



HAL
open science

Embodied Evolution in Collective Robotics: A Review

Nicolas Bredeche, Evert Haasdijk, Abraham Prieto

► **To cite this version:**

Nicolas Bredeche, Evert Haasdijk, Abraham Prieto. Embodied Evolution in Collective Robotics: A Review. *Frontiers in Robotics and AI*, 2018, 5, pp.12. 10.3389/frobt.2018.00012 . hal-03313845

HAL Id: hal-03313845

<https://hal.sorbonne-universite.fr/hal-03313845>

Submitted on 4 Aug 2021

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.



Embodied Evolution in Collective Robotics: A Review

Nicolas Bredeche^{1*}, Evert Haasdijk² and Abraham Prieto³

¹Sorbonne Université, CNRS, Institute of Intelligent Systems and Robotics, ISIR, Paris, France, ²Computational Intelligence Group, Department of Computer Science, Vrije Universiteit, Amsterdam, Netherlands, ³Integrated Group for Engineering Research, Universidade da Coruna, Ferrol, Spain

This article provides an overview of evolutionary robotics techniques applied to online distributed evolution for robot collectives, namely, embodied evolution. It provides a definition of embodied evolution as well as a thorough description of the underlying concepts and mechanisms. This article also presents a comprehensive summary of research published in the field since its inception around the year 2000, providing various perspectives to identify the major trends. In particular, we identify a shift from considering embodied evolution as a parallel search method within small robot collectives (fewer than 10 robots) to embodied evolution as an online distributed learning method for designing collective behaviors in swarm-like collectives. This article concludes with a discussion of applications and open questions, providing a milestone for past and an inspiration for future research.

OPEN ACCESS

Edited by:

Elio Tuci,
Middlesex University,
United Kingdom

Reviewed by:

Yara Khaluf,
Ghent University, Belgium
Aparajit Narayan,
Aberystwyth University,
United Kingdom

*Correspondence:

Nicolas Bredeche
nicolas.bredeche@sorbonne-
universite.fr

Specialty section:

This article was submitted to
Multi-Robot Systems,
a section of the journal
Frontiers in Robotics and AI

Received: 30 November 2017

Accepted: 29 January 2018

Published: 22 February 2018

Citation:

Bredeche N, Haasdijk E and Prieto A
(2018) Embodied Evolution in
Collective Robotics: A Review.
Front. Robot. AI 5:12.
doi: 10.3389/frobt.2018.00012

Keywords: embodied evolution, online distributed evolution, collective robotics, evolutionary robotics, collective adaptive systems

1. INTRODUCTION

This article provides an overview of evolutionary robotics research where evolution takes place in a population of robots in a continuous manner. Ficici et al. (1999) coined the phrase *embodied evolution* for evolutionary processes that are distributed over the robots in the population to allow them to adapt autonomously and continuously. As robotics technology becomes simultaneously more capable and economically viable, individual robots operated at large expense by teams of experts are increasingly supplemented by collectives of robots used cooperatively under minimal human supervision (Bellingham and Rajan, 2007), and embodied evolution can play a crucial role in enabling autonomous online adaptivity in such robot collectives.

The vision behind embodied evolution is one of the collectives of truly autonomous robots that can adapt their behavior to suit varying tasks and circumstances. Autonomy occurs at two levels: not only the robots perform their tasks without external control but also they assess and adapt—through evolution—their behavior without referral to external oversight and so learn autonomously. This adaptive capability allows robots to be deployed in situations that cannot be accurately modeled *a priori*. This may be because the environment or user requirements are not fully known, or it may be due to the complexity of the interactions among the robots as well as with their environment effectively rendering the scenario unpredictable. Also, onboard adaptivity intrinsically avoids the reality gap (Jakobi et al., 1995) that results from inaccurate modeling of robots or their environment when developing controllers before deployment because controllers continue to develop *after* deployment. A final benefit is that embodied evolution can be seen as parallelizing the evolutionary process because it distributes the evaluations over multiple robots. Alba (2002) has shown that such parallelism can provide substantial benefits, including superlinear speedups. In the case of robots,

this has the added benefit of reducing the amount of time spent executing poor controllers per robot, reducing wear and tear.

Embodied evolution's online nature contrasts with "traditional" evolutionary robotics research. Traditional evolutionary robotics employs evolution in the classical sequential centralized optimization paradigm: parent and survivor selection are centralized and consider the entire population. The "robotics" part entails a series of robotic trials (simulated or not) in an evolution-based search for optimal robot controllers (Nolfi and Floreano, 2000; Bongard, 2013; Doncieux et al., 2015). In terms of task performance, embodied evolution has been shown to outperform alternative evolutionary robotic techniques in some setups such as surveillance and self-localization with flying UAVs (Schut et al., 2009; Prieto et al., 2016), especially regarding convergence speed.

To provide a basis for a clear discussion, we define embodied evolution as a paradigm where evolution is implemented in multirobotic (two or more robots) system. Two robots are already considered a multirobotic system since it is still possible to distribute an algorithm among them. These systems exhibit the following features.

Decentralized: There is no central authority that selects parents to produce offspring or individuals to be replaced. Instead, robots assess their performance, exchange, and select genetic material autonomously on the basis of locally available information.

Online: Robot controllers change on the fly, as the robots go about their proper actions: evolution occurs during the operational lifetime of the robots and in the robots' task environment. The process continues after the robots have been deployed.

Parallel: Whether they collaborate in their tasks or not, the population consists of multiple robots that perform their actions and evolve concurrently, in the same environment, interacting frequently to exchange genetic material.

The decentralized nature of communicating genetic material implies that the selection is executed locally, usually involving only a part of the whole population (Eiben et al., 2007), and that it must be performed by the robots themselves. This adds a third opportunity for selection in addition to parent and survivor selection as defined for classical evolutionary computing. Thus, embodied evolution extends the collection of operators that define an evolutionary algorithm (i.e., evaluation, selection, variation, and replacement (Eiben and Smith, 2008)) with *mating* as a key evolutionary operator.

Mating: An action where two (or more) robots decide to send and/or receive genetic material, whether this material will or will not be used for generating new offspring. When and how this happens depends not only on predefined heuristics but also on evolved behavior, the latter determining to a large extent whether robots ever meet to have the opportunity to exchange genetic material.

In the past 20 years, online evolutionary robotics in general and embodied evolution in particular have matured as research fields. This is evidenced by the growing number of relevant publications in respected evolutionary computing venues such as in conferences (e.g., ACM GECCO, ALIFE, ECAL, and EvoApplications), journals (e.g., *Evolutionary Intelligence's* special issue on Evolutionary Robotics (Haasdijk et al., 2014b)), workshops (PPSN 2014 ER workshop, GECCO 2015 and 2017

Evolving collective behaviors in robotics workshop), and tutorials (ALIFE 2014, GECCO 2015 and 2017, ECAL 2015, PPSN 2016, and ICDL-EPIROB 2016). A Google Scholar search of publications citing the seminal embodied evolution paper by Watson et al. (2002) illustrates this growing trend. Since 2009, the paper has attracted substantial interest, more than doubling the yearly number of citations since 2008 (approximately 20 citations per year since then).¹

To date, however, a clear definition of what embodied evolution is (and what it is not) and an overview of the state of the art in this area are not available. This article provides a definition of the embodied evolution paradigm and relates it to other evolutionary and swarm robotics research (Sections 2 and 3). We identify and review relevant research, highlighting many design choices and issues that are particular to the embodied evolution paradigm (Sections 4 and 5). Together this provides a thorough overview of the relevant state-of-the-art and a starting point for researchers interested in evolutionary methods for collective autonomous adaptation. Section 6 identifies open issues and research in other fields that may provide solutions, suggests directions for future work, and discusses potential applications.

2. CONTEXT

Embodied evolution considers collectives of robots that adapt online. This section positions embodied evolution vis à vis other methods for developing controllers for robot collectives and for achieving online adaptation.

2.1. Offline Design of Behaviors in Collective Robotics

Decentralized decision-making is a central theme in collective robotics research: when the robot collective cannot be centrally controlled, the individual robots' behavior must be carefully designed so that global coordination occurs through local interactions.

Seminal works from the 1990s such as Mataric's Nerd Herd (Mataric, 1994) addressed this problem by hand-crafting behavior-based control architectures. Manually designing robot behaviors has since been extended with elaborate methodologies and architectures for multirobot control (see Parker (2008) for a review) and with a plethora of bioinspired control rules for swarm-like collective robotics (see Nouyan et al. (2009) and Rubenstein et al. (2014) for recent examples involving real robots and Beni (2005), Brambilla et al. (2012), and Bayindir (2016) for discussions and recent reviews).

Automated design methods have been explored with the hope of tackling problems of greater complexity. Early examples of this approach were applied to the robocup challenge for learning coordination strategies in a well-defined setting. See the study by Stone and Veloso (1998) for an early review and Stone et al. (2005) and Barrett et al. (2016) for more recent work in this vein. However, Bernstein et al. (2002) demonstrated that solving even the simplest multiagent learning problem is intractable in

¹See <https://plot.ly/~evertwh/17/> for more details and the underlying data.

polynomial time (actually, it is NEXP-complete), so obtaining an optimal solution in reasonable time is currently infeasible. Recent works in reinforcement learning have developed theoretical tools to break down complexity by operating a move from considering many agents to a collection of single agents, each of which being optimized separately (Dibangoye et al., 2015), leading to theoretically well-founded contributions, but with limited practical validation involving very few robots and simple tasks (Amato et al., 2015).

Lacking theoretical foundations, but instead based on the experimental validation, swarm robotics controllers have been developed with black-box optimization methods ranging from brute-force optimization using a simplified (hence tractable) representation of a problem (Werfel et al., 2014) and evolutionary robotics (Hauert et al., 2008; Trianni et al., 2008; Gauci et al., 2012; Silva et al., 2016).

The methods vary, but all the approaches described here (including “standard” evolutionary robotics) share a common goal: to design or optimize a set of control rules for autonomous robots that are part of a collective *before* the actual deployment of the robots. The particular challenge in this kind of work is to design individual behaviors that lead to some required global (“emergent”) behavior without the need for central oversight.

