

RESEARCH ARTICLE

Evolution of Collective Behaviors for a Real Swarm of Aquatic Surface Robots

Miguel Duarte^{1,2,3*}, Vasco Costa^{1,2,3}, Jorge Gomes^{1,2,4}, Tiago Rodrigues^{1,2,3}, Fernando Silva^{1,2,4}, Sancho Moura Oliveira^{1,2,3}, Anders Lyhne Christensen^{1,2,3}

1 BioMachines Lab, Lisbon, Portugal, **2** Instituto de Telecomunicações, Lisbon, Portugal, **3** Instituto Universitário de Lisboa (ISCTE-IUL), Lisbon, Portugal, **4** BiolSI, Faculdade de Ciências, Lisbon, Portugal

* miguel_duarte@iscte.pt



OPEN ACCESS

Citation: Duarte M, Costa V, Gomes J, Rodrigues T, Silva F, Oliveira SM, et al. (2016) Evolution of Collective Behaviors for a Real Swarm of Aquatic Surface Robots. PLoS ONE 11(3): e0151834. doi:10.1371/journal.pone.0151834

Editor: Long Wang, Peking University, CHINA

Received: October 29, 2015

Accepted: March 5, 2016

Published: March 21, 2016

Copyright: © 2016 Duarte et al. This is an open access article distributed under the terms of the [Creative Commons Attribution License](https://creativecommons.org/licenses/by/4.0/), which permits unrestricted use, distribution, and reproduction in any medium, provided the original author and source are credited.

Data Availability Statement: The raw onboard logs from our experiments are available online at the following URL: <http://dx.doi.org/10.5281/zenodo.45570>.

Funding: MD acknowledges support from Fundação para Ciência e Tecnologia (FCT), of which he is a doctoral researcher (grant number SFRH/BD/76438/2011). JG acknowledges support from Fundação para Ciência e Tecnologia (FCT), of which he is a doctoral researcher (grant number SFRH/BD/89095/2012). FS acknowledges support from Fundação para Ciência e Tecnologia (FCT), of which he is a doctoral researcher (grant number SFRH/BD/89573/2012). The work presented in this paper was

Abstract

Swarm robotics is a promising approach for the coordination of large numbers of robots. While previous studies have shown that evolutionary robotics techniques can be applied to obtain robust and efficient self-organized behaviors for robot swarms, most studies have been conducted in simulation, and the few that have been conducted on real robots have been confined to laboratory environments. In this paper, we demonstrate for the first time a swarm robotics system with evolved control successfully operating in a real and uncontrolled environment. We evolve neural network-based controllers in simulation for canonical swarm robotics tasks, namely homing, dispersion, clustering, and monitoring. We then assess the performance of the controllers on a real swarm of up to ten aquatic surface robots. Our results show that the evolved controllers transfer successfully to real robots and achieve a performance similar to the performance obtained in simulation. We validate that the evolved controllers display key properties of swarm intelligence-based control, namely scalability, flexibility, and robustness on the real swarm. We conclude with a proof-of-concept experiment in which the swarm performs a complete environmental monitoring task by combining multiple evolved controllers.

1 Introduction

Swarm robotics systems (SRS) are a promising approach to collective robotics, in which large groups of relatively simple and autonomous robots are able to display collectively intelligent behavior [1]. Control in a SRS is decentralized, meaning that each individual robot operates based on its local observations of the environment and coordination with neighboring robots. During task execution, the swarm-level behavior emerges from the interactions between neighboring robots, and from the interactions between robots and the environment. The swarm robotics approach has the potential to incite a number of important properties in multirobot systems [1–3], namely: (i) robustness, the ability to cope with the faults of individual robots; (ii) flexibility, the ability to operate in a variety of different environments and to perform various tasks; and (iii) scalability, the ability to generally maintain the group behavior regardless of the size of the swarm. Due to these properties, SRS have an enormous potential in several real-

supported by Fundação para Ciência e Tecnologia (FCT) with the projects CORATAM (grant number EXPL/EEI-AUT/0329/2013) and HANCAD (grant number UID/EEA/50008/2013).

Competing Interests: The authors have declared that no competing interests exist.

world domains, such as search and rescue, exploration, surveillance, and clean up [1, 2]. Despite the potential of swarm robotics for real-world tasks, no demonstrations of swarm behaviors outside of controlled laboratory conditions have been performed so far. As discussed by Brambilla *et al.* [1] and Bayındır [3] in recent reviews of the field of swarm robotics, all real-robot experiments presented in the literature covered have, in fact, been performed in controlled laboratory environments such as enclosed arenas in which the relevant conditions are defined by the experimenter. That is, although SRS are ultimately intended to operate in potentially unstructured, real-world environments, experiments have been conducted mainly in simulation, and at the current state of development no study has been able to demonstrate the benefits of SRS in real-world environments.

One of the key challenges in SRS is the synthesis of behavioral control. Manual design of control for each robot requires the decomposition of the macroscopic, swarm-level behavior into microscopic behavioral rules that determine the interactions between neighboring robots, and between the robots and the environment so that the global, collectively self-organized behavior emerges. The designer thus needs to be able to understand the relation between local robot interactions and emerging swarm-level properties. There is, however, no approach to derive the microscopic behavioral rules based on a desired global behavior or task description for the general case [4]. In this respect, the use of evolutionary computation techniques has been studied as a promising alternative to traditional control approaches, due to its capability to automatically synthesize self-organized control based only on the definition of a high-level performance metric (see [5–7] for examples).

SRS are typically associated with ground robots [1, 8] and, more recently, with aerial robots [9, 10]. In this paper, we show that SRS also have the potential to be applied to maritime tasks, which are usually expensive to carry out because they rely on manned vehicles with large operational crews. While significant effort has been made to adapt unmanned vehicle technology to maritime tasks, such systems are currently relatively expensive to acquire and operate, and only a single or a few robots are typically deployed [11]. The use of SRS at sea is advantageous because: (i) several maritime tasks, such as environmental monitoring, sea-life localization, and sea-border patrolling, require distributed sensing at a high spatial and temporal resolution, and are therefore challenging to carry out with only a single or a few robots; and (ii) robots operating at sea need to display a high degree of fault-tolerance and robustness [11], which can be provided by SRS approaches [12].

In this paper, we show for the first time a SRS system with evolved, distributed control performing self-organized behaviors in real-world environments. Our study relies on a swarm of up to ten small, simple, and inexpensive aquatic surface robots (see Fig 1). First, we demonstrate our SRS performing four canonical tasks in an uncontrolled aquatic environment: (i) homing, (ii) dispersion, (iii) clustering, and (iv) monitoring. Second, we perform a set of experiments to demonstrate the scalability and robustness of the evolved controllers when operating in a real environment. Finally, we show how multiple evolved swarm behaviors can be combined to produce more sophisticated controllers, which we use to perform an environmental monitoring task where the robots have to sample the water temperature in a given area. In summary, our experiments demonstrate that a swarm of robots with evolved control can operate successfully in a real-world environment.

2 Related work

2.1 Swarm Robotics

The field of swarm robotics, as well as the more general field of swarm intelligence, takes inspiration from the observation of social insects, such as ants, bees, wasps, and termites [12]. In



Fig 1. The robotic swarm with 10 units performing a homing task.

doi:10.1371/journal.pone.0151834.g001

these animal societies, relatively simple units rely on self-organization to display collectively intelligent behavior. The self-organized behavior of social insects confers the swarms a high level of scalability, versatility, and robustness to individual failures [13]. The key motivation behind swarm robotics is to harness these properties to build large-scale decentralized multi-robot systems [12, 14]. Due to their properties, SRS have significant potential to take on a number of applications requiring distributed sensing and/or action. In a recent review, Bayindir [3] discusses the categories of behaviors that have been achieved with SRS, such as aggregation, flocking, foraging, object clustering and sorting, navigation, path formation, swarm self-deployment, collaborative manipulation, and task allocation.

Despite the advances in the field, SRS have, however, up to this point not leveraged their full potential, and have been confined to controlled laboratory environments [1], still far from real-world applications. There are multiple reasons for the absence of real-world SRS, one of the most prevalent being the difficulty in designing behavioral control for the individual robots that results in the desired swarm-level behavior. This problem is exacerbated when trying to achieve behaviors capable of dealing with the complexity of real tasks and uncontrolled environments, as opposed to the abstract tasks and highly controlled environments that have been used in the majority of previous works [1, 3].

2.2 Evolutionary Synthesis of Control for SRS

Evolutionary robotics (ER) [5] is the research field that studies the application of evolutionary computation to the synthesis of robot controllers and/or morphology. ER has become a popular alternative to manual programming in a variety of SRS applications [1]. Given the specification of a task, an evolutionary algorithm evaluates and optimizes controllers in a holistic manner, thereby facilitating the emergence of self-organized behavior [15, 16]. That is, based on the performance of the resulting swarm-level behavior, the evolutionary algorithm iteratively fine-tunes the parameters and the microscopic rules governing each individual robot, thus eliminating the need for manual and detailed specification of low-level control [17]. Optimization of genomes is based on Darwinian evolution, namely blind variations and survival of the fittest.

Evolutionary approaches have been used to solve a large number of swarm robotics tasks such as coordinated motion [18, 19], chain formation [20], aggregation [21–23], flocking [24], hole avoidance [25], aerial vehicles communication [26], categorization [27], group transport

[28], social learning [29], sharing of an energy recharging station [23, 30], and patrolling of an area [31]. The majority of these studies were, however, performed exclusively in a simulated environment. The limited number of studies that used real robots have been conducted mostly with small (up to 15 cm) wheeled robots [32], such as the *Khepera* [33], the *s-bot* [34] and the *e-puck* [35]. These studies relied on abstract tasks in highly controlled environments: relatively small and enclosed arenas in laboratories.

The use of simulation in ER is widespread due to the large number of evaluations required by an evolutionary process to find adequate controllers. Transferring controllers evolved in simulation to real robots is currently one of the biggest challenges in ER [36], since the evolutionary process tends to exploit simulation-specific features that are not present in the real world, resulting in the real robots behaving differently and obtaining a lower task performance than in simulation. This issue is usually referred to as the *reality gap* [37].

Several approaches have been proposed to try to overcome the reality gap. Miglino et al. [38] suggested three complementary techniques: (i) using samples from the real robots' sensors to more accurately model them in simulation, (ii) introducing a conservative form of noise to promote the evolution of robust controllers, and (iii) continuing evolution in real hardware to tune controllers to the differences between simulation and reality. Jakobi [37] proposed the use of *minimal simulations*, in which the simulation model is based on a small set of features deemed necessary for successful synthesis of controllers, while remaining features are masked by an envelope of noise. Koos et al. [39] proposed *the transferability approach*, a multi-objective formulation of ER in which controllers are evaluated both by their performance in simulation and their performance on real robots. The approach requires periodical transfer of control during the evolutionary process in order to update a surrogate model, which estimates how well candidate solutions will transfer to the real world. Lehman et al. [40] proposed a different approach where the evolutionary process explicitly rewards behaviors that display a high reactivity, as measured by the mutual information between the magnitude of changes in a robot's sensors and actuators. The study showed that controllers with a reactive disposition transfer better to real robots than those not rewarded for reactive behavior, regardless of the noise conditions with which the controllers were evolved.

Another challenge in ER is how to drive the evolutionary process towards high-quality solutions, avoiding potential local optima. When the task to which solutions are sought reaches a certain level of complexity, traditional evolutionary approaches are prone to suffer from bootstrap problems and deception [41, 42]. One approach to deal with these issues is to employ diversity maintenance techniques, such as novelty search [43]. The goal is to allow an evolutionary algorithm to explore multiple different paths through the search space [23, 44, 45]. A different approach is to directly assist the evolutionary process with domain-specific knowledge, by using techniques such as incremental evolution [46, 47], behavioral decomposition [48, 49], or semi-interactive human-in-the-loop-techniques [50, 51]. In Section 7, we demonstrate a simple form of behavioral decomposition: we show how different evolved behaviors can be combined to produce a more complex controller, capable of performing an entire mission.

