

Engineering the Evolution of Self-Organising Behaviours in Swarm Robotics: A Case Study

Vito TRIANNI[†]

`vtrianni@ulb.ac.be`

Stefano NOLFI[‡]

`stefano.nolfi@istc.cnr.it`

[†] IRIDIA-CoDE, ULB, Avenue F. Roosevelt, n. 50, CP 194/6, 1050 Brussels, Belgium

[‡] ISTC-CNR, via San Martino della Battaglia 44, 00185 Rome, Italy

Corresponding author: Vito Trianni[†], phone: +3226502712, fax: +3226502715

Abstract

Evolutionary Robotics (ER) is a powerful approach for the automatic synthesis of robot controllers, as it requires little a priori knowledge about the problem to be solved in order to obtain good solutions. This is particularly true for collective and swarm robotics, in which the desired behaviour of the group is an indirect result of the control and communication rules followed by each individual. However, the experimenter must make several arbitrary choices in setting up the evolutionary process, in order to define the correct selective pressures that can lead to the desired results. In some cases, only a deep understanding of the obtained results can point to the critical aspects that constrain the system, which can be later modified in order to re-engineer the evolutionary process towards better solutions. In this paper, we discuss the problem of engineering the evolutionary machinery that can lead to the desired result in the swarm robotics context. We also present a case study about self-organising synchronisation in a swarm of robots, in which some arbitrarily chosen properties of the communication system hinder the scalability of the behaviour to large groups. We show that by modifying the communication system, artificial evolution can synthesise behaviours that properly scale with the group size.

Keywords: Evolutionary Robotics, Swarm Robotics, Self-Organisation, Engineering Emergence

1 Introduction

The synthesis of controllers for autonomous robots is a complex problem that has been faced with a large number of different techniques [32]. Among the various possibilities, Evolutionary Robotics (ER) represents a viable approach for the automatic synthesis of robot controllers requiring little a priori knowledge about the solution of a given problem [26]. In the swarm robotics context, ER can be very useful to synthesise efficient self-organising behaviours and to obtain systems with desired emergent properties, such as robustness to individual failures, flexibility and adaptivity to environmental changes, or scalability to different group size. In fact, ER operates by introducing variations at the level of the fine-grained characteristics of the robots, and by retaining or discarding these variations on the basis of the effect they have on the global behaviour exhibited by the swarm. This bottom-up approach contrasts with classic top-down approaches (e.g., divide-and-conquer), which require an a priori arbitrary decomposition of the group behaviour into individual behaviours and interaction rules. Such decomposition is difficult to be performed, due to the indirect relationships between the rules executed by the robots—which determine the individual reaction to perceived sensory states—and the overall group behaviour. ER solves this *design problem* by synthesising self-organisation from the bottom up, therefore relieving the experimenter from the need to guess both the appropriate behaviour that each individual robot should produce, and the control mechanisms that can lead to the exhibition of such behaviour [39, 36]. More importantly, in some cases the designer might be unable to identify the behaviour that should be exhibited by each individual robot that, in interaction with the other robots, can lead to the desired swarm behaviour [12].

On the other hand, ER does not completely exempt the experimenter from the design problem since she/he still has to face the problem of designing the setup of the evolutionary experiment, together with those characteristics of the robot and of the environment that are predetermined and fixed. In this case, the design effort concerns the preconditions that make the evolution of the self-organising process possible, along with its emergent properties. Although designing an experimental setup that enables the evolution of effective solutions should be in principle simpler than designing the solutions themselves, the lack of systematic methodologies that can guide such an experimental design may prevent an effective exploitation of the potential advantages of the ER approach, or might condition such exploitation to the intuition of the experimenter. We believe that there is a strong necessity of formalising an engineering approach for the evolution of robot behaviours, above all within a relatively novel and complex domain like swarm robotics.

In this paper, we make a step in this direction by discussing which are the relevant choices that must be made when setting up an evolutionary experiment in a swarm robotics context. In particular, we focus on the robot sensory-motor system, the genotype-to-phenotype mapping, the fitness function and the ecology in which the evolutionary process is carried out. Moreover, we suggest that, like in any engi-

neering exercise, an iterative approach can guide the design process through incremental improvements. In this iterative process, a fundamental role is played by testing and evaluation of the obtained results, an activity that can convey useful information about the critical aspects of the system that must be improved. In this respect, we discuss a case study in which self-organising synchronisation behaviours are evolved in a robotic system. The goal of the experiment is to understand which are the minimal behavioural and communication strategies that allow a group of robots to synchronise their individual periodic behaviour [38]. In particular, we are interested in the scalability property of the evolved behaviours to large groups. By analysing the initial results, we discovered that the arbitrary choice made in the communication system was hindering the evolved behaviours to suitably scale to large groups. This finding allowed us to re-engineer the characteristics of the robots by identifying a new communication channel, and to run further evolutionary experiments that resulted in properly scalable behaviours.

Engineering emergence in ER can be broadly linked to recent studies on guided self-organisation [28, 29]. In these studies, various techniques are proposed to obtain desired behaviours through guidance rather than control. In fact, self-organising systems are complex and strongly non-linear, and attempts to obtain desired effects through direct control are often ill-suited [14]. It is therefore more convenient to steer self-organisation toward desirable outcomes, exploiting the dynamics of the self-organising process itself. Therefore, on the one hand self-organisation is promoted in a task-independent fashion, and on the other hand guidance is given to the self-organising process in order to steer it towards desired patterns. Recent studies investigated different approaches to guide self-organisation. Information-driven evolution is an ER technique that allows to produce robot controllers exploiting a task-independent fitness function, usually based on information theoretic measures, such as Shannon’s entropy or mutual information [30, 16, 15, 33]. The task independent metric provides intrinsic selective pressures that favour sensory-motor coordination or correlations among different agents’ behaviours. However, some constraint is necessary to channel evolution towards the desired behaviour. Constraints can be given in the form of explicit reward functions to be merged with the task-independent ones, or can be implicit in the robot configuration or in the robot ecology. A similar approach, but on a ontogenetic time-scale, is homeokinetic learning [17]. Here, the spontaneous generation of behavioural patterns through homeokinesis is guided by learning through reinforcement signals [18]. Homeokinesis refers to the ability to produce and maintain a definite kinetic regime. In the robotic domain, this is achieved through learning an internal representation of the current behaviour (a self-model) and to adapt the behaviour in order to minimise the difference with the learned model [8]. The interactions between model and behaviour produce a spontaneous discovery and maintenance of behavioural patterns [7]. However, reinforcement signals are necessary to channel the spontaneous generation of behaviours toward desired outcomes [17].