2.2. Lifelong Learning in Evolutionary Robotics

It has long been argued that deploying robots in the real world may benefit from continuing to acquire new capabilities *after* initial deployment (Thrun and Mitchell, 1995; Nelson and Grant, 2006), especially if the environment is not known beforehand. Therefore, the question we are concerned with in this article is *how to endow a collective robotics system with the capability to perform lifelong learning*. Evolutionary robotics research into this question typically focuses on individual autonomous robots. Early works in evolutionary robotics that considered lifelong learning explored learning mechanisms to cope with minor environmental changes (see the classic book by Nolfi and Floreano (2000) and Urzelai and Floreano (2001) and (Tonelli and Mouret, 2013) for examples and Mouret and Tonelli (2015) for a nomenclature). More recently, Bongard et al. (2006) and Cully et al. (2015) addressed

resilience by introducing fast online re-optimization to recover from hardware damage.

Bredeche et al. (2009), Christensen et al. (2010), and Silva et al. (2012) are some examples of online versions of evolutionary robotics algorithms that target the fully autonomous acquisition of behavior to achieve some predefined task in individual robots. Targeting agents in a video game rather than robots, Stanley et al. (2005) tackled the online evolution of controllers in a multiagent system. Because the agents were virtual, the researchers could control some aspects of the evaluation conditions (e.g., restarting the evaluation of agents from the same initial position). This kind of control is typically not feasible in autonomously deployed robotic systems.

Embodied evolution builds on evolutionary robotics to implement lifelong learning in robot *collectives*. Its clear link with traditional evolutionary robotics is exemplified by work such as by Usui and Arita (2003), where a traditional evolutionary algorithm is encapsulated on each robot. Individual controllers are evaluated sequentially in a standard time sharing setup, and the robots implement a communication scheme that resembles an island model to exchange genomes from one robot to another. It is this communication that makes this an instance of embodied evolution.

3. ALGORITHMIC DESCRIPTION

This section presents a formal description of the embodied evolution paradigm by means of generic pseudocode and a discussion about its operation from a more conceptual perspective.

The pseudocode in **Algorithm 1** provides an idealized description of a robot’s control loop as it pertains to embodied evolution. Each robot runs its own instance of the algorithm, and the evolutionary process emerges from the interaction between the robots. In embodied evolution, there is no entity outside the robots that oversees the evolutionary process, and there is typically no synchronization between the robots: the replacement of genomes is asynchronous and autonomous.

Some steps in this generic control loop can be implicit or entwined in particular implementations. For instance, robots may continually broadcast genetic material over short range, so that other robots that come within this range receive it

ALGORITHM 1 | An individual robot’s control loop for embodied evolution.

```

initialize robot;
for ever do
  Sense - Act cycle (depends on robotic paradigm);
  perf ← calculate performance;
  if mating? then
    transmit my genome;
    g ← receive mate’s genome;
    store(g);
  end
  if replacement? then
    parents ← select parents;
    offspring ← variation(parents);
    activate(offspring) //Time-sharing: control is handed over to the new candidate controller
  end
end
end

```

automatically. In such a case, the *mating* operation is implicitly defined by the selected broadcast range. Similarly, genetic material may be incorporated into the currently active genome as it is received, merging the mating and replacement operations. Implicitly defined or otherwise, the steps in this algorithm are, with the possible exception of performance calculation, necessary components of any embodied evolution implementation.

The following list describes and discusses the steps in the algorithm in detail.

Initialization: The robot controllers are typically initialized randomly, but it is possible that the initial controllers are developed offline, be it through evolution or handcraft (e.g., see the work by Hettiarachchi et al. (2006)).

Sense-act cycle: This represents “regular,” i.e., not related to the evolutionary process, robot control. The details of the sense-act cycle depend on the robotic paradigm that governs robot behavior; this may include planning, subsumption, or other paradigms. This may also be implemented as a separate parallel process.

Calculate performance: If the evolutionary process defines an objective function, the robots monitor their own performance. This may involve measurements of quantities such as speed, number of collisions, or amount of collected resources. Whatever their nature, these measurements are then used to evaluate and compare genomes (as fitness values in evolutionary computation). The possible discrepancy between the individual’s objective function and the population welfare will be discussed further in Section 6.2.

Mating: This is the essential step in the evolutionary process where robots exchange genetic material. The choice to mate with another robot may be purely based on the environmental contingencies (e.g., when robots mate whenever they are within communication range), but other considerations may also play a part (e.g., performance, genotypic similarity). The pseudocode describes a symmetric exchange of genomes (both with a transmit and a receive operation), but this may be asymmetrical for particular implementations. In implementations such as that of Schwarzer et al. (2011) or Haasdijk et al. (2014a), for instance, robots suspend normal operation to collect genetic material from other, active robots. Mating typically results in a pool of candidate parents that are considered in the parent selection process.

Replacement: The currently active genome is replaced by a new individual (the offspring), implying the removal of the current genome. This event can be triggered by a robot’s internal conditions (e.g., running out of time or virtual energy, reaching a given performance level) or through interactions with other robots (e.g., receiving promising genetic material (Watson et al., 2002)).

Parent selection: This is the process that selects which genetic information will be used for the creation of new offspring from the received genetic information through mating events. When an objective is defined, the performance of the received genome is usually the basis for selection, just as in regular evolutionary computing. In other cases, the selection among received genomes can be random or depend on non-performance related heuristics (e.g., random, genotypic proximity). In the absence of objective-driven selection pressure, individuals are still competing with

respect to their ability to spread their own genome within the population, although that cannot be explicitly captured during parent selection. This will be further discussed in Section 5.2.

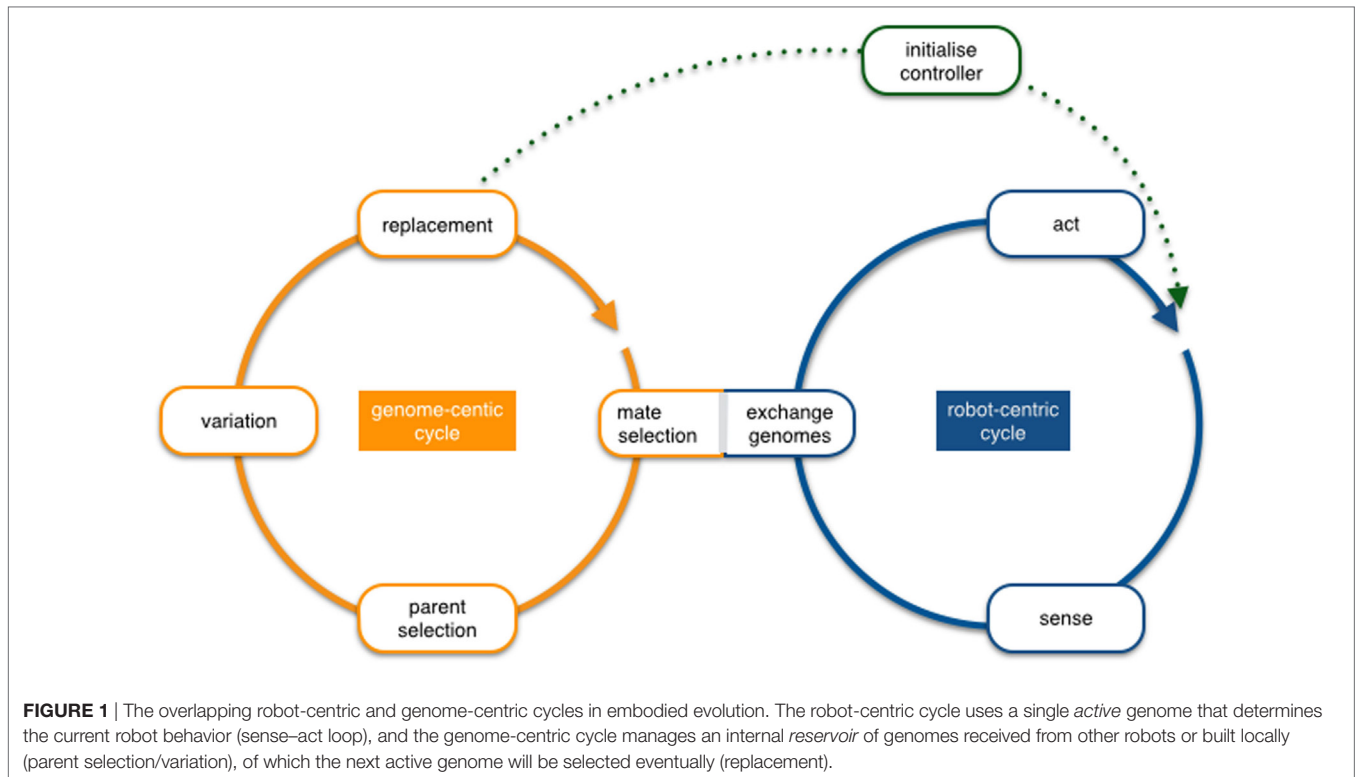
Variation: A new genome is created by applying the variation operators (mutation and crossover) on the selected parent genome(s). This is subsequently activated to replace the current controller.

From a conceptual perspective, embodied evolution can be analyzed at two levels that are represented by two as depicted in **Figure 1**.

The robot-centric cycle is depicted on the right in **Figure 1**. It represents the physical interactions that occur between the robot and its environment, including interactions with other robots and extends this sense-act loop commonly used to describe real-time control systems by accommodating the exchange and activation of genetic material. At this particular point, the genome-centric and robot-centric cycles overlap. The cycle operates as follows: each robot is associated with an *active* genome, and the genome is interpreted into a set of features and control architecture (the phenotype), which produces a behavior that includes the transmission of its own genome to some other robots. Each robot eventually switches from an active genome to another, depending on a specific event (e.g., minimum energy threshold) or duration (e.g., fixed lifetime), and consequently changes its active genome, probably impacting its behavior.