2.3 Aquatic Robots

In recent years, significant development has been conducted in the areas of autonomous underwater vehicles and autonomous surface vehicles [52], focusing on individual robots with a high degree of hardware and software complexity. Such systems have been applied to a variety of scenarios, such as environmental monitoring [53] and search and rescue [54]. Nonetheless, they are expensive, and restricted in terms of the tasks they can undertake as they have a

limited degree of autonomy, such as basic waypoint navigation capabilities. There have been numerous conceptual studies on systems composed of multiple robots with decentralized control for aquatic environments. However, only few such systems have been realized. Scerri *et al.* [55], for instance, made use of groups of up to five autonomous robots for distributed environmental monitoring missions. Their robots, however, need to be controlled by a central computer, which limits deployment in remote locations, and reduces their potential to scale to larger groups. Other studies in the domain of swarms of underwater robots have been performed in the scope of the CoCoRo project [56]. Regarding the evolution of control for aquatic robots, significantly fewer studies have approached the subject. Examples include the evolution of station keeping abilities for a single, simulated underwater robot [57], and the evolution of *centralized* control for a team of four underwater robots in a simulated predator-prey task [58].

3 Methodology

In this section, we present the methodology that we followed to produce controllers for the four tasks, and the steps we took to ensure a successful transfer of the controllers to the real robotic swarm. The methodology is summarized in Fig 2.

1. Definition of the simulation platform The evolutionary process used to synthesize behavior for the swarm relied exclusively on simulation for the evaluation of the candidate solutions. It is therefore necessary to implement a simulator in which the performance of the controllers can be assessed. We used a two-dimensional simulation environment, where the robots were abstracted as circular objects with a certain heading and position. The general principle behind the simulation was to model the motion of the robots based on real measurements taken in the water, but without including physics simulation and fluid dynamics, which would have resulted in a complex and computationally expensive simulation environment. Maintaining a low computational complexity is essential for making the evolutionary robotics process feasible, as a large number of simulations have to be conducted for each evolutionary run.

We developed a middle-layer that is shared by both the on-board software and the simulation platform. This enabled the same code-base and controllers to be executed on both the real robots and on robots in simulation, thereby minimizing the implementation differences between simulation and reality.

2. Configuration of the robots' sensors and actuators The processing of raw sensory data into values that can be fed to the neural controller is an essential step to enable the evolution of control. To this end, the robots (both real and simulated) relied on sensors that return the distance to objects of interest in equal-sized segments around the robot. These sensor values are obtained by pre-processing the GPS locations of entities in the environment of which the robot is currently aware, and the current heading and position of the robot. It should be noted that the sensors and actuators used by the real robots are inherently associated with a high degree of stochasticity. To account for these phenomena in simulation, we added noise to the sensors and actuators of the simulated robots, based on the observations made on the real robots. This conservative noise approach [38] is simple to implement, requires only basic assumptions and domain knowledge, and has proven effective in previous studies conducted in controlled conditions in our lab [49]. The noise included (i) the variation of robot parameters at the beginning of each simulation trial (e.g. maximum motor speed), (ii) noise injected at every simulation step (e.g. errors in sensory readings), and (iii) environmental noise (water currents).

3. Definition of the fitness evaluation For each task, we specified a fitness function that translated the task objective that needed to be achieved. The task setup chosen for the evaluation of controllers for each task approximated the setup in which the real robots were going to be used, meaning that the environment had approximately the same size, the number of robots

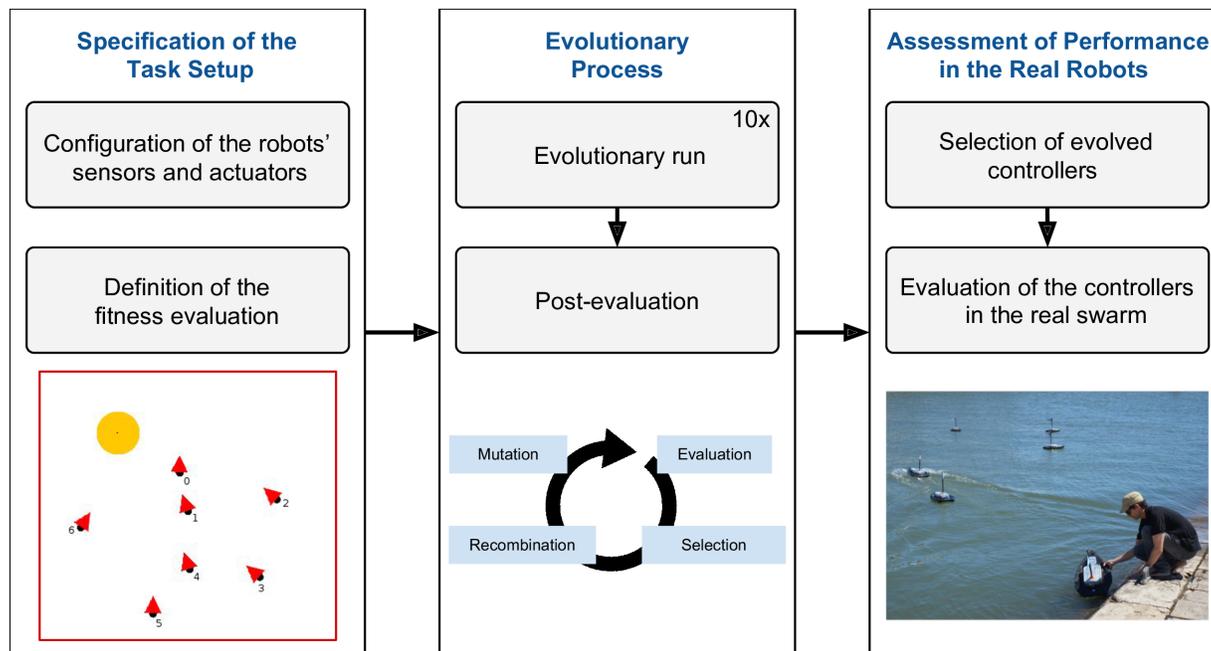


Fig 2. The control synthesis and performance assessment process for each task.

doi:10.1371/journal.pone.0151834.g002

was within the same range, and the robots had the same capabilities. To evaluate each candidate solution (controller), ten independent simulation trials were conducted, and the fitness of the controller was the mean score obtained in these simulations. To avoid overfitting the controllers to specific task conditions, we introduced random variations in the simulation trials during the evolutionary process. In every trial, we varied the number of robots (between five and ten), their starting position and orientation, and the location of other spatial entities (when applicable).

4. Evolutionary process The evolutionary algorithm optimized the neural network-based controllers for the robots. We used the NEAT algorithm [59] to evolve the neural network controllers for all tasks. NEAT is a widely used generational evolutionary algorithm that simultaneously optimizes the connection weights and the topology of artificial neural networks. The default NEAT parameters were used for all different tasks studied in this paper. Each evolutionary setup was repeated for ten independent evolutionary runs. After each evolutionary run, we post-evaluated the top controller of each generation in 100 simulations to obtain an accurate estimate of the controllers' quality.

5. Selection of evolved controllers For each task, we identified the highest-performing controller of each evolutionary run based on the post-evaluation results. Out of these ten best-of-run controllers, we then selected the three with the highest fitness scores, as evolutionary algorithms are stochastic, and separate evolutionary runs are likely to produce different controllers.

6. Assessment of performance in the real robots We finally assessed the real-robot performance of the top three controllers for each task, and compared their performance to that obtained in simulation. Each selected controller was evaluated in three independent experiments. For every experiment, the initial positions of the robots were randomized: a set of positions was generated inside a task-specific region, and each robot navigated to its assigned starting position. When all robots were in place, the experiment started and the controllers were simultaneously activated in all robots. The experiment ended after a fixed amount of

time. When evaluating the three different controllers for a given task, the same set of initial positions was used. To establish a fair comparison with controller performance in simulation, the same set of initial conditions were also used, i.e., the initial positions logged by robots were used as the initial robot positions for the simulation trials.

4 Experimental Setup

4.1 Robotic Platform

We designed and produced a total of ten simple, small (60 cm in length) and inexpensive (\approx 300 eur/unit) robots, using widely available, off-the-shelf hardware. Each robot model is a differential drive mono-hull boat (see Fig 3). The maximum speed of the robot is 1.7 m/s (3.3 kts), and the maximum angular velocity is 90°/s. The on-board control of each robot was executed by a Raspberry Pi 2 single-board computer. The robots form a distributed network without any central coordination or single point of failure through a IEEE 802.11g (Wi-Fi) ad-hoc wireless network, and communicate with one another by broadcasting UDP datagrams. Each robot is equipped with a GPS and a compass, and broadcasts its position to neighboring robots every second. See S1 Text for additional information on the hardware configuration of the robots.

The robotics simulation was performed on JBotEvolver [60]—an open-source framework that has been used in a large number of evolutionary robotics studies, and is available at <https://github.com/BioMachinesLab/jbotevolver> under GNU GPL license. The robot's dynamic model was implemented based on measurements taken from a single robot. We measured the robot's movement with multiple combinations of motor speeds, which allowed us to characterize the motion dynamics, friction, and inertia. We then tuned the model's parameters by comparing simulated and real motion patterns.

Each robot was controlled by an artificial neural network. The inputs of the neural network were the normalized readings of the sensors, and the outputs of the network controlled the robot's actuators. The sensor readings and actuation values were updated every 100 ms. The neural network controlling each robot had two output neurons which controlled the linear speed and the angular velocity of the robot. These two outputs were then converted to left and right speeds and applied to the robot's motors. We implemented three sensors for the detection



Fig 3. The robot is an autonomous surface vehicle equipped with Wi-Fi for communication, and a compass and GPS for navigation. It has a length of 60 cm and can move at speeds of up to 1.7 m/s.

doi:10.1371/journal.pone.0151834.g003

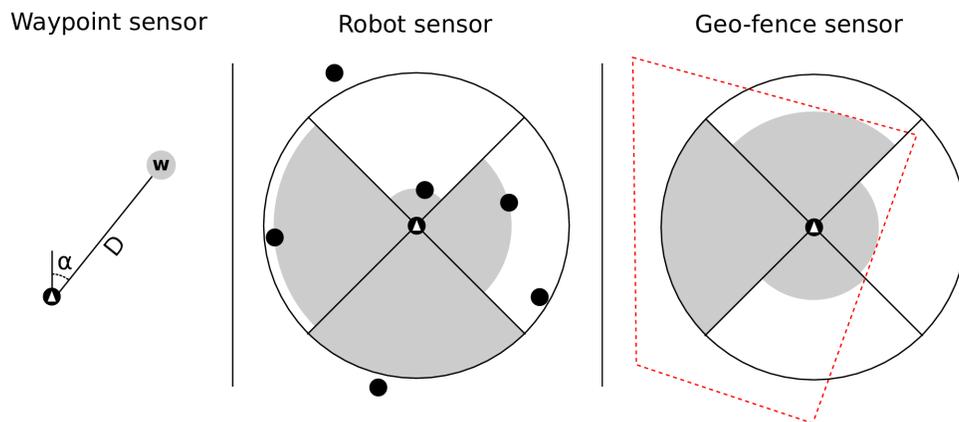


Fig 4. Illustration of the three types of sensors.

doi:10.1371/journal.pone.0151834.g004

of points and objects of interest in the task environment, based on GPS locations and on information shared by neighboring robots (see Fig 4):

- **Waypoint sensor:** This sensor is used to locate waypoints in the environment, if any. The sensor returns two values: (i) the relative angle from the robot to the waypoint, and (ii) the distance from the robot to the waypoint.
- **Robot sensor:** This sensor measures the distance to nearby robots. The area around the robot is divided into four equally-sized slices. The sensor returns four values, one per slice, linearly proportional to the distance between the robot and the closest neighbor in that slice, or the maximum value if there is none within the sensor range.
- **Geo-fence sensor:** This sensor is used to detect the geo-fence. Similarly to the robot sensor, the geo-fence sensor returns four values, one for each slice around the robot, indicating the minimum distance to the fence, or the maximum sensor value if the fence is outside the sensing range. One additional value is returned, indicating whether the robot is inside the fence or not (the geo-fence is a polygon).

