Maintaining Diversity in Robot Swarms with Distributed Embodied Evolution

Iñaki Fernández Pérez, Amine Boumaza, François Charpillet

Maintaining Diversity in Robot Swarms with Distributed Embodied Evolution

Iñaki Fernández Pérez, Amine Boumaza, François Charpillet
University of Toulouse, IRIT, UMR 5505, France
Université de Lorraine, CNRS, Inria, LORIA, F-54000 Nancy, France
inaki.fernandez-perez@irit.fr
amine.boumaza@loria.fr
francois.charpillet@loria.fr

October 2, 2018

1 Introduction

Diversity in an evolving population, as a measure of how different its individuals are, is crucial for effective evolutionary adaptation. In artificial evolution and evolutionary robotics, diversity has been investigated either to analyze the dynamics of the evolutionary process, or to explicitly promote the search for diverse or novel individuals [17, 12]. An adequate level of diversity through evolution allows to better search, balancing between exploration, to find promising areas, and exploitation, to refine good solutions. This is even more necessary when the search space is deceptive, *i.e.* it is rugged, with valleys and many local optima, which corresponds to difficult optimization problems. A very active research topic in Evolutionary Computation concerns the explicit promotion of diversity, where diversity measures are used as an auxiliary objective to be maximized: searching for diverse solutions to the problem [12, 5]. Diversity measures can also be used to monitor and analyze the evolutionary process, better understand its dynamics, and trigger specific events depending on the diversity in the population (*e.g.* restarting an evolutionary process to enhace exploration, or stop evolution when the diversity gets too low). Typically, work on diversity in evolutionary robotics is restricted to evolving single-robot behaviors with a centralized evolutionary algorithm. The work by Gomes [10] is an exception, where the authors evolve behaviors for multirobot and swarm robotic systems using a novelty-based centralized algorithm. On the other hand, in distributed Embodied Evolution (dEE), [18, 1] robots in a swarm locally communicate with each other to build their respective local populations. This entails different evolutionary dynamics to the global process, compared to centralized algorithms, due to local interactions between robots. Here, we analyze the influence on the diversity of the evolved behaviors of the distributed nature of dEE algorithms and the intensity of local selection pressure. Our experiments aim at answering the following questions: a) does distributed Embodied Evolution for robot swarms intrinsically maintains more diversity than centralized evolution?, and b) does local selection pressure influence diversity in distributed EE as it does in centralized algorithms? We first describe related work on dEE, and approaches to measure diversity in single and multirobot systems. Then, we describe the distributed EE algorithm used in our experiments, and our proposed generic diversity metric, that we compute at two levels, *i.e. global* (over the swarm) and *local* (on each local population). Finally, we detail our experiments, discuss the results, conclude and provide further research questions.

2 Related Work

A particularity of dEE is that selection is decentralized, with each robot of the swarm selecting over its local population, which is progressively built over the evaluation of controllers: robots exchange their active controllers and their respective fitness value when meeting. As such, local populations on different robots are different, and selection pressure applied over such subpopulations has different dynamics as compared to more classical centralized EAs. In [2], the authors investigate the influence of the environment on the behaviors evolved by mEDEA, a dEE algorithm that does not use a fitness measure to perform selection: selection is performed at random inside the local population of each robot. As such, the algorithm does not apply any task-driven selection pressure: it is rather the environmental selection pressure to reproduce and spread their genes that pushes evolution toward behaviors adapted to the environment that maximize the opportunities to meet other robots and mate. In [6], the authors evaluate the impact on the performance of the swarm of the intensity of selection pressure of the local selection operator in a dEE algorithm. The authors evolve neurocontrollers in a swarm of robots using different intensities of selection pressure, and conclude that the higher the selection pressure, the higher the performance, as opposed to classical centralized evolutionary algorithms, in which a lower intensity of selection pressure is usually preferred to maintain diversity in the population. This could indicate that distributed EE algorithms maintain such a diversity, necessary for the search to escape local minima.

