

Embodied Evolution: Distributing an Evolutionary Algorithm in a Population of Robots

Richard A. Watson, Sevan G. Ficici, Jordan B. Pollack

*Dynamical and Evolutionary Machine Organization
Volen Center for Complex Systems, Brandeis University
Waltham Massachusetts USA
www.demo.cs.brandeis.edu*

Abstract

We introduce Embodied Evolution (EE) as a new methodology for evolutionary robotics (ER). EE uses a population of physical robots that autonomously reproduce with one another while situated in their task environment. This constitutes a fully-distributed evolution algorithm embodied in physical robots. Several issues identified by researchers in the evolutionary robotics community as problematic for the development of ER are alleviated by the use of a large number of robots being evaluated in parallel. Particularly, EE avoids the pitfalls of the simulate-and-transfer method and allows the speed-up of evaluation time by utilizing parallelism. The more novel features of EE are that the evolutionary algorithm is entirely decentralized, which makes it inherently scalable to large numbers of robots, and that it uses many robots in a shared task environment, which makes it an interesting platform for future work in collective robotics and Artificial Life. We have built a population of eight robots and successfully implemented the first example of embodied evolution by designing a fully-decentralized, asynchronous evolutionary algorithm. Controllers evolved by EE outperform a hand-designed controller in a simple application. We introduce our approach and its motivations, detail our implementation and initial results, and discuss the advantages and limitations of EE.

Key words: Evolutionary Robotics, Artificial Life, Evolutionary Algorithms, Distributed Learning, Collective Robotics

1 Introduction

1.1 Vision

Our work is inspired by the following vision. A large number of robots freely interact with each other in a shared environment, attempting to perform some task—say the collection of objects representing food or energy. The robots mate with each other, i.e., exchange genetic material, producing (offspring) control programs that become resident in other members of the robot population. Naturally, the likelihood of a robot producing offspring is regulated by its ability to perform the task or collect ‘energy.’ Further, there is no need for human intervention either to evaluate, breed, or reposition the robots for new trials.

This vision, to our knowledge first described in [17], aspires to an ideal where the robot population evolves in a completely hands-free and autonomous manner; in so doing, it offers intriguing possibilities for the future of evolutionary robotics. Nevertheless, many substantial technological demands are made by this conception of robot evolution, and considerable algorithmic detail must be added before it is implementable.

We have developed this vision into a working implementation and have termed our methodology *embodied evolution* (EE) [48,9]. We define EE as evolution taking place within a population of real robots where evaluation, selection, and reproduction are carried out by and between the robots in a distributed, asynchronous, and autonomous manner. Thus, as Figure 1 shows, we distinguish embodied evolution from ER methods that use either simulated trials, or embodied trials (the evaluation of candidate controllers on real robots) where a) only one or a small number of real robots are used [15,10,11,33], or b) the evolutionary algorithm maintains and manipulates the specifications of individual agents in a centralized manner. We wish to create a population of physical robots that evolve autonomously (as well as perform their tasks autonomously), using a distributed evolutionary algorithm.

1.2 Motivation

1.2.1 An Artificial Life Experiment

The vision of embodied evolution described above is largely inspired by experiments in Artificial Life. In natural evolution the adaptive mechanism is completely decentralized and distributed. Evaluation is implicit and reproduction is carried out autonomously by the agents in the population. This con-

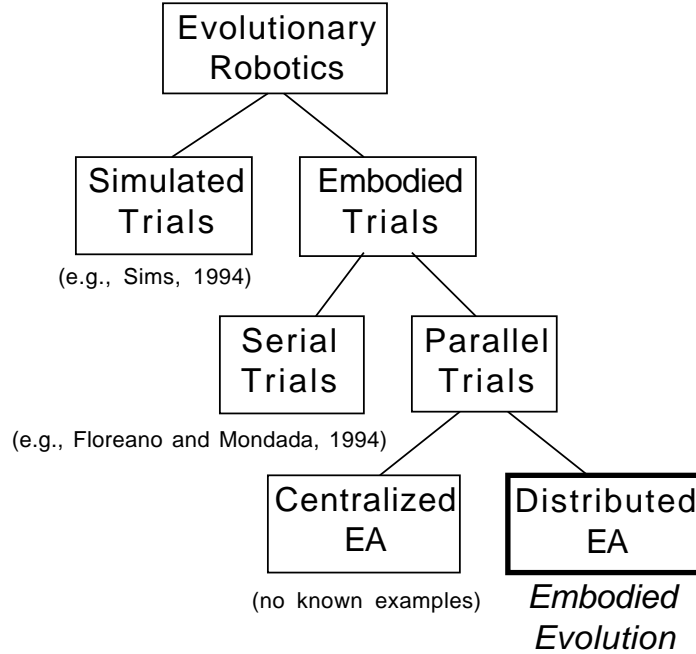


Fig. 1. Embodied evolution is an evolutionary robotics methodology that embodies a distributed evolutionary algorithm within a population of real robots.

trasts with the common practice of artificial evolution, where reproduction is performed by some centralized authority. The artificial life literature provides several examples of simulated systems where agent behavior and reproductive activity are integrated together [49,12,36,47]. In these systems, agent behavior either impacts reproduction directly, or, in some cases, is synonymous with reproduction. These experiments enable researchers to explore the critical effects that result from the merging of reproductive behavior with other behaviors.

In contrast, ER experiments that use physical robots have not been able to integrate reproduction with other autonomous behaviors. Although some evolutionary robotics has used real robots for evaluation of individuals, the evolving *population* is virtual—a set of controllers centrally stored either off-board or on-board—and so reproduction cannot occur between two robots. A significant motive for our EE research is to implement, in a population of real robots, artificial evolution using the distributed and autonomous properties of natural evolution. We wish to employ autonomy and distributed control not only in the task behavior of robots, but in their adaptive mechanism as well.

1.2.2 An Evolutionary Robotics Methodology

Mataric and Cliff [28] identify several key problems with existing ER methods, such as simulator fidelity, evaluation time, and hardware robustness. We

argue below that one way to ameliorate these three particular concerns is to use what we will term *embodied trials in parallel* (ETP)—that is, the parallel evaluation of a large number of evolving controllers on a large number of physical robots. The parallelization of the evaluation procedure converts costs in fidelity, time, and robustness to costs in hardware duplication and coordination. The magnitudes of these different costs determine whether ETP provides a net gain.

Embodied evolution uses the ETP approach: EE presumes a large number of robots, each simultaneously evaluating a controller. But, additionally, EE employs a distributed evolutionary algorithm. Though embodied trials in parallel are, to some extent, feasible using a centralized evolutionary algorithm, scaling issues regarding coordination persist in principle; moreover, centralized ETP is not fundamentally distinct from conventional ER approaches. Specifically, centralized ETP is merely the marriage of existing ER work that uses a small number of real robots with existing EA work that farms the evaluation of individuals to parallel processors. In contrast, parallelized embodied trials coupled with an evolutionary algorithm that is distributed throughout the population of robots—that has no centralized component—is both inherently scalable and methodologically novel. Embodied evolution is a case in point.

1.2.3 Opportunities for Interactive and Collective Tasks

Our most intriguing, yet currently unexplored, motivation for EE stems from the novelty of the approach. Since EE uses a distributed EA in a population of real robots, it may provide an effective way to investigate interactive and collective tasks. Pursuit and evasion, and collective foraging are but two examples of domains that would naturally be amenable to the use of EE. Also, EE provides interesting opportunities for experiments that investigate the interaction of reproductive behaviors and task behaviors. We develop these points in later sections.