4.2 Evolutionary Setup and Task Configurations

In this paper, we study the evolution of controllers for a real swarm of aquatic robots operating in a large uncontrolled environment. Our study relies on four cooperative tasks, which have been previously studied in simulation and laboratory environments [3]: (i) homing in group while avoiding collisions, (ii) dispersion, (iii) clustering, and (iv) monitoring of a pre-defined area. In this section, we describe the evolutionary setup that was used to synthesize control for each task, including the fitness functions and relevant parameters. All evolutionary and task parameters are listed in [S1 Text](#).

Homing. Navigation and obstacle avoidance is an essential feature for autonomous robots operating in the real world. This type of tasks was among the first to be studied in the field of evolutionary robotics [61]. A variation of the task, known as *homing*, involves one or more robots moving towards a point of interest in the environment, sometimes with obstacles in the way [62].

In the homing task used in our experiments, the swarm of robots had to navigate to a given waypoint while avoiding collisions between the robots. To evolve solutions for this task, we rewarded controllers for minimizing the average distance of the swarm to a waypoint,

according to the following equation:

$$f_{\text{homing}} = \left(\frac{1}{T} \sum_{t=1}^T \frac{1}{R} \sum_{r=1}^R \frac{\text{startingDist}_r - \text{dist}_{r,t}}{\text{startingDist}_r} \right) \times S, \quad (1)$$

where T is the maximum number of time steps, R is the number of robots, startingDist_r is the initial distance from robot r to the waypoint, and $\text{dist}_{r,t}$ is the distance of robot r to the waypoint at time t .

Collisions between robots are undesirable in any task, as they can damage the real robots. To avoid collisions, we introduced a *safety coefficient* (S) when assessing the fitness in all task setups, which penalized solutions where robots get too close to one another (less than 3 m). The safety coefficient S is in the range [0.1, 1], and is inversely proportional to the minimum distance between any two robots in the current simulation trial (minDist). The safety coefficient is defined according to the following equation:

$$S = 0.1 + \frac{\max(0, \min(3, \text{minDist}))}{3} \times 0.9, \quad (2)$$

Dispersion. In a dispersion task, each individual robot in the swarm has to maintain a pre-defined distance (*target distance*) from its nearest neighbor. Robots should cover a large area without risking losing contact with the rest of the swarm, which can be an issue in unbounded environments. Examples in the literature include solutions that rely on preprogrammed behaviors [63], potential fields [64], automatically synthesized state-machines [65], and artificially evolved ANNs [31].

For this task, we chose a target distance of 20 m (half of the robots' communication range). During evolution, controllers were rewarded for minimizing the difference between the current distance to the nearest neighbor and the target distance, according to the following equation:

$$f_{\text{dispersion}} = \left(\frac{1}{T} \sum_{t=1}^T \frac{1}{R} \sum_{r=1}^R 1 - |\text{minDist}_{r,t} - \text{targetDist}| \right) \times S, \quad (3)$$

where T is the maximum number of time steps, R is the number of robots, $\text{minDist}_{r,t}$ is the distance of robot r to the nearest neighbor at time t , and targetDist is the distance at which the robots should disperse.

Clustering. Clustering, also known as *aggregation*, is a canonical task in swarm robotics [1]. Clustering is a challenging task because it combines several aspects of multirobot tasks [66], including distributed individual search, coordinated movement, and cooperation. Furthermore, clustering plays an important role in robotics because it is a precursor of other collective behaviors such as group transport of heavy objects [67].

In our instance of the clustering task, the robots started randomly spread in a square-shaped area with a side-length of 100 m, and had the objective of finding one another so as to form a single cluster. However, since the communication range was limited and the environment was unbounded, certain initial conditions could lead to the formation of more than one cluster. During evolution, controllers were scored based on the number of clusters that they formed. Two robots belonged to the same cluster if the distance between them was inferior to 7 m. The fitness function was defined based on a weighted sum of the number of clusters throughout

time:

$$f_{clustering} = \frac{\sum_{t=1}^T t \times \frac{R-c_t}{R-1}}{\sum_{t=1}^T t} \times S, \tag{4}$$

where T is the total number of time steps in each trial, R is the number of robots, and c_t is the number of clusters formed by the robots at step t .

Area Monitoring. Swarms of robots are ideal for tasks where large areas need to be covered for monitoring or surveillance purposes. This task category has been explored in the past with several different control techniques, such as pheromone traces [68], odor sources [69] and evolved artificial neural networks [70].

In our area monitoring task, a geo-fence is defined to delimit an area of interest, and the robots should coordinate to continuously cover as much of the area as possible. During evolution, we tested the controllers in a variety of randomly generated geo-fences with different shapes, in order to obtain general behaviors. For the purpose of this task, we considered that each robot covers a circular area with a radius of 5 m (V) around it. The challenge in this task is to find a general movement pattern that takes into account many different shapes of the monitoring area and varying number of robots.

In order to assess the performance of the controllers, we divided the monitoring area into a grid with a fixed cell size of 1 m². Each robot could visit cells within the coverage radius V , setting its value to 1. The value of previously visited cells decayed linearly over a time frame of 100 s, down to 0. Controllers were scored based on how much of the grid was covered over time, according to the following equation:

$$f_{monitor} = \left(\frac{1}{T} \sum_{t=1}^T \frac{1}{C} \sum_{c=1}^C val(c_t) \right) \times S, \tag{5}$$

$$val(c_t) = \begin{cases} 0 & , t = 1 \\ \max(0, val(c_{t-1}) - \Delta) & , minDist_{c,t} > V \\ 1 & , minDist_{c,t} \leq V \end{cases}$$

where T is the maximum number of time steps, C is the number of cells, $minDist_{c,t}$ is the minimum distance of any robot to cell c at time t , and Δ is the decay for any cell (0.001 per time step).

4.3 Evolutionary Results

The results of the evolutionary algorithm for all tasks can be found in Fig 5. Solutions for homing, dispersion, and monitoring were found within 100 generations, while clustering required up to 400 generations for effective solutions to evolve. This can be explained by the challenge in forming a single cluster from random initial positions: since the environment is unbounded, some robots might find themselves outside of each other’s communication range, making the task more difficult. Section 5 covers the results in detail, including comparisons between performance in simulation and on the real robots.

5 Transferring Evolved Controllers to Real Robots

In this section, we assess the evolved controllers on real robots and compare their performance on real robots to their performance in simulation. The experiments reported in this section were conducted over the course of four days, at Parque das Nações, Lisbon, Portugal, in a semi-

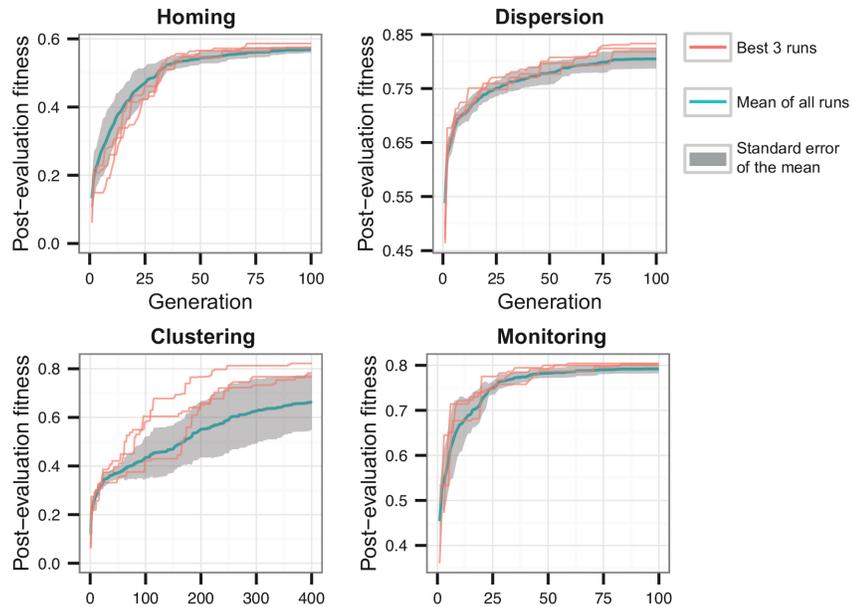


Fig 5. Fitness plot for the four different tasks. The plot shows the highest fitness scores found so far at each generation. The red lines depict the three highest-scoring evolutionary runs, while the blue line depicts the average of the ten runs, with the respective standard deviation shown in gray.

doi:10.1371/journal.pone.0151834.g005

enclosed waterbody with an area of 300 m × 190 m which is connected to the Tagus river, see Fig 6. The wind speed ranged from 15 to 40 km/h. The average drift speed of the robots was 0.2 m/s, twice as high as we had used in simulation (see SI Text). The experiments at Parque das Nações were carried out with the permission and cooperation of the nautical center Marina Parque das Nações, and preliminary experiments at the Lisbon Naval Base were carried out with the permission and cooperation of Arsenal do Alfeite and the Portuguese Navy. As described in Section 3, the top three controllers for each task were tested (referred to as Controllers 1, 2 and 3 in this section). Each controller was evaluated in three independent experiments (referred to as samples A, B and C).

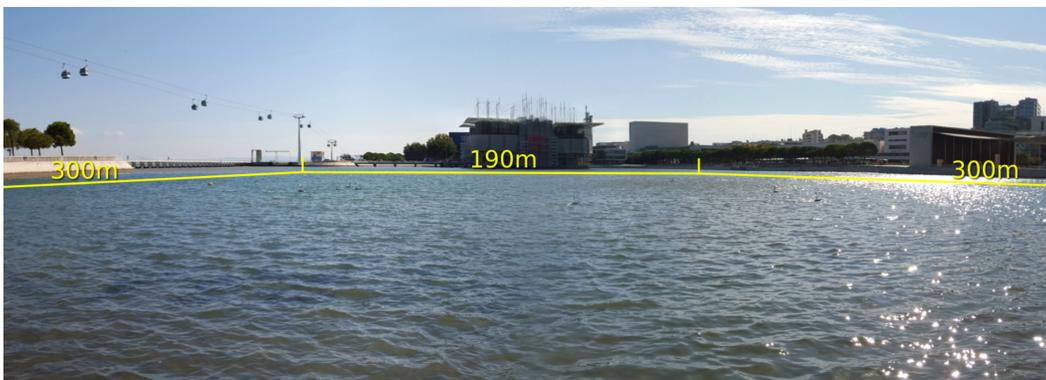


Fig 6. Panoramic photo of the location the experiments at Parque das Nações, Lisbon, Portugal.

doi:10.1371/journal.pone.0151834.g006

5.1 Homing

We setup a scenario where the swarm had to navigate sequentially to four waypoints to assess the performance of the evolved homing controllers. Each waypoint was placed at a distance of 40 m from the previous one, and the active waypoint was changed periodically every 60 s. Each controller was tested once in this scenario, for a total of 4 minutes. All three controllers tested in the real robots were able to navigate successfully to the waypoint while avoiding collisions with the neighboring robots. [Fig 7](#) (top) shows the average distance to the active waypoint, compared to the same controller in the simulation environment. The controllers display a similar behavior and level of performance in reality and in simulation, arriving at the intermediate waypoints at approximately the same time, and keeping the same average distance to them (around 10 m). Once the swarm arrived at one of the waypoints, the robots would start moving around the waypoint at slow speeds in concentric circles, see [Fig 7](#) (bottom).

5.2 Dispersion

In the real-robot dispersion tests, eight robots started in a cluster, and had to disperse to a distance of 20 m from their nearest neighbor. The robots were initially placed randomly in a square-shaped area with a side-length of 28 m, at a minimum distance of 5 m from the nearest neighbor. Each experiment lasted for 90 seconds.

