

Embodied Evolutionary Robotics with Large Number of Robots

Nicolas Bredeche^{1,2}

¹Sorbonne Universités, UPMC Univ Paris 06, UMR 7222, ISIR, F-75005, Paris, France

²CNRS, UMR 7222, ISIR, F-75005, Paris, France

nicolas.bredeche@upmc.fr

Abstract

Embodied evolutionary robotics is a particular flavour of evolutionary robotics, where the evolutionary optimization of behaviours is achieved in an on-line and distributed fashion (Watson et al., 2002). The question asked in this paper is: does population size play a role in the evolution of particular behaviours? We experimentally demonstrate that varying the number of robots and the size of the environment can lead to very different outcomes in terms of evolved behaviours.




Figure 1: Experimental setup: a population of robots with 8 infra-red (IR) sensors (shown in blue) is deployed in an environment where 10 (yellow) landmarks are randomly placed. The robots are modelled after the famous e-puck robot, and communication between robots is achieved through the IR devices. The red tail is visible to the user only (used for identifying directions).

The meDEA algorithm

For this study, we use a particular algorithm that does not address any objective function, but rather perform

environment-driven evolution. Similar to other works in artificial life, such as Tierra (Ray, 1991), meDEA favors individuals that are capable to spread their genomes in other individuals, without any explicit consideration for any user-defined task. meDEA has been extensively described elsewhere (Bredeche et al., 2012), and has been tested with up to 20 real robots.

meDEA considers a set of robots with limited communication capabilities. Evolution relies on the diffusion of genomes from peer to peer, and genomes compete with one another to spread over the population of robots. To do so, each robot carries a genome, which defines its behaviour during a predefined amount of time (its “lifetime”). Whenever two robots are close enough, each transmits a mutated copy of its current genome to the other, and store the incoming genome in a list for further use. At the end of the robot’s lifetime, the list of previously received genomes is emptied except for one arbitrarily selected genome, which is then mutated (using gaussian mutation) and replaces the genome used so far. Because the new genome is arbitrarily selected, there is pressure towards genomes that are able to drive their “vehicles” (i.e., the robots) to spread themselves over the population, favouring behavioural strategies better fitted to the environment.

Experiments

In a previous work (Bredeche et al., 2012), we showed that the average number of encounters is critical to survival, and is related to the number of robots, the communication radius and the size of the environment. That is: if robots get lost and end up with no new genomes to carry on evolution, the algorithm fails (ie. the robots do not survive).

The question addressed here is quite the opposite: given the population is large enough for the algorithm to work (ie. each robot will encounter at least one other robot during its lifetime), what happens when we consider larger populations? To answer this question, we consider experiments with different setups (units are given in pixels, a robot is 5x5 pixels): (1) Setup 1: 100 robots in a 400x400 environment; (2) Setup 2: 200 robots in a 1000x1000 environment; (3)




Figure 2: Trajectories from the last generations of three different setups (typical runs): (a) 100 robots, 400x400 environment ; (b) 200 robots, 1000x1000 environments ; (c) 4000 robots, 4000x4000 robots.

Setup 3: 4000 robots in a 4000x4000 environment.

The setup we consider is similar to what is shown in Figure 1 (see caption for description), except for the environment size and number of robots which vary wrt. the setups. We consider that each robot lives for 400 steps, during which it spreads its genome to every robots encountered. Each robot senses its surrounding with 12 infra-red sensors with limited range, which are also used for local communication (sending/receiving genomes), as well as a particular sensor which gives the direction and distance to the closest landmark (infinite range). There are 10 landmarks in the environment (initially placed at random locations). Landmarks provide no direct advantage nor disadvantage to nearby robots (ie. landmarks are *a priori* useless). Each robot is controlled by an Elman recurrent neural network with 5 neurons in the hidden layer, and the network weights are read from the genome.

Fig. 2 shows three typical examples from each of the three setups (at least 10 replications per setup have been done, each leading to similar results). These images show the trajectories of each of the robots over the last few generations of one run. Several conclusions can be drawn just from experimental observations.

Setup 1 displays randomly moving robots. This is an efficient behavioural strategy as each robot gets the opportunity to spread its own genome, while not being able to distinguish itself from others by doing so. Though a tendency to orbit around landmarks can be seen, this is far from the only behaviours observed, and behaviours from all robots are very similar with no clear clusters identified;

Setup 2 is rather different as three different behaviours can be observed: (1) some robots gather in corners (2) a group of robots is closely orbiting around one of the 10 landmarks and (3) some (few) robots appears to be travelling in straight lines.

Setup 3 provides a clear vision of what Setup 2 merely sketched: robots are either orbiting around landmarks, or

travelling in straight lines away from the landmarks. Also, very small groups are found orbiting around landmarks, but at some distance, avoiding interaction with the robots which are positioned right on the landmark (e.g. the landmark in the center).

These experiments show that increasing the number of robots and size of the environment leads to more diversity in the behaviours observed. However, the most compelling observation comes from robots travelling in straight lines, seemingly avoiding the landmarks. While it was not predicted *a priori*, this behaviour can be understood from observations: robots keep the farther away from any landmarks, resulting in clusters of robots travelling on the exact frontier between two landmarks. It is also interesting, if not useful, to note that this type of behaviour produces some kind of distributed voronoi tessellation, using landmarks as seeds.

Conclusion

The take-home message is that population size may dramatically change the results of evolution. In particular, the emergence of what resembles distributed Voronoi tessellation shows a striking example of what can only be observed when scaling up the population size.

One current limitation of this study it that population size *and* environment size were changed simultaneously, without considering the *density* of robots. This is currently under investigation.

Acknowledgements

The author wishes to thank Leo Cazenille, Arthur Bernard and Jean-Baptiste Mouret for comments.

References

- Bredeche, N., Montanier, J.-M., Wenguo, L., and Winfield, A. F. (2012). Environment-driven Distributed Evolutionary Adaptation in a Population of Autonomous Robotic Agents. *Mathematical and Computer Modelling of Dynamical Systems*, 18(1).
- Ray, T. (1991). Is it alive, or is it ga? In *Proceedings of International Conference on Genetic Algorithms*.
- Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, 39(1):1–18.