1.3 Paper Outline

Our first motive, outlined above, is essentially an issue of curiosity, but the latter two motives warrant additional detail—this follows in sections 2 and 3. Section 4 summarizes the suitability and utility of the EE approach. Section 5 describes our implementation of EE. Section 6 describes our experiments and results that show the first proof-of-concept for EE. Finally, Section 7 outlines some caveats of our approach, followed by conclusions.

2 Parallelized Embodied Trials: A Response to Challenges in Evolutionary Robotics

As an alternative to the hand design of robotic controllers, evolutionary robotics (ER) [15,7,16] has generated a great deal of exuberance. While expected by many to emerge as an important technology [30], ER is currently a new research area in which a number of potentially serious problems are known to exist. Mataric and Cliff’s [28] comprehensive summary of the challenges that face evolutionary robotics attracted considerable attention in the ER community. The issues they raise ultimately question whether ER techniques can produce a net savings of human effort when applied to complex robotic domains.

Our continued investment in developing ER techniques, and the EE methodology in particular, is done with the belief that problem domains of interest will soon be, if they are not already, too difficult for the hand design of solutions; indeed, the seminal work in evolutionary robotics was done with this realization in mind [16]. That evolutionary techniques have the potential to find unexpected, yet effective solutions is the real promise of evolutionary techniques.

Nonetheless, for further progress to occur in ER, we must address the issues Mataric and Cliff discuss. We argue below that one way to alleviate the issues they raise is to use embodied trials in parallel.

2.1 Simulation

- “The difficulty of accurately simulating physical systems is well known in robotics.” [28, p. 76]
- “As the complexity of robotic systems grows and the gap between the simulation and the real system widens, the question of the value of investing in a specialized simulation will become increasingly important.” [28, p. 76]

Mataric and Cliff, in agreement with Brooks [6], raise the concern that a lack of simulation fidelity can lead to problems of *transference*, where robotic controllers that evolve in simulation are unable to perform effectively when transferred to real robots because they exploit features of the simulator that are not found in the physical world. They argue that this problem is magnified as the system to be modeled becomes more complex. Realistic physical models represent significant engineering effort [22,29,31]. Even virtual worlds that are not intended to faithfully model the real-world can become complex and difficult to create [42,21].

More recently, a method has been devised to provably eliminate transference

risks from the design of a simulator, but only provided that the environmental factors responsible for correct behavior are known *a priori* [18,19]. This allows the environmental factors that are not salient to the robot to be approximated, and those that cannot be modeled accurately to be masked with noise, which minimizes the computational requirements of the simulation and prevents the evolution of behaviors that depend upon inaccurate features of the simulation. However, as environments continue to become more complex (for example, if they involve a multiplicity of robots, or robots with high-resolution sensory apparatus such as vision), the critical environmental factors will become more difficult to ascertain, more difficult to model, and more computationally expensive to simulate, arguably to the extent that simulation is prohibited.

One approach to alleviate the concerns of simulator fidelity is to not use a simulator—to conduct evaluations on a real robot [10,11]. Obviously, embodied trials have perfect fidelity and the problem of transference is side-stepped completely. However, evaluations on real robots can be time consuming.

2.2 Time

- “Evolution on physical systems takes prohibitively long.” [28, p. 76]
- “[Consider a problem of]...five free parameters, each of which is tested with four different values, [and where]...each trial takes 15 s [seconds] ... With a population of size 100...the 100th generation will finish in roughly five years.” [28, p. 80]

The issue of time raised by Mataric and Cliff is in contrast to the (supposed) speed of simulation. Though simplifying assumptions and stochastic approximations minimize simulator complexity, they do not eliminate it; a sufficiently complex environment can still cause simulation to run slower than real time. Trials on real robots, although they may be slower than approximate models of simple domains, are never slower than real time and they have perfect fidelity.

Nevertheless, serial evaluation of candidate robot controllers on a single robot can take a very long time indeed. Hardware parallelization can alleviate this concern for both simulated and embodied approaches. Even in Mataric and Cliff’s example scenario above, which they characterize as “exaggerated,” a population of 100 physical robots reduces five years to 18 days (5 years = 1825 days; dividing by population size of 100 gives 18). However, where virtual robots perform interactive tasks in a simulated shared environment, parallelization is not trivial: for example, collision detection and mutual-sightedness can easily incur high computation costs. Parallelization in physical robots, however, accommodates shared environments and interactive tasks naturally.

Thus, concerns over simulation fidelity and evaluation time combine to direct us towards the use of embodied trials in parallel. As we point out above, the time savings gained by the ETP approach obviously come at the cost of significant hardware duplication. Yet, as large-scale multi-agent domains become increasingly important arenas of research, large numbers of robots will become commonplace; ETP is a straight-forward way to exploit the ubiquity of hardware as it becomes available.

2.3 *Power*

- “The unavoidable need to recharge robot batteries further slows down the experimental procedure.” [28, p. 76]

The issue of power pervades robotics. Battery power is able to sustain a robot only for a period on the order of hours, often no more than two or three [6]. Longer periods of uninterrupted power can be achieved by either tethering a robot directly to a power source [32], or by providing battery recharge stations for the robot to visit periodically. Nevertheless, tethers easily tangle with only a few robots, and recharge stations cannot be made transparent with respect to the robotic task, as they force robots to interrupt their activity for non-trivial amounts of time.

However, though the issue of power continues to raise problems, we do not view it as a fundamental restriction to ER. There are, at least in some cases, technological solutions to this issue. For example, in our experiments, we use an electrified floor to provide power to our robots. While building our floor, we learned of two other research groups that have built floors of similar construction [25,20]. These parallel achievements attest to the viability and utility of this power supply approach. Other approaches [1], like earlier prototypes of our own, use a floor-and-ceiling “bumper-car” style set-up. Together these examples demonstrate acceptable solutions for laboratory conditions, if not some industrial settings as well. Given these technologies we see that electrical power need not be a restriction to the use of parallelized embodied trials. While technologies such as the powered floor are important to the implementation of our experiments, we do not consider them to be an intrinsic part of the EE methodology.

2.4 *Robustness*

- “...a robotic system cannot survive the necessary continuous testing...” [28, p. 76]

Robots used in research are rarely endowed with the robustness that is engineered into industrial robots, usually for reasons of economy of development

time or expense. As a result, research robots demand almost constant care and attention to keep them in operational order. The robots we built for our experiments are no exception in this respect.

Nevertheless, the population of robots that ETP uses is a valuable source of redundancy, which allows the performance of the evolutionary system to degrade gracefully with the number of robot failures. While we concede that physical failures are inevitable for long running times, especially when many robots are involved in physical interaction, hardware parallelization reduces the amount of run-time per robot by a factor equal to the size of the robot population. Thus, the ETP approach appears to provide a feasible way forward for ER.

2.5 *Other Issues*

We mention here, for completeness, that Mataric and Cliff raise several other issues that pertain to any method of machine learning, or even to hand-design methods. While our EE methodology is silent on these points, they are important to recognize. One such point in particular concerns the creation of an effective metric of agent success; all automated learning methods, including embodied evolution, require feedback to function. The question of how researchers are to construct good metrics of behavior for autonomous robots, especially as environments become more complex and interactive, is undeniably important and will continue to require special attention.