The results in [Fig 8](#) show some differences in the performance of controllers on the real robots and in simulation. While performance was relatively homogeneous across all controllers in the simulation environment, with an error close to 1 m to the target distance, no controller was able to achieve the same level of performance on real robots. Controller 3 achieved an average error of 2 m on the real robots, but the other two controllers performed considerably worse, despite all of them displaying similar performance in simulation. Upon inspection of the behaviors, we observed that the robots in simulation relied on moving in circles at very low speeds to maintain their positions after dispersion. The same movement pattern was observed on the real robots. However, the circles performed by the real robots were larger, meaning that they tended to move away from their position, which is explained by the disparities between the motion model in simulation and the motion of the real robots.

5.3 Clustering

The clustering controllers were tested by randomly placing the real robots in an area of 100x100 m, up to a maximum of 40 m from the nearest robot. Each experiment lasted for 180 seconds. Results for the clustering experiments greatly varied depending on the initial conditions, but this variation was also observed in the simulation environment (see [Fig 9](#), top). If one or more robots were initially positioned far away from the remaining robots (and therefore not within communication range), the swarm would sometimes aggregate in more than one cluster, see example in [Fig 9](#) (bottom, Controller 2). The results show that the controllers were generally able to cross the reality gap successfully. In the case of Controller 1, it had a better performance in the real-robot experiments than in simulation, by successfully aggregating into a single cluster in two of the three samples. Controller 2, on the other hand, performed worse than in simulation, and Controller 3 displayed a similar level of performance.

5.4 Area Monitoring

We assessed each of the three highest-performing area monitoring controllers in real robots in three areas: square, L-shaped, and rectangular. Each of the areas covered a total of 10,000 m². It should be noted that the evolutionary process did not optimize the controllers for areas with

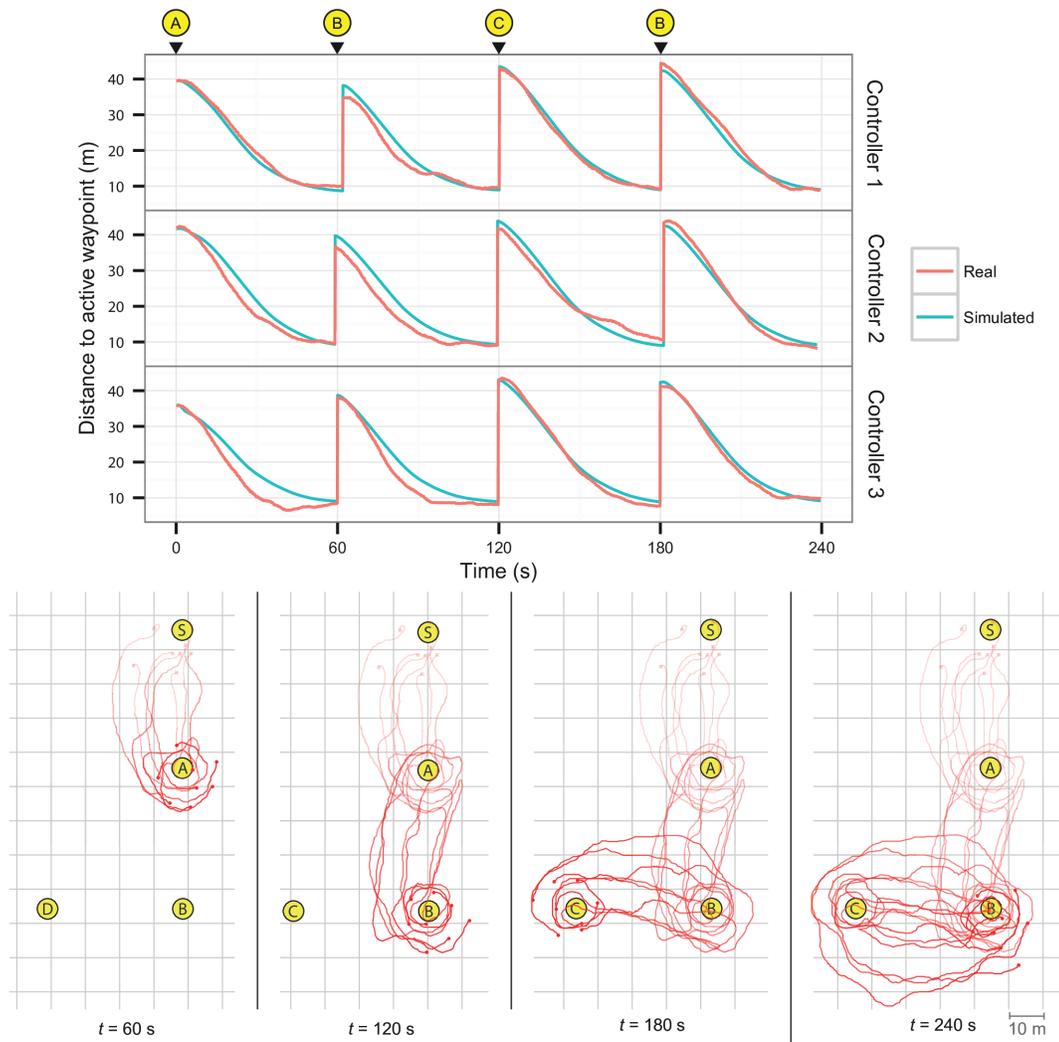


Fig 7. Real-world homing experiments with eight robots. The robots started around S. The active waypoint was then changed at 60 second intervals, in the order A→B→C→B, for a total of four minutes per experiment. Top: comparison between the real and simulated robots, showing the average distance to the active waypoint, for similar conditions. The top of the figure shows the current active waypoint. Bottom: trajectories of the real robots for Controller 3. The waypoints are marked with yellow circles.

doi:10.1371/journal.pone.0151834.g007

these specific shapes, but rather for randomly generated shapes in order to promote the evolution of general behaviors (see Section 4.2). Only the highest-performing controller was tested on real robots, since we did not observe any significant behavioral or performance differences in the other best controllers. Each experiment lasted for a total of five minutes.

The controllers successfully monitored the three different areas, performing significantly better on the real robots than in simulation (see Fig 10, top). These results are explained by speed differences between the simulated robots and some of the real robots. The speed differences directly influence the coverage measure, since a faster robot will be able to cover a larger area in the same amount of time. In terms of behavior (Fig 10, bottom), it is possible to see that the robots were able to effectively cover the square and rectangular areas, staying within the boundaries and moving away from the other nearby robots. The L-shaped area was also covered reasonably well, with the robots passing over all regions, but the intensity of the coverage was not uniform across the whole area. Although some robots would eventually move outside

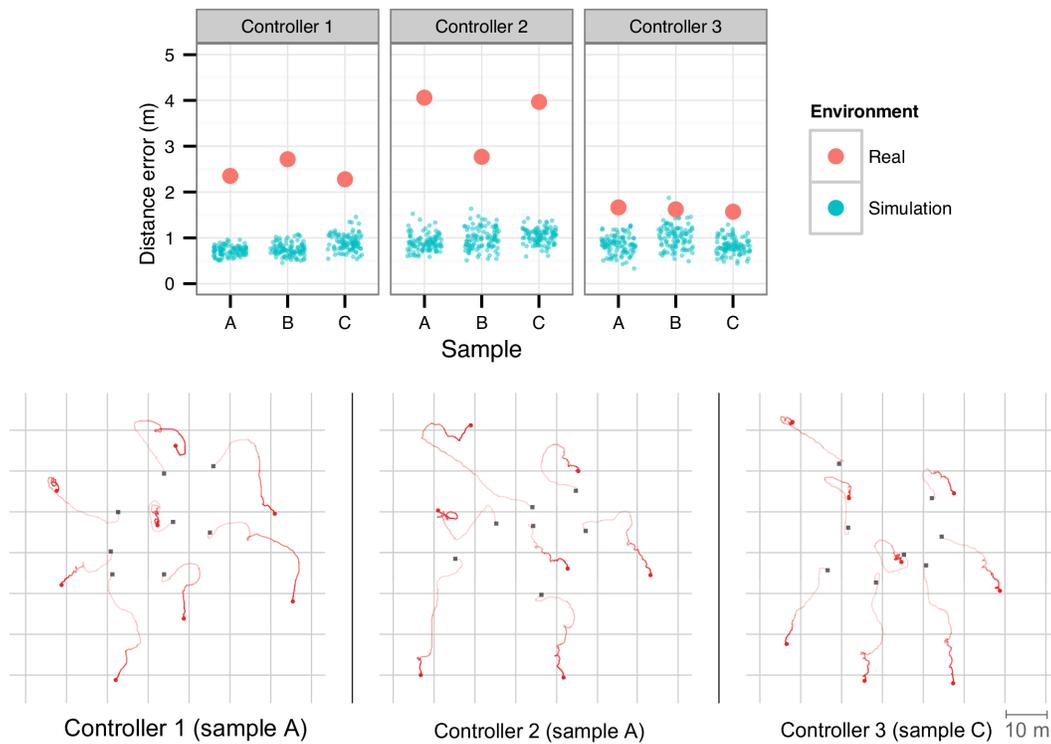


Fig 8. Real-world dispersion experiments with eight robots, one for each controller tested, over a period of 90 seconds. Top: average error to target distance (20 m) of the nearest robot in the last 10 s of each dispersion experiment. Bottom: trajectories of the real robots. The black squares mark the starting positions, and the red circles mark the final positions.

doi:10.1371/journal.pone.0151834.g008

of the pre-defined boundaries (both in simulation and on the real robots) either because of wind, currents, GPS inaccuracies, or too many robots nearby, they would quickly return to the area inside the boundaries and resume monitoring.

6 Scalability and Robustness

Desirable characteristics of swarm behaviors include scalability and robustness [12]. Performance should increase as more units are added (until issues such as congestion and overcrowding become a limiting factor), and failures in individual units should not compromise the performance of the rest of the swarm. To evaluate if our swarm of real robots with evolved control would display these properties, we ran an additional series of experiments: (i) scalability experiments for the dispersion and clustering tasks with four and six robots, (ii) robustness experiments for the dispersion task where a second group of robots is added to the swarm after a period of time, and (iii) robustness experiments for the monitoring task where robots are physically removed during the task executing (simulating faults), and then reintroduced after a period of time. All experiments were conducted in the real environment with the controllers that showed the highest performance in real hardware in the previous experiments.

6.1 Scalability

The dispersion and clustering controllers were tested with four and six robots (in addition to the first experiments with eight robots) to determine if the evolved controllers were able to perform well independently of the number of robots in the swarm. Fig 11 summarizes the results for these experiments.

