Information Sciences 312 (2015) 55–77

Contents lists available at ScienceDirect

Information Sciences

journal homepage: www.elsevier.com/locate/ins

Towards the standardization of distributed Embodied Evolution

A. Prieto, F. Bellas *, P. Trueba, R.J. Duro

Integrated Group for Engineering Research, Universidade da Coruña, Spain

ARTICLE INFO

Article history:
Received 2 August 2014
Received in revised form 13 March 2015
Accepted 22 March 2015
Available online 31 March 2015

Keywords:
Embodied Evolution
Evolutionary computation
Distributed optimization
Collective intelligence
Multi-robot system

ABSTRACT

The Embodied Evolution (EE) paradigm arose in the early 2000s as a response to the automatic design of distributed control systems in real time for teams of autonomous robots. The interest for this type of evolutionary approach has been increasing steadily, not only in its native field of robotics, but also in other fields related to distributed optimization problems since previous works have shown its capability to outperform traditional evolutionary techniques when the scenario requires an on-line coordination of the team. Most of the activity in this research field has been eminently practical, meaning that authors have focused their efforts on developing EE algorithms and variations adapted to solve very specific practical cases. The problem that arises is that, on one hand, all these dissimilar variations of the basic EE structure produce an unclear state of the art and, on the other, that there is a high dependence between the performance obtained by the algorithms and the specific problems where they have been tested, which complicates extrapolating conclusions to different scenarios. As a consequence, this work has two main objectives, namely, designing and implementing a standard EE algorithm that captures the more general principles of this paradigm and that can be applied to any distributed optimization problem, and analyzing how its parameters influence the performance of the algorithms in a set of theoretical representative problems so that objective and reliable conclusions about the behavior of EE can be obtained. At the same time, this work presents an analysis of the evaluation criteria required for coordination tasks when using decentralized distributed approaches, which has influenced both, the definition of the algorithm and the selection of experimental set to test it.

1. Introduction

In 1999 Ficici et al. presented the Embodied Evolution (EE) paradigm [13] as "evolution taking place in a population of robots". Their original motivation for developing a new evolutionary methodology for multi-robot systems was, on one hand, to avoid the reality gap problem of evolutionary robotics when transferring simulated controllers to real robots, and on the other, enabling scalability and increasing robustness in multi-robot systems by moving from a centralized evolution strategy to a decentralized one. Hence, EE arose as an alternative evolutionary methodology for dealing with the design of controllers in real robotic systems made up of several components, one of the most complex scenarios one can face in autonomous robotics.

* Corresponding author at: Escuela Politécnica Superior, Universidade da Coruña, Mendizábal s/n, 15403 Ferrol, A Coruña, Spain. Tel.: +34 981 337400x3886; fax: +34 981 337410.

E-mail addresses: abprieto@udc.es (A. Prieto), francisco.bellas@udc.es (F. Bellas), pedro.trueba@udc.es (P. Trueba), richard@udc.es (R.J. Duro).

http://dx.doi.org/10.1016/j.ins.2015.03.044

EE is inspired by Artificial Life experiments [24,33]. The individuals that make up the population are embodied and situated in an environment where they interact in a local, decentralized and asynchronous fashion. According to this inspiration, evolution in EE is open-ended, leading to a paradigm that is intrinsically adaptive and highly suitable for real time learning in distributed dynamic problems (see [29] for a comparison of an EE technique and traditional evolutionary approaches). The main potential of EE comes from the management of the multiple interactions that occur between individuals, leading to emergent solutions in highly complex domains, like that of real multi-robot systems. On the other hand, its main drawback has to do with the high parametric sensitivity it presents, usually implying a highly problem-dependent hand tuning stage for obtaining stable solutions [11].