In the context of evolutionary swarm robotics, self-organisation need to be guided at both the phylogenetic and ontogenetic time-scale. That is, we need to specify the correct selective pressures in order

to observe the evolution of the desired group behaviours. Moreover, we need to setup the robotic system in order to obtain desired emergent properties of the evolved self-organising behaviours. In Section 2 we discuss which are the relevant choices that need to be taken in setting up an evolutionary experiment within a swarm robotic context. We consider four main components: the robot sensory-motor system in Section 2.1, the genotype-to-phenotype mapping in Section 2.2, the fitness function in Section 2.3 and the ecology in Section 2.4. In Section 3, we discuss our iterative approach in engineering and re-engineering the evolution of self-organising behaviours. An instance of this process is presented in the case-study of Section 4, which is discussed in detail in order to show how the proposed methodology can be applied. Section 5 concludes the paper.

2 Engineering emergence in evolutionary swarm robotics

As mentioned above, the design of an evolutionary experiment in a swarm robotics context requires to make a number of choices. First of all, the experimenter should determine which are the characteristics of the robots that are predetermined and fixed (e.g., the available sensors and actuators and/or the robot body structure), and which are the characteristics subjected to the evolutionary process. Secondly, the experimenter should determine the way in which these evolved features are represented in the artificial genotype—in other words, how the genotype maps to the phenotype, that is, to the robotic system. Thirdly, the experimenter should devise the performance metric or fitness function to evaluate the evolving genotypes, which determines the selective pressures that allow to progressively evolve the desired solution. Finally, the experimenter should determine the ecology (i.e., the characteristics of the environment in which the robots are evaluated), which in turns introduces ecological selective pressures to which evolving robots are exposed.

Guidance to the evolution of self-organisation can be obtained through suitable choices at the moment of the design of the experiment. In this section, we discuss the above issues under an engineering perspective. That is, we consider that the robotic hardware and the task to be performed by the robots are specified in advance, and the focus of the experiment is the synthesis of a controller for the robotic group.¹ We limit our discussion to the swarm robotics domain, and we refer the reader to the ER literature for the topics not covered here [19, 26, 13, 9]. A schematic view of the relevant design choices for the evolution of self-organising behaviours that we consider in this work is given in Table 1.

Place Table 1 about here.

¹A thorough discussion of all possible instances of evolutionary robotic experiments is out of the scope of this paper. Nevertheless, this discussion can be easily generalised to issues that are not covered in these pages.

2.1 Sensory-motor system

The body and the sensory-motor system of the robots crucially constrain the way in which the robots can interact with the physical and social environment, and set the preconditions for the evolution of the desired global behaviour and for the emergence of group-level properties. When the objective of an experiment is, as in the case considered here, to enable a group of robots to perform a certain function or to solve a certain task, the characteristics of the body and of the overall sensory-motor system are predetermined and fixed. The experimenter, however, should determine a subset of the available sensors and actuators that will be used. Moreover, it must be defined how the chosen sensors and actuators are interfaced with the robot control system. The selection of the appropriate sensors and actuators is usually straightforward. However, it can be desirable and sometimes necessary to process the raw sensor data to obtain a more compact or better usable information. This pre-processing can be as simple as a linear scaling of the sensor reading, but can even be a complex function of many sensory inputs. For instance, a coloured camera provides a very rich information, and some feature extraction algorithm is necessary in order to recognise useful patterns within the image. The choices of the quantity and quality of the information extracted by the camera, and in general by a pre-processing of raw data, can be of fundamental importance for the evolvability of the system, and should be carefully taken into account. A similar discussion holds for communication. Robots may be provided with different communication devices, which enable to chose the type of messages exchanged (e.g., implicit communication through infrared sensors detecting the other robots body [31] or torque sensors affected by the movement of other robots [3], sound communications through microphones and speakers [37], or light communication through coloured LEDs and cameras [6, 22, 23]). The definition of the communication channel might significantly affect the evolutionary process, as it has a strong influence on the ability of the robots to interact with each other, and therefore to self-organise. In Section 4, we present a case study in which an ill-suited communication system severely limits the scalability of the evolved behaviours.

Overall, the important remark is that the selection of the appropriate sensory-motor system can be extremely relevant, as it sets the preconditions for the evolution of the desired group behaviour. Not only this choice defines the actual capabilities of the individual robots, but it also states which are the interaction abilities among the robots, that is, through which channels information is exchanged among the robots, and which is the capacity of these information channels.

2.2 Genotype-to-phenotype mapping

In evolutionary computing methods (including ER), a genotype is usually a string of bits or real numbers that encodes a potential solution to a given problem. In ER, the genotype specifies the characteristics of a robotic system that should be able to display a desired behaviour. The experimenter therefore has to specify the genotype-to-phenotype mapping, i.e., the rules or the processes that determine the relation

between the genotype—the string of bits or numbers—and the phenotype—the robotic system [19, 9].