The genome-centric cycle deals with the events that directly affect the genomes existing in the robot population and therefore also the evolution *per se*. Again, the mating and the replacement are the events that overlap with the robot-centric cycle. The operation from the genome cycle perspective is as follows: each robot starts with an initial genome, either initialized randomly or *a priori* defined. While this genome is *active*, it determines the phenotype of the robot, hence its behavior. Afterward, when replacement is triggered, some genomes are selected from the reservoir of genomes previously received according to the parent selection criteria and later combined using the variation operators. This new genome will then become part of the population. In the case of fixed-size population algorithms, the replacement will automatically trigger the removal of the old genome. In some other cases, however, there is a specific criterion to trigger the removal event producing populations of individuals that change their size along the evolution.

The two circles connect on two occasions, first by the “exchange genomes” (or mating) process, which implies the transmission of genetic material, possibly together with additional information (fitness if available, general performance, genetic affinity, etc.) to modulate future selection. Generally, the received information is stored to be used (in full or in part) to replace the active genome in the later parent selection process. Therefore, the event is triggered and modulated by the robot cycle, but it impacts on the genomic cycle. Also, the decentralized nature of the paradigm enforces that these transmissions occur locally, either one-to-one or to any robot in a limited range. There are several ways in which mate selection can be implemented, for instance, individuals may send and receive genomic information indiscriminately within a certain location range or the frequency of transmission can depend on the task performance. The second overlap between



the two cycles is the activation of new genomic information (replacement). The activation of a genome in the genomic cycle implies that this new genome will now take control of the robot and therefore changes the response of the robot in the scenario (in evolutionary computing terms, this event will mark the start of a new individual evaluation). This aspect is what creates the online character of the algorithm that, together with the locality constraints, implies that the process is also asynchronous.

This conceptual representation matches what has been defined as *distributed* embodied evolution by Eiben et al. (2010). The authors proposed a taxonomy for online evolution that differentiates between encapsulated, distributed, and hybrid schemes. Most embodied evolution implementations are distributed, but this schematic representation also covers hybrid implementations. In such cases, the robot locally maintains a population that is augmented through mating (rather like an island model in parallel evolutionary algorithms). It should be noted that encapsulated implementations (where each robot runs independently of the others) are not considered in this overview.

4. EMBODIED EVOLUTION: THE STATE OF THE ART

In this section, we identify the major research topics from the works published since the inception of the domain, all summarized in **Table 1**. **Table 1** provides an overview of published research on embodied evolution with robot collectives. Each entry describes a contribution, which may cover several papers. The entries are described in terms of their implementation details, the robot

behavior, experimental settings, mating conditions, selection, and replacement schemes. The glossary in **Table 2** explains these features in more detail.

First, we distinguish between works that consider embodied evolution as a parallel search method for optimizing *individual* behaviors and works where embodied evolution is employed to craft *collective* behavior in robot populations. The latter trend, where the emphasis is on collective behavior, has emerged relatively recently and since then has gained importance (32 papers since 2009).

Second, we consider the homogeneity of the evolving population; borrowing definitions from biology, we use the term *monomorphic* (resp. *polymorphic*) for a population containing one (resp. more than one) class of genotype, for instance, to achieve specialization. A monomorphic population implies that individuals will behave in a similar manner (except for small variations due to minor genetic differences). On the contrary, polymorphic populations host multiple groups of individuals, each group with its particular genotypic signature, possibly displaying a specific behavior. Research to date shows that cooperation in monomorphic populations can be easily achieved (e.g. (Prieto et al., 2010; Schwarzer et al., 2010; Montanier and Bredeche, 2011, 2013; Silva et al., 2012)), while polymorphic populations (e.g., displaying genetic-encoded behavioral specialization) require very specific conditions to evolve (e.g., Trueba et al. (2013); Haasdijk et al. (2014a); Montanier et al. (2016)).

A notable number of contributions employ real robots. Since the first experiments in this field, the intrinsic online nature of embodied evolution has made such validation comparatively straightforward (Ficici et al., 1999; Watson et al., 2002). “Traditional” evolutionary robotics is more concerned with

TABLE 1 | Overview of Embodied Evolution research.

	Implementation				Robot behavior				Experimental settings		Mating conditions			Selection scheme		Replacement scheme					
	Distributed	Hybrid	Simulation	Real robots	Monomorphic	Polymorphic	Individual behavior	Cooperation	Division of labor	Task	No of robots	Pannictic	Proximity	Other	Performance	Random	Genotypic distance	Fixed lifetime	Variable lifetime	Event based	Limited lifetime
Ficici et al. (1999); Watson et al. (2002)	•			•	•		•			Phototaxis	8		•		•				•		
Simões and Dimond (2001)	•			•	•		•			Obstacle avoidance	6	•			•			•			
Usui and Arita (2003)		•		•	•		•			obstacle avoidance	6		•		•			•			
Bianco and Nolfi (2004)	•		•		•		•			Self-assembly	64		•			•					•
Hettiarachchi et al. (2006); Hettiarachchi and Spears (2009)	•		•		•		•			Navigation with obstacle avoidance	60		•		•						•
Wischmann et al. (2007)	•		•		•		•			Foraging ^a	3		•		•				•		
Perez et al. (2008)		•	•		•		•			Obstacle avoidance	5	•			•			•			
König and Schmeck (2009); König et al. (2009)	•		•		•		•			Obstacle avoidance with gate passing	26; 30		•		•			•			
Pugh and Martinoli (2009)		•	•	•	•		•			Obstacle avoidance	1–10	•	•		•			•			
Prieto et al. (2009); Trueba et al. (2011, 2012)	•		•			•		•	•	Surveillance, foraging, construction	20		•		•					•	•
Bredeche and Montanier (2010, 2012) Bredeche (2014)	•		•	•	•			•		None	20; 4000		•			•					•
Prieto et al. (2010)	•			•		•		•	•	Surveillance	8		•		•					•	•
Schwarzer et al. (2010)	•		•	•		•		•		None	Up to 40		•			•			•		
Schwarzer et al. (2011)	•		•		•		•			Phototaxis	4		•		•			•			
Montanier and Bredeche (2011, 2013)	•		•		•			•		None	100		•			•	•				•
Huijsman et al. (2011)	•	•	•		•		•			Obstacle avoidance	4–400	•		•	•			•			

(Continued)

TABLE 1 | Continued

	Implementation				Robot behavior				Task	Experimental settings No of robots	Mating conditions			Selection scheme		Replacement scheme				
	Distributed	Hybrid	Simulation	Real robots	Monomorphic	Polymorphic	Individual behavior	Cooperation			Division of labor	Panmictic	Proximity	Other	Performance	Random	Genotypic distance	Fixed lifetime	Variable lifetime	Event based
Karafotias et al. (2011)		•	•		•	•			•	Obstacle avoidance, phototaxis, and patrolling	10			•	•					
Silva et al. (2012, 2013, 2015, 2017)		•	•	•		•	•	•		Navigation, aggregation, surveillance, and phototaxis	2–20		•		•				•	
Weel et al. (2012a,b)		•	•		•			•		Foraging	10; 50		•		•				•	
García-Sánchez et al. (2012)		•	•		•	•				Obstacle avoidance	4–36		•	•	•				•	
Haasdijk and Bredeche (2013) Haasdijk et al. (2013, 2014a); Noskov et al. (2013); Haasdijk and Eigenhuis (2016); Bangel and Haasdijk (2017); Kemeling and Haasdijk (2017)	•		•			•	•	•		Foraging	100		•		•				•	
Trueba et al. (2013) Trueba (2017)	•		•	•	•			•	•	Synthetic mapping, gathering, self-location	40; 20; 9	•	•		•	•			•	•
O'Dowd et al. (2014)		•		•	•	•				Foraging	10		•		•				•	
Fernandez Pérez et al. (2014)	•		•		•			•		Foraging	50		•		•	•			•	
Fernandez Pérez et al. (2015)	•		•		•			•		Foraging	100		•		•				•	
Hart et al. (2015); Steyven et al. (2016)	•		•		•			•		Foraging	100		•		•				•	
Heinerman et al. (2015, 2016)		•	•	•	•	•				Obstacle avoidance	6	•			•				•	
Montanier et al. (2016); Bredeche et al. (2017)	•		•		•			•		Foraging	100; 500		•		•	•			•	
Fernandez Pérez et al. (2017)	•		•		•			•		Foraging	200		•		•				•	

^aAs a proxy for predator avoidance.

TABLE 2 | Glossary.

Field	Comment
Implementation	Distributed implementations have one genome for each robot, and an offspring is created only as the result of a mating event or by mutating the current genome. Hybrid implementations have multiple genomes per robot, and offspring can be created from this internal pool and from genomes “imported” through mating events. As stated earlier, the encapsulated scheme is not considered embodied evolution as there is no exchange of genomes between robots in this case. The experiments can use real robots or simulation .
Robot behavior	A monomorphic population contains individuals with similar genotypes (with variations due to mutation). A polymorphic population is divided into two (or more) subgroups of genetically similar individuals, and different genotypic signatures from one group to the other, e.g., to achieve specialization. We distinguish between experiments that target efficient individual behavior vs. collective behaviors (i.e., social behaviors, incl. cooperation)
Experimental settings	Identifies the task(s) considered in the experiment, e.g., obstacle avoidance, foraging, ... None indicates that there is no user-defined task and that consequently, selection pressure results from the environment only. The number of robots used is also included. $n_1 - n_2$ indicates the interval for one experiment and $n1; n2$ gives numbers for two experiments.
Mating conditions	Mating can be based on proximity : two robots can mate whenever they are physically close to each other (e.g., in infrared communication range). In panmictic systems, robots can mate with all other robots, regardless of their location. Other comprises systems where robots maintain an explicit list of potential mates (a social network), which may be maintained through gossiping.
Selection scheme	Parents are selected from the received and internal genomes on the basis of their performance if a task is defined. Random parent selection implies only environment-driven selection. Currently, the only examples of other selection schemes use genotypic distance, but this category also covers metrics such as novelty.
Replacement scheme	Genomes can have a fixed lifetime , variable lifetime , or limited lifetime (similar to variable lifetime, but with an upper bound). Event-based replacement schemes do not depend on time but on events such as reception of genetic material (e.g., in the microbial GA used by Watson et al. (2002)).

robustness at the level of the evolved *behavior* (mostly caused from the reality gap that exists between simulation and the real world) than is embodied evolution, which emphasizes the design of robust *algorithms*, where transfer between simulation and real world may be less problematic. In the contributions presented here, simulation is used for extensive analysis that could hardly take place with real robots due to time or economic constraint. Still, it is important to note that many researchers who use simulation have also published works

with real robots, thus including real-world validation in their research methodology.