2.6 *Scalability: The Transition from ETP to EE*

We have examined a number of core issues raised by Mataric and Cliff and have argued that embodied trials in parallel on a large number of robots may be one approach to alleviate their concerns. If we are to take the ETP approach seriously, however, we must address a new concern that does not arise in the methods considered by Mataric and Cliff, namely the *scalability* of the evolutionary algorithm. The feasibility of ETP using a centralized evolutionary algorithm is limited, for any centralized adaptive algorithm is susceptible to scaling problems. For example, centralized coordination of evaluation and reproduction can potentially result in communication bottlenecks [25] as the number of robots increases.

Thus we arrive at the suggestion of an ER methodology that uses a large population of robots, performing evaluations in parallel (i.e., ETP), but that also implements a distributed evolutionary algorithm. This is our definition of embodied evolution. The use of an evolutionary algorithm that has no central-

ized component—that is distributed throughout the population of robots—is inherently scalable, enabling the use of hundreds or thousands of robots. As we state in the introduction, we wish to create a population of physical robots that evolve autonomously as well as perform their tasks autonomously. In this way, we adhere to the goals of autonomy and distributed control now common in the field of robotics, not only in the task behavior of the robots, but in their adaptive mechanism as well.

3 Opportunities for Interactive and Collective Tasks

The previous section asserts the benefits of embodied evolution from the perspective of an evolutionary robotics methodology. In this section, we emphasize potential advantages of EE that are neither directly motivated by the issues that Mataric and Cliff raise, nor by the need for scalability. Specifically, since EE requires the use of a large number of robots in a shared task environment, it offers an appropriate substrate for experiments in which robot interaction is of interest (for example, pursuit and evasion tasks). Similarly, it offers a substrate for collective problem solving where robots must work together to solve some task (for example, collective foraging or box pushing).

Distributed robotics systems pose serious challenges to established controller-design methods. Distributed control is easy to achieve if the decomposition of a problem is known and the problem sub-parts are neatly separable into independent tasks; in such a case, we may build an independent autonomous agent for each sub-problem (using either hand design or machine learning). The structures of most real-world problems, however, are neither known *a priori*, nor composed of neatly separable sub-parts. As a result, much work to-date in collective robotics focuses on restricted cases, such as systems that are composed of homogeneous and independent sub-systems, for example flocking and foraging. Typically, agents in such experiments have used hand-built (non-learning) controller architectures [4,3,38,8]. Work that does involve learning typically occurs in simulation [44,23,39,2], or in relatively simple physical domains/environments [24,26,27,34,46].

The difficulty of accomplishing highly coordinated multi-robot behavior in complex interactive domains provides an additional area of motivation for the development of EE. To date, evolutionary robotics has not addressed collective tasks in real robots because of the many technical and engineering challenges involved, such as the need for continuous power and the difficulty of coordinating multiple robots. As robot populations become larger and deployed in more complex environments these challenges become more acute.

Although the experiments we describe below do not involve an interactive

task, nor collective behaviors, we suggest that EE may provide an appropriate substrate for future work in these areas. As an intrinsically population-based method where robots adapt in the task environment, embodied evolution appears potentially well-suited for the study of emergent group behavior and mechanisms that adaptively discover problem decomposition. This research direction returns to a theme from our ALife-inspired motivations, specifically, the integration of reproductive and task-oriented behaviors. Behaviors that modify reproduction can be used as a tool to enable speciation and heterogeneous teams/collective behaviors. This can be enabled by placing reproductive behaviors explicitly under the adaptive control of the robots, allowing certain environmental cues to trigger reproductive events. Alternatively, segregated reproduction may occur via some subtle interaction with task behaviors. This is not as esoteric as it may seem. The reproductive events in our implementation of EE (which we shall detail shortly), are subject to the co-locations of the robots, and thus an experimental set-up not much more complicated than ours may enable reproductive ‘preference’ to be implemented simply by a behavioral disposition to move to one area of the task environment more often than another area.

Although much can be learned about these topics from studies that use simulation, the caveats regarding simulation fidelity, originally raised with respect to the task domain, apply to the “reproductive domain,” as well. For task behaviors, subtleties afforded by the physical substrate may be precluded by the use of a simplified or approximate simulation. Similarly, there exists the potential for subtle forms of sexual segregation to emerge in a physical substrate that may be precluded by simulation. Further, if the robots’ choices for mating are expressed and determined by their behaviors, then a centralized reproductive algorithm that determines which robots reproduce with which other robots is excluded: reproductive and task-oriented behaviors are no longer categorically distinct, and a centralized EA would require an interpretive process to disentangle reproductive from non-reproductive behaviors, or else be reduced to a proxy for the robot’s reproductive choices.

4 Suitability of Embodied Evolution

This section summarizes the strengths and suitability of embodied evolution according to the issues we have described above. We suggest that EE is suitable in any of the following circumstances. Taken together, the points below provide a strong motivation for a distributed, embodied approach:

- *Where a simulator for the task domain is impossible, unavailable, or insufficiently accurate.*

The use of embodied trials, as in EE, allows the world to act as ‘its own

best model' and avoids the need for a simulator.

- *Where a centralized, globally-coordinated adaptive algorithm is not implementable or is unavailable, or where coordination of parallelized embodied trials is difficult.*

The distributed architecture of EE's evolutionary algorithm is intrinsically scalable with respect to the number of robots.

In addition, the EE approach offers possibilities for future work in the following circumstances, the last two of which are applications we have not previously discussed.

- *Where we are interested in evolving interactive or collective behaviors.*

The shared task environment utilized in EE provides a natural substrate for interactive and collective tasks. EE may be especially suitable where the interactions involve subtle complexities that are difficult to simulate.

- *Where we are interested in the interaction between task behaviors and reproductive behaviors.*

Studies in Artificial Life that involve the interaction of task and reproductive behaviors, for example through implicit fitness metrics, may be implemented using the EE approach. Again, EE may be especially suitable where the interactions involve subtle complexities that are difficult to simulate.

- *Where the agents must learn "in the field."*

One can easily imagine applications where learning must occur after the agents have been deployed in the actual task domain, for example if the agents are to operate in a remote region (e.g., Mars or perhaps a Micro Electro-Mechanical Systems (MEMS) substrate). In such a case, a centralized coordinator of agent learning would be both difficult to design and perhaps precarious to use, as it would give the learning system a single point of failure. The distributed nature of EE provides a robust method for adaptation in the field.

- *Where agents have been evolved in simulation and must now adjust their control programs to cross the "reality gap."*

Embodied evolution can be applied as the final stage of a hybrid ER process, as described in [35]. Agents are first evolved in simulation, then downloaded into a population of physical robots where EE is used to fine-tune their control programs in the real world.

5 Implementing EE

5.1 A Distributed Evolutionary Algorithm

The principal components of any evolutionary algorithm are *evaluation*, *selection*, and *reproduction*, and all of these are carried out autonomously by

and between the robots in a distributed fashion according to our definition of embodied evolution.

Because the process of evaluation is carried out autonomously by each robot, some metric must be programmed into the robots with which they can measure their performance. This can be quite implicit, for example, where failing to maintain adequate power results in “death” [32]. Or, it can be explicitly hard-coded, for example, where fitness is a function of objects collected and time. Whatever metric is used, performance against it must be monitored by the robot itself, as no external observer exists to measure a robot’s ability explicitly.

Reproduction in EE must also be both distributed and asynchronous. Assuming that we can not really create new robots spontaneously, the offspring must be implemented using (other) robots of the same population. And, if we do not have structurally reconfigurable bodies, reproduction must simply mean the exchange of control program code.