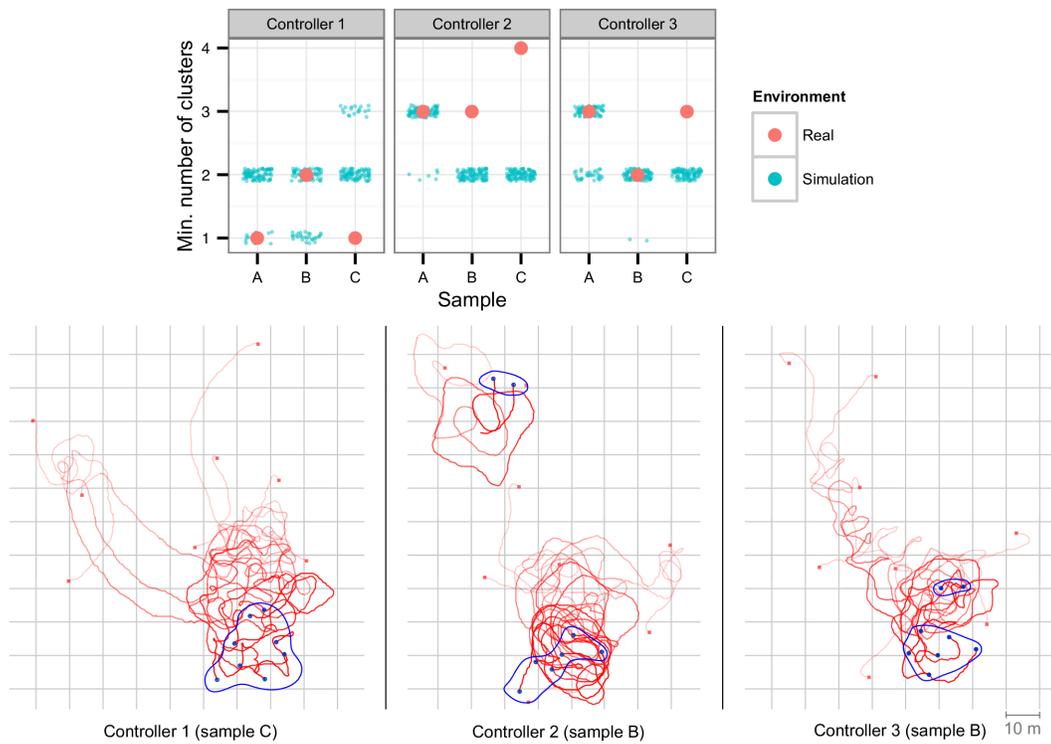


Fig 9. Real-world clustering experiments with eight robots, over a period of 180 seconds. Top: minimum number of clusters obtained in each sample. Bottom: trajectories of the real robots. The final clusters are highlighted in blue.

doi:10.1371/journal.pone.0151834.g009

In the dispersion experiments, while the controllers were able to perform the task successfully with swarms of different sizes, the results show a slight performance decrease in the experiments with six robots. It must be noted, however, that the error obtained is still relatively low, when compared to the target distance of 20 m. In the clustering experiments, the swarm was able to aggregate into a single cluster in all three samples with four robots, and in two out of three samples with six and eight robots. Since the communication range is limited, subsets of robots may initially be outside communication range or lose contact during an experiment, which can result in two or more isolated clusters.

6.2 Robustness

We setup a dispersion task that starts with a group of four robots (G_a). In the beginning of the experiment, G_a start dispersing. After 60 s, a group of four additional robots (G_b) starts moving towards the center of the swarm. When all robots of G_b are at the center ($t = 130$ s), their presence disturbs the dispersion of G_a , forcing G_a to separate even further. After the robots in G_b have been in the center for 60 s, they start the dispersion behavior along with the robots in G_a (at $t = 180$ s).

Fig 12 shows the results for the robustness experiments. When G_b moves to the center of G_a , the robots in G_a are forced to move away from one another since they try to stay at a fixed distance from the closest robot. When G_b starts dispersing, robots from G_a move further away to accommodate the movement pattern of G_b . When the whole swarm is dispersing, the distance error decreases with time as the robots continue adjusting their positions. The results demonstrate that the controllers were able to adjust to conditions that they were not exposed to during

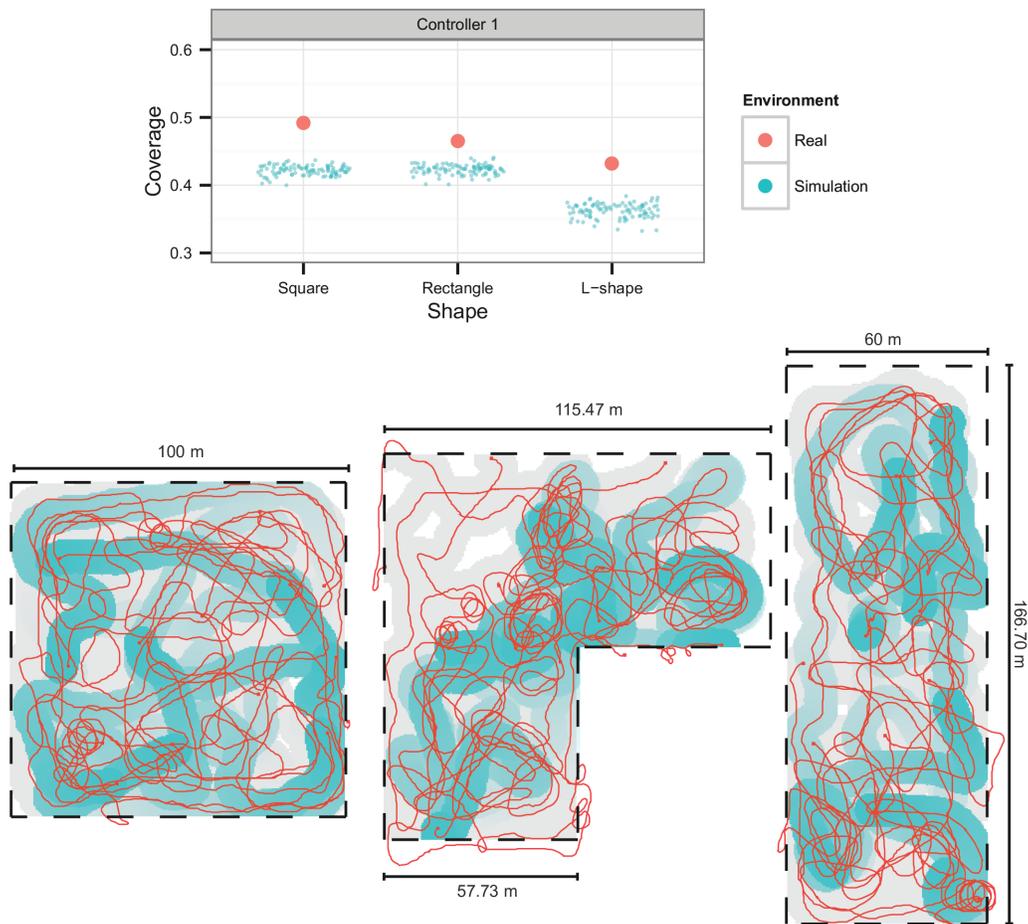


Fig 10. Real-world monitoring experiments with eight robots for Controller 1, over a period of five minutes. Top: coverage of the three different monitoring areas. Bottom: coverage maps in the experiments with the real swarm. The coverage of the area is presented in blue, and has a decay of 100 s. Trajectories for the full duration of the task are presented in red, and all the areas visited by the robots are filled in gray.

doi:10.1371/journal.pone.0151834.g010

evolution, particularly the addition of new robots during task execution, and still carry out the task successfully.

In an additional experiment, we physically removed and added units during the execution of the monitoring task. The area to monitor was a square with 100×100 m. The monitoring

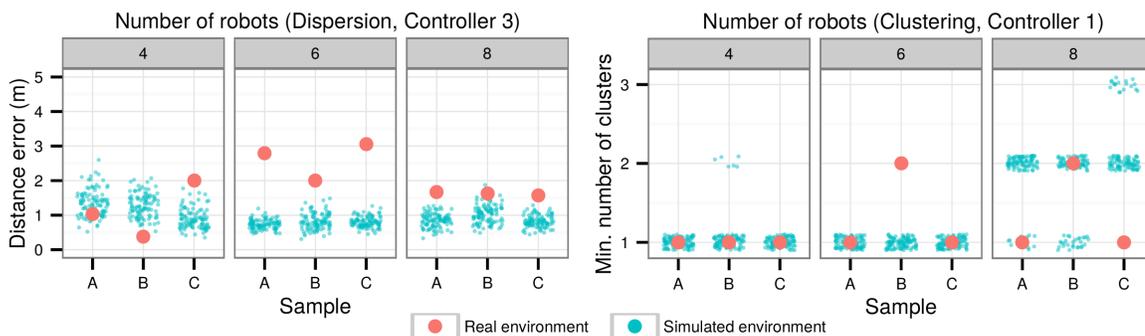


Fig 11. Scalability experiments with dispersion (Controller 3, left) and clustering (Controller 1, right) controllers. In each task, the same controller was used in a swarm of four, six, and eight robots, with three samples for each setup.

doi:10.1371/journal.pone.0151834.g011

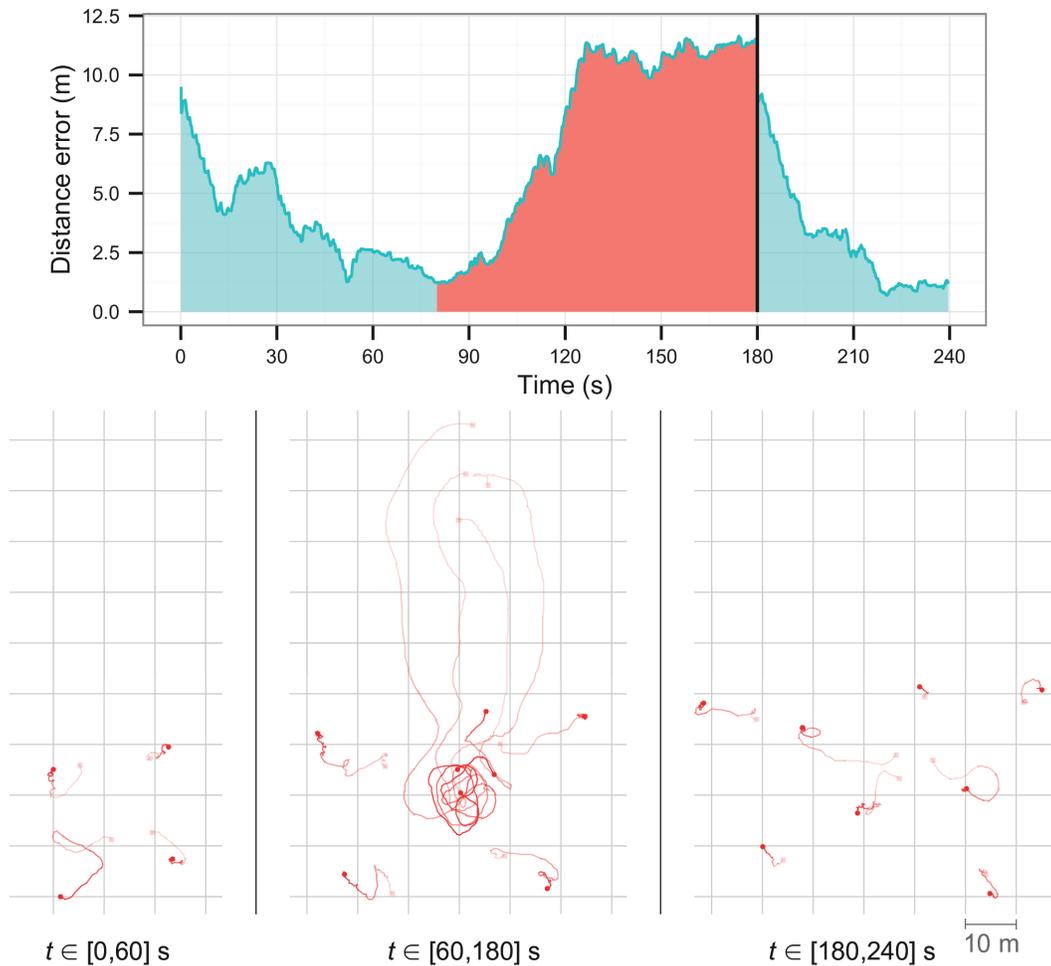


Fig 12. Robustness experiments with Controller 3 of the dispersion behavior. The red area represents the period where the robots of G_b are disturbing the dispersion of G_a , and the black vertical line at $t = 180$ s indicates the point where the robots in G_b start dispersing, and where the distance error starts being measured for all eight robots.

doi:10.1371/journal.pone.0151834.g012

controller was executed for a total of 15 minutes, starting with eight robots. At the 5-minute mark, we removed four of the robots, and at the 10-minute mark, two new robots were added. Fig 13 shows the coverage of the monitoring area at any given time during the experiment. The performance decrease after the 5-minute (300 s) mark, and the consequent increase at the 10 minute (600 s) mark show that the performance of the evolved controllers tends to scale linearly with the number of robots executing the task (see mean coverage in Fig 13). These results indicate that the evolved controller is robust to variations in the composition of the swarm during task execution. A video of the experiment can be found in the supplemental information (S1 Video).