Measuring diversity has been a topic of interest in the literature, and typically aims at two non-exclusive goals: understanding the dynamics of an evolutionary algorithm (diversity *analysis*, *e.g.* [13]), and reinjecting diversity measures into the EA, *e.g.* for diversity *promotion* (*e.g.* Novelty Search [12]), to evolve a diverse set of individuals (*e.g.* Quality-Diversity algorithms [16]), to restart the algorithm [8], or to maintain a population able to adapt to unforeseen changes [14]. Generally, when investigating diversity in Evolutionary Robotics it is measured based on *behaviors*, instead of genotypic or phenotypic diversity. A behavioral descriptor must be defined (task-specific or task-agnostic, *i.e.* generic, based on sensorimotor values) to capture adequate features of the behavior resulting from a controller. These are then used by distance functions to compute diversity metrics. In [4], the authors propose four different behavioral diversity measures as a auxiliary objectives to evolve single-robot behaviors, which help circumvent the deceptiveness of the chosen task. In [9], the authors propose two diversity measures specifically designed for swarms of robots by capturing features of the joint behavior of a swarm, instead of features of single-robot behaviors. In their paper, the authors use these measures as novelty objective, linearized with fitness values into a single objective, for a centralized novelty-based EA to evolve diverse behaviors for robot swarms. In this paper, we measure behavioral diversity as a postanalysis measure to provide insights on the internal dynamics of distributed evolution. Specifically, we propose a generic behavioral diversity metric for distributed Embodied Evolution, taken at two levels (*global*, over the swarm, and *local* diversity, on the local population of each robot). While the algorithm on each robot can only rely on local information, since the diversity measures are not used by the robots, but used to analyze how diverse the behaviors are, this does not contradict the decentralized nature of the approach. Since we focus on characterizing diversity between individual robot behaviors, either among local populations or in the swarm, and not joint swarm behaviors, we chose to use mono-robot behavioral diversity measures, closer to [4], instead of basing our study on the diversity measures for swarm robotics in [9].

3 Methods and Experiments

The algorithm used in our experiments corresponds mEDEA with task-driven selection pressure [2, 6]. Each robot in the swarm runs an independent instance of the algorithm. At every moment, a robot carries an active genome corresponding to its current neurocontroller, which is randomly initialized at the beginning of each experiment. A robot executes its controller for some time T_e, while estimating its fitness and continuously broadcasting the active genome and its current fitness estimate to other nearby robots (and vice versa). Once T_e timesteps are elapsed, the robot stops and selects a parent genome using a given selection operator. The selected genome is mutated and replaces the active genome (no crossover is used), the local population $\mathbf{l}$ is emptied, and a new generation begins. We designed a parameterized tournament selection operator, that, given a parameter $\theta_{sp} \in [0, 1]$ and a local population, selects the genome with the best fitness in a random θ_{sp} fraction of the population. The parameter θsp influences selection pressure by determining the actual tournament size, and the higher the tournament size, the stronger the selection pressure. If $\theta_{sp} = 0$, the fitness is disregarded and selection is random, while if $\theta_{sp} = 1$, the best genome in the population is selected (maximal selection pressure). Each experiment consists in running this algorithm for a given task, with a given θ_{sp}, and either with selection operating on local populations (*distributed*), or on the global one (*centralized*), *i.e.* the set of all active genomes in the swarm. At each generation], in addition to measuring the swarm's average fitness, we measure behavioral diversity using our proposed metric of dispersion among a set of behaviors $\mathbf{b}$:

$$Div(\mathbf{b}) = \frac{2}{|\mathbf{b}| \cdot (|\mathbf{b}| - 1)} \sum_{i=0}^{|\mathbf{b}-1|} \sum_{j=i+1}^{|\mathbf{b}|} d(b_i, b_j), \tag{1}$$

where $\mathbf{b}$ is a set of behavioral descriptors b_i, and $d(\cdot, \cdot)$ is a distance function between two behavioral descriptors. We aim at defining a diversity measure as generic as possible while still capturing differences in functional features of the corresponding neurocontrollers. In our approach, a behavioral descriptor for a given robot controller is defined as the list of motor outputs corresponding to an input dataset I, sampled at the beginning of each run, $I = [in^1, in^2, \ldots, in^N]$. Each in^k is a random vector of the size of the inputs of the controllers, uniformly sampled in the corresponding value range. To compute the behavioral descriptor of a controller c_i, the entries in the input dataset are fed to the controller, and the corresponding outputs ($o_i^k = c_i(in^k)$ are recorded, serving as the behavioral descriptor for c_i, *i.e.* $b_i = [o_i^1, o_i^2, \ldots, o_i^N]$. The distance between two behaviors, b_i and b_j, is computed as the average Euclidian distance between all their paired elements from b_i and b_j. In other words, the distance measures how different are the motor outputs computed by two neurocontrollers when confronted with the same set of inputs, and the global diversity, $Div(\cdot)$, is then computed as the average functional distance between each pair of behaviors in $\mathbf{b}$. We use our proposed diversity metric to evaluate at each generation how diverse are the behaviors at the *global* level of the swarm ($Div(\mathbf{b}^g_{swarm})$, where $\mathbf{b}^g_{swarm}$ is the set of behavioral descriptors of the active robot controllers in the swarm at generation g), and at the *local* level of the local populations (for each robot r, $Div(\mathbf{b}^g_r)$, where $\mathbf{b}^g_r$ is the set of behavioral descriptors of the local population of r at generation g; we report the average over the swarm).