Since 2010, there have been a number of experiments that employ large (≥ 100) numbers of (simulated) robots, shifting toward more swarm-like robotics where evolutionary dynamics can be quite different (Huijsman et al., 2011; Bredeche, 2014; Haasdijk et al., 2014b). Recent works in this vein focus on the nature of selection pressure, emphasizing the unique aspect of embodied evolution that selection pressure results from both the environment (which impacts mating) and the task. Bredeche and Montanier (2010, 2012) showed that environmental pressure alone can drive evolution toward self-sustaining behaviors. Haasdijk et al. (2014a) showed that these selection pressures can to some extent be modulated by tuning the ease with which robots can transmit genomes. Steyven et al. (2016) showed that adjusting the availability and value of energy resources results in the evolution of a range of different behaviors. These results emphasize that tailoring the environmental requirements to transmit genomes can profoundly impact the evolutionary dynamics and that understanding these effects is vital to effectively develop embodied evolution systems.

5. ISSUES IN EMBODIED EVOLUTION

What sets embodied evolution apart from classical evolutionary robotics (and, indeed, from most evolutionary computing) is the fact that evolution acts as a force for continuous adaptation, not (just) as an optimizer before deployment. As a continuous evolutionary process, embodied evolution is similar to some evolutionary systems considered in artificial life research (e.g. Axelrod (1984); Ray (1993), to name a few). The operations that implement the evolutionary process to adapt the robots' controllers are an integral part of their behavior in their task environment. This includes mating behavior to exchange and select genetic material, assessing one's own and/or each other's task performance (if a task is defined) and applying variation operators such as mutation and recombination.

This raises issues that are particular to embodied evolution. The research listed in the previous section has identified and investigated a number of these issues, and the remainder of this section discusses these issues in detail, while Section 6.2 discusses issues that so far have not benefited from close attention in embodied evolution research.

5.1. Local Selection

In embodied evolution, the evolutionary process is generally implemented through local interactions between the robots, i.e., the mating operation introduced above. This implies the concept of a neighborhood from which mates are selected. One common way to define neighborhood is to consider robots within communication range, but it can also be defined in terms of other distance measures such as genotypic or phenotypic distance. Mates are selected by sampling from this neighborhood, and a new individual is created by applying variation operators to the sampled genome(s). This local interaction has its origin in constraints that derive from communication limitations in some distributed robotic scenarios. Schut et al. (2009) showed it to

be beneficial in simulated setups as an exploration/exploitation balancing mechanism.

Embodied evolution, with chance encounters providing the sampling mechanism, has some similarities with other flavors of evolutionary computation. Cellular evolutionary algorithms (Alba and Dorronsoro, 2008) consider continuous random rewiring of a network topology (in a grid of CPUs or computers) where all elements are evaluated in parallel. In this context, locally selecting candidates for reproduction is a recurring theme that is shared with embodied evolution (e.g., García-Sánchez et al. (2012); Fernández Pérez et al. (2014)).

5.2. Objective Functions vs Selection Pressure

In traditional evolutionary algorithms, the optimization process is guided by a (set of) objective function(s) (Eiben and Smith, 2008). Evaluation of the candidate solutions, i.e., of the genomes in the population, allows for (typically numerical) comparison of their performance. Beyond its relevance for performance assessment, the evaluation process *per se* has generally no influence on the manner in which selection, variation, and replacement evolutionary operators are applied. This is different in embodied evolution, where the behavior of an individual can directly impact the likelihood of encounters with others and so influence selection and reproductive success (Bredeche and Montanier, 2010). Evolution can not only improve task performance but can also develop mating strategies, for example, by maximizing the number of encounters between robots if that improves the likelihood of transmitting genetic material.

It is therefore important to realize that the *selection pressure* on the robot population does not only derive from the specified *objective function(s)* as it traditionally does in evolutionary computation. In embodied evolution, the environment, including the mechanisms that allow mating, also exert selection pressure. Consequently, evolution experiences selection pressure from the aggregate of objective function(s) and environmental particularities. Steyven et al. (2016) researched how aspects of the robots' environment influence the emergence of particular behaviors and the balance between pressure toward survival and task. The objective may even pose requirements that are opposed to those by the environment. This can be the case when a task implies risky behaviors or because a task requires resources that are also needed for survival and mating. In such situations, the evolutionary process must establish a tradeoff between objective-driven optimization and the maintenance of a viable environment where evolution occurs, which is a challenge in itself (Haasdijk, 2015).

5.3. Autonomous Performance Evaluation

The decentralized nature of the evolutionary process implies that there is no omniscient presence who knows (let alone determines) the fitness values of all individuals. Consequently, when an objective function is defined, it is the robots themselves that must gauge their performance and share it with other robots when mating: each robot must have an evaluation function that can be computed onboard and autonomously. Typical examples of such evaluation functions are the number of resources collected,

the number of times a target has been reached, or the number of collisions. The requirement of autonomous assessment does not fundamentally change the way one defines fitness functions, but it does impact their usage as shown by Nordin and Banzhaf (1997), Walker et al. (2006), and Wolpert and Tumer (2008).

Evaluation time: The robots must run a controller for some time to assess the resultant behavior. This implies a *time sharing* scheme where robots run their current controllers to evaluate their performance. In many similar implementations, a robot runs a controller for a fixed evaluation time; Haasdijk et al. (2012) showed that this is a very important parameter in encapsulated online evolution, and it is likely to be similarly influential in embodied evolution.

Evaluation in varying circumstances: Because the evolutionary machinery (mating, evaluating new individuals, etc.) is an integral part of robot behavior, which runs in parallel with the performance of regular tasks, there can be no thorough re-initialization or re-positioning procedure between genome replacements. This implies a noisy evaluation: a robot may undervalue a genome starting in adverse circumstances and vice versa. As Nordin and Banzhaf (1997) (p. 121) put it: *“Each individual is tested against a different real-time situation leading to a unique fitness case. This results in ‘unfair’ comparison where individuals have to navigate in situations with very different possible outcomes. However, [...] experiments show that over time averaging tendencies of this learning method will even out the random effects of probabilistic sampling and a set of good solutions will survive.”* Bredeche et al. (2009) proposed a re-evaluation scheme to address this issue: seemingly efficient candidate solutions have a probability to be re-evaluated to cope with possible evaluation noise. A solution with a higher score *and* a lower variance will then be preferred to one with a higher variance. While re-evaluation is not always used in embodied evolution, the evaluation of relatively similar genomes on different robots running in parallel provides another way to smooth the effect of noisy evaluations.

Multiple objectives: To deal with multiple objectives, evolutionary computation techniques typically select individuals on the basis of Pareto dominance. While this is eminently possible when selecting partners as well, Pareto dominance can only be determined vis à vis the population sample that the selecting robot has acquired. It is unclear how this affects the overall performance and if the robot collective can effectively cover the Pareto front. Bangel and Haasdijk (2017) investigated the use of a “market mechanism” to balance the selection pressure over multiple tasks in a concurrent foraging scenario, showing that this at least prevents the robot collective from focusing on single tasks, but that it does not lead to specialization in individual robots.

6. DISCUSSION

The previous sections show that there is a considerable and increasing amount of research into embodied evolution, addressing issues that are particular to its autonomous and distributed nature. This section turns to the future of embodied evolution research, discussing potential applications and proposing a

research agenda to tackle some of the more relevant and immediate issues that so far have remained insufficiently addressed in the field.

6.1. Applications of Embodied Evolution

Embodied evolution can be used as a design method for engineering, as a modeling method for evolutionary biology, or as a method to investigate evolving complex systems more generally. Let us briefly consider each of these possibilities.

Engineering: The online adaptivity afforded by embodied evolution offers many novel possibilities for deployment of robot collectives: exploration of unknown environments, search and rescue, distributed monitoring of large objects or areas, distributed construction, distributed mining, etc. Embodied evolution can offer a solution when robot collectives are required to be versatile, since the robots can be deployed in and adapt to open and *a priori* unknown environments and tasks. The collective is comparatively robust to failure through redundancy and the decentralized nature of the algorithm because the system continues to function even if some robots break down. Embodied evolution can increase autonomy because the robots can, for instance, learn how to maintain energy while performing their task without intervention by an operator.

Currently, embodied evolution has already provided solutions to tasks such as navigation, surveillance, and foraging (see **Table 1** for a complete list), but these are of limited interest because of the simplicity of the tasks considered in research to date. The research agenda proposed in Section 6.2 provides some suggestions for further research to mitigate this.

Evolutionary biology: In the past 100 years, evolutionary biology benefited from both experimental and theoretical advances. It is now possible, for instance, to study evolutionary mechanisms through methods such as gene sequencing (Blount et al., 2012; Wisser et al., 2013). However, *in vitro* experimental evolution has its limitations: with evolution in “real” substrates, the time scales involved limit the applicability to relatively simple organisms such as *Escherichia coli* (Good et al., 2017). From a theoretical point of view, population genetics (see Charlesworth and Charlesworth (2010) for a recent introduction) provides a set of mathematically grounded tools for understanding evolution dynamics, at the cost of many simplifying assumptions.

Evolutionary robotics has recently gained relevance as an individual-based modeling and simulation method in evolutionary biology (Floreano and Keller, 2010; Waibel et al., 2011; Long, 2012; Mitri et al., 2013; Ferrante et al., 2015; Bernard et al., 2016), enabling the study of evolution in populations of robotic individuals in the physical world. Embodied evolution enables more accurate models of evolution because it is possible to embody not only the physical interactions but also the evolutionary operators themselves.

