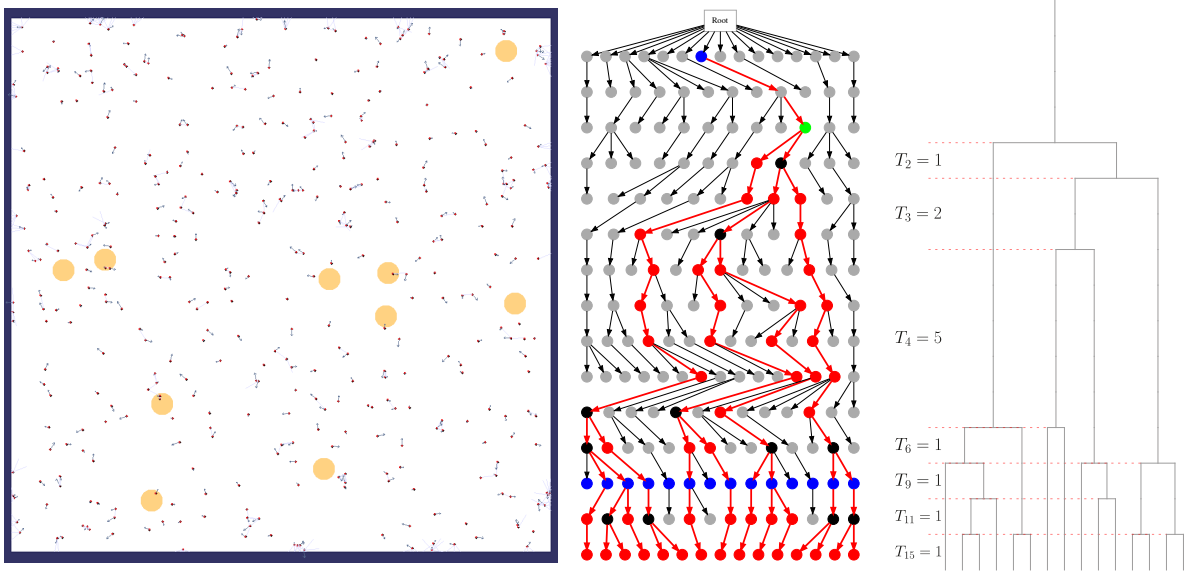


Figure 3: A screenshot of the simulator (left) where yellow circles represent the landmarks the rest are agents. An example of a phylogenetic tree for 15 agents and 14 generations (middle) and the corresponding coalescence tree (right). Gray nodes represent active genomes, black nodes represent coalescence events. The coalescence tree is a summarized representation of the sub-tree that is colored in red.

3.2 Measures

We devise two sets of experiments, each with a different population size: a small one with $\lambda = 400$, and a large one with $\lambda = 1000$. Both population evolve in the same environment, and we consider a long evolution time of 1000 generations. For each experiment we test both schemes (open-ended and novelty) compare them on different aspects.

The number of active robots indicates how many were successful at gathering genomes and managed to survive. To assess if robots were successful at spreading their genomes, we measure the number of transmissions and the number of inseminations. The difference between these two values is that transmission is the number of times a robot transmitted its genome (possibly multiple times to the same robots), whereas the number of inseminations is the number of robots that received a given agent's genome i.e. the number of distinct copies of its genome in the population. The ratio between the two values indicate the success of the agent at spreading its genome. Finally, a simple quantity that contributes in assessing agents' behaviors, is the range or the maximum distance traveled from the starting position during one generation.

To observe the evolved behaviors, we record at different times of the evolution, the trajectories of the robots over the course of a few generations. These trajectories help to observe the behaviors visually. In addition, we estimate their complexity in terms of information (Kolmogorov complexity) using a lossless compression² algorithm (Gomez, 2009). Rich and diverse trajectories with tend to have a larger complexity than ones with more regularity.