7 Multi-controller Mission

In this section, we show how relatively simple swarm behaviors can be combined to produce more sophisticated behaviors that could be used in realistic applications. We integrate the highest-performing behaviors described in the previous section to accomplish a mission of sampling the water temperature over a given area of interest. The robots all started near a base

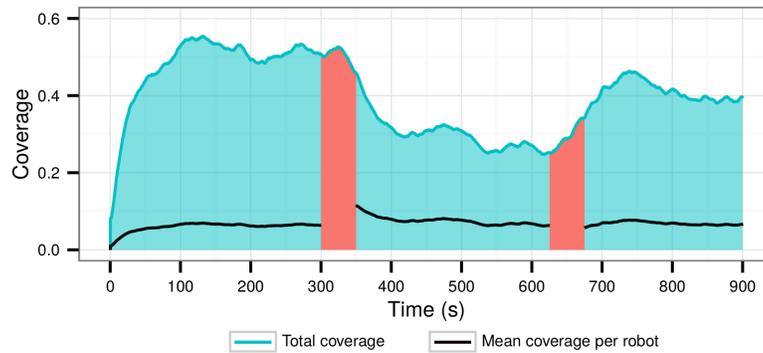


Fig 13. Robustness experiments with Controller 1 of the monitoring behavior. The time regions highlighted in red correspond to the periods when robots were either entering or leaving the monitoring area.

doi:10.1371/journal.pone.0151834.g013

station, outside the area of interest. The mission was decomposed in five sequential sub-tasks: (i) collectively navigate from the base station to the center of the area of interest (homing); (ii) disperse; (iii) monitor the area of interest; (iv) aggregate, using the clustering behavior; and finally (v) return to the base station (homing). The robots start collecting temperature data after reaching the center of the monitoring area, before sub-task (ii). The different behaviors were triggered sequentially to produce a complete mission, with each behavior lasting for a pre-defined amount of time.

The trajectories of the real robots executing the multi-controller mission and the temperature data collected can be seen in Fig 14. The uncertainty of the temperature estimate decreases as the robots cover the area, resulting in a more accurate profile. At the end of the monitoring period ($t = 360$ s), the estimated standard deviation of the error was < 0.14 for the whole area except in the corners where it was slightly higher (< 0.24). A video of the experiment is available in S2 Video.

8 Discussion

Our experiments showed that evolutionary computation is a viable approach for the synthesis of control for swarms of robots operating outside of controlled laboratory environments. Below, we discuss the main challenges we faced in evolving control for the SRS, and the properties observed in the SRS with evolved control.

8.1 Evolutionary Process

For the most part, the definition and configuration of the evolutionary process was straightforward. The main challenge in configuring the evolutionary algorithm was the definition of the fitness functions. The fact that we used common swarm robotics tasks facilitated the definition of the fitness functions, however, once we started considering more complex tasks (as the temperature monitoring task in Section 7), it became hard to design effective fitness function that could express the task objectives. Moreover, previous works have shown that using standard evolutionary techniques for such complex and sequential tasks can be inadequate due to bootstrapping and deception issues [46, 72–74].

One solution is to synthesize the various *building blocks* a-priori, and then combine them in order to solve complex tasks [49]. We have explored a variant of this line of research in this paper with the multi-controller mission (Section 7), where a sequence of different behaviors were triggered to produce a complete mission, with each behavior executing for a predefined amount of time. However, as we have successfully shown in previous work, more complex decision

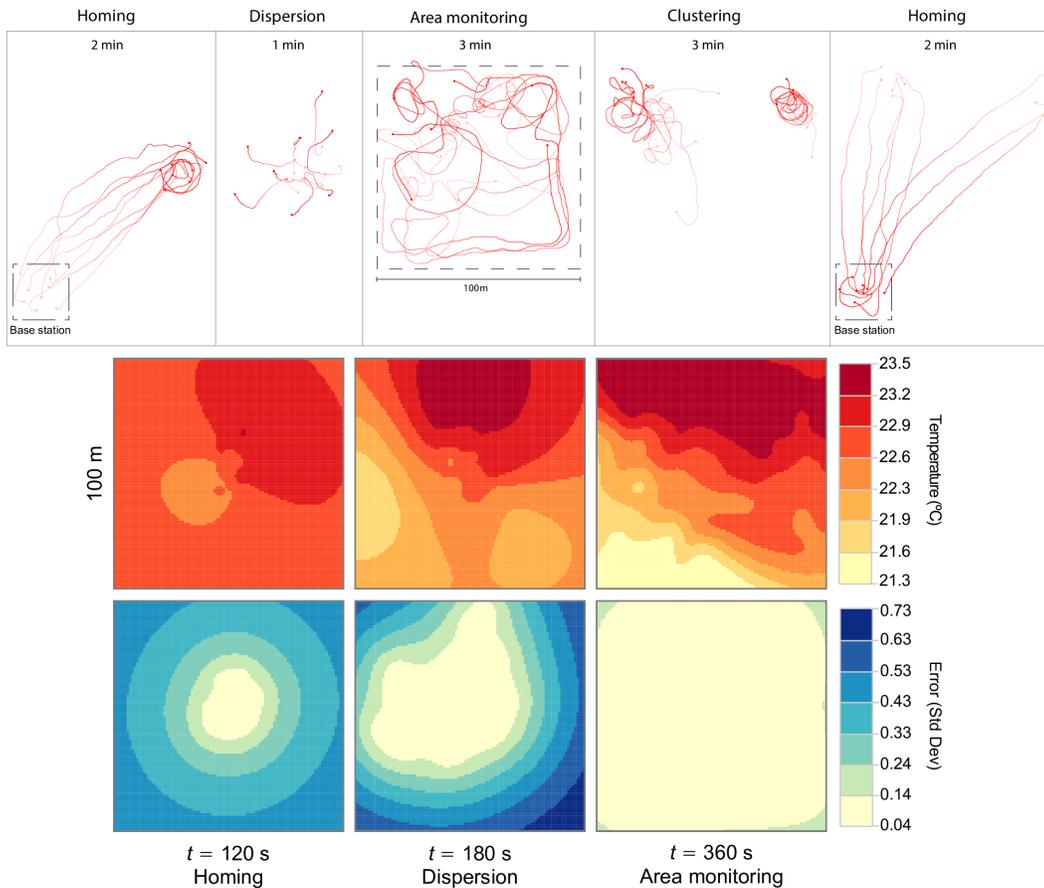


Fig 14. Results for the multi-controller mission. Top: robot trajectories for each sub-task. Middle and bottom: temperatures in the monitoring area. Measurements taken by the robots' temperature sensors were spatially interpolated using kriging [71]. Data collection started after the robots arrived at the waypoint ($t = 100$ s). The middle row shows the predicted temperatures, while the bottom row shows the estimated error of the predictions.

doi:10.1371/journal.pone.0151834.g014

schemes can be used to articulate different behaviors [31, 49, 75]. In recent work [76], we have taken the first steps towards the demonstration of hierarchical control in a real robotic swarm, and demonstrate how this control approach can produce modular and flexible behaviors.

8.2 Transferability of Controllers

One of the challenges of synthesizing control for aquatic surface robots is dealing with the unpredictable environmental conditions and unexpected circumstances associated with this environment. Many of the environmental variables affecting the robots can be extremely difficult, or even impossible, to model accurately. The simulation model was therefore a substantial simplification of reality. To compensate for the differences between the motion model and the real dynamics of the robots, as well as for other potential sources of stochasticity, we introduced a conservative amount of noise in the sensors, actuators, and environment, during evolution (Section 3). The introduction of noise is a computationally-effective way of promoting the evolution of robust controllers [38]. In addition to noise, each controller was evaluated in multiple trials with variations in the task setup, such as the number of robots in each sample, or the initial positions of the robots and environmental features, which also contributed to the increase of transferability of the evolved controllers.

Our experimental results showed that the evolved controllers were generally able to cross the reality gap. In a number of experiments, a speed difference between some of the real robots and the simulated robots (up to 20%, in some cases) contributed to differences in terms of performance. While the behaviors of the real robots were generally similar to those observed in simulation, some motion patterns did not transfer well. With the dispersion controllers, for instance, in simulation the robots started moving in small circles to maintain their current position. In reality, however, these circles were larger, which resulted in a lower performance. This difference highlights the limitations of the simplified physics model we adopted in the simulation environment. We are considering two different approaches to address these issues in future work: (i) further improve the simulated movement dynamics, thereby decreasing the mismatch between simulation and reality, and (ii) select for controllers that rely on motion patterns that transfer well from simulation to reality, as proposed in previous works [39, 40, 77].

8.3 Swarm Behaviors

In this paper, we relied on four canonical swarm tasks (homing, dispersion, clustering, and area monitoring) to validate the basic properties of swarm robotics in real conditions. All the collective tasks presented in this paper can be classified as swarm robotics tasks, according to the classification proposed by Brambilla et al. [1]: (i) robots are autonomous; (ii) robots are situated in the environment; (iii) robots' sensing and communication capabilities are local (up to 40 m, while the robots can be over 200 m apart from each other); (iv) robots do not have access to centralized control and/or to global information; and (v) robots cooperate to solve a given task. Additionally, Şahin [12] states three properties as motivations for swarm robotics, which we were able to verify with the experiments described in the paper:

Robustness: We demonstrated robustness explicitly with the monitoring task, by removing robots during task execution (see Section 6.2). In all tasks, there were also sporadic hardware failures, which resulted in individual failures that did not impact the behavior of the group (Section 6.2).

Flexibility: The flexibility of the swarm could be observed in the different coordination strategies that were adopted by the swarm, when using controllers for different tasks. Furthermore, we showed the behavioral modularity of the system by assembling a high-level controller that relied on the combination of multiple previously evolved controllers (see Section 7).

Scalability: We studied scalability explicitly using the clustering and dispersion behaviors (Section 6.1). We tested the same controller with different swarm sizes (from four to eight robots), and the performance and behavior of the swarm was similar regardless of the swarm size. In the monitoring task, we showed that the performance of the swarm (amount of space covered) was proportional to the number of robots (Section 6.2).

8.4 Robustness to Hardware Faults

The use of robots in real-world conditions will often imply hardware malfunctions, caused by prolonged use and environmental factors. In our experiments, the main source of faults were the robots' motors, which sporadically and temporarily stopped, causing the robot to stop completely or start moving in circles. Other problems included the decrease of motor power with lower battery levels, and often erroneous GPS readings. These faults occurred often during the execution of all experiments reported in this paper. Even though these faults were not contemplated in simulation, and therefore were not taken into account during the evolutionary process, they did not compromise the overall performance of the swarm, highlighting the inherent robustness of swarm behaviors. In future work, we will study the fault-tolerance

properties of the system more systematically by injecting different types of faults [78], and by measuring the impact of these faults in the performance of the swarm.

An additional source of uncertainty, related to the hardware, was the heterogeneity among the robots of the swarm. Robots within the swarm often had different battery levels and motors with slight performance differences. We measured this heterogeneity in the monitoring task, where all robots display similar behaviors. In the real swarm, the difference between the slowest and the fastest robot was up to 0.6 m/s at certain times, while for the same task, in simulation, the difference between the slowest and the fastest robot was just 0.1 m/s. This suggests that the real swarm had to deal with a much higher degree of heterogeneity than what was contemplated during the evolutionary process.

8.5 Towards Real World Applications

In this study, we have demonstrated that swarm robotics systems can effectively operate in real-world conditions. There are, however, a number of control-related issues to be addressed before the practical use of swarms of robots in real-world applications. Prominent challenges include:

Integration of mission-specific sensors: In the experiments presented in this paper, the robot controllers relied mostly on their position and on the position of the neighboring robots for the behavioral decision process. The integration of data gathered from onboard sensors such as cameras or sonars, may however be an essential input for the controllers in other tasks. The challenge of evolving controllers capable of dealing with such data is twofold: (i) the data might not be trivial to decompose into meaningful information that can be fed to neural controllers; and (ii) the sensors would have to be accurately modeled in simulation.