Synthetic approach: Embodied evolution can also be used to “understand by design” (Pfeifer and Scheier, 2001). As Maynard Smith (1992), a prominent researcher in evolutionary biology, advocated in a famous (Maynard Smith, 1992)’s Science paper (originally referring to Tierra (Ray, 1993)): “so far, we have been able to study only one evolving system and we cannot wait for

interstellar flight to provide us with a second. If we want to discover generalizations about evolving systems, we have to look at artificial ones.”

This *synthetic approach* stands somewhere between biology and engineering, using tools from the latter to understand mechanisms originally observed in nature and aiming at identifying general principles not confined to any particular (biological) substrate. Beyond improving our *understanding* of adaptive mechanisms, these general principles can also be used to improve our ability to *design* complex systems.

6.2. Research Agenda

We identify a number of open issues that need to be addressed so that embodied evolution can develop into a relevant technique to enable online adaptivity of robot collectives. Some of these issues have been researched in other fields (e.g., credit assignment is a well-known and often considered topic in reinforcement learning research). Lessons can and should be learned from there, inspiring embodied evolution research into the relevance and applicability of findings in those other fields.

In particular, we identify the following challenges.

Benchmarks: The pseudocode in Section 3 provides a clarification of embodied evolution’s concepts by describing the basic building blocks of the algorithm. This is only a first step toward a theoretical and practical framework for embodied evolution. Some authors have already taken steps in this direction. For instance, Prieto et al. (2015) propose an abstract algorithmic model to study both general and specific properties of embodied evolution implementations. Montanier et al. (2016) described “vanilla” versions of embodied evolution algorithms that can be used as practical benchmarks. Further exploration of abstract models for theoretical validation is needed. Also, standard benchmarks and test cases, along with systematically making the source code available, would provide a solid basis for empirical validation of individual contributions.

Evolutionary dynamics: Embodied evolution requires new tools for analyzing the evolutionary dynamics at work. Because the evolutionary operators apply *in situ*, the dynamics of the evolutionary process are not only important in the context of understanding or improving an optimization procedure, but they also have a direct bearing on how the robots behave and change their behavior when deployed.

Tools and methodologies to characterize the dynamics of evolving systems are available. The field of population genetics has produced techniques for estimating the selection pressure compared to genetic drift possibly occurring in finite-sized populations (see, for instance, Wakeley (2008) and Charlesworth and Charlesworth (2010) for a comprehensive introduction). Similarly, tools from adaptive dynamics (Geritz et al., 1998) can be used to investigate how particular solutions spread within the population. Finally, embodied evolution produces phylogenetic trees that can be studied either from a population genetics viewpoint (e.g., coalescence theory to understand the temporal structure of evolutionary adaptation) and graph theory (e.g., to characterize the particular structure of the inheritance graph). Boumaza (2017) shows an interesting first foray into using this technique to analyze embodied evolution.

Credit assignment: In all the research reviewed in this article that considers robot tasks, the fitness function is defined and implemented at the level of the individual robot: it assesses its own performance independently of the others. However, collectively solving a task often requires an assessment of performance at group level rather than individual level. This raises the issue of estimating each individual's contribution to the group's performance, which is unlikely to be completely captured by a fitness function (e.g. all individuals going toward the single larger food patch may not always be the best strategy if one aim to bring back the largest amount of food to the nest).

Closely related to our concern, Stone et al. (2010) formulated the *ad hoc teamwork* problem in multirobot systems, involving robots that each must “collaborate with previously unknown teammates on tasks to which they are all individually capable of contributing as team members.” As stated by Wolpert and Tumer (2008), this implies devoting substantial attention to the problem of estimating the *local utility* of individual agents with respect to the *global welfare* of the whole group and how to make a tradeoff between individual and group performance (e.g., Hardin (1968); Arthur (1994)).

While a generally applicable method to estimate an individual's local utility in an online distributed setting has so far eluded the community, it is possible to provide an exact assessment in controlled settings. Methods from cooperative game theory, such as computing the Shapley value (Shapley, 1953), could be used in embodied evolution but are computationally expensive and require the ability to replay experiments. However, replaying experiments is possible only with simulation and/or controlled experimental settings. While these methods cannot apply when robots are deployed in the real world, they at least provide a method to *compare* the outcome of candidate solutions to estimate individuals' marginal contributions and choose which should be deployed.

Social complexity: Section 4 shows that embodied evolution so far demonstrated only a limited set of social organization concepts: simple cooperative and division of labor behaviors. To address more complex tasks, we must first gain a better understanding of the mechanisms required to achieve complex collective behaviors. This raises two questions. First, there is an *ethological* question: what are the behavioral mechanisms at work in complex collective behaviors? Some of them, such as positive and negative feedback between individuals, or indirect communication through the environment (i.e., *stigmergy*), are well known from examples found both in biology (Camazine et al., 2003) and theoretical physics (Deutsch et al., 2012). Second, there is a question about the origins and stability of behaviors: what are the key elements that make it possible to evolve collective behaviors, and what are their limits? Again, evolutionary ecology provides relevant insights, such as the interplay between the level of cooperation and relatedness between individuals (West et al., 2007). The literature on such phenomena in biological systems may provide a good basis for research into the evolution of social complexity in embodied evolution.

A first step would be to clearly define the nature of social complexity that is to be studied. For this, evolutionary game

theory (Maynard Smith, 1992) has already produced a number of well-grounded and well-defined “games” that capture many problems involving interactions among individuals, including thorough analysis of the evolutionary dynamics in simplified setups. Of course, results obtained on abstract models may not be transferable within more realistic settings (as Bernard et al. (2016) showed for mutualistic cooperation), but the systematic use of a formal problem definition would greatly benefit the clarity of contributions in our domain.

Open-ended adaptation: As stated in Section 2, embodied evolution aims to provide continuous adaptation so that the robot collective can cope with changes in the objectives and/or the environment. Montanier and Bredeche (2011) showed that embodied evolution enables the population to react appropriately to changes in the regrowth rate of resources, but generally this aspect of embodied evolution has to date not been sufficiently addressed.

We reformulate the goal of continuous adaptation as providing *open-ended* adaptation, i.e., having the ability to continually keep exploring new behavioral patterns, constructing increasingly complex behaviors as required. Bedau et al. (2000), Soros and Stanley (2014), and Taylor et al. (2016) and others identified open-ended adaptation in artificial evolutionary systems as one of the big questions of artificial life. Open-ended adaptation in artificial systems, in particular in combination with learning relevant task behavior, has proved to be an elusive ambition.

A possible avenue to achieve this ambition may lie in the use of quality diversity approaches in embodied evolution. Recent research has considered *quality diversity* measures as a replacement (Lehman and Stanley, 2011) or additional (Mouret and Doncieux, 2012) objective to improve the population diversity and consequently the efficacy of evolution. To date, such research has focused on the evolution of behavior for particular tasks with task-specific metrics of behavioral diversity that must be tailored for each application. To be able to exploit quality diversity in unknown environments and for arbitrary tasks, generic measures of behavioral diversity must be developed.

Another avenue of research would be to take inspiration from the behavior of a passerine bird, the great tit (*parus major*), as recently analyzed by Aplin et al. (2017). It appears that great tits combine collective and individual learning with varying intensity as they age and that the motivations to pursue behaviors also vary with age. Reward-based learning occurs primarily in young birds and is often individual, while adult birds engage mostly in social learning to copy the behavior that is most common, regardless of whether it produces more or less rewards than alternative behavior. This combination of conformist and payoff-sensitive reinforcement allows individuals and populations both to acquire adaptive behavior and to track environmental change.

Combining embodied evolution, individual reinforcement learning with task-based and diversity-enhancing objectives may yield similar behavioral plasticity for collectives of robots.

Safety and robot ethics: To deploy the kind of adaptive technology that embodied evolution aims for responsibly, one must ensure that the adaptivity can be controlled: autonomous

adaptation carries the risk of adaptation developing in directions that do not meet the needs of human users or that they even may find undesirable. Even so, the adaptive process should be curtailed as little as possible to allow effective, open-ended, learning. The user cannot be expected to monitor and closely control the robot's behavior and learning process; this may in fact be impossible in exactly those scenarios where robotic autonomy is most beneficial and adaptivity most urgently required. There is growing awareness that it may be necessary to endow robots with innately ethical behavior (e.g., Moor (2006); Anderson and Anderson (2007); Vanderelst and Winfield (2018)), where the systems select actions based on a "moral arithmetic" (Bentham, 1878), often informed by casuistry, i.e., generalizing morality on the basis of example cases in which there is agreement concerning the correct response (Anderson and Anderson, 2007). Moral reasoning along these lines could conceivably be enabled in embodied evolution as well, in which case interactive evolution to develop surrogate models of user requirements may offer one possible route to allow user guidance.

Additional open issues and opportunities will no doubt arise from advances in this and other fields. A relevant recent development, for instance, is the possibility of evolvable morphofunctional machines that are able to change both their software and hardware features (Eiben and Smith, 2015) and replicate through 3D printing (Brodbeck et al., 2015). This would allow embodied evolution holistically to adapt the robots' morphologies as well as their controllers. This can have profound consequences for embodied evolution implementations that exploit these developments: it would, for instance, enable dynamic population sizes, allowing for more risky behavior as broken robots could be replaced or recycled.