To measure the genetic diversity in the population, we take a gene perspective and study their genealogy independently of the agents. At each generation, we record the descendants of the genes and construct the phylogenetic tree of the population (Figure 3, center). Nodes in this tree represent active genomes (those that got selected by agents), and edges represent parenthood relations (child nodes are one mutation away from their parents). Edges connect genomes regardless of their vehicles; i.e., parents and offspring may

²The compression algorithm is a variant of the LZ77 (Lempel-Ziv) implemented in the zlib library

Table 1: Simulation parameters.

Arena size	1000 ² pixels
Nb. Landmarks	10
Sens. & comm. range	24 pixels
Robot diameter	6 pix.
Max trs. vel.	2 pix./tic
Max rot. vel.	30 deg/tic
Init. weights	$[-1, 1]$
Behav. maps size	10 ²
Fake inputs	128
λ	$\{400, 1000\}$
t_{\max}, g_{\max}	400 tics, 1000 gen
$\sigma^0, (\sigma^-, \sigma^+)$	0.25, (0.01, 0.5)
γ	0.35
τ	0.2

have existed on different agents. From the phylogenetic tree, we extract the corresponding coalescence tree, which synthesizes the genealogy of the last surviving genomes (the population at the last generation). 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).

From the coalescence tree, we measure the branch lengths which gives a sense of how much history genomes share. Genomes would share the least history if they come from a common ancestor far back in time and then evolved along distinct lineages. Assuming there are $r \leq \lambda$ survivors at the last generation, the branch length is defined as:

$$L = \sum_{k=2}^r k \cdot T_k, \quad (1)$$

where T_k is the number of generations there were k distinct lineages in the tree (Figure 3, right). Said differently, L is the sum of the lengths of all the vertical bars shown in the figure. To put this measure into perspective, the value can fall between two extreme cases: one where there exists only one lineage from the start to the end ($L = g_{\max}$) and one where there exists as many lineages as agents ($L = \lambda g_{\max}$). The former corresponds to the case where the entire population is one niche and all agents select the same genome. The later corresponds to the case where each agent is its own niche and selects its own genome all the time.

3.2.1 Statistical Significance

To compare both settings, each instance was executed 32 times, and we compare the above measures using median values and percentiles. Furthermore, we perform a pairwise comparison using a Mann-Whitney U test with the null hypothesis being “*the samples of the results of both instances are drawn from the same distribution*”. We establish that instance “A” outperforms instance “B” on a given measure, if the median value for instance “A” is better than the one for “B” and there is significant statistical difference between the distributions. The level of confidence at which we reject the null hypothesis is indicated by the number of asterisks on the figures, when applicable, where n asterisks indicate a p -value $< 10^{-n}$. In all the figures, curves report median values over the independent runs (solid line) and the shaded range between the 25th and the 75th percentiles (resp. the 5th and the 95th with slightly lighter shade).




Figure 4: Snapshots of agents trajectories from a typical run with $\lambda = 1000$. From left to right beginning, middle and end of evolution. Random selection (top row), behavioral diversity (bottom row). Colors have no meaning except to distinguish agents. Black dots represent landmarks for that run.

4 Results and discussions

To save space, we discuss only the results of the large experiment ($\lambda = 1000$)³. Each figure compares the two settings open-ended evolution (random selection) and novelty based selection (diversity).

We start with the trajectories realized by the agents. Figure 4 shows snapshots at various times. We clearly see that in the early stages of evolution, both selection schemes evolve similar behaviors. Agents tend to gather at landmarks, and close to walls. They also tend to orbit around landmarks and travel between them. There are what appear to be lines between landmarks, akin of a Voronoi tessellation. These behaviors are consistent with those described by (Bredecche, 2014). When we advance in time, things change a bit. With random selection, agents reinforce their gathering behavior to end up clustered and tightly packed towards the end. On the other hand, with novelty, things change, and agents seem to travel all around the arena. In terms of complexity, Figure 5 shows that indeed the information contained in the trajectories generated in the case of diversity selection is higher than that of random selection. In other words, there is more randomness and less regularity. The behavior evolve in the open-ended environment converge whereas it is not the case with novelty.