A common approach in ER consists in a direct mapping between genotype and controller parameters, in which each free parameter of the controller is encoded into a corresponding part (gene) of the genotype. This implies that artificial evolution operates on the parameters that regulate the fine-grained interactions between the robot and the environment, which in turns determine the behaviour exhibited by the robot, within the possibilities and the limits imposed by the hardware characteristics of the robots and by the architecture of the robot controllers that are designed by the experimenter. A widely used approach in the literature consists in encoding into the genotype a fixed number of parameters of a neural network, while the neural network size, topology and connectivity remain fixed. The use of neural controllers allows the evolutionary process to operate on the network parameters that regulate the fine-grained interactions between the robot and the environment, which in turns determine the behaviour exhibited by the robot. In this way, the evolutionary process shapes the behaviour exhibited by the robot in detail. The use of a fixed architecture and of a direct encoding has the advantage of being simple, but requires to choose the architecture of the controller beforehand. This again might turn out to be critical since the characteristics of the architecture (i.e., the number of neurons and/or the presence of recurrent connections) might strongly influence the obtained result. Moreover, it is difficult for the experimenter to predict the type of behaviours that the robot should exhibit and consequently the characteristics that the architecture should have in order to enable the production of such behaviours. A different approach consists in the indirect encoding, for which the genotype “develops” in a corresponding phenotype according to some developmental rules that are encoded in the genotype. This approach is followed to mimic the embryogenesis, and offers the possibility to specify complex control structures with compact genetic encodings (therefore reducing the search space), and to avoid fixing the controller structure beforehand. However, these advantages come at some cost: it is necessary to define a suitable set of developmental rules, and to ensure that such rules have the necessary expressive power.

In collective and swarm robotics, whether a direct or indirect encoding is used, it is anyway necessary to specify the characteristics of the whole robotic group. The relevant choice here concerns the genetic relatedness between the individuals forming the group, that is, whether they are *genetically homogeneous* (i.e., they are clones) or *heterogeneous* (i.e., they differ from each other). In a homogeneous group, the genotype usually encodes the parameters of a single controller, which is duplicated several times and embodied in all the robots taking part to the experiment (i.e., all individuals forming the swarm are generated from the same genotype). This method has the advantage of being simple and compact with respect to the number of free parameters, since the parameters of all individuals are encoded in the same genes. Moreover, it eliminates the problems related to the identification of the individual contributions to the performance produced by the entire swarm, and removes the conflict of interest that might arise between genetically different individuals, as discussed in the following section.

In heterogeneous group, the swarms is constituted by individuals generated by different genotypes or by different parts of the same genotype. The use of heterogeneous groups might be advantageous when the robots forming the group should play well differentiated roles. In this case, however, the different roles must be somehow encoded into the genotype. The simplest approach consists in a priori defining how many roles are necessary (at most, one role per robot in the group), and encoding in a single genotype all the parameters of all controllers. Alternatively, the robots might be genetically identical, but might express different parts of their genome [4]. These approaches allow to evolve tightly cooperating teams, at the cost of substantially increasing the search space for the evolutionary algorithm. In order to reduce the search space, heterogeneous teams can be obtained from controllers evolved in different populations, which are updated in parallel. Each population is therefore dedicated to a specific role, and teams are formed by drawing from the different population with a certain strategy. Eventually, the best individual of each population is the representative of the corresponding role. A similar approach can be instantiated with a single population of genotypes: here, different roles are drawn from the same population. However, in this case, a strong convergence of the population would result in rather homogeneous teams. It would be required to use some technique to maintain enough diversity in the population, which would result in niches adapted to the required roles. In both cases, however, it is challenging to identify an effective way to assign the fitness to the different genotypes forming a team, as discussed below.

2.3 Behavioural selective pressures: the fitness function

The definition of the performance metric that rewards the desired behaviour is usually task-dependent (i.e., a function that estimates the extent to which the swarm solves the given task).² There are multiple ways to define a fitness function for a given problem [25]. To discriminate between different types, Floreano and Urzelai proposed the usage of a three-dimensional *fitness space*, in which the different dimensions refer to important features of a fitness function [11]:

functional vs. behavioural: a functional fitness rewards a particular working modality (i.e., gives an indication on the actuators outputs), while a behavioural fitness measures the quality of the behaviour (i.e., gives an indication about the outcome of a sequence of actions);

external vs. internal: an external fitness is computed through variables that are available to an external observer (i.e., the absolute position of the robot in the environment), while an internal fitness is computed through variables available to the robot (i.e., the sensor readings). While external fitness functions can be easier to deploy, internal ones may reduce possible biases introduced by the designer;

²For alternative approaches based on task-independent fitness function or combination of the two, see [30, 16, 15, 33], already mentioned in Section 1.

explicit vs. implicit: an explicit fitness measures the way in which a goal is achieved (i.e., the trajectory performed to get close to a light source), and therefore puts constraints on the displayed behaviours. An implicit fitness function, instead, measures the level of attainment of a goal (i.e., how close to the light source the robot ends). An implicit fitness gives more freedom to explore the solution space, therefore allowing to find solutions that are not a priori envisioned by the experimenter.

In swarm robotics, the indirect relationship between individual actions and group organisation makes it difficult to devise functional measures. A functional measure, in fact, is directly related to the causes of the observed behaviour, which are a priori unknown to the experimenter. Similarly, internal fitness functions may be more difficult to devise, given that they require the evaluation of the group behaviour from the perspective of the individual robots. However, a common approach is to devise an internal fitness function that is measured on each robot taking part to the experiment, obtaining individual fitness values that are aggregated either averaging over the group or selecting the best or the worst performing robot. In this way, it is possible to obtain a group measure starting from internal variables. This can be done only if the individual measure is directly related to the global organisation. Finally, implicit measures should be preferred when the relationship between the individual control rules and the group behaviour is indirect or unknown, as they pose less constraints on the way the desired collective behaviour is achieved.

In evolutionary swarm robotics, it is often the case that a homogeneous group of robots should present a self-organising behaviour. In this case, it is useful to evaluate the group level properties through external fitness functions, rather than looking at the individual actions. It is also useful to evaluate the group organisation (i.e., the spatio-temporal pattern), rather than the way in which the organisation is achieved, resorting to implicit metrics. In fact, external and implicit fitness functions pose less constraints on the way in which the problem should be solved. On the contrary, if we consider the case in which a heterogeneous group should display a teamwork with highly specific roles, internal and explicit metrics could be preferred, as they may allow to develop the implementation details of specific solutions beforehand identified.