1.1. Encapsulated and distributed Embodied Evolution

Due to the promising results obtained by Watson and Ficici in their original work [33], several authors have continued with the development of EE algorithms for real robots [3,15,21,27]. They have followed two different approaches. The first one is found in the original EE algorithm, where each individual in the population only carries its own genotype, as in the case of natural evolution and Alife simulations, and any genotypic change occurs strictly through interactions with other individuals. This approach has been called distributed Embodied Evolution (dEE) [10]. In the second approach, called encapsulated Embodied Evolution (eEE) [10], each individual carries a population of genotypes and their evolution is carried out partially or entirely as independent processes associated to each individual.[1] This leads to a sort of on-line parallel evolutionary algorithm with "islands" of evolutionary niches running in real time and where the interactions between niches are asynchronous, occurring only when the evaluation of the performance of a group of individuals is carried out. eEE arose due to practical implementation issues, based on the premise that dEE requires a large number of robots in the team to provide the algorithm with sufficient genotypic variety to avoid premature convergente. Therefore, it was thought that dEE was not suitable when using small robot teams [12]. In fact, eEE has been successfully applied in different robotic experiments involving a small number of robots and requiring a low degree of coordination. For example, Elfwing et al. [12] used two Cyber Rodent robots in an open-ended mating and foraging experiment, and six Khepera robots were coordinated by Usui and Arita [32] in a simple patrolling task. In [15], Haasdijk et al. analyzed the $(\mu + 1)$ ON-LINE evolutionary algorithm in a simulated task consisting in patrolling an arena with walls using a single e-puck robot.

One of the main problems that researchers in eEE must face is that of how to evaluate the candidate genotypes for each embodied individual, as this must be done in real-time by means of time-sharing strategies [12,16], which makes the outcome highly strategy-dependent, adds noise to the evaluation, and thus slows down evolution [19]. This establishment of a set of arbitrary evaluation criteria places this approach nearer to a classical evolutionary algorithm for non-collective optimization. Moreover, eEE introduces a second level of evolution in the individuals during their lifetime that is far from the bio-inspired background of Alife simulations, on which the original EE algorithm [33] was based.

On the other hand, as commented above, the main feature of the distributed Embodied Evolution (dEE) algorithm comes from the dynamics of the interactions that occur among the components of the population during their collective evaluation, which enables the emergence of self-adaptive cooperative behaviors. This property is attenuated when the population size is reduced, as in the case of the eEE approaches commented above. In this sense, the tendency in dEE has been quite different, addressing tasks involving larger robotic teams[2] and where coordination was a key aspect. This has implied having to face typical problems in collective intelligence and distributed optimization like self-organization, adaptation, emergence of specialization, etc. For instance, Bredeche et al. [3] apply the mEDEA algorithm, an implementation of dEE, using 100 simulated e-pucks in a survival task to analyze environment-driven adaptation and robustness in the face of environmental change, and 20 real e-puck robots in a two-sun experiment to study the emergence of consensus and specialization. Another remarkable example is [34], where the authors apply the original EE algorithm (PGTA) developed by Watson et al. [33] to a prey-predator task with 9 simulated Khepera robots. The interplay between evolution and learning is studied in this case using a highly dynamic environment with two co-evolving populations. The authors of the current paper applied in [27] the ASiCo algorithm, another implementation of dEE, to a collective cleaning task with 8 real e-puck robots, and also with a set of 30 simulated ones. In this case, the self-organization capabilities of dEE were studied together with the emergence of specialization within the population in the case of requiring different sub-tasks. A collective gathering experiment was designed and analyzed in [31] with 20 simulated Khepera robots, focusing the discussion on the parameterization of the ASiCo algorithm. Finally, it must be highlighted that dEE has also been extended to other domains different from robotics. An example is found in [28] where a shipping freight optimization task is successfully solved in a highly dynamic setup.

[1] Some authors [10,18] also consider a third type or *hybrid* EE approach where information transfer between the niches of the encapsulated approach is allowed and so the evolution is carried out in a partially independent process for each individual [10,12,32].

[2] In [21] Nehmzow presented the PEGA algorithm, a physically embedded algorithm that follows principles similar to those of dEE, and which was tested in real cases using only two robots. This approach, however, cannot be considered as Embodied Evolution because mating is not guided by the evolutionary process but by a predefined behavior that is forced after evaluation.