We measure the fitness and behavioral diversity over time when a swarm of robots uses this algorithm to adapt to two classical benchmark tasks for swarm robotics: navigation and item collection. For each task, we perform 10 variants, with 5 levels of selection pressure, $\theta_{\text{sp}} \in \{0, 0.25, 0.5, 0.75, 1\}$, with either robots locally exchanging genomes (distributed), or selecting on the global population (centralized). The experiments with selection on the global population do not comply with the distributed nature of swarm systems, and are used as control experiments to test if dEE intrinsically maintain more diversity than when selection is performed on the global population. In each experiment, a swarm of robotic agents is deployed in a simulated environment (Figure 1), containing food items in the collection task. Our experiments are run using the RoboRobo simulator [3], which is a fast simulator for collective robotics. For the navigation task, each robot has 8 proximity sensors evenly spaced around the robot, which detect walls and other robots, with 8 additional item proximity sensors in the collection task. Each robot is controlled by a fully-connected perceptron with a bias neuron and no hidden layers, and maps sensory inputs to motor outputs (left and right wheel speed). The genome corresponds to a real-valued vector containing the weights of the controller (18 for navigation, and 34 for collection), adapted by either the distributed algorithm, or the centralized version.

The fitness for navigation rewards moving fast, straight and avoiding obstacles [15], while in item collection it is the number of items collected by a robot. To evaluate the impact of distributed evolution on swarm performance and diversity, at every generation of each experiment, we measure the swarm fitness (average fitness over all the robots), and the global and local diversity. We compare the results (swarm fitness and diversity) of distributed evolution to centralized evolution, and the impact of the intensity of selection pressure in both cases.

4 Results and Conclusion

To compare diversity (either global or local) between centralized and distributed evolution, we use 2D histograms represented as heatmaps, where the x-axis and the y-axis correspond to the diversity in the distributed variant and in the centralized variant, respectively. Each datapoint is then the pair of diversity values corresponding to the same generation g in a distributed and a centralized run (randomly paired), *i.e.* (Div^g_D, Div^g_C) for each pair of runs. The density of each bin in the histogram corresponds to the number of generations across all the runs when the pair of diversity values from the distributed variant and the centralized falls into that bin. If a plot is denser under the diagonal, it means that, overall, distributed evolution maintains more diversity, and vice versa. When comparisons are made between swarm fitness values, difference is reported *iff* Mann-Whitney tests yield $p < 0.05$. Figure 2 (resp. Figure 3) show the fitness of the swarm over generations for the navigation and the collection task (resp. the global and local

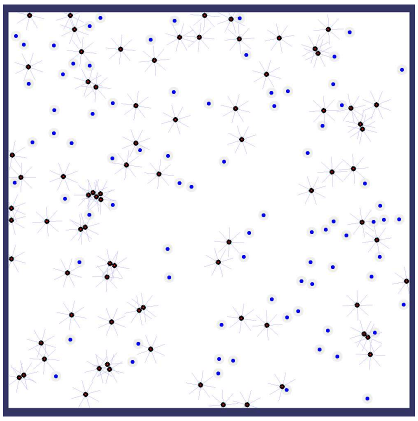

Figure 1: Simulated environment: enclosed square arena containing a swarm of robots and items (black and blue circles).

# Robots	80
# Items	80
Envir. size	$1000 \times 1000px$
Sensor range	$30px$
# runs	30
Generations	200
T_e	800 *steps*
σ	0.1

Table 1: T_e and σ are the evaluation time and std. dev. of the Gaussian mutation.

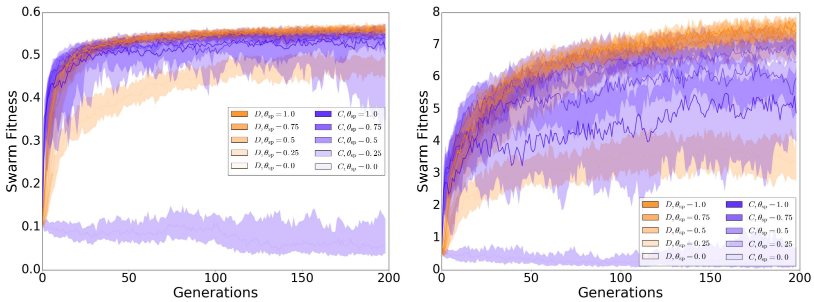

Figure 2: Swarm fitness over generations for navigation (left) and item collection (right). Blue curves represent centralized evolution (C), while orange curves represent distributed evolution (D). θ_sp is the intensity of selection pressure.