The use of heterogeneous swarms constituted by individuals generated from different genotypes pose additional constraint on the definition of the fitness function since it requires to estimate the contribution of genetically different individuals to the overall ability of the group. Without going too deep into the details, it is important here to notice that this problem does not affect homogeneous groups and heterogeneous teams generated from a single genotype. In these cases, in fact, the group performance can be directly assigned to the single genotype that generated the entire swarm. Whenever the group members correspond to different genotypes, it is necessary to deploy a fitness function that directly measures the individual contribution. When this is not possible, the fitness of a single genotype must be evaluated by forming multiple groups, choosing the teammates randomly or with a specific strategy in order to have a good estimate of the individual contribution to the group performance. The latter, however, is a complex

and time-consuming procedure, which adds further uncertainty over the estimation of the genotype fitness in varying environmental conditions, as discussed below. It is also worth mentioning that the use of heterogeneous groups constituted by genetically different individuals and the use of fitness functions that estimate the performance at the level of the single individuals tend to cause conflicts of interest between the individuals forming the group, which might prevent the evolution of stable coordinated/cooperative behaviour [10, 20, 40].

2.4 Ecology

The ability of a swarm to perform a certain task does not depend only on the characteristics of the robot, but is also influenced by the characteristics of the environment (e.g., variability in space and time) and by the relation between the robots and the environment (e.g., the set of possible initial positions and orientations of the robots within the environment). The behaviour of the swarm should be robust with respect to possible variations of the environment and of other parameters that contribute to define the *ecological niche* in which the behaviour is evolved. Indeed, the definition of the ecological niche of the robotic system can significantly affect the results of the evolutionary process [27]. A precise computation of the fitness would require testing the behaviour systematically for every possible environmental condition in which the robot may find itself. This is normally not feasible, and therefore it is necessary to sample the space of the possible ecological conditions in an appropriate way, in order to obtain a reasonable fitness estimate. In a collective robotics setup, the problem is worsened by the presence of multiple robots, which increase the variability of the ecological niche. Interaction among individuals, physical interferences and collisions among robots may be very relevant to the accomplishment of the task, requiring the definition of experimental conditions that can let the group experience the interaction patterns relevant for obtaining a robust behaviour.

It is important to notice that indirect selective pressures may be created through the definition of the ecological niche and through the sampling employed to estimate the fitness. Given that the group is evaluated for presenting a robust behaviour within the parameter space of the ecological niche, the choice of the sampling may influence the evolutionary path. For instance, in [2] communication and cooperation emerge solely due to ecological selective pressures, as the fitness function does not contain any indication about cooperative strategies. Thus, the ecological niche and the sampling of the parameter space must be appropriately defined in order to account for robust group behaviour and to take into account implicit selective pressures.

A final issue to consider concerns symmetry breaking, which pertain many collective phenomena. Symmetry breaking refers to the situation in which a system passes from a disordered condition (which is symmetric in the sense that small changes do not change the overall appearance) to a more ordered one, characterised by some structure or pattern. For instance, a group of robots may pass from a disordered

(symmetric) condition in which all robots are randomly oriented to a ordered one in which all robots have the same orientation. Symmetric conditions in a collective robotic system must be carefully identified: symmetry breaking may not be possible exploiting the inherent randomness of the robotic system, and therefore suitable behavioural strategies may be required. The evolutionary machinery needs to encounter such conditions often in order to synthesise the collective behaviour necessary to break the symmetry. For this reason, it is necessary to force the system into symmetric conditions, as well as into asymmetric ones, to evolve robust behaviours. An example of systematic testing in symmetric and asymmetric conditions for a two-robot system is given in [1].

3 The engineering design cycle in evolutionary robotics

An attentive setup of the evolutionary experiments does not guarantee the achievement of the desired results. In fact, some choices, although reasonable without any a priori knowledge of the evolutionary dynamics, may turn out to be too constraining, or may prevent the system to display certain properties. However, negative results should be exploited to acquire information on the system dynamics and re-engineer hand-designed characteristics accordingly. In fact, by understanding the properties of unsuccessful results it may be possible to recognise which are the critical aspects that constrain the system in sub-optimal solutions. As a consequence, it is possible to proceed with an informed re-engineering of the evolutionary experiments.

This iterative process is not any different from any engineering design cycle. We can highlight three main phases: (i) *setup* of the evolutionary experiment, (ii) the actual *evolution* and (iii) the *analysis* of the results. In the setup phase, all the relevant choices are taken to devise the evolutionary experiment, as discussed above. The setup phase is fundamental within the design cycle. Here, the expertise of the experimenter is of utmost importance to orient the evolutionary design in the right direction. The main goal of this phase is recognising which are the free parameters of the system that must be specified beforehand, and which are the ones that will be subjected to evolution (i.e., through variation and selective reproduction). In doing so, there are anyway many choices that are performed somehow arbitrarily, especially when there is no a priori knowledge available about the system and evolutionary dynamics.

In the evolution phase, a sufficient number of instances (or replications) of the evolutionary experiment are performed, each starting with a different random population of genotypes. On the one hand, multiple instances increase the chances of finding an optimal solution to the given problem. On the other hand, the quality of the design choices can be verified by looking at the percentage of instances that produced good results. In the best case, all instances of the experiment produce acceptable results. Instead, a small percentage of successful instances might indicate that there is room for improvement (i.e., the setup of the evolutionary experiments can be modified in order to increase the quality of the evolved solutions).

Also, a sub-optimal solution may reveal that some features of the robotic system defined in the setup phase hinder the production of the desired behaviour.