In general, selection in an evolutionary algorithm may be realized by having more-fit individuals supply genes (i.e., be parents) or by having less-fit individuals lose genes (i.e., be replaced by the offspring) or by a combination of both. The Microbial GA [14] uses this observation to simplify the steady-state genetic algorithm; rather than pick two (above-average fitness) parents and produce an offspring from the combination of their genes to replace a (below-average) third, the Microbial GA selects two individuals at random and overwrites some of the genes of the less fit (of the two) with those from the more fit. In effect, the less fit of the two becomes the offspring. To achieve decentralized and asynchronous reproduction in EE, we have developed a probabilistic version of the Microbial GA.

5.2 Probabilistic Gene Transfer Algorithm (PGTA)

The Probabilistic Gene Transfer Algorithm (PGTA) is our name for a probabilistic version of the Microbial GA. This algorithm requires minimal inter-agent communication, and eliminates the need to coordinate the communication of each reproduction event.

In the PGTA, reproduction is concurrent with task behavior—there is no “reproduction mode” as such. Each robot maintains a virtual energy level, which reflects the robot’s performance at the task, and each robot probabilistically broadcasts genetic information on its local-range communication channel at a rate proportional to this energy level. Each broadcast contains a mutated version of one randomly-selected gene from the robot’s genome (i.e., one parameter from the robot’s control specification). If another robot receives the

broadcast, that robot may allow the received gene value to overwrite its own corresponding gene. The receiving robot will accept the broadcast gene with a probability inversely related to its own energy level. Robots with higher energy thus attempt to reproduce, and resist the reproductive attempts of others, more frequently than do robots with lower energy. Nevertheless, because sending and receiving is probabilistic, and genes are picked at random, the PGTA does not guarantee that a fitter robot will transfer all its genes to a less fit robot. On average robots are left with a mixture of genes in proportion to their relative energy levels. This implements a fitness-proportionate recombinative evolutionary algorithm.

Using the PGTA, each reproduction event requires only one-way communication—there is no need for robots to coordinate reproductive acts, for a robot to know the fitness or identity of another robot, or even to know that any robot received its broadcast. The PGTA thus allows the complete decentralization of selection and reproduction. Though the PGTA provides an interesting mechanism for EAs in general, its robustness to genetic information “dropped” in communication makes the PGTA particularly advantageous for implementation in a population of real robots. In a later section, we describe how the maintenance of energy levels for the PGTA is integrated into our experimental task.

5.3 A Population of Robots

Embodied evolution requires a larger number of robots than that used in any evolutionary robotics work to-date. The short-term proof-of-concept experiments (described in the next section) require only minimal capabilities of each robot. Similarly, the long-term objectives of EE emphasize the interaction of robots rather than the sophistication of individual robots. Accordingly, we have built a population of simple robots of our own design that are quite minimal in their individual capacity yet have the necessary capabilities for EE. Our robots employ the “Cricket” micro-controller board, supplied by the MIT Media Laboratory [37], which uses a PIC micro-controller. Shown in Figure 2, each robot measures 12cm in diameter, has two light sensor inputs, two motor outputs, as well as local-range omnidirectional infra-red (IR) communication. The IR range is such that a robot may communicate with any other robot when the peripheries of their bodies are less than about 4cm apart.

5.4 Continuous Power Technology

We have developed and refined a powered-floor technology that transparently provides continuous, untethered power to our robots, without the use of




Fig. 2. (Left) The robot design used in our initial EE experiments. A: Infra-red transmit/receive. The directional infra-red diodes are directed vertically downwards and use reflectance off the floor to achieve local omnidirectional communication. B: PIC micro-controller. C: Lego™ motor. D: Tupperware™ body. E: Rechargeable cell. F: Recharge circuit. (Right) Robot underside showing the two light sensors and four contact points that collect power from the floor.

recharging stations. Our powered floor is surfaced with strips of stainless-steel tape, which are alternately connected to the positive and negative poles of a DC power source. Each robot has four contact points on the underside of its body, with which it draws power from the floor. The geometry of the contacts, shown in Figure 3, guarantees that at least one point can make contact with each pole of the DC power supply, regardless of the rotation or translation of the robot on the floor. The strip a contact touches may be either polarity, according to the position of the robot on the floor, and so the power is rectified before being delivered to the robot’s controller and motors. It transpires that this method of power delivery was described by Claude Shannon in an internal memo during his tenure at the AT&T Bell Laboratories in 1950 [43][p. 678]. We add a rechargeable cell to cover intermittency in contact.

6 Experiments and Results

6.1 A Phototaxis Task

Our first embodied evolution environment employs eight of our robots. The behavior of a robot is controlled by a simple artificial neural-network architecture, the weights of which are evolved to perform phototaxis similar to that described in [5]. The task environment consists of a 130cm by 200cm pen with a lamp located in the middle, visible from all positions on the floor plane, as

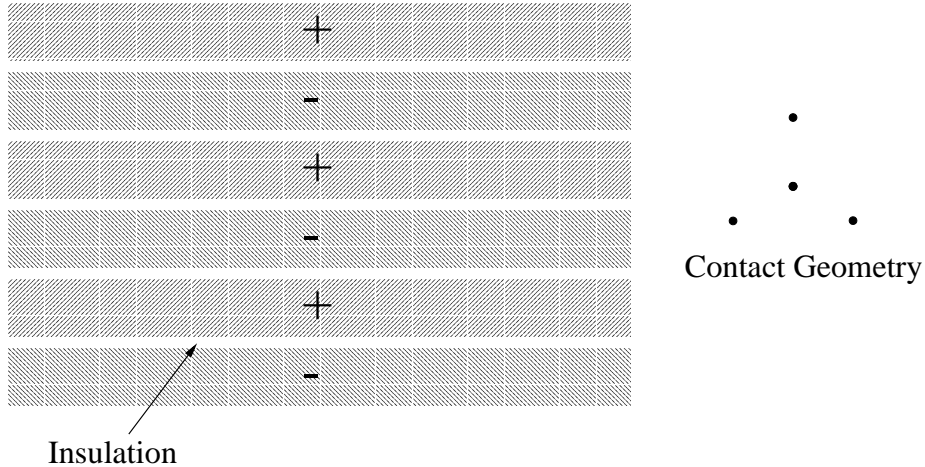


Fig. 3. Shannon’s design for power delivery [43]. The floor is constructed of conductive strips (separated by a thin insulator) that are alternately connected to the positive and negative poles of a DC power supply. Power is picked up from the floor by four contacts arranged in an equilateral triangle with one in the center. This geometry guarantees contact with both poles of the power supply regardless of its rotation and translation on the floor. Since an individual contact point may be on either pole of the supply, power drawn through the contacts must be rectified before being delivered to its ultimate destination.

seen in Figure 4. The robot task is to reach the light from any starting point in the pen. An infra-red beacon mounted above the light emits a signal that can be detected by a robot when it reaches the light source. This beacon signal triggers a built-in reset behavior that moves the robot to a random position and orientation along the periphery of the pen, from where the robot recommences its light-seeking behavior. If a robot’s sensor values do not change for some period of time, it is assumed that the robot is stuck against a wall (or another robot) and the robot invokes a second built-in behavior designed to free the robot by rotating it a random amount. Both of these built-in behaviors operate independently of the evolving neural-network controller. Because the pen contains a multitude of robots, the *de facto* environment also includes some amount of robot-to-robot interference [41]; therefore, the task implicitly requires that each robot also successfully overcome this interference.