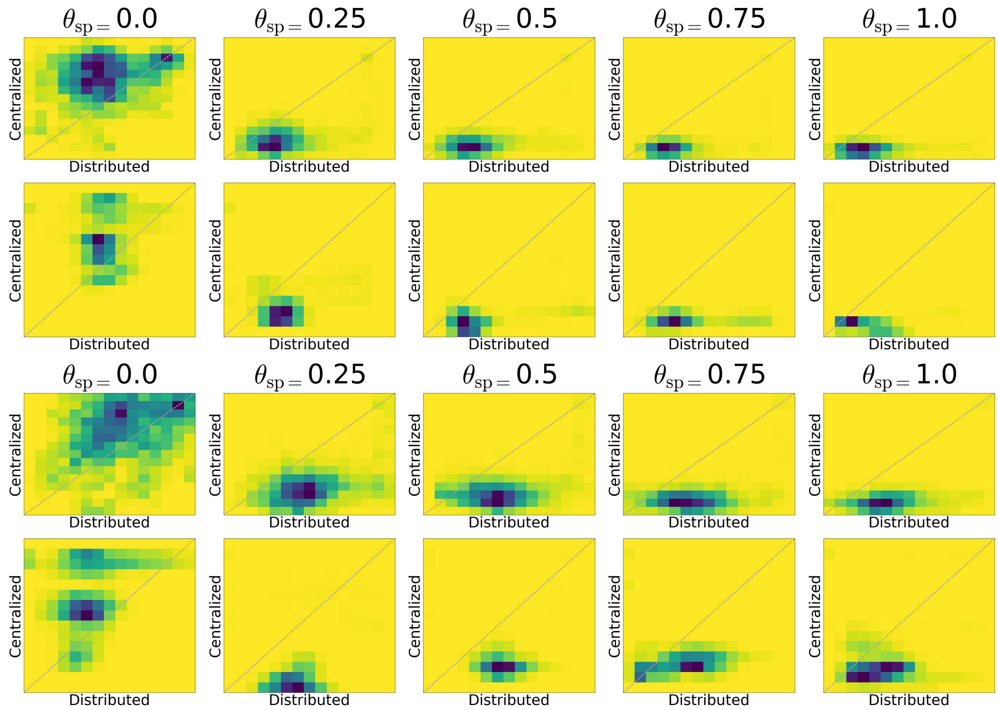

Figure 3: Heatmap for comparing global and local diversity between centralized and distributed experiments in navigation (top 2 rows) and collection (bottom 2 rows).

behavioral diversity heatmaps). In both tasks, robots adapt solve the task, reaching high fitness in all the experiments except for the centralized experiment with $\theta_{sp} = 0.0$, which corresponds to random search in the entire population. The distributed variants with $\theta_{sp} \neq 0.0$ reach slightly higher values with lower variance than the centralized variants, especially in the more challenging collection task. Regarding item collection, the intensity of selection pressure seems to have little impact on the fitness in the distributed case, while in the centralized case, the highest performance is obtained when $\theta_{sp} = 0.25$ or $\theta_{sp} = 0.5$. On the other hand, when $\theta_{sp} = 0.75$, and especially when $\theta_{sp} = 1$, the swarm fitness is lower. This could be due to a possible loss of diversity when selection pressure is strong in the centralized case. Search could stagnate in local minima, being unable to escape, and thus yielding lower fitness, especially since item collection is arguably more difficult to evolve than navigation: the search space is bigger, and information from sensors of different nature needs to be integrated. In the case of distributed evolution with $\theta_{sp} = 0.0$, which corresponds to mEDEA algorithm, there is also an improvement, although slower, even in the absence of task-driven selection pressure. This is due to environmental selection pressure pushing toward behaviors that maximize mating chances by navigating the environment, and collecting items by chance in the item collection task. Figure 3 show that, when there is selection pressure ($\theta_{sp} \neq 0.0$), distributed evolution maintains more diversity, both local and global (denser areas under the diagonal). In the case of $\theta_{sp} = 0.0$, centralized evolution yields higher diversity than distributed evolution: the centralized case corresponds to random search, and, even if a diversity of behaviors is maintained, those behaviors do not provide any fitness, as shown before.