Many ER practitioners just cycle between the setup and evolution phases, exploiting intuition and experience in adapting the evolutionary setup to obtain better results. Less frequently, an attentive analysis of the negative results obtained in previous iterations is performed to identify the critical aspects of the evolutionary design. We believe that the analysis phase should be always conducted, and suitable methodological tools developed, in order to support the re-engineering of the evolutionary experiment toward promising directions. Two possible analyses should be conducted to guide the evolutionary process towards the desired results. First, the analysis of the selective pressures and of the evolutionary dynamics can inform the experimenter about the limits of the system in producing better and better solution starting from scratch. Second, the evolved sub-optimal solutions should be analysed to acquire knowledge about the characteristics of the robotic system and of the produced behaviour, in order to understand whether or not it can display the desired properties. In this case, generalisation tests should be performed to verify the ability of the system to work in conditions not experienced during the evolutionary phase. In fact, during evolution is not always possible to test all the conditions that can be faced by the robotic system. Therefore, it is important to test the robustness to failures, flexibility to varying environmental conditions, or the scalability to larger groups. In the following section, similar analyses are performed to understand the properties of the evolved system with respect to the scalability to larger groups. In this case study, the design cycle is closed by re-engineering the communication channel in order to obtain better scalability and better performance, which result also in more robust evolutionary dynamics.

4 Case study: self-organised synchronisation

Self-organised synchronisation is a common phenomenon observed in many natural and artificial systems: simple coupling rules at the level of the individual components of the system result in an overall coherent behaviour [34]. A well known synchronisation phenomenon is the flashing behaviour of some firefly species in South-East Asia, which aggregate at dusk and engage in massively synchronous displays [5]. Models of this behaviour describe fireflies as a population of pulse-coupled oscillators with equal or very similar frequencies. These oscillators can influence each other by emitting a pulse that shifts or resets their oscillation phase. The numerous interactions among the individual oscillator-fireflies are sufficient to explain the synchronisation of the whole population (for more details, see [5, 21, 35]). This model has been often exploited to engineer systems capable of synchronous behaviour, also in collective and swarm robotics [41, 6]. In this study, we have investigated which are the minimal behavioural and Communicative conditions that can lead to synchronisation in a group of robots, in which each individual presents a periodic behaviour. For this purpose, we chose to provide robots with simple reactive controllers and

basic communication abilities. The period and the phase of the individual behaviour are defined by the sensory-motor coordination of the robot, that is, by the dynamical interactions with the environment that result from the robot embodiment. We show that such dynamical interactions can be exploited for self-organised synchronisation, allowing to keep a minimal complexity of both the behavioural and the communication level (for more details, see [38]).

4.1 Phase 1: Evolutionary Setup

We make use of a simple evolutionary algorithm that works on a population of 100 binary genotypes, which are randomly generated. Each genotype is a string of bits that encode some parameters of the robotic system. Once the fitness of each genotype in the population has been evaluated, a new population is produced by a combination of selection with elitism and mutation. Recombination is not used. At each generation, the four best individuals of the population—i.e., the *elite*—are retained in the subsequent generation. The remainder of the new population is generated by the 20 individuals of the previous generation that scored the highest fitness. Each selected genotype reproduces at most 5 times by applying mutation with 3% probability of flipping a bit. The evolutionary process runs for 500 generations.

The evolutionary experiments are performed in simulation, using a simple kinematic model of the *s-bot* robot (see Fig. 1a,b and refer to [24] for details), and the results are afterwards validated on the physical platform. The experimental scenario for the evolution of self-organising synchronisation requires that each robot in the group displays a simple periodic behaviour, which should be entrained with the periodic behaviour of the other robots present in the arena. The individual periodic behaviour consists in oscillations along the y direction of a rectangular arena (see Figure 1c). Oscillations are possible through the exploitation of a symmetric gradient in shades of grey painted on the ground. The gradient presents a white stripe for $|y| < 0.2 m$, and black stripe for $|y| > 1 m$.

Place Figure 1 about here.

4.1.1 Sensory-motor system

For the purpose of engineering the evolutionary system, both the characteristics of the arena and the capabilities of the robots give several constraints to the experimental setup. According to these constraints, we select among the various possibilities the minimal set of sensors and actuators that are required to accomplish the task, that is, individual periodic oscillations over the grey gradient and synchronisation of the oscillation phase. Certainly, the controller needs access to the wheels' motors, and we set $\omega_M \approx 4.5 s^{-1}$ as the maximum angular speed of the wheels. The grey gradient of the arena can be perceived by the robots through four infrared sensors placed under their chassis (ground sensors), which are appropriately scaled to encode the grey-level in the range $[0, 1]$, where 0 corresponds to black and 1 to

white. The perception of the gradient through these sensors provides the robot with enough information to perform oscillations along the y axis. Additionally, robots need to use the infrared proximity sensors placed around their cylindrical body, in order to avoid collisions with walls or with other robots. These choices, which are mainly constrained by the arena setup and by the features of the physical robot, are sufficient for the individual behaviour.

For what concerns the group behaviour, instead, we need to provide the robots with suitable interaction modalities that can lead to synchronisation of their movements. The choice of the communication system is the aspect we focus on in this paper. In fact, the *s-bot* platform features various communication devices, and we need to select among them the one that best fits our experimental scenario. Robots are provided with speakers and microphones for sound communication. Moreover, robots can exploit coloured LEDs positioned around their turret to display a colour pattern that can be perceived through the omni-directional camera. Finally, robots have wireless communication abilities. Therefore, there is a large freedom in choosing the communication system. In order to maintain a minimal configuration, we decided to provide the robots with a *global* and *binary* communication system:

$$s(t) = \max_r S_r(t), \quad (1)$$

where $S_r(t) \in \{0, 1\}$ is the binary signal emitted by robot r at time t , and $s(t) \in \{0, 1\}$ is the binary signal perceived by all robots. In other words, each robot r can produce a signal $S_r(t)$. Signals produced by different robots cannot be distinguished, and result in a single signal $s(t)$ perceived by every robot in the arena, including the signalling one. Signals are perceived in a binary way: either there is someone signalling in the arena, or there is no one. This communication channel is probably the poorest one in terms of the amount of information that can be conveyed. However, this is sufficient for our purposes, as we will see in the following. Note that this communication channel can be easily implemented with sound signals: a robot can emit a single frequency tone with an intensity high enough to be perceived everywhere in the arena. Additionally, it is worth mentioning that, differently from the other sensors and actuators, the choice of the communication system is not constrained by the robotic hardware or by other aspects of the experimental setup, but is only dictated by the communication system we have chosen.