It is interesting to observe, at least in the case of random selection, that agents end up gathering at different places in the environment. In fact, this is an efficient strategy for survival where agents are guaranteed a longtime supply of genetic material, but the choice of meeting points could seem surprising.

A likely explanation could be the following. In the absence of all other information except for what can be perceived (obstacles, agents, landmarks), finding and agreeing on gathering places is not an easy task. The fact that this strategy evolves suggests that agents end up taking advantage of what their environment has to offer. The existence of these specific affordances created selection pressure

³The results on the smaller population size are qualitatively similar.




Figure 5: The complexity of the trajectories over time. Light blue corresponds to random selection and purple to novelty.

towards these meeting places. Agents are not aware that landmarks and alike are good meeting places, their environment induced a selection pressure that made this behavioral trait survive (Reed, 1996). As suggested earlier, affordances do not cause behavior but make it possible, as they can be considered as opportunities for behavior.

To take that opportunity, agents need to perceive what in their environment provides it. Landmarks are perceived from faraway, and in the absence of any other features (obstacle or robot), they are the only activity recorded on the sensors. Furthermore, landmarks are the only constant stimulus, whereas, obstacles and other agents are intermittent. They provide the only exploitable gradient of information (in the sens of Gibson’s optical flow). This may explain, why landmarks or information about them, plays an important role in shaping the agents’ behavior. Landmarks do not provide any advantage, but evolution found an effective use for them as meeting places, the same goes for corners and walls. Environmental pressure is overcome by this strategy, and in the absence of other forces, the strategy persists.

The “subjective” visual observations above, are further supported by a few objective quantities. In the case of random selection, the number of active robots indicates that survival is achieved, except for a negligible few. Furthermore, the distance traveled also suggests that agents do not move very far. On the other hand, this is not the case with diversity selection, agents travel longer distances (Figure 6) and somehow manage to survive even better (with less variance) than in the case of random selection (Figure 7).

This suggests that selection pressure (from the selection scheme) prevented the fixation of gathering behaviors and drove evolution to discover other strategies, e.g. travel and search for partners. In a sense, the selection pressure induced by the novelty selection scheme was stronger than the one induces by the affordances of the environment. This is further supported by the number of transmissions and the number of inseminations Figure 8 and Figure 9. In the case of random selection, agents have higher values for both measures, than with diversity selection, which can be explained by their gathering strategy. But again, the fact that agents survive in the case of diversity, suggests that they evolved strategies that achieve what is necessary, in terms of mating, to survive; they overcame the pressure from the environment. This is consistent with results on energy foraging, where agents only forage what is needed to survive (Bredeche et al., 2012). We may even further argue that they evolved a more efficient mating strategy. Indeed, if




Figure 6: The maximum distance traveled from the generation starting point. Light blue corresponds to random selection and purple to novelty.

we look at the ratio of the number of inseminations over the number of transmissions (Figure 10), we see that they are twice as much more efficient at mating than in the random case.

As mentioned previously, gathering tends to create clusters of agents that exchange genetic material extensively, which generally leads to a loss of genetic diversity. Clusters create niches that evolve independently and in isolation. Although gathering at meeting points is efficient for survival, it can be disastrous for genetic diversity. To support this claim, we look at the coalescence tree and the length of its branches (Figure 11). Here, the branches are longer in the case of random selection, which suggests the existence of niches that do not share a common genetic history. Indeed, many branches can be traced all the way back to the beginning of evolution, implying that cluster formation begins early and is reinforced over time (Figure 12).

The loss of genetic diversity, may also explain why gathering strategies tend to fixate over time. On the other hand, we also notice long branches, although to a lesser extent, in the case of diversity selection, which suggest the existence of niches. This could be the case if the robots, which already move extensively, do not move far enough and stay orbiting in limited areas, as if the space were partitioned into territories among the niches.