In this paper, our main hypothesis is that such algorithms intrinsically maintain diversity, since the genomes on the local repositories of the robots are built through local exchanges between robots when meeting, and are therefore different. To test such a hypothesis, we perform a set of experiments where a swarm of robots adapts to given tasks using a distributed EE algorithm. We test 5 intensities of selection pressure, in the distributed algorithm and in a control experiment with selection on the global population. We measure both the performance on the tasks and a proposed diversity measure designed for distributed evolution in robot swarms, both from local and global perspectives, and we conclude that, when there is selection pressure in our experiments, this approach systematically maintains more diversity, compared to centralized evolution, allowing to reach slightly higher performances, especially in the item collection task. This work opens questions on how to exploit such diversity measures: they could help regulating evolutionary operators, including the *mating* operator that defines genome migration between robots in distributed evolution [1]: mating could be restricted to robots with similar behaviors, a form of reproductive isolation, which might favor the evolution of specialized niches. On the other hand, diversity measures could be used as novelty objectives. Searching for novelty in distributed evolution has recently received attention [7, 11], and we believe that our proposed diversity measures could be used to guide search in robot swarms.

References

[1] Nicolas Bredeche, Evert Haasdijk, and Abraham Prieto. Embodied evolution in collective robotics: A review. *Frontiers in Robotics and AI*, 5:12, 2018.

[2] Nicolas Bredèche and Jean-Marc Montanier. Environment-driven Embodied Evolution in a Population of Autonomous Agents. In *Parallel Problem Solving from Nature, PPSN 2010*, pages 290–299, Krakow, Poland, 2010.

[3] Nicolas Bredèche, Jean-Marc Montanier, Berend Weel, and Evert Haasdijk. Roborobo! a fast robot simulator for swarm and collective robotics. *CoRR*, abs/1304.2888, 2013.

[4] Stéphane Doncieux and Jean-Baptiste Mouret. Behavioral diversity measures for Evolutionary Robotics. In *Congress on Evolutionary Computation (CEC)*, pages 1303–1310, Espagne, 2010.

[5] Stephane Doncieux and Jean-Baptiste Mouret. Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence*, 7(2):71–93, 2014.

[6] Iñaki Fernández Pérez, Amine Boumaza, and François Charpillet. Comparison of selection methods in on-line distributed evolutionary robotics. In *Proceedings of the International Conference on the Synthesis and Simulation of Living Systems (Alife'14)*, pages 282–289, New York, July 2014. MIT Press.

[7] Marco Galassi, Nicola Capodieci, Giacomo Cabri, and Letizia Leonardi. Evolutionary strategies for novelty-based online neuroevolution in swarm robotics. In *Systems, Man, and Cybernetics (SMC), 2016*, pages 002026–002032. IEEE, 2016.

[8] Farzad Ghannadian, Cecil Alford, and Ron Shonkwiler. Application of random restart to genetic algorithms. *Information Sciences*, 95(1-2):81–102, 1996.

[9] Jorge Gomes and Anders L Christensen. Generic behaviour similarity measures for evolutionary swarm robotics. In *Proceedings of the 15th annual conference on Genetic and evolutionary computation*, pages 199–206. ACM, 2013.

[10] Jorge Gomes, Paulo Urbano, and Anders Lyhne Christensen. Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*, 7(2-3):115–144, 2013.

[11] Emma Hart, Andreas S.W. Steyven, and Ben Paechter. Evolution of a functionally diverse swarm via a novel decentralised quality-diversity algorithm. 2018.

[12] Joel Lehman and Kenneth O. Stanley. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*, 19(2):189–223, June 2011.

[13] Ronald W Morrison and Kenneth A De Jong. Measurement of population diversity. In *International Conference on Artificial Evolution*, pages 31–41. Springer, 2001.

[14] Trung Thanh Nguyen, Shengxiang Yang, and Juergen Branke. Evolutionary dynamic optimization: A survey of the state of the art. *Swarm and Evolutionary Computation*, 6:1–24, 2012.

[15] Stefano Nolfi and Dario Floreano. *Evolutionary Robotics*. MIT Press, 2000.

[16] Justin K. Pugh, Lisa B. Soros, and Kenneth O. Stanley. Quality diversity: A new frontier for evolutionary computation. *Frontiers in Robotics and AI*, 3:40, 2016.

[17] Rasmus K Ursem. Diversity-guided evolutionary algorithms. In *International Conference on Parallel Problem Solving from Nature*, pages 462–471. Springer, 2002.

[18] Richard A. Watson, Sevan G. Ficici, and Jordan B. Pollack. Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Syst.*, 2002. Elsevier.