Robustness to faults: Another significant challenge is endowing controllers with the capability to cope with faults within the swarm. Although our experiments confirmed that swarm behaviors inherently have some degree of tolerance for individual faults, in real-world applications it may be necessary to guarantee that faulty individuals will not compromise the whole swarm. Approaches for fault detection [78], fault tolerance [79], and online adaptation [77] must therefore be implemented in the system.

Safety: The robots should have the capability of dealing with external factors, such as nearby vessels, in a way that guarantees the safety of all parties and conforms with the existing regulations. Combining such fail-safe behaviors with the mission controllers may require the combination of evolved and manually programmed behaviors [75].

Behavior flexibility and modularity: One concern when using evolutionary robotics is the difficulty in rapidly producing new behaviors or modify existing ones. One possible solution for this problem is to compose behaviors for new tasks based on existing building blocks [49]. In this study, we showed how multiple evolved behaviors could be sequentially arranged in order to perform an environmental monitoring task. Future work should explore the usage of more capable techniques for combining controllers, such as hierarchical composition [49, 76].

9 Conclusion

In this paper, we present for the first time a swarm robotics system with evolved control operating in a real-world uncontrolled environment. Our experiments relied on a swarm of small and inexpensive aquatic surface robots that used distributed self-organized control and local communication. In a simulation environment, we evolved controllers for four common collective tasks: homing, clustering, dispersion, and area monitoring.

The evolved controllers were then systematically evaluated in the real robots in a large and uncontrolled aquatic environment. Our results showed that the controllers displayed similar

behaviors and levels of performance in the real swarm as those observed in simulation. Moreover, we verified key desirable properties of swarm robotics systems, namely robustness, flexibility, and scalability. We finally showed how simple evolved behaviors can be combined to produce a complete mission, by carrying out a complete environmental monitoring task.

Past studies in the field of swarm robotics and evolutionary robotics have been confined either to simulation (the vast majority), or to highly-controlled laboratory conditions. Our study contrasts with the previous work in these fields, as it relies on an evolved swarm of robots operating in a realistic and uncontrolled environment. Overall, our research demonstrates the potential of both swarm robotics systems and evolutionary robotics techniques as viable approaches for controlling multirobot systems in real-world scenarios.

Supporting Information

S1 Text. Experimental Details. Document detailing the aquatic robotic platform and the experimental parameters of both the simulated and real experiments, in some cases repeating and extending the information given in the main manuscript.

(PDF)

S1 Video. Adaptive monitoring experiment. Traces of the robots' positions when performing the experiment where the robustness of the swarm is assessed using the area monitoring task, as described in Section 6.2.

(MP4)

S2 Video. Multi-controller mission. Video of the experiment detailed in Section 7, where the different swarm behaviors are composed sequentially in order to produce control for a temperature-sampling mission.

(MP4)

Author Contributions

Conceived and designed the experiments: MD JG SMO ALC. Performed the experiments: MD VC JG TR FS SMO ALC. Analyzed the data: MD JG. Wrote the paper: MD JG FS ALC. Developed the robotic platform: MD VC TR.

References

1. Brambilla M, Ferrante E, Birattari M, Dorigo M. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*. 2013; 7(1):1–41. doi: [10.1007/s11721-012-0075-2](https://doi.org/10.1007/s11721-012-0075-2)
2. Bayındır L, Şahin E. A Review of Studies in Swarm Robotics. *Turkish Journal of Electrical Engineering & Computer Sciences*. 2007; 15(2):115–147.
3. Bayındır L. A review of swarm robotics tasks. *Neurocomputing*. 2015; In press.
4. Dorigo M, Trianni V, Şahin E, Groß R, Labella TH, Baldassarre G, et al. Evolving self-organizing behaviors for a swarm-bot. *Autonomous Robots*. 2004; 17(2–3):223–245. doi: [10.1023/B:AURO.0000033973.24945.f3](https://doi.org/10.1023/B:AURO.0000033973.24945.f3)
5. Nolfi S, Floreano D. *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge, MA; 2000.
6. Bongard JC. Evolutionary robotics. *Communications of the ACM*. 2013; 56(8):74–83. doi: [10.1145/2492007.2493883](https://doi.org/10.1145/2492007.2493883)
7. Lipson H, Pollack J. Automatic design and manufacture of robotic lifeforms. *Nature*. 2000; 406(6799):974–978. doi: [10.1038/35023115](https://doi.org/10.1038/35023115) PMID: [10984047](https://pubmed.ncbi.nlm.nih.gov/10984047/)
8. Dorigo M, Floreano D, Gambardella LM, Mondada F, Nolfi S, Baaboura T, et al. Swarmanoid: a novel concept for the study of heterogeneous robotic swarms. *IEEE Robotics & Automation Magazine*. 2013; 20(4):60–71. doi: [10.1109/MRA.2013.2252996](https://doi.org/10.1109/MRA.2013.2252996)

9. Lindsey Q, Mellinger D, Kumar V. Construction with quadrotor teams. *Autonomous Robots*. 2012; 33(3):323–336. doi: [10.1007/s10514-012-9305-0](https://doi.org/10.1007/s10514-012-9305-0)
10. Basiri M, Schill F, Floreano D, Lima PU. Audio-based localization for swarms of micro air vehicles. In: *Proceedings of the 2014 IEEE International Conference on Robotics and Automation*. IEEE Press, Piscataway, NJ; 2014. p. 4729–4734.
11. Yan Rj, Pang S, Sun Hb, Pang Yj. Development and missions of unmanned surface vehicle. *Journal of Marine Science and Application*. 2010; 9(4):451–457. doi: [10.1007/s11804-010-1033-2](https://doi.org/10.1007/s11804-010-1033-2)
12. Şahin E. Swarm robotics: From sources of inspiration to domains of application. In: *Proceedings of the International Workshop on Swarm Robotics*. Springer, Berlin, Germany; 2005. p. 10–20.
13. Camazine S, Deneubourg JL, Franks NR, Sneyd J, Theraulaz G, Bonabeau E. *Self-organization in biological systems*. Princeton University Press, Princeton, NJ; 2003.
14. Jones CV, Mataric MJ. Behavior-based Coordination in Multi-robot Systems. In: Ge SS, Lewis FL, editors. *Autonomous Mobile Robots: Sensing, Control, Decision Making and Applications*. CRC Press, Boca Raton, FL; 2006. p. 549–569.
15. Nolfi S. Evolutionary robotics: Exploiting the full power of self-organization. *Connection Science*. 1998; 10(3–4):167–184. doi: [10.1080/095400998116396](https://doi.org/10.1080/095400998116396)
16. Trianni V, Nolfi S. Engineering the evolution of self-organizing behaviors in swarm robotics: A case study. *Artificial Life*. 2011; 17(3):183–202. doi: [10.1162/artl_a_00031](https://doi.org/10.1162/artl_a_00031) PMID: [21554112](https://pubmed.ncbi.nlm.nih.gov/21554112/)
17. Floreano D, Keller L. Evolution of Adaptive Behaviour in Robots by Means of Darwinian Selection. *PLoS Biology*. 2010; 8(1):1–8. doi: [10.1371/journal.pbio.1000292](https://doi.org/10.1371/journal.pbio.1000292)
18. Baldassarre G, Trianni V, Bonani M, Mondada F, Dorigo M, Nolfi S. Self-Organized Coordinated Motion in Groups of Physically Connected Robots. *IEEE Transactions on Systems, Man, and Cybernetics*. 2007; 37(1):224–239. doi: [10.1109/TSMCB.2006.881299](https://doi.org/10.1109/TSMCB.2006.881299) PMID: [17278574](https://pubmed.ncbi.nlm.nih.gov/17278574/)
19. Sperati V, Trianni V, Nolfi S. Evolving coordinated group behaviours through maximisation of mean mutual information. *Swarm Intelligence*. 2008; 2(2–4):73–95. doi: [10.1007/s11721-008-0017-1](https://doi.org/10.1007/s11721-008-0017-1)
20. Sperati V, Trianni V, Nolfi S. Self-organised path formation in a swarm of robots. *Swarm Intelligence*. 2011; 5(2):97–119. doi: [10.1007/s11721-011-0055-y](https://doi.org/10.1007/s11721-011-0055-y)
21. Trianni V, Groß R, Labella TH, Şahin E, Dorigo M. Evolving aggregation behaviors in a swarm of robots. In: *Proceedings of the 7th European Conference on Artificial Life*. Springer, Berlin, Germany; 2003. p. 865–874.
22. Soysal O, Bahçeci E, Sahin E. Aggregation in swarm robotic systems: Evolution and probabilistic control. *Turkish Journal of Electrical Engineering & Computer Sciences*. 2007; 15(2):199–225.
23. Gomes J, Urbano P, Christensen AL. Evolution of swarm robotics systems with novelty search. *Swarm Intelligence*. 2013; 7(2–3):115–144. doi: [10.1007/s11721-013-0081-z](https://doi.org/10.1007/s11721-013-0081-z)
24. Baldassarre G, Nolfi S, Parisi D. Evolving mobile robots able to display collective behaviors. *Artificial Life*. 2003; 9(3):255–267. doi: [10.1162/106454603322392460](https://doi.org/10.1162/106454603322392460) PMID: [14556687](https://pubmed.ncbi.nlm.nih.gov/14556687/)
25. Trianni V, Nolfi S, Dorigo M. Cooperative hole avoidance in a swarm-bot. *Robotics and Autonomous Systems*. 2006; 54(2):97–103. doi: [10.1016/j.robot.2005.09.018](https://doi.org/10.1016/j.robot.2005.09.018)
26. Hauert S, Zufferey JC, Floreano D. Evolved swarming without positioning information: an application in aerial communication relay. *Autonomous Robots*. 2009; 26(1):21–32. doi: [10.1007/s10514-008-9104-9](https://doi.org/10.1007/s10514-008-9104-9)
27. Ampatzis C, Tuci E, Trianni V, Dorigo M. Evolution of Signaling in a Multi-Robot System: Categorization and Communication. *Adaptive Behavior*. 2008; 16(1):5–26. doi: [10.1177/1059712307087282](https://doi.org/10.1177/1059712307087282)
28. Groß R, Dorigo M. Evolution of solitary and group transport behaviors for autonomous robots capable of self-assembling. *Adaptive Behavior*. 2008; 16(5):285–305. doi: [10.1177/1059712308090537](https://doi.org/10.1177/1059712308090537)
29. Pini G, Tuci E. On the design of neuro-controllers for individual and social learning behaviour in autonomous robots: an evolutionary approach. *Connection Science*. 2008; 20(2–3):211–230. doi: [10.1080/09540090802092014](https://doi.org/10.1080/09540090802092014)
30. Gomes J, Christensen AL. Generic Behaviour Similarity Measures for Evolutionary Swarm Robotics. In: *Genetic and Evolutionary Computation Conference (GECCO)*. ACM Press; 2013. p. 199–206.
31. Duarte M, Oliveira SM, Christensen AL. Hybrid Control for Large Swarms of Aquatic Drones. In: *Proceedings of the 14th International Conference on the Synthesis & Simulation of Living Systems*. MIT Press, Cambridge, MA; 2014. p. 785–792.
32. Khaldi B, Cherif F. An Overview of Swarm Robotics: Swarm Intelligence Applied to Multi-robotics. *International Journal of Computer Applications*. 2015; 126(2):31–37. doi: [10.5120/ijca2015906000](https://doi.org/10.5120/ijca2015906000)
33. Mondada F, Franzi E, Guignard A. The development of khepera. In: *Experiments with the Mini-Robot Khepera*, Proceedings of the First International Khepera Workshop; 1999. p. 7–14.