4.1.2 Genotype-to-phenotype mapping

As mentioned below, in self-organising synchronisation all individuals are (nearly) identical and display the same behaviour. Synchronisation emerges from the interactions among individuals which adapt their behaviour in order to entrain with each other. This feature constrains the genotype-to-phenotype mapping to produce a homogeneous group of robots. To do so, the genotype is mapped into a single control structure, which is cloned and downloaded onto all the robots taking part in the experiment.

The control structure chosen is a fully connected, feed forward neural network—a perceptron network. The choice of this type of neural network stems from our working hypothesis, that is, searching the minimal complexity of the control and communication system that can support synchronisation behaviours. For this purpose, we choose a controller that directly transforms the sensory input into the motor output, without recurrent connections or internal states. The neural network has 11 sensory neurons directly connected to 3 motor neurons. The sensory neurons are simple relay units and the output neurons are sigmoid units whose activation is computed as follows:

$$O_j = \sigma \left(\sum_i w_{ij} I_i + \beta_j \right), \quad \sigma(z) = \frac{1}{1 + e^{-z}}, \quad (2)$$

where I_i is the activation of the i^{th} input unit, β_j is the bias term, O_j is the activation of the j^{th} output unit, w_{ij} is the weight of the connection between the input neuron i and the output neuron j , and $\sigma(z)$ is the sigmoid function. Six sensory neurons— I_1 to I_6 —receive input from a subset of the infrared proximity sensors evenly distributed around the *s-bot*'s turret. Four sensory neurons— I_7 to I_{10} —are dedicated to the readings of the four ground sensors. The state of all infrared and ground sensors is linearly scaled to the range $[0.0, 1.0]$. A simulated uniform noise within 5% of the input range is also added. The last sensory neuron I_{11} receives a binary input corresponding to the perception of sound signals. The activation states of the first two motor neurons— O_1 and O_2 —is scaled onto the range $[-\omega_M, +\omega_M]$, where ω_M is the maximum angular speed of the wheels ($\omega_M \approx 4.5 \text{ s}^{-1}$). The third motor neuron controls the speaker in such a way that a sound signal is emitted whenever the activation state O_3 is greater than 0.5. The connection weights w_{ij} and bias terms β_j are the genetically encoded parameters. Each parameter is represented with a 8-bit binary code mapped onto a real number ranging in $[-10, +10]$. In other words, we have a direct encoding of the genotype into the phenotype, as there is a bijective function that relates the genotype to the phenotype.

4.1.3 The fitness function

The performance of a genotype is evaluated by a 2-components function: $F = 0.5 \cdot F_{\mathcal{M}} + 0.5 \cdot F_{\mathcal{S}} \in [0, 1]$. The movement component $F_{\mathcal{M}}$ simply rewards robots that move along the y direction within the arena at maximum speed. With respect to the taxonomy introduced in Section 2.3, this component is *behavioural*, *external* and *implicit*. In fact, it rewards the movements of the robot from the observer perspective, without explicitly indicating how to perform a periodic behaviour: the oscillatory behaviour derives from the fact that the arena is surrounded by walls, so that oscillations during the whole trial are necessary to maximise $F_{\mathcal{M}}$. The second fitness component $F_{\mathcal{S}}$ rewards synchrony among the robots as the cross-correlation coefficient between the distance of the robots from the x axis. Also this component is *behavioural*, *external* and *implicit*: it is related to the group behaviour, and measures a quantity—the

cross-correlation—that is available only to the observer. In addition to the fitness computation described above, two indirect selective pressures are present. First of all, a trial is stopped when a robot moves over the black-painted area, and we assign to the trial a performance $F = 0$. In this way, robots are rewarded to exploit the information coming from the ground sensors to perform the individual oscillatory movements. Secondly, a trial is stopped when a robot collides with the walls or with another robot, and also in this case we set $F = 0$. In this way, robots are evolved to efficiently avoid collisions. For more details on the fitness computation, refer to [38].

4.1.4 The ecology

The ecology in which the group behaviour evolves need also to be defined. The arena is a rectangle of 6×3 m sides, completely surrounded by walls. The ground is painted in white for $|y| < 0.2$ m, and linearly changes to black until $|y| = 1$ m. For larger distances, the arena is painted in black. We are interested in behavioural and communication strategies that scale well with the group size. However, it is not possible to test every possible group size, and we therefore decided to fix the number of robots in the group to just three robots. This low number of robots allows to obtain fast simulations, and still support the evolution of group coordination. In order to have a good estimate of the group performance, the genotype fitness is the average of the group fitness computed over 10 different trials. In each trial, we vary the initial positions and orientations of the three robots by choosing them uniformly random within the arena. This allows to let the robots experience many different initial conditions, and should result in robust and efficient synchronisation behaviours.

4.1.5 Design and evolution

Before presenting the obtained results, it is useful to discuss which are the features that are fixed by the experimenter, and those that are adaptively set by the evolutionary process. We have defined an experimental scenario that is intrinsically cooperative, because robots are homogeneous and are explicitly rewarded to display a desired group behaviour. We have also fixed the sensory-motor configuration and the controller architecture. In particular, we have fixed the interaction modality between different robots, which mainly happens through the binary and global communication signal. Notwithstanding this, the motor and communicative behaviour is not at all pre-determined, but it is the result of the evolutionary process. The individual behaviour and the synchronisation mechanisms are completely determined by the parameters of the neural controller (i.e., connection weights and biases). Individual behaviour and communication signals co-evolve and mutually influence: the individual behaviour determines how the robot moves and experience the environment, which influences the signals emitted. In turns, perceived signals change the way in which the robot reacts to the environment. During evolution, the group behaviour is shaped in order to maximise the user-defined utility metric, within the constraints imposed

by the pre-determined features. In the following, we will see how the communication channel we have chosen influences the obtained results.