6.2 Control Architecture

Our initial experiments use a simple artificial neural-network control architecture to serve as the evolving substrate, depicted in Figure 5. The weights of the network are evolved. The network consists of two output nodes, one for each of the two motors, one binary-valued input node, which indicates which of the robot’s two light sensors is receiving more light, and one bias node. Being a fully-connected feed-forward architecture, there are four weights. Each




Fig. 4. The robot pen for the phototaxis experiments. Eight robots, the power-floor, and the light in the center are shown. The unique ID of a robot is collected when it reaches the light (via infra red receivers on the overhead beam). This data is time-stamped and stored for monitoring experiment progress.

weight has an integer value in the range $[-8, 7]$. The values sent to the output nodes (controlling motor speed and direction) are the weighted sum of the input nodes; no sigmoid function is used. This network is simple enough to be computable by the PIC micro-controller in real time, yet provides a non-trivial search space of 16^4 network weight configurations.

As no individual learning takes place in our experiments, robots only get new weight values from other robots during reproduction, which is performed via local broadcasts (omni-directional signals that can be received by any robot within a short-range) on the robots' infra-red communications channel. This limited communication range (approximately one body width in radius) combined with the freely mixed population of our shared environment essentially implements random selection of mates, as appropriate for the PGTA.

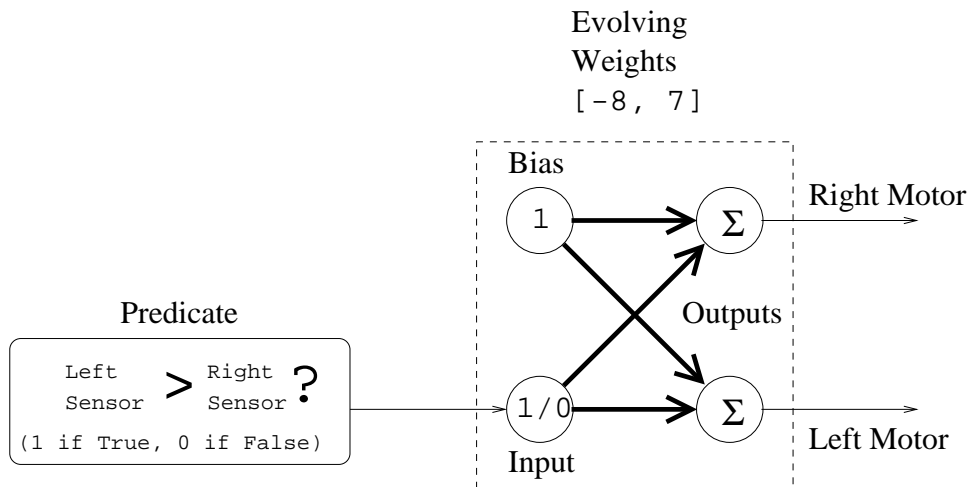


Fig. 5. Control architecture for phototaxis experiment. The one-bit input is 1 if left sensor is brighter than right sensor, 0 otherwise; the bias node has constant activation of 1.

6.3 *Maintaining Reproductive Energy Levels*

Energy levels regulate reproduction events and should reflect the robots performance at the task. The virtual energy level maintained by a robot is updated as follows: whenever a robot reaches the light its energy is increased by a fixed amount, up to a maximum energy value; whenever a robot sends a gene for reproduction (regardless of whether another robot receives or accepts the gene) its energy is decreased by a small fixed amount, down to a minimal energy value. Since the robot’s rate of sending genes is proportional to its energy level and decrements occur with each send, the rate of broadcasting decays exponentially over the time from its most recent visit to the light. The energy level thus approximates a leaky integral of the robot’s performance at its task (i.e., the frequency with which it reaches the light). Figure 6 provides an overview of how the reproductive energy levels are maintained in our experiments and how the PGTA is integrated with the robots’ other behaviors.

Note that the energy level maintained by a robot is an odd representation of its performance compared to the usual meaning of “fitness.” In our implementation, the energy level is not reset in a robot when it receives a new gene during a reproductive event—and so, the energy level is related to the performance of the various controllers that have been resident on that robot. In contrast, one would not normally expect the fitness of previously-resident controllers to affect the current fitness of a robot. However, assuming that the offspring is similar to the parent, our method of using inherited energy potentially shortens the evaluation process that must take place before the fitness measure of a new controller is reliable.

6.4 *Results*

Figure 7 shows the frequency with which the light is successfully reached by the robot population over time in each of three experiments. The main experiment evolves the neural-network weights to perform the light-seeking task. The initial condition for the networks is that all weights have a value of zero (this configuration produces no output to the motors and provides a neutral starting point). The other two experiments are controls where the robots do not evolve; in one case the robots’ weights are random values, in the other the robots use weights of a hand-designed solution.

As Figure 7 shows, the two controls delimit a broad range of possible performance levels and provide useful references against which to judge the success of the trials where evolution takes place. We see that embodied evolution al-

```

define embodied_evolve
  initialize_genes[]
  energy = min_energy
  repeat forever
    if (excited?)
      send(genes[random(num_genes)] + mutation)
    if (receptive? and received?)
      genes[indexof(received)] = valof(received)
    do_task_specific_behavior
    energy = limit(energy + reward - penalty)
  endrepeat
enddefine

```

Fig. 6. Pseudo-code of control program that implements the Probabilistic Gene Transfer Algorithm (PGTA). This code is run on every robot. No methods for synchronizing or coordinating the robots, nor any centralized elements, are used in the PGTA. The predicates `excited?` and `receptive?` are probabilistic functions of `energy`. `send` takes a gene value and broadcasts it on local infrared (wrapped with gene locus). `received?` is *true* if any gene received on infrared. `indexof` and `valueof` return the locus and value of received gene, respectively. `limit` bounds the energy value between `min_energy` and `max_energy`. `random` returns an integer in the range of its argument. `task_specific_behaviour` includes monitoring performance at the task and setting the values of `reward` and `penalty`. In our phototaxis experiments, `min_energy` is 10; `max_energy` is 255. `excited?` returns *true* if `energy > random(max_energy)`, *false* otherwise; `receptive?` returns *true* if `energy < random(max_energy)`, *false* otherwise. Each gene, `genes[1..4]`, is a weight value for the network. `initialize_genes` sets all genes to 0. `mutation` returns $\{0, 1, -1\}$ with uniform probability. `task_specific_behavior` includes reading sensor values, updating network outputs, setting motor speeds/directions accordingly, monitoring sensor readings and performing random turn if robot appears to be stuck, and monitoring for arrival at the beacon. `reward` is set to 127, if the robot detects the beacon, 0 otherwise, and `penalty` is set to 1 whenever the robot broadcasts a gene, 0 otherwise.

lows the population of robots to achieve performance favorably comparable to that of our hand-designed solution—the Wilcoxon rank-sum test ¹ indicates that EE outperforms the designed solution with probability $p = 0.935$. These results provide the first evidence that a fully decentralized, asynchronous evolutionary algorithm, can operate effectively in a population of physical robots and provide high-quality control programs.

There are several additional points of interest in this result. First, though the robots learn to approach the light in a multi-robot environment, they are able to perform effectively in isolation, as well (not shown here). Second, despite its minimal structure, the artificial neural-network control architecture used

¹ The Wilcoxon test is preferred over the Student t-test in this case because the variances of the designed and evolved runs are very different from each other.