REFERENCES

- Alba, E. (2002). Parallel evolutionary algorithms can achieve super-linear performance. *Inf. Process. Lett.* 82, 7–13. doi:10.1016/S0020-0190(01)00281-2
- Alba, E., and Dorronsoro, B. (2008). *Cellular Genetic Algorithms*. Berlin, Heidelberg, New York: Springer-Verlag.
- Amato, C., Konidaris, G. G., Cruz, G., Maynor, C. A., How, J. P., and Kaelbling, L. P. (2015). "Planning for decentralized control of multiple robots under uncertainty," in *IEEE International Conference on Robotics and Automation (ICRA) (IEEE)*, 1241–1248.
- Anderson, M., and Anderson, S. L. (2007). Machine ethics: creating an ethical intelligent agent. *AI Mag.* 28, 15–26. doi:10.1609/aimag.v28i4.2065
- Aplin, L. M., Sheldon, B. C., and McElreath, R. (2017). Conformity does not perpetuate suboptimal traditions in a wild population of songbirds. *Proc. Natl. Acad. Sci. U.S.A.* 114, 7830–7837. doi:10.1073/pnas.1621067114
- Arthur, W. (1994). Inductive reasoning and bounded rationality. *Am. Econ. Rev.* 84, 406–411. doi:10.2307/2117868
- Axelrod, R. (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Bangel, S., and Haasdijk, E. (2017). "Reweight rewards in embodied evolution to achieve a balanced distribution of labour," in *Proceedings of the 14th European Conference on Artificial Life ECAL 2017* (Cambridge, MA: MIT Press), 44–51.
- Barrett, S., Rosenfeld, A., Kraus, S., and Stone, P. (2016). Making friends on the fly: cooperating with new teammates. *Artif. Intell.* 242, 1–68. doi:10.1016/j.artint.2016.10.005
- Bayindir, L. (2016). A review of swarm robotics tasks. *Neurocomputing* 172, 292–321. doi:10.1016/j.neucom.2015.05.116
- Bedau, M. A., McCaskill, J. S., Packard, N. H., Rasmussen, S., Adami, C., Green, D. G., et al. (2000). Open problems in artificial life. *Artif. Life* 6, 363–376. doi:10.1162/106454600300103683

7. CONCLUSION

This article provides an overview of embodied evolution for robot collectives, a research field that has been growing since its inception around the turn of the millennium. The main contribution of this article is threefold. First, it clarifies the definitions and overall process of embodied evolution. Second, it presents an overview of embodied evolution research conducted to date. Third, it provides directions for future researches.

This overview sheds light on the maturity of the field: while embodied evolution was mostly used as a parallel search method for designing individual behavior during its first decade of existence, a trend has emerged toward its collective aspects (i.e., cooperation, division of labor, specialization). This trend goes hand in hand with a trend toward larger, swarm-like, robot collectives.

We hope this overview will provide a stepping stone for the field, accounting for its maturity and acting as an inspiration for aspiring researchers. To this end, we highlighted possible applications and open issues that may drive the field's research agenda.

AUTHOR CONTRIBUTIONS

All authors listed have made a substantial, direct, and intellectual contribution to the work equally and approved it for publication.

ACKNOWLEDGMENTS

The authors gratefully acknowledge the support from the European Union's Horizon 2020 research and innovation program under grant agreement No 640891. The authors would also like to thank A.E. Eiben and Jean-Marc Montanier for their support during the writing of this paper.

- Bellingham, J. G., and Rajan, K. (2007). Robotics in remote and hostile environments. *Science* 318, 1098–1102. doi:10.1126/science.1146230
- Beni, G. (2005). From swarm intelligence to swarm robotics. *Robotics* 3342, 1–9. doi:10.1007/978-3-540-30552-1_1
- Bentham, J. (1878). *Introduction to the Principles of Morals and Legislation*. Oxford, UK: Clarendon.
- Bernard, A., André, J.-B., and Bredeche, N. (2016). To cooperate or not to cooperate: why behavioural mechanisms matter. *PLoS Comput. Biol.* 12:e1004886. doi:10.1371/journal.pcbi.1004886
- Bernstein, D. S., Givan, R., Immerman, N., and Zilberstein, S. (2002). The complexity of decentralized control of Markov decision processes. *Math. Oper. Res.* 27, 819–840. doi:10.1287/moor.27.4.819.297
- Bianco, R., and Nolfi, S. (2004). Toward open-ended evolutionary robotics: evolving elementary robotic units able to self-assemble and self-reproduce. *Connect. Sci.* 16, 227–248. doi:10.1080/09540090412331314759
- Blount, Z. D., Barrick, J. E., Davidson, C. J., and Lenski, R. E. (2012). Genomic analysis of a key innovation in an experimental *Escherichia coli* population. *Nature* 489, 513–518. doi:10.1038/nature11514
- Bongard, J., Zykov, V., and Lipson, H. (2006). Resilient machines through continuous self-modeling. *Science* 314, 1118–1121. doi:10.1126/science.1133687
- Bongard, J. C. (2013). Evolutionary robotics. *Commun. ACM* 56, 74–83. doi:10.1145/2492007.2493883
- Boumaza, A. (2017). "Phylogeny of embodied evolutionary robotics," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion on – GECCO '17* (New York, NY: ACM Press), 1681–1682.
- Brambilla, M., Ferrante, E., Birattari, M., and Dorigo, M. (2012). Swarm robotics: a review from the swarm engineering perspective. *Swarm Intell.* 7, 1–41. doi:10.1007/s11721-012-0075-2

- Bredeche, N. (2014). "Embodied evolutionary robotics with large number of robots," in *Artificial Life 14: Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, 272–273.
- Bredeche, N., Haasdijk, E., and Eiben, A. E. (2009). On-line, on-board evolution of robot controllers. *Lect. Notes Comput. Sci.* 5975, 110–121. doi:10.1007/978-3-642-14156-0_10
- Bredeche, N., and Montanier, J.-M. (2010). "Environment-driven embodied evolution in a population of autonomous agents," in *Proceedings of the 11th International Conference on Parallel Problem Solving from Nature: Part II* (Berlin, Heidelberg: Springer-Verlag), 290–299.
- Bredeche, N., and Montanier, J.-M. (2012). "Environment-driven open-ended evolution with a population of autonomous robots," in *Evolving Physical Systems Workshop*, 7–14.
- Bredeche, N., Montanier, J.-M., and Carrignon, S. (2017). "Benefits of proportionate selection in embodied evolution: a case study with behavioural specialization," in *Proceedings of the Genetic and Evolutionary Computation Conference Companion*, 1683–1684.
- Brodbeck, L., Hauser, S., and Iida, F. (2015). Morphological evolution of physical robots through model-free phenotype development. *PLoS ONE* 10:e0128444. doi:10.1371/journal.pone.0128444
- Camazine, S., Deneubourg, J.-L., Franks, N., Sneyd, J., Theraulaz, G., and Bonabeau, E. (2003). *Self-Organization in Biological Systems*. Princeton University Press.
- Charlesworth, B., and Charlesworth, D. (2010). *Elements of Evolutionary Genetics*. Greenwood Village: Roberts and Company Publishers.
- Christensen, D. J., Spröwitz, A., and Ijspeert, A. J. (2010). "Distributed online learning of central pattern generators in modular robots," in *From Animals to Animats 11: Proceedings of the 11th international conference on Simulation of Adaptive Behavior*, 402–412.
- Cully, A., Clune, J., Tarapore, D., and Mouret, J. B. (2015). Robots that can adapt like animals. *Nature* 521, 503–507. doi:10.1038/nature14422
- Deutsch, A., Theraulaz, G., and Vicsek, T. (2012). Collective motion in biological systems. *Interface Focus* 2, 689–692. doi:10.1098/rsfs.2012.0048
- Dibangoye, J. S., Amato, C., Buffet, O., and Charpillet, F. (2015). "Exploiting separability in multiagent planning with continuous-state MDPs," in *IJCAI International Joint Conference on Artificial Intelligence*, 4254–4260.
- Doncieux, S., Bredeche, N., Mouret, J.-B., and Eiben, A. (2015). Evolutionary robotics: what, why, and where to. *Front. Robot. AI* 2:4. doi:10.3389/frobt.2015.00004
- Eiben, A., and Smith, J. (2008). *Introduction to Evolutionary Computing*. Springer.
- Eiben, A. E., Haasdijk, E., and Bredeche, N. (2010). "Embodied, on-line, on-board evolution for autonomous robotics," in *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, Vol. 5.2, eds P. Levi and S. Kernbach (Springer), 361–382.
- Eiben, A. E., Schoenauer, M., Laredo, J. L. J., Castillo, P. A., Mora, A. M., and Merelo, J. J. (2007). "Exploring selection mechanisms for an agent-based distributed evolutionary algorithm," in *Proceedings of the 9th Annual Conference Companion on Genetic and Evolutionary Computation*, 2801–2808.
- Eiben, A. E., and Smith, J. (2015). From evolutionary computation to the evolution of things. *Nature* 521, 476–482. doi:10.1038/nature14544
- Fernandez Pérez, I., Boumaza, A., and Charpillet, F. (2014). "Comparison of selection methods in on-line distributed evolutionary robotics," in *Proceedings of the Fourteenth International Conference on the Synthesis and Simulation of Living Systems*, 1–16.
- Fernandez Pérez, I., Boumaza, A., and Charpillet, F. (2015). "Decentralized innovation marking for neural controllers in embodied evolution," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 161–168.
- Fernandez Pérez, I., Boumaza, A., and Charpillet, F. (2017). "Learning collaborative foraging in a swarm of robots using embodied evolution," in *Proceedings of the 14th European Conference on Artificial Life ECAL 2017* (Cambridge, MA: MIT Press), 162–161.
- Ferrante, E., Turgut, A. E., Duéñez-Guzman, E., Dorigo, M., and Wenseleers, T. (2015). Evolution of self-organized task specialization in robot swarms. *PLoS Comput. Biol.* 11:e1004273. doi:10.1371/journal.pcbi.1004273
- Ficci, S. G., Watson, R. A., and Pollack, J. B. (1999). "Embodied evolution: a response to challenges in evolutionary robotics," in *Proceedings of the Eighth European Workshop on Learning Robots*, eds J. L. Wyatt and J. Demiris 14–22.
- Floreano, D., and Keller, L. (2010). Evolution of adaptive behaviour in robots by means of Darwinian selection. *PLoS Biol.* 8:e1000292. doi:10.1371/journal.pbio.1000292
- García-Sánchez, P., Eiben, A. E., Haasdijk, E., Weel, B., and Merelo-Guervós, J.-J. (2012). "Testing diversity-enhancing migration policies for hybrid on-line evolution of robot controllers," in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 7248 (LNCS), 52–62.
- Gauci, M., Chen, J., Dodd, T. J., and Groß, R. (2012). "Evolving aggregation behaviors in multi-robot systems with binary sensors," in *Int. Symposium on Distributed Autonomous Robotic Systems (DARS 2012)* (Springer).
- Geritz, S., Kisdi, E., Meszner, G., and Metz, J. (1998). Evolutionarily singular strategies and the adaptive growth and branching of the evolutionary tree. *Evol. Ecol.* 12, 35–57. doi:10.1023/A:1006554906681
- Good, B. H., McDonald, M. J., Barrick, J. E., Lenski, R. E., and Desai, M. M. (2017). The dynamics of molecular evolution over 60,000 generations. *Nature* 551, 45–50. doi:10.1038/nature24287
- Haasdijk, E. (2015). "Combining conflicting environmental and task requirements in evolutionary robotics," in *IEEE 9th International Conference on Self-Adaptive and Self-Organizing Systems*, 131–137.
- Haasdijk, E., and Bredeche, N. (2013). "Controlling task distribution in MONEE," in *Advances in Artificial Life, ECAL 2013* (MIT Press), 671–678.
- Haasdijk, E., Bredeche, N., and Eiben, A. E. (2014a). Combining environment-driven adaptation and task-driven optimisation in evolutionary robotics. *PLoS ONE* 9:e98466. doi:10.1371/journal.pone.0098466
- Haasdijk, E., Bredeche, N., Nolfi, S., and Eiben, A. E. (2014b). Evolutionary robotics. *Evol. Intell.* 7, 69–70. doi:10.1007/s12065-014-0113-7
- Haasdijk, E., and Eigenhuis, F. (2016). "Increasing reward in biased natural selection decreases task performance," in *Artificial Life {XV}: Proceedings of the 15th International Conference on Artificial Life*, 314–322.
- Haasdijk, E., Smit, S. K., and Eiben, A. E. (2012). Exploratory analysis of an on-line evolutionary algorithm in simulated robots. *Evol. Intell.* 5, 213–230. doi:10.1007/s12065-012-0083-6
- Haasdijk, E., Weel, B., and Eiben, A. E. (2013). "Right on the MONEE," in *Proceeding of the Fifteenth Annual Conference on Genetic and Evolutionary Computation Conference*, ed. C. Blum (New York, NY: ACM), 207–214.
- Hardin, G. (1968). The tragedy of the commons. *Science* 162, 1243–1248. doi:10.1126/science.162.3859.1243
- Hart, E., Steyven, A., and Paechter, B. (2015). "Improving survivability in environment-driven distributed evolutionary algorithms through explicit relative fitness and fitness proportionate communication," in *Proceedings of the 2015 Annual Conference on Genetic and Evolutionary Computation*, 169–176.
- Hauert, S., Zufferey, J.-C., and Floreano, D. (2008). Evolved swarming without positioning information: an application in aerial communication relay. *Auton. Robot.* 26, 21–32. doi:10.1007/s10514-008-9104-9
- Heinerman, J., Drupsteen, D., and Eiben, A. E. (2015). "Three-fold adaptivity in groups of robots: the effect of social learning," in *Proceedings of the 17th Annual Conference on Genetic and Evolutionary Computation, GECCO '15*, ed. S. Silva (ACM), 177–183.
- Heinerman, J., Rango, M., and Eiben, A. E. (2016). "Evolution, individual learning, and social learning in a swarm of real robots," in *Proceedings – 2015 IEEE Symposium Series on Computational Intelligence, SSCI 2015* (IEEE), 1055–1062.
- Hettiarachchi, S., and Spears, W. M. (2009). Distributed adaptive swarm for obstacle avoidance. *Int. J. Intell. Comput. Cybern.* 2, 644–671. doi:10.1108/17563780911005827
- Hettiarachchi, S., Spears, W. M., Green, D., and Kerr, W. (2006). "Distributed agent evolution with dynamic adaptation to local unexpected scenarios," in *Innovative Concepts for Autonomic and Agent-Based Systems, Volume LNCS 3825*, eds M. G. Hinchey, P. Rago, J. L. Rash, C. A. Rouff, R. Sterritt, and W. Truszkowski (Greenbelt, MD: Springer), 245–256.
- Huijsman, R.-J., Haasdijk, E., and Eiben, A. E. (2011). "An on-line on-board distributed algorithm for evolutionary robotics," in *Artificial Evolution, 10th International Conference Evolution Artificielle, Number 7401* in LNCS, eds J.-K. Hao, P. Legrand, P. Collet, N. Monmarché, E. Lutton, and M. Schoenauer (Springer), 73–84.
- Jakobi, N., Husbands, P., and Harvey, I. (1995). Noise and the reality gap: the use of simulation in evolutionary robotics. *Lect. Notes Comput. Sci.* 929, 704–720. doi:10.1007/3-540-59496-5_337
- Karafotias, G., Haasdijk, E., and Eiben, A. E. (2011). "An algorithm for distributed on-line, on-board evolutionary robotics," in *GECCO '11: Proceedings of the 13th*