4.2 Phase 2: Evolution

In order to test the robustness of the evolutionary setup we devised, we decided to run 20 different instances of the experiment. Each instance was initialised with a different randomly generated population of genotypes, and lasted 500 generations. At the end of the evolutionary process, to assess the quality of the evolved behaviours, we select a single genotype per instance to be chosen among the best individuals of the final generation. To do so, we evaluate the performance of the 20 best individuals of the final generation in 500 different trials, and we choose the individual with highest average fitness. In the remainder of the paper, we refer to the best controllers evolved in replication i as $c_i, i = 1, \dots, 20$. The performance of these controllers over the 500 post-evaluation trials, sorted according to decreasing median values, is shown in Figure 2. The obtained results show that in most replications the performance obtained is in average within the interval $[0.7, 0.9]$, which indicates that robots are able to maximise both the movement fitness component, F_M , and the synchronisation component, F_S .

In order to assess the difference in performance among the controllers evolved in different evolutionary replications, we used the performance data recorded over 500 trials to perform a series of pairwise Wilcoxon tests among all possible controller couples. The results are plotted in Figure 2 as vertical lines spanning over the controller numbers having a performance that is not statistically different (at 99% confidence). So, for example, controllers c_{13} and c_{15} are not statistically different from the performance point of view. Similarly, controller c_1 has a performance equivalent to c_{18} and c_{19} , but it performs worse than controllers c_{10} and c_{14} . As can be seen in Figure 2, controller c_8 outperforms all other controllers.

The fitness values just tells us that the system is able to maximise the performance metric we devised. It does not tell us if and how synchronisation is achieved, and whether or not the system presents desired properties like robustness or scalability. For this purpose, we performed behavioural and scalability analyses, which are discussed in the following section.

Place Figure 2 about here.

4.3 Phase 3: Analysis

By looking at the behaviour exhibited by a group of three robots, we observed that each instance of the evolutionary experiment produced a successful synchronisation behaviour, in which robots display oscillatory movements along the y direction and synchronise with each other, according to the requirements of the devised fitness function. In general, it is possible to distinguish two phases in the evolved

behaviours: an initial transitory phase during which robots achieve synchronisation, and a subsequent synchronised phase. The transitory phase may be characterised by physical interferences between robots due to collision avoidance, if robots are initialised close to each other. The collision avoidance behaviour performed in this condition eventually leads to a separation of the robots in the environment, so that further interferences to the individual oscillations are limited and synchronisation can be achieved. The synchronous phase is characterised by a stable synchronous oscillation of all robots, and small deviations from synchrony are immediately compensated.

The individual ability to perform oscillatory movements is based on the perception of the gradient painted on the arena floor, which gives information about the direction parallel to the y axis and about the point where to perform a U-turn and move back towards the x axis, therefore avoiding to end up into the black painted area. Each evolved controller produces a signalling behaviour that varies while the robots oscillate. The main role of the evolved signalling behaviour is to provide a coupling between the oscillating robots, in order to achieve synchronisation. In response to a perceived signal, robots react by moving in the environment, changing the trajectory of their oscillations. This results in a modulation of the oscillation amplitude and frequency, which allows the robots to reduce the phase difference among each other, and eventually synchronise. In a previous work [38], we developed a mathematical model and exploited dynamical systems theory to thoroughly analyse the synchronisation behaviour. We invite the reader to refer to that work for further details on the synchronisation mechanism, which are out of the scope of the present paper.

Once analysed the synchronisation behaviours evolved using three robots only, we tested scalability to large groups. To do so, we compared the performance of the evolved behaviour varying the group size. To avoid overcrowding, we performed the scalability analysis in larger arenas, ensuring a constant density of robots across the different settings. By ensuring a constant initial density we limit the negative effects of overcrowding and we are able to compare the performance of robotic systems with varying group size. In order to keep a constant robot density equal to the one used in the evolutionary experiments, we lengthened the arena in the x direction, trying to keep a uniform initial density of 0.25 robots per square meter. Despite the increased arena length, we still keep the same communication system, that is, communication continues to be binary and global, with all robots affecting each other. This choice allows us to evaluate the scalability of a behaviour as it was evolved, without modifying the features of the communication channel. We evaluated all best evolved controllers 100 times using six different group sizes (3, 6, 12, 24, 48 and 96 robots). The obtained results are presented in the top part of Figure 3. It is possible to notice that most of the best evolved controllers have a good performance for groups composed of 6 robots. Performance degrades for larger group sizes and only few controllers produce scalable behaviours up to groups formed by 96 robots. The main problem that reduces the scalability of the evolved controllers is given by the physical interactions among robots. Despite the constant initial

density we introduced in order to limit the disruptive effect of collision avoidance, physical interactions nevertheless occur with a higher probability per time step, as the group size increases. Every collision avoidance action provokes a temporary desynchronisation of at least two robots, which have to adjust their movements in order to re-gain synchronous oscillations with other robots. In such cases, the whole group is influenced by the attempt of few robots to re-gain synchronisation, due to the global and binary communication.

Place Figure 3 about here.

To summarise, the above analysis showed that physical interactions and collision avoidance have a disruptive effect on the synchronisation ability of the robots, and this effect is more and more visible as the group size increases. However, the synchronisation mechanism evolved may scale with the group size if we ignore physical interactions. To test this hypothesis, we performed an identical scalability analysis, but in this case we ignore the physical interactions among the robots, as if each robot was placed in a different arena and perceived the other robots only through communication signals. The obtained results are plotted in the bottom part of Figure 3. Differently from what was observed above, in this case many controllers present good scalability, with only a slight decrease in performance due to the longer time required by larger groups to perfectly synchronise (namely, controllers evolved in replication number 2, 8, 10, 12, 14, 18 and 19). This result confirms the analysis about the negative impact of physical interferences and collisions among robots. In fact, removing the necessity to avoid collisions leads to scalable self-organising behaviours.