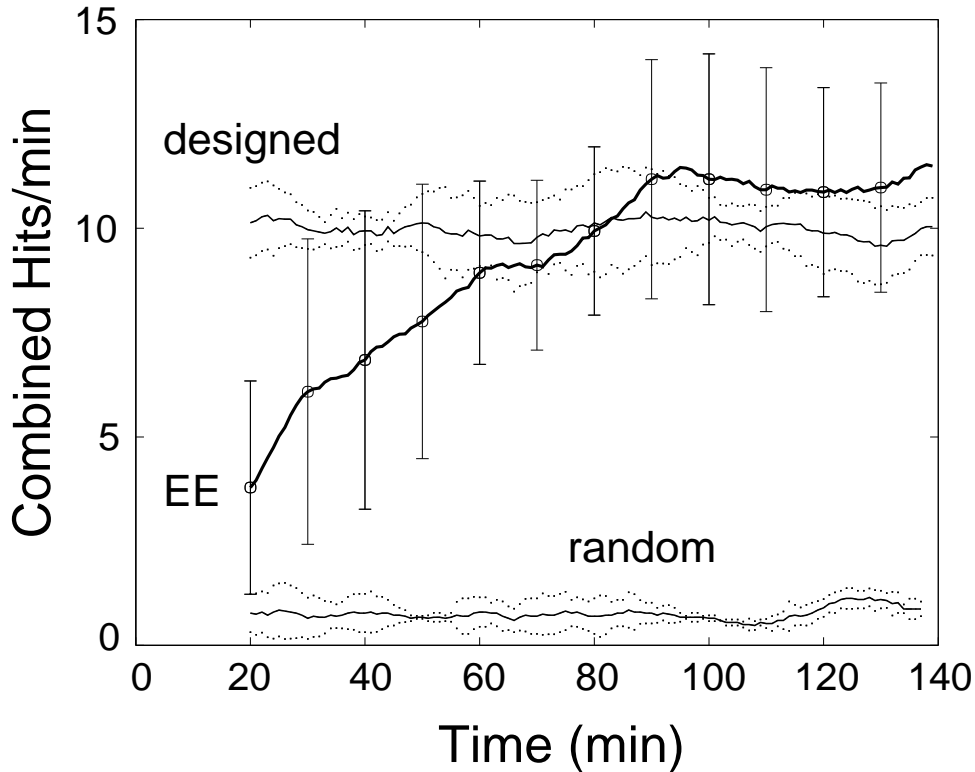


Fig. 7. Performance Over Time. Upon detecting the beacon’s signal, a robot sends a reply signal to the beacon. These replies are then delivered to a desktop computer where they are time-stamped. Three solid curves show the performance of the robot population using hand-designed (non-evolved), evolved, and random (non-evolved) networks. The data from the hand-designed and evolved experiments are averaged over six runs, while the data from the random-solution experiment are averaged over two runs. Each run lasts 140 minutes and uses eight robots. The vertical axis represents the average rate (in hits per minute) at which the group of robots reaches the light. A time window of 20 minutes is used to compute the instantaneous hit rate for each data point on the graph (hence the first data points appear at Time = 20 minutes). Vertical bars on the evolved run, shown every 10 minutes, and the dotted lines on the control experiments, show \pm one standard deviation. Though the evolved solutions begin with network weights of zero, we see that the robots achieve an average performance of four hits per minute within the first twenty minutes of the experiment and eventually exceed the hand-designed hit rate (the Wilcoxon rank-sum test gives $p = 0.935$).

in the robots allows a surprising variety of solutions to be discovered by the evolutionary process. Interestingly, the best evolved solutions exhibit behaviors that are qualitatively different from our hand-designed solution; evolution appears to favor a “looping” solution, whereas, with our hand-designed solution, the robot “swaggers” to the light, as shown in Figure 8. The reasons for this are not known and we intend to address this in future work.

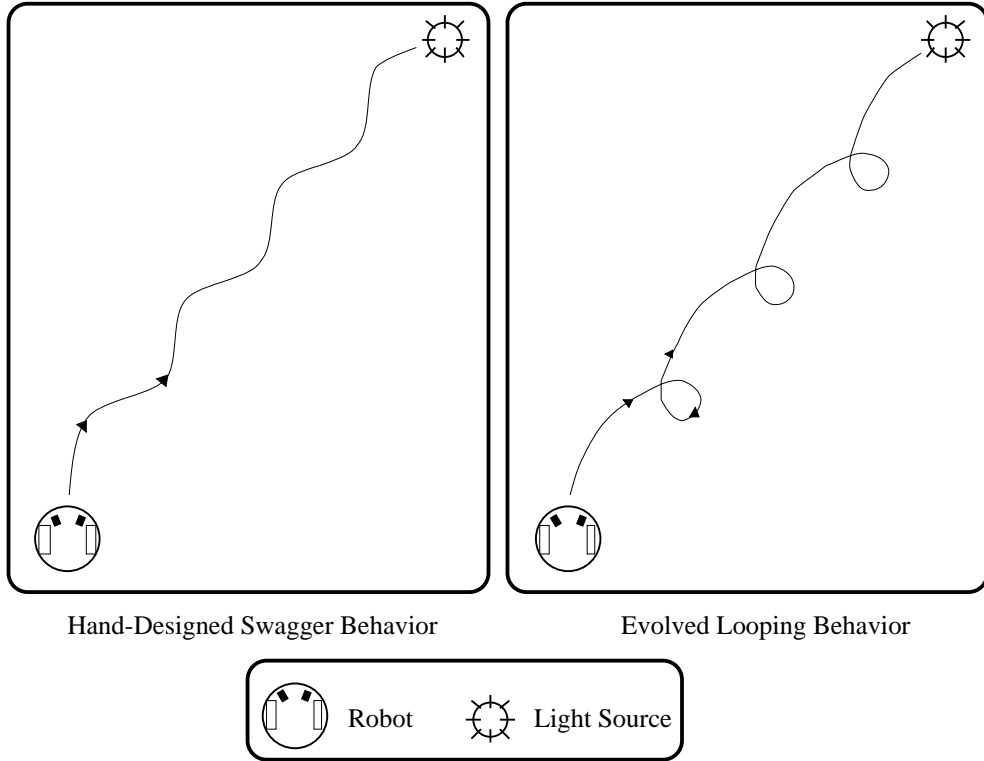


Fig. 8. Trajectories of light-seeking solutions.

7 Caveats

In spite of our enthusiasm for the embodied evolution approach, we must recognize that it is still a developing methodology, and although EE offers solutions for some issues, it also introduces new difficulties.

- An environment that contains a multitude of robots also includes some amount of robot-to-robot interference [41]; e.g., our phototaxis environment implicitly requires that each robot also successfully overcome such interference. Hence we suggest that EE is more suitable where interaction is native to the task domain.
- Our experiments to date have concerned only simple tasks for which many other learning approaches have also been effective. Though we suggest EE is suited to multi-agent systems, we are only just now designing experiments that involve explicitly interactive tasks.
- Our powered-floor is constrained to research or industrial environments; other applications will require different power technologies.
- Because we eschew centralization, the jobs of monitoring our experiments and collecting data are made more awkward. With our current robotic hardware we are unable to monitor reproductive activity, for example.
- Because reproduction is based upon the principle of locality in our particular experimental setup, our implementation of EE is susceptible to failure if the

agents become physically, and therefore reproductively, isolated. Moreover, the need to prevent reproductive isolation produces a selective pressure that may interfere with the objective of the task. And if, as we suggest, the reproductive mechanism is modified to allow speciation, we can imagine that reproductive behaviors could become quite elaborate, worsening this interference.