- Annual Conference on Genetic and Evolutionary Computation* (ACM Press), 171–178.
- Kemeling, M., and Haasdijk, E. (2017). “Incorporating user feedback in embodied evolution,” in *Proceedings of the Genetic and Evolutionary Computation Conference Companion on – GECCO ’17* (New York, NY: ACM Press), 1685–1686.
- König, L., Mostaghim, S., and Schmeck, H. (2009). Decentralized evolution of robotic behavior using finite state automata. *Int. J. Intell. Comput. Cybern. 2*, 695–723. doi:10.1108/17563780911005845
- König, L., and Schmeck, H. (2009). “A completely evolvable genotype-phenotype mapping for evolutionary robotics,” in *SASO 2009 – 3rd IEEE International Conference on Self-Adaptive and Self-Organizing Systems*, 175–185.
- Lehman, J., and Stanley, K. O. (2011). Abandoning objectives: evolution through the search for novelty alone. *Evol. Comput. 19*, 189–223. doi:10.1162/EVCO_a_00025
- Long, J. (2012). *Darwin’s Devices: What Evolving Robots Can Teach Us about the History of Life and the Future of Technology*. Basic Books.
- Mataric, M. J. (1994). *Interaction and Intelligent Behavior*. PhD thesis.
- Maynard Smith, J. (1992). Evolutionary biology. Byte-sized evolution. *Nature 355*, 772–773. doi:10.1038/355772a0
- Mitri, S., Wischmann, S., Floreano, D., and Keller, L. (2013). Using robots to understand social behaviour. *Biol. Rev. Camb. Philos. Soc. 88*, 31–39. doi:10.1111/j.1469-185X.2012.00236.x
- Montanier, J.-M., and Bredeche, N. (2011). “Surviving the tragedy of commons: emergence of altruism in a population of evolving autonomous agents,” in *European Conference on Artificial Life* (Paris, France), 550–557.
- Montanier, J.-M., and Bredeche, N. (2013). “Evolution of altruism and spatial dispersion: an artificial evolutionary ecology approach,” in *Advances in Artificial Life, ECAL 2013*, 260–267.
- Montanier, J.-M., Carrignon, S., and Bredeche, N. (2016). Behavioural specialization in embodied evolutionary robotics: why so difficult? *Front. Robot. AI 3:38*. doi:10.3389/frobot.2016.00038
- Moor, J. H. (2006). The nature, importance, and difficulty of machine ethics. *IEEE Intell. Syst. 21*, 18–21. doi:10.1109/MIS.2006.80
- Mouret, J. B., and Doncieux, S. (2012). Encouraging behavioral diversity in evolutionary robotics: an empirical study. *Evol. Comput. 20*, 91–133. doi:10.1162/EVCO_a_00048
- Mouret, J. B., and Tonelli, P. (2015). Artificial evolution of plastic neural networks: a few key concepts. *Stud. Comput. Intell. 557*, 251–261. doi:10.1007/978-3-642-55337-0_9
- Nelson, A. L., and Grant, E. (2006). Using direct competition to select for competent controllers in evolutionary robotics. *Rob. Auton. Syst. 54*, 840–857. doi:10.1016/j.robot.2006.04.010
- Nolfi, S., and Floreano, D. (2000). *Evolutionary Robotics: The Biology, Intelligence, and Technology*. Cambridge, MA: MIT Press.
- Nordin, P., and Banzhaf, W. (1997). An on-line method to evolve behavior and to control a miniature robot in real time with genetic programming. *Adapt. Behav. 5*, 107–140. doi:10.1177/105971239700500201
- Noskov, N., Haasdijk, E., Weel, B., and Eiben, A. E. (2013). “MONEE: using parental investment to combine open-ended and task-driven evolution,” in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 7835 LNCS, 569–578.
- Nouyan, S., Gross, R., Bonani, M., Mondada, F., and Dorigo, M. (2009). Teamwork in self-organized robot colonies. *IEEE Trans. Evol. Comput. 13*, 695–711. doi:10.1109/TEVC.2008.2011746
- O’Dowd, P. J., Studley, M., and Winfield, A. F. T. (2014). The distributed co-evolution of an on-board simulator and controller for swarm robot behaviours. *Evol. Intell. 7*, 95–106. doi:10.1007/s12065-014-0112-8
- Parker, L. E. (2008). “Multiple mobile robot systems,” in *Handbook of Robotics*, Chap. 40, ed. S. B. Heidelberg (Springer), 921–941.
- Perez, A. L. F., Bittencourt, G., and Roisenberg, M. (2008). “Embodied evolution with a new genetic programming variation algorithm,” in *Proceedings – 4th International Conference on Autonomic and Autonomous Systems, ICAS 2008* (Los Alamitos, CA: IEEE Press), 118–123.
- Pfeifer, R., and Scheier, C. (2001). *Understanding Intelligence*. MIT Press.
- Prieto, A., Becerra, J., Bellas, F., and Duro, R. (2010). Open-ended evolution as a means to self-organize heterogeneous multi-robot systems in real time. *Rob. Auton. Syst. 58*, 1282–1291. doi:10.1016/j.robot.2010.08.004
- Prieto, A., Bellas, F., and Duro, R. J. (2009). Adaptively coordinating heterogeneous robot teams through asynchronous situated coevolution. *Lect. Notes Comput. Sci. 5864(PART 2)*, 75–82. doi:10.1007/978-3-642-10684-2_9
- Prieto, A., Bellas, F., Trueba, P., and Duro, R. J. (2015). Towards the standardization of distributed embodied evolution. *Inf. Sci. 312*, 55–77. doi:10.1016/j.ins.2015.03.044
- Prieto, A., Bellas, F., Trueba, P., and Duro, R. J. (2016). Real-time optimization of dynamic problems through distributed embodied evolution. *Integr. Comput. Aided Eng. 23*, 237–253. doi:10.3233/ICA-160522
- Pugh, J., and Martinoli, A. (2009). Distributed scalable multi-robot learning using particle swarm optimization. *Swarm Intell. 3*, 203–222. doi:10.1007/s11721-009-0030-z
- Ray, T. S. (1993). An evolutionary approach to synthetic biology: Zen and the art of creating life. *Artif. Life 1*, 179–209. doi:10.1162/artl.1993.1.1_2.179
- Rubenstein, M., Cornejo, A., and Nagpal, R. (2014). Programmable self-assembly in a thousand-robot swarm. *Science 345*, 795–799. doi:10.1126/science.1254295
- Schut, M. C., Haasdijk, E., and Prieto, A. (2009). “Is situated evolution an alternative for classical evolution?” in *IEEE Congress on Evolutionary Computation, CEC 2009*, 2971–2976.
- Schwarzer, C., Hösler, C., and Michiels, N. (2010). “Artificial sexuality and reproduction of robot organisms,” in *Symbiotic Multi-Robot Organisms: Reliability, Adaptability, Evolution*, eds P. Levi and S. Kernbach (Berlin, Heidelberg, New York: Springer-Verlag), 384–403.
- Schwarzer, C., Schlachter, F., and Michiels, N. K. (2011). Online evolution in dynamic environments using neural networks in autonomous robots. *Int. J. Adv. Intell. Syst. 4*, 288–298.
- Shapley, L. S. (1953). A value for n-person games. *Ann. Math. Stud. 28*, 307–317. doi:10.1515/9781400881970-018
- Silva, F., Correia, L., and Christensen, A. L. (2013). “Dynamics of neuronal models in online neuroevolution of robotic controllers,” in *Lecture Notes in Computer Science (Including Subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, Vol. 8154, 90–101.
- Silva, F., Correia, L., and Christensen, A. L. (2017). Evolutionary online behaviour learning and adaptation in real robots. *R. Soc. Open Sci. 4*, 160938. doi:10.1098/rsos.160938
- Silva, F., Duarte, M., Correia, L., Oliveira, S. M., and Christensen, A. L. (2016). Open issues in evolutionary robotics. *Evol. Comput. 24*, 205–236. doi:10.1162/EVCO_a_00172
- Silva, F., Urbano, P., Correia, L., and Christensen, A. L. (2015). odNEAT: an algorithm for decentralised online evolution of robotic controllers. *Evol. Comput. 23*, 421–449. doi:10.1162/EVCO_a_00141
- Silva, F., Urbano, P., Oliveira, S., and Christensen, A. L. (2012). odNEAT: an algorithm for distributed online, onboard evolution of robot behaviours. *Artif. Life 13*, 251–258. doi:10.7551/978-0-262-31050-5-ch034
- Simões, E. D. V., and Dimond, K. R. K. (2001). Embedding a distributed evolutionary system into a population of autonomous mobile robots. *IEEE Int. Conf. Syst. Man Cybern. 2*, 1069–1074. doi:10.1109/ICSMC.2001.973061
- Soros, L. B., and Stanley, K. K. O. (2014). “Identifying necessary conditions for open-ended evolution through the artificial life world of chromaria,” in *Proc. of Artificial Life Conference (ALife 14)*, 793–800.
- Stanley, K. O., Bryant, B. D., and Miikkulainen, R. (2005). Real-time neuroevolution in the NERO video game. *IEEE Trans. Evol. Comput. 9*, 653–668. doi:10.1109/TEVC.2005.856210
- Steyven, A., Hart, E., and Paechter, B. (2016). “Understanding environmental influence in an open-ended evolutionary algorithm,” in *Parallel Problem Solving from Nature – PPSN XIV: 14th International Conference, Edinburgh, UK, September 17–21, 2016, Proceedings*, eds J. Handl, E. Hart, P. R. Lewis, M. López-Ibáñez, G. Ochoa, and B. Paechter (Cham: Springer), 921–931.
- Stone, P., Kaminka, G. A., Kraus, S., and Rosenschein, J. S. (2010). “Ad hoc autonomous agent teams: collaboration without pre-coordination,” in *The Twenty-Fourth Conference on Artificial Intelligence (AAAI)*.
- Stone, P., Sutton, R. S., and Kuhlmann, G. (2005). Reinforcement learning for RoboCup-soccer keep away. *Adapt. Behav. 13*, 165–188. doi:10.1177/105971230501300301
- Stone, P., and Veloso, M. (1998). Layered approach to learning client behaviors in the RoboCup soccer server. *Appl. Artif. Intell. 12*, 165–188. doi:10.1080/088395198117811