One could argue that gathering strategies might be simple and easily attainable behaviors, and to some extent, can be considered as local minima in the space of behaviors. In our experiments, both selection schemes started evolving strategies that tended in that direction; however, in the case of novelty, these strategies were not retained. The additional force from the selection pressure pushed agents to diverge and breakaway from clusters (which are genetically close) and search for other strategies that were still as efficient in terms of survival. If as we said previously, niche formation begins at an early stage of evolution, how come in the majority of niches, agents evolved similar behaviors? These niches are physically and genetically separated, they only share one thing the availability of certain affordances.




Figure 7: The median number of active robots at the end of the lifetime.

5 Conclusions

In this work we revisit open-ended evolution in EER and attempt to shed some light on which mechanisms are required to address particular environmental conditions. We focus on the features provided by the environment that favor the evolution of certain types of behaviors. We argued that those features, which represent affordances, provide a selection pressure on the behavior, and in the absence of any other selection pressure, the population converges on those behaviors.

We presented experiments in two settings, one in which there is no selection pressure on the agents, and a second one with a strong selection pressure towards behavioral novelty. The behaviors that evolved in the open-ended case tend to be similar, achieving just what is required to survive, while exploiting the affordances in the environment, on the other hand, novelty avoided the fixation of behaviors and evolved rich and complex behaviors that ignored those same affordances.

Our hypothesis remains to be tested further on different settings in other richer environments where different kinds of affordances are available (energy resources, obstacles etc.). If in our experiment, evolution stumbled easily on strategies that exploit the affordances of the environment, could that always be the case? And if not, what would be the necessary conditions for that? Could we design affordances that drive open-ended evolution? Investigating these questions seems to be very interesting.

References

- Boumaza, A. (2023). Promoting originality in online swarm robotics. In *Proceedings of the Companion Conference on Genetic and Evolutionary Computation, GECCO '23 Companion*, page 115–118, New York, NY, USA. Association for Computing Machinery.
- Bredeche, N. (2014). Embodied evolutionary robotics with large number of robots. In *the 14th international conference on the synthesis and simulation of living systems (ALIFE'14)*, pages 1–2, New York, USA.
- Bredeche, N., Haasdijk, E., and Prieto, A. (2018). Embodied evolution in collective robotics: A review. *Front. in Robo. and AI*, 5:12.




Figure 8: The number of transmissions for $\lambda=1000$.

- Bredeche, N. and Montanier, J.-M. (2010). Environment-driven Embodied Evolution in a Population of Autonomous Agents. In *PPSN 2010*, pages 290–299, Krakow.
- Bredeche, N., Montanier, J.-M., Liu, W., and Winfield, A. (2012). Environment-driven Distributed Evolutionary Adaptation in a Population of Autonomous Robotic Agents. *Mathematical and Computer Modelling of Dynamical Systems*, 18(1):101–129.
- Bredeche, N., Montanier, J.-M., Weel, B., and Haasdijk, E. (2013). Roborobo! a fast robot simulator for swarm and collective robotics. *CoRR*, abs/1304.2888.
- Endres, D. and Schindelin, J. (2003). A new metric for probability distributions. *IEEE Transactions on Information Theory*, 49(7):1858–1860.
- Ficici, S., Watson, R., and Pollack, J. (1999). Embodied evolution: A response to challenges in evolutionary robotics. In *Proc. of the 8th Europ. Wrks. on Learning Robots*.
- Gibson, J. J. (2015). *The Ecological Approach to Visual Perception*. Psychology Press.
- Gomez, F. J. (2009). Sustaining diversity using behavioral information distance. In *Proceedings of the 11th Annual Conference on Genetic and Evolutionary Computation*, page 113–120, New York, NY, USA. Association for Computing Machinery.
- Lehman, J. and Stanley, K. O. (2011). Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology*. MIT Press, Cambridge, MA, USA.
- Reed, S. E. (1996). *Encountering the World, Toward an Ecological Psychology*. Oxford University Press.




Figure 9: The number of inseminations.




Figure 10: The ratio of inseminations over transmissions. Light blue corresponds to random selection and purple to novelty.




Figure 11: The branch length of the coalescence trees L in eq. Equation (1)




Figure 12: Typical coalescence trees for $\lambda=1000$. Left, random selection and right diversity selection.