- Though the shared adaptive environment used in EE has interesting possibilities for team tasks, the precise manner in which EE should be applied to team evolution is unclear. Any requirement of locality (for reproduction) seems to suggest that an awkward overlapping of multiple teams is needed for reproduction to take place, and the mechanism for organizing games between teams is problematic.

Thus we see that a great deal of research effort is still required to meet the long-term goals of EE, and although EE provides advantages in some domains it is not suitable for all applications.

8 Future Work and Conclusions

8.1 Future Work

There exist a number of control experiments that will help us map the parameter space of the PGTA. Through these controls, we expect to refine the PGTA and understand more precisely the dynamics of the algorithm and the settings that provide the most robust operation. For example, a simulation built by Miguel Schneider-Fontán [40] suggests that good solutions are not stable in the population if we remove the robots' ability to resist the reproductive attempts of others. The resistance model we use, while effective, is not known to be optimal. Other parameters to investigate include the rates at which a robot's energy is increased and decreased as it reaches the light and attempts reproduction, respectively. We are in the process of developing more complex task environments and control architectures for our experiments, beginning with a recurrent version of the network architecture that will operate on raw sensor inputs. Though the phototaxis task described in this report is simple and does not involve explicit robot interaction, the transparency of this domain allows us to investigate the strong *implicit* interactive forces within the EE approach. For example, the reproductive process and physical robot-to-robot interference are two types of interaction that we are currently investigating before moving to a more complex task.

One question we hope to answer is why the looping, which intuitively we would expect to be less efficient, is in practice more effective than the swagger. One

hypothesis is that the looping behavior overcomes the physical interference caused by the other robots in the pen more efficiently than does our hand-designed solution. Another hypothesis is that looping is more robust to the inevitable hardware variances that exist between the robots. Or, perhaps, we will find the cause is more mundane, such as our choice of algorithm parameters.

As stated, a long-term goal of distributed robotics is a method for the automatic discovery of problem decomposition and balancing local autonomy with group coordination. By employing a large number of robots together in the task environment and allowing them to evolve interactive behaviors, we avoid introducing preconceptions about how a problem should be decomposed, how many robots should be assigned to each task/sub-task, or how many groups/sub-groups will be needed. Potentially, we allow the robots to discover appropriate working groups and interactive behaviors that reflect the nature and structure of the task at hand. Achieving this will require that we address many critical issues: credit assignment, the balance of cooperation and competition, homogeneity and heterogeneity, encapsulation and modularity.

Finally, our future work will address advanced robotic platforms, including evolvable hardware [45] and evolvable morphology [22,13].

8.2 *Conclusions*

We have introduced embodied evolution as a new methodology for evolutionary robotics. Embodied evolution is a new methodology for evolutionary robotics. EE uses a population of robots that evolve together while situated in the task environment. The adaptive mechanism is distributed in the population using robot-to-robot reproduction that is carried out autonomously by the robots. Evolutionary adaptation is seamlessly integrated with the robot's task behavior. Our experiments in EE have employed a population of eight robots that are supplied continuous power via an electrified floor. We have developed an evolutionary algorithm that operates via the probabilistic transfer of genetic information between robots on local-range communication. This PGTA is entirely distributed and is robust in ways that make it effective for implementation in a population of robots.

Our experiments provide the first proof-of-concept for embodied evolution. We have successfully applied EE to a simple phototaxis task. The neural-network control architecture, though minimal, has a non-trivial search space and provides surprisingly novel solutions for phototaxis. Results show solutions evolved with EE to perform better than our best hand-designed solutions.

In good part, our motivation for the development of EE was simply to see

if the artificial evolution of a population of robots could be implemented in the distributed and autonomous manner of natural evolution. But, EE also addresses several significant issues identified by ER researchers as problematic. We must be careful to separate technological issues from fundamental ones when considering the long term prospects for evolutionary robotics. Our implementation of embodied evolution reveals some concerns, such as power, to be merely technological. Other issues, such as combinatorics and the limits of simulation, are fundamental in nature and can not be solved technologically. Nevertheless, some of these fundamental issues are ameliorated by the embodied parallelism that EE employs. Specifically, ETP eliminates problems of simulation transference, alleviates the slow running time of numerous embodied trials, and provides robustness. But, the most important contributions of the embodied evolution approach are twofold. First, it provides the inherent scalability of a distributed EA. Second, it provides a substrate for future work exploring the evolution of interactive and collective behaviors that are particularly difficult to simulate.

Acknowledgments

The industrious contributions of several people were essential to this paper. Miguel Schneider-Fontán built an EE simulator, which allowed us to conduct the embodied experiments on real robots with confidence. Giovanni Motta designed and built the power backup and battery recharge circuit for our robots. Prem Melville administered many of the experiments and maintained the robots. Paaras Kumar designed and built an early prototype of the power backup circuit as well as the IR communication circuit for future work. Greg Hornby, Hod Lipson, and other members of DEMO challenged us with many insightful questions. We are also grateful for careful review and suggestions from Dave Cliff, and to Fred Martin of the MIT Media Lab, who supplied the Cricket micro-controllers.

References

- [1] Applied AI Systems, Inc., Kanata, Ontario, Canada. *Continuous Power Supply for Khepera*, 1998. Product literature.
- [2] T. Balch. Learning roles: Behavioral diversity in robot teams. In *1997 AAAI Workshop on Multiagent Learning*. AAAI, 1997.
- [3] T. Balch and R. Arkin. Motor schema-based formation control for multiagent robot teams. In *Proceedings of the First International Conference on Multi-Agent Systems ICMAS-95*, pages 10–16. AAAI Press, 1995.

- [4] R. Beckers, O. Holland, and J. Deneubourg. From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks and P. Maes, editors, *Artificial Life IV*, pages 181–189. MIT Press, 1994.
- [5] V. Braitenberg. *Vehicles: experiments in synthetic psychology*. MIT Press, 1984.
- [6] R. Brooks. Artificial life and real robots. In F. Varela and P. Bourguine, editors, *Proceedings of the First European Conference on Artificial Life*, pages 3–10. MIT Press, 1992.
- [7] D. Cliff, I. Harvey, and H. P. Explorations in evolutionary robotics. *Adaptive Behavior*, 2(1):73–110, 1993.
- [8] B. Donald, J. Jennings, and D. Rus. Minimalism + distribution = supermodularity. *Journal on Experimental and Theoretical Artificial Intelligence*, 9(2–3):293–321, 1997.
- [9] S. G. Ficici, R. A. Watson, and J. B. Pollack. Embodied evolution: A response to challenges in evolutionary robotics. In J. L. Wyatt and J. Demiris, editors, *Proceedings of the Eighth European Workshop on Learning Robots*, pages 14–22, 1999.
- [10] D. Floreano and F. Mondada. Automatic creation of an autonomous agent: Genetic evolution of a neural-network driven robot. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *From Animals to Animats 3*, pages 421–430. MIT Press, 1994.
- [11] D. Floreano and F. Mondada. Evolution of homing navigation in a real mobile robot. *IEEE Transactions on Systems, Man, and Cybernetics—Part B: Cybernetics*, 26(3):396–407, 1996.
- [12] W. Fontana. Algorithmic chemistry. In C. Langton, C. Taylor, J. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 159–209. Addison-Wesley, 1991.
- [13] P. Funes and J. Pollack. Computer evolution of buildable objects. In P. Husbands and I. Harvey, editors, *Proceedings of the Fourth European Conference on Artificial Life*, pages 358–367. MIT Press, 1997.
- [14] I. Harvey. The microbial genetic algorithm. Unpublished manuscript, 1996.
- [15] I. Harvey, P. Husbands, and D. Cliff. Issues in evolutionary robotics. In J.-A. Meyer, H. Roitblat, and S. Wilson, editors, *From Animals to Animats 2*, pages 364–373. MIT Press, 1993.
- [16] P. Husbands and I. Harvey. Evolution versus design: Controlling autonomous robots. In *Proceedings of the Third Annual Conference on Artificial Intelligence, Simulation and Planning*, pages 139–146. IEEE Press, 1992.
- [17] P. Husbands, I. Harvey, and D. Cliff. Central issues in evolutionary robotics. Unpublished manuscript presented at ALife III, Sante Fe, 1992.
- [18] N. Jakobi. Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behavior*, 6(2):325–368, 1997.