- Taylor, T., Bedau, M., Channon, A., Ackley, D., Banzhaf, W., Beslon, G., et al. (2016). Open-ended evolution: perspectives from the OEE workshop in York. *Artif. Life* 22, 408–423. doi:10.1162/ARTL
- Thrun, S., and Mitchell, T. M. (1995). Lifelong robot learning. *Rob. Auton. Syst.* 15, 25–46. doi:10.1016/0921-8890(95)00004-Y
- Tonelli, P., and Mouret, J. B. (2013). On the relationships between generative encodings, regularity, and learning abilities when evolving plastic artificial neural networks. *PLoS ONE* 8:e79138. doi:10.1371/journal.pone.0079138
- Trianni, V., Nolfi, S., and Dorigo, M. (2008). “Evolution, self-organization and swarm robotics,” in *Swarm Intelligence. Natural Computing Series*, eds C. Blum and D. Merkle (Berlin, Heidelberg: Springer).
- Trueba, P. (2017). “Embodied evolution versus cooperative,” in *Proceeding GECCO '17, Proceedings of the Genetic and Evolutionary Computation Conference Companion*, Berlin, Germany (New York, NY, USA: ACM).
- Trueba, P., Prieto, A., Bellas, F., Caamaño, P., and Duro, R. J. (2011). “Task-driven species in evolutionary robotic teams,” in *International Work-Conference on the Interplay between Natural and Artificial Computation* (Springer), 138–147.
- Trueba, P., Prieto, A., Bellas, F., Caamaño, P., and Duro, R. J. (2012). “Self-organization and specialization in multiagent systems through open-ended natural evolution,” in *Lecture Notes in Computer Science, Volume 7248 LNCS of Lecture Notes in Computer Science* (Berlin, Heidelberg: Springer), 93–102.
- Trueba, P., Prieto, A., Bellas, F., Caamaño, P., and Duro, R. J. (2013). Specialization analysis of embodied evolution for robotic collective tasks. *Rob. Auton. Syst.* 61, 682–693. doi:10.1016/j.robot.2012.08.005
- Urzelai, J., and Floreano, D. (2001). Evolution of adaptive synapses: robots with fast adaptive behavior in new environments. *Evol. Comput.* 9, 495–524. doi:10.1162/10636560152642887
- Usui, Y., and Arita, T. (2003). “Situated and embodied evolution in collective evolutionary robotics,” in *Proceedings of the 8th International Symposium on Artificial Life and Robotics, Number c*, 212–215.
- Vanderelst, D., and Winfield, A. (2018). An architecture for ethical robots inspired by the simulation theory of cognition. *Cogn. Syst. Res.* 48, 56–66. doi:10.1016/j.cogsys.2017.04.002
- Waibel, M., Floreano, D., and Keller, L. (2011). A quantitative test of Hamilton's rule for the evolution of altruism. *PLoS Biol.* 9:e1000615. doi:10.1371/journal.pbio.1000615
- Wakeley, J. (2008). *Coalescent Theory, an Introduction*. Roberts & Company Publishers.
- Walker, J. H., Garrett, S. M., and Wilson, M. S. (2006). The balance between initial training and lifelong adaptation in evolving robot controllers. *IEEE Trans. Syst. Man Cybern. B Cybern.* 36, 423–432. doi:10.1109/TSMCB.2005.859082
- Watson, R. A., Ficici, S. G., and Pollack, J. B. (2002). Embodied evolution: distributing an evolutionary algorithm in a population of robots. *Rob. Auton. Syst.* 39, 1–18. doi:10.1016/S0921-8890(02)00170-7
- Weel, B., Haasdijk, E., and Eiben, A. E. (2012a). “The emergence of multi-cellular robot organisms through on-line on-board evolution,” in *Applications of Evolutionary Computation, Volume 7248 of Lecture Notes in Computer Science*, ed. C. Di Chio (Springer), 124–134.
- Weel, B., Hoogendoorn, M., and Eiben, A. E. (2012b). “On-line evolution of controllers for aggregating swarm robots in changing environments,” in *PPSN XII: Proceedings of the 12th International Conference on Parallel Problem Solving from Nature, LNCS 7492*, 245–254.
- Werfel, J., Petersen, K., and Nagpal, R. (2014). Designing collective behavior in a termite-inspired robot construction team. *Science* 343, 754–758. doi:10.1126/science.1245842
- West, S. A., Griffin, A. S., and Gardner, A. (2007). Social semantics: altruism, cooperation, mutualism, strong reciprocity and group selection. *J. Evol. Biol.* 20, 415–432. doi:10.1111/j.1420-9101.2006.01258.x
- Wischmann, S., Stamm, K., and Florentin, W. (2007). “Embodied evolution and learning: the neglected timing of maturation,” in *ECAL 2007: Advances in Artificial Life*, Vol. 4648 (Springer), 284–293.
- Wiser, M. J., Ribbeck, N., and Lenski, R. E. (2013). Long-term dynamics of adaptation in asexual populations. *Science* 342, 1364–1367. doi:10.1126/science.1243357
- Wolpert, D. H., and Tumer, K. (2008). *An Introduction to Collective Intelligence*. Technical Report, NASA.

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2018 Bredeche, Haasdijk and Prieto. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.