34. Mondada F, Pettinaro GC, Guignard A, Kwee IW, Floreano D, Deneubourg JL, et al. SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*. 2004; 17(2–3):193–221. doi: [10.1023/B:AURO.0000033972.50769.1c](https://doi.org/10.1023/B:AURO.0000033972.50769.1c)
35. Mondada F, Bonani M, Raemy X, Pugh J, Cianci C, Klaptocz A, et al. The e-puck, a robot designed for education in engineering. In: *Proceedings of the 9th conference on autonomous robot systems and competitions*. IPCB: Instituto Politécnico de Castelo Branco; 2009. p. 59–65.
36. Doncieux S, Bredeche N, Mouret JB, Eiben AE. Evolutionary robotics: what, why, and where to. *Frontiers in Robotics and AI*. 2015; 2(4).
37. Jakobi N. Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive Behavior*. 1997; 6(2):325–368. doi: [10.1177/105971239700600205](https://doi.org/10.1177/105971239700600205)
38. Miglino O, Lund HH, Nolfi S. Evolving Mobile Robots in Simulated and Real Environments. *Artificial Life*. 1996; 2(4):417–434. doi: [10.1162/artl.1995.2.4.417](https://doi.org/10.1162/artl.1995.2.4.417)
39. Koos S, Mouret JB, Doncieux S. The transferability approach: Crossing the reality gap in evolutionary robotics. *IEEE Transactions on Evolutionary Computation*. 2013; 17(1):122–145. doi: [10.1109/TEVC.2012.2185849](https://doi.org/10.1109/TEVC.2012.2185849)
40. Lehman J, Risi S, D'Ambrosio D, Stanley KO. Encouraging reactivity to create robust machines. *Adaptive Behavior*. 2013; 21(6):484–500. doi: [10.1177/1059712313487390](https://doi.org/10.1177/1059712313487390)
41. Doncieux S, Mouret JB. Beyond black-box optimization: a review of selective pressures for evolutionary robotics. *Evolutionary Intelligence*. 2014; 7(2):71–93. doi: [10.1007/s12065-014-0110-x](https://doi.org/10.1007/s12065-014-0110-x)
42. Silva F, Duarte M, Correia L, Oliveira SM, Christensen AL. Open Issues in Evolutionary Robotics. *Evolutionary Computation*. 2016; In press.
43. Lehman J, Stanley KO. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary Computation*. 2011; 19(2):189–223. doi: [10.1162/EVCO_a_00025](https://doi.org/10.1162/EVCO_a_00025) PMID: [20868264](https://pubmed.ncbi.nlm.nih.gov/20868264/)
44. Mouret JB, Doncieux S. Encouraging behavioral diversity in evolutionary robotics: An empirical study. *Evolutionary Computation*. 2012; 20(1):91–133. doi: [10.1162/EVCO_a_00048](https://doi.org/10.1162/EVCO_a_00048) PMID: [21838553](https://pubmed.ncbi.nlm.nih.gov/21838553/)
45. Gomes J, Mariano P, Christensen AL. Novelty-driven Cooperative Coevolution. *Evolutionary computation*. 2016; In press.
46. Mouret JB, Doncieux S. Incremental evolution of animats' behaviors as a multi-objective optimization. In: *Proceedings of the 10th International Conference on Simulation of Adaptive Behaviour*. Springer, Berlin, Germany; 2008. p. 210–219.
47. Christensen AL, Dorigo M. Incremental Evolution of Robot Controllers for a Highly Integrated Task. In: *Proceedings of the 9th International Conference on Simulation of Adaptive Behavior*. Springer, Berlin, Germany; 2006. p. 473–484.
48. Lee WP. Evolving complex robot behaviors. *Information Sciences*. 1999; 121(1):1–25. doi: [10.1016/S0020-0255\(99\)00078-X](https://doi.org/10.1016/S0020-0255(99)00078-X)
49. Duarte M, Oliveira SM, Christensen AL. Evolution of Hybrid Robotic Controllers for Complex Tasks. *Journal of Intelligent and Robotic Systems*. 2015; 78(3–4):463–484. doi: [10.1007/s10846-014-0086-x](https://doi.org/10.1007/s10846-014-0086-x)
50. Bongard J, Beliveau P, Hornby G. Avoiding local optima with interactive evolutionary robotics. In: *Proceedings of the 14th Genetic and Evolutionary Computation Conference*. ACM Press, New York, NY; 2012. p. 1405–1406.
51. Celis S, Hornby GS, Bongard J. Avoiding local optima with user demonstrations and low-level control. In: *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE Press, Piscataway, NJ; 2013. p. 3403–3410.
52. Manley JE. Unmanned surface vehicles, 15 years of development. In: *Proceedings of OCEANS 2008*. IEEE Press, Piscataway, NJ; 2008. p. 1–4.
53. Pinto E, Marques F, Mendonça R, Lourenço A, Santana P, Barata J. An Autonomous Surface-Aerial Marsupial Robotic Team for Riverine Environmental Monitoring: Benefiting from Coordinated Aerial, Underwater, and Surface Level Perception. In: *Proceedings of the IEEE International Conference on Robotics and Biomimetics*. IEEE Press, Piscataway, NJ; 2014. p. 443–450.
54. De Cubber G, Doroftei D, Serrano D, Chintamani K, Sabino R, Ourevitch S. The EU-ICARUS project: developing assistive robotic tools for search and rescue operations. In: *IEEE international symposium on Safety, Security, and Rescue Robotics (SSRR)*. IEEE Press, Piscataway, NJ; 2013. p. 1–4.
55. Scerri P, Velagapudi P, Kannan B, Valada A, Tomaszewski C, Dolan J, et al. Real-world testing of a multi-robot team. In: *Proceedings of the International Conference on Autonomous Agents & Multiagent Systems*. IFAAMAS, Richland, SC; 2012. p. 1213–1214.
56. Schmickl T, Thenius R, Moslinger C, Timmis J, Tyrrell A, Read M, et al. CoCoRo—The Self-Aware Underwater Swarm. In: *Proceedings of the 5th IEEE Conference on Self-Adaptive and Self-Organizing Systems Workshops*. IEEE Press, Piscataway, NJ; 2011. p. 120–126.

57. Moore JM, Clark AJ, McKinley PK. Evolution of station keeping as a response to flows in an aquatic robot. In: Proceedings of the 15th Genetic and Evolutionary Computation Conference. ACM Press, New York, NY; 2013. p. 239–246.
58. Praczyk T. Using augmenting modular neural networks to evolve neuro-controllers for a team of underwater vehicles. *Soft Computing*. 2014; 18(12):2445–2460. doi: [10.1007/s00500-014-1221-0](https://doi.org/10.1007/s00500-014-1221-0)
59. Stanley K, Miikkulainen R. Evolving neural networks through augmenting topologies. *Evolutionary Computation*. 2002; 10(2):99–127. doi: [10.1162/106365602320169811](https://doi.org/10.1162/106365602320169811) PMID: [12180173](https://pubmed.ncbi.nlm.nih.gov/12180173/)
60. Duarte M, Silva F, Rodrigues T, Oliveira SM, Christensen AL. JBotEvolver: A Versatile Simulation Platform for Evolutionary Robotics. In: Proceedings of the 14th International Conference on the Synthesis & Simulation of Living Systems. MIT Press, Cambridge, MA; 2014. p. 210–211.
61. Floreano D, Mondada F. Evolution of Homing Navigation in a Real Mobile Robot. *IEEE Transactions on Systems, Man, and Cybernetics*. 1996; 26(3):396–407. doi: [10.1109/3477.499791](https://doi.org/10.1109/3477.499791) PMID: [18263042](https://pubmed.ncbi.nlm.nih.gov/18263042/)
62. Christensen AL, Dorigo M. Evolving an integrated phototaxis and hole-avoidance behavior for a swarm-bot. In: Proceedings of the 10th International Conference on the Simulation and Synthesis of Living Systems. MIT Press, Cambridge, MA; 2006. p. 248–254.
63. Batalin MA, Sukhatme GS. Spreading out: A local approach to multi-robot coverage. In: Proceedings of the 5th International Symposium on Distributed Autonomous Robotic Systems. Springer, Tokyo, Japan; 2002. p. 373–382.
64. Reif JH, Wang H. Social potential fields: A distributed behavioral control for autonomous robots. *Robotics and Autonomous Systems*. 1999; 27(3):171–194. doi: [10.1016/S0921-8890\(99\)00004-4](https://doi.org/10.1016/S0921-8890(99)00004-4)
65. Francesca G, Brambilla M, Brutschy A, Garattoni L, Miletitch R, Podevijn G, et al. An experiment in automatic design of robot swarms. In: Proceedings of the 9th International Conference on Swarm Intelligence. Springer, Berlin, Germany; 2014. p. 25–37.
66. Silva F, Urbano P, Oliveira S, Christensen AL. odNEAT: An Algorithm for Distributed Online, Onboard Evolution of Robot Behaviours. In: Proceedings of the 13th International Conference on the Simulation & Synthesis of Living Systems. MIT Press, Cambridge, MA; 2012. p. 251–258.
67. Groß R, Dorigo M. Towards group transport by swarms of robots. *International Journal of Bio-Inspired Computing*. 2009; 1(1–2):1–13.
68. Wagner I, Lindenbaum M, Bruckstein AM, et al. Distributed covering by ant-robots using evaporating traces. *IEEE Transactions on Robotics and Automation*. 1999; 15(5):918–933. doi: [10.1109/70.795795](https://doi.org/10.1109/70.795795)
69. Marques L, Nunes U, de Almeida AT. Particle swarm-based olfactory guided search. *Autonomous Robots*. 2006; 20(3):277–287. doi: [10.1007/s10514-006-7567-0](https://doi.org/10.1007/s10514-006-7567-0)
70. Haasdijk E, Atta-ul Qayyum A, Eiben AE. Racing to improve on-line, on-board evolutionary robotics. In: Proceedings of the 13th Genetic and Evolutionary Computation Conference. ACM Press, New York, NY; 2011. p. 187–194.
71. Stein ML. *Interpolation of Spatial Data: Some Theory for Kriging*. Springer-Verlag, New York, NY; 1999.
72. Nelson AL, Barlow GJ, Doitsidis L. Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*. 2009; 57(4):345–370. doi: [10.1016/j.robot.2008.09.009](https://doi.org/10.1016/j.robot.2008.09.009)
73. Gomez F, Miikkulainen R. Incremental evolution of complex general behavior. *Adaptive Behavior*. 1997; 5(3–4):317–342. doi: [10.1177/105971239700500305](https://doi.org/10.1177/105971239700500305)
74. Silva F, Correia L, Christensen AL. R-HybrID: Evolution of Agent Controllers with a Hybrisation of Indirect and Direct Encodings. In: Proceedings of the International Conference on Autonomous Agents & Multiagent Systems. IFAAMAS, Richland, SC; 2015. p. 735–744.
75. Duarte M, Oliveira SM, Christensen AL. Evolution of Hierarchical Controllers for Multirobot Systems. In: Proceedings of the 14th International Conference on the Synthesis & Simulation of Living Systems. MIT Press, Cambridge, MA; 2014. p. 657–664.
76. Duarte M, Gomes J, Costa V, Oliveira SM, Christensen AL. Hybrid Control for a Real Swarm Robotic System in an Intruder Detection Task. In: Proceedings of the 18th European Conference on the Applications of Evolutionary Computation (EvoStar). Springer, Berlin, Germany; 2016. In press.
77. Cully A, Clune J, Tarapore D, Mouret JB. Robots that can adapt like animals. *Nature*. 2015; 521(7553):503–507. doi: [10.1038/nature14422](https://doi.org/10.1038/nature14422) PMID: [26017452](https://pubmed.ncbi.nlm.nih.gov/26017452/)
78. Christensen AL, O’Grady R, Birattari M, Dorigo M. Fault detection in autonomous robots based on fault injection and learning. *Autonomous Robots*. 2008; 24(1):49–67. doi: [10.1007/s10514-007-9060-9](https://doi.org/10.1007/s10514-007-9060-9)
79. Tarapore D, Lima PU, Carneiro J, Christensen AL. To err is robotic, to tolerate immunological: fault detection in multirobot systems. *Bioinspiration & Biomimetics*. 2015; 10(1):016014. doi: [10.1088/1748-3190/10/1/016014](https://doi.org/10.1088/1748-3190/10/1/016014)