- [19] N. Jakobi. Half-baked, ad hoc, and noisy: minimal simulations for evolutionary robotics. In P. Husbands and I. Harvey, editors, *Fourth European Conference on Artificial Life*, pages 348–357. MIT Press, 1997.
- [20] D. Keating. Personal communication, 1998.
- [21] M. Komosiński and S. Ulatowski. Framsticks: Towards a simulation of a nature-like world, creatures and evolution. In D. Floreano, J.-D. Nicoud, and F. Mondada, editors, *Proceedings of the Fifth European Conference on Artificial Life*, pages 261–265, 1999.
- [22] H. Lipson and J. B. Pollack. Evolving physical creatures. In M. Bedau, J. McCaskill, N. Packard, and S. Rasmussen, editors, *Artificial Life VII*. MIT Press, 2000.
- [23] M. Littman. Markov games as a framework for multi-agent reinforcement learning. In *Proceedings of the International Machine Learning Conference*, pages 157–163, 1994.
- [24] S. Mahadevan and H. Connell. Automatic programming of behavior-based robots using reinforcement learning. In *Proceedings of the Ninth National Conference on Artificial Intelligence (AAAI '91)*, pages 8–14, 1991.
- [25] A. Martinoli, E. Franzi, and O. Matthey. Towards a reliable set-up for bio-inspired collective experiments with real robots. In A. Casals and A. de Almeida, editors, *Proceedings of the Fifth Symposium on Experimental Robotics ISER-97*, pages 597–608. Springer Verlag, 1997.
- [26] M. Mataric. Learning to behave socially. In D. Cliff, P. Husbands, J.-A. Meyer, and S. Wilson, editors, *From Animals to Animats 3*, pages 453–462. MIT Press, 1994.
- [27] M. Mataric. Reward functions for accelerated learning. In W. Cohen and H. Hirsh, editors, *Proceedings of the Eleventh International Conference on Machine Learning*, pages 181–189. Morgan Kaufman, 1994.
- [28] M. Mataric and D. Cliff. Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems, Special Issue on Evolutional Robotics*, 19(1):67–83, 1996.
- [29] Mathengine. <http://www.mathengine.com>.
- [30] J.-A. Meyer, P. Husbands, and I. Harvey. Evolutionary robotics: A survey of applications and problems. In P. Husbands and J.-A. Meyer, editors, *Evolutionary Robotics: First European Workshop, EvoRobot98*, pages 1–21. Springer-Verlag, 1998.
- [31] B. Mirtich. *Impulse-Based Dynamic Simulation of Rigid Body Systems*. PhD thesis, University of California, Berkeley, 1996.
- [32] F. Mondada and D. Floreano. Evolution and mobile autonomous robotics. In E. Sanchez and M. Tommasini, editors, *Towards Evolvable Hardware*, pages 221–249. Springer Verlag, Berlin, 1996.

- [33] S. Nolfi. Evolving non-trivial behaviors on real robots: A garbage collecting robot. *Robotics and Autonomous Systems*, 22(3-4):187-198, 1997.
- [34] L. Parker. Task-oriented multi-robot learning in behavior-based systems. *Advanced Robotics, Special Issue on Selected Papers from IROS '96*, 11(4):305-322, 1997.
- [35] J. B. Pollack, H. Lipson, S. G. Ficici, P. Funes, G. Hornby, and R. A. Watson. Evolutionary techniques in physical robotics. In J. Miller, editor, *Proceedings of the Third International Conference on Evolvable Systems(ICES 2000)*, pages 175-186. Springer, 2000.
- [36] T. Ray. An approach to the synthesis of life. In C. Langton, C. Taylor, J. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 371-408. Addison-Wesley, 1991.
- [37] M. Resnick, R. Berg, M. Eisenberg, and S. Turkle. Beyond black boxes: Bringing transparency and aesthetics back to scientific instruments, 1997. MIT project funded by the National Science Foundation (1997-1999).
- [38] D. Rus, B. Donald, and J. Jennings. Moving furniture with teams of autonomous robots. In *Proceedings of IEEE/RSJ IROS'95*, pages 235-242, 1995.
- [39] G. Saunders and J. Pollack. The evolution of communication schemes of continuous channels. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, and S. Wilson, editors, *From Animals to Animats IV*, pages 580-589. MIT Press, 1996.
- [40] Schneider-Fontán. Personal communication, 1998. Brandeis University.
- [41] M. Schneider-Fontán and M. Mataric. A study of territoriality: The role of critical mass in adaptive task division. In P. Maes, M. Mataric, J.-A. Meyer, J. Pollack, and S. Wilson, editors, *From Animals to Animats IV*, pages 553-561. MIT Press, 1996.
- [42] K. Sims. Evolving 3d morphology and behavior by competition. In *Artificial Life IV*, pages 28-39. MIT Press, 1994.
- [43] N. J. A. Sloane and A. D. Wyner, editors. *Claude Elwood Shannon: Collected Papers*. IEEE Press, 1993.
- [44] M. Tan. Multi-agent reinforcement learning: independent vs. cooperative agents. In *Proceedings of the Tenth International Machine Learning Conference*, pages 330-337, 1993.
- [45] A. Thompson. *Automatic Design of Electronic Circuits in Reconfigurable Hardware by Artificial Evolution*. Springer, 1998.
- [46] E. Uchibe, M. Asada, and K. Hosoda. Cooperative behavior acquisition in multi mobile robots environment by reinforcement learning based on state vector estimation. In *Proceedings of International Conference on Robotics and Automation*, pages 1558-1563, 1998.

- [47] J. Ventrella. Attractiveness vs efficiency (how mate preference affects locomotion in the evolution of artificial swimming organisms). In C. Adami, R. Belew, H. Kitano, and C. Taylor, editors, *Artificial Life VI*, pages 178–186. MIT Press, 1998.
- [48] R. A. Watson, S. G. Ficici, and J. B. Pollack. Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In P. Angeline, Z. Michalewicz, M. Schoenauer, X. Yao, and A. Zalzal, editors, *1999 Congress on Evolutionary Computation*, pages 335–342. IEEE Press, 1999.
- [49] G. M. Werner and M. G. Dyer. Evolution of communication in artificial organisms. In C. Langton, C. Taylor, J. Farmer, and S. Rasmussen, editors, *Artificial Life II*, pages 659–687. Addison-Wesley, 1991.