Nevertheless, many other controllers present a strange behaviour (namely, controllers evolved in replication number 3, 4, 7, 9, 11, 13, 15, 16, 17, 20). It is possible to notice that the performance presents a high variability up to a certain group size. The variable performance indicates that in some cases the robots are able to synchronise, and in other cases not. With larger group sizes, the performance stabilises to a low, constant value, independent from the initial conditions and the number of robots used. This value, which is characteristic of each non-scaling controller, represents the performance of the robotic system trapped into the basin of an *incoherent attractor*. In other words, the robotic system always converges into a dynamical condition in which no robot can synchronise with any other. By observing the actual behaviour produced by these controllers, we realised that the incoherent condition is caused by a communicative interference problem: the signals emitted by different robots overlap in time and are perceived as a constant signal (signals are global and are perceived in a binary way, preventing a robot from recognising different signal sources). If the perceived signal does not vary in time, it does not bring enough information to be exploited for synchronisation, and the system remains desynchronised. This result is confirmed by the dynamical system analysis that we performed, which revealed how the individual signalling behaviour is responsible for producing such communicative interference, allowing

also to predict which controllers present scalability just looking at the individual behaviour (for more details, see [38]).

4.4 Re-engineering for scalability

The analysis of the unsuccessful controllers revealed that scalability cannot be always obtained, due to the physical and communicative interferences among robots. In particular, the communication channel we selected has a strong impact on the scalability of the system. In fact, communication is global and binary, that is, the signal emitted by a robot is perceived by any other robot everywhere in the arena. Moreover, from the robot point of view, there is no difference between a single robot and a thousand signalling at the same time. Therefore, a single robot can influence the whole group. This has no negative effect as long as robots are synchronous, but can have severe consequences when a robot modifies its behaviour due to collision avoidance following some physical interaction with other robots. Furthermore, the binary communication channel generates the communicative interference we described above, which prevents the group from synchronising in certain conditions.

The main problems are therefore related to the absence of *locality*—i.e., signals are perceived everywhere in the arena—and of *additivity*—i.e., signals overlap without adding, preventing to recognise how many robots are contemporaneously signalling. The lack of locality and additivity is the main cause of failure for the scalability of the evolved synchronisation mechanisms.³ We therefore decided to re-engineer our evolutionary experiments changing the communication system, which was arbitrarily chosen in the first place. Given that we are interested in studying global synchronisation, we decided to re-engineer our experiments focusing only on the additivity of the communication system. This allows us to make only minor changes to the experimental setup and directly compare the effects of the re-engineering approach.

We evolved self-organising synchronisation behaviours exploiting exactly the same setup as above, but changing the way robots signal and perceive emitted signals. Specifically, we change the binary communication system with a continuous one:

$$\tilde{s}(t) = \frac{1}{N} \sum_{r=1}^N \tilde{S}_r(t), \quad (3)$$

Now, robots always emit a signal $\tilde{S}_r(t) \in [0, 1]$, encoding a number in a continuous range. The emitted signals are perceived as the average $\tilde{s}(t)$ among all the perceived signals. By doing so, the influence of an individual robot on the global perceived signal—which is equal for all robots in the arena—depends on the signalling behaviour of the whole group: the bigger the group, the smaller the influence of the single individual. This communication channel can be easily implemented on the *s-bots*. For instance, signals could be sent as messages over the wireless network containing a real number in $[0, 1]$. On the basis of

³However, as we have seen, this problem affects only some of the analysed controllers.

the analysis performed so far, we expect that self-organising synchronisation behaviours can be evolved with such a communication system, and that these are more scalable.

Place Figure 4 about here.

Also in this case, we executed 20 instances of the evolutionary experiments, using groups of three robots. All replications were successful, and produced synchronisation behaviours that are qualitatively similar to those obtained with the binary communication system: robots perform oscillations over the painted gradient and react to the perceived signal by modifying the individual behaviour, in order to synchronise with other robots. We therefore executed a scalability analysis, which was performed with the same modalities as described above, and the obtained results are presented in Figure 4.

In the upper plot, scalability is tested including physical interactions. Also in this case, we notice that collisions prevent the scalability of some controllers, in which a good avoidance behaviour was not evolved. Recall that when a collision is detected, the group scores a null performance. However, it is possible to notice that the usage of an additive communication system leads to better performance even with large groups. Most controllers present good scalability for every tested group size, and only collisions substantially reduce the performance. Here, differently from what was observed before, physical interactions and collision avoidance do not have a severe impact on the performance of the whole group. In fact, the signals of few non-synchronous robots are averaged with those emitted by the rest of the group. As a consequence, the influence on the group of a single robot attempting to synchronise decreases with increasing group size. This leads to a quick convergence to synchrony and to an improved group performance.

To better understand the effects of the re-engineering approach, we also performed a scalability analysis for the evolved synchronisation mechanisms, again removing the physical interactions among robots. The results plotted in the lower part of Figure 4 show that all evolved synchronisation mechanisms perfectly scale, and they do not suffer from the communicative interference observed with binary signals. In fact, the perceived signal brings information about the average signalling behaviour of all robots. As a consequence, synchronisation is always achieved, no matter the group size. Notice also that all controllers present a linear decrease in performance in correspondence to an exponential growth of the group size. This observation suggests that the self-organising synchronisation mechanism is very efficient, and is only slightly affected by the group size.

5 Discussions and conclusions

Evolutionary Robotics can be a very useful methodology for the automatic synthesis of controllers for robotic systems. This is especially true in the swarm robotics domain, which is characterised by an

indirect, non-linear relationship between the control rules followed by an individual and the group behaviour. However, ER cannot be simply applied blindfold and effortlessly. Most importantly, ER does not exclude arbitrary choices in setting up an experiments, which can strongly influence the outcome of the evolutionary process. The advantage given by ER is that, despite such arbitrary choices, it can find good or even optimal solutions for a given problem. However, much as in conventional engineering methods, multiple design loops may be needed to find such optimal results. In this paper, we suggest that a structured, formal approach to the design of experiments in evolutionary swarm robotics can be very advantageous. On the one hand, we identified and discussed which are the relevant choices to be performed in designing the setup of an evolutionary experiment in the context of swarm robotics. On the other hand, we discussed the appropriateness of an engineering design cycle and showed its advantages in a practical case study.

The case study we presented was about the evolution of self-organising synchronisation in a robotic system. We showed that in setting up the experiments, some characteristics of the system were chosen somehow arbitrarily, given that no a priori knowledge was available about the possible solutions to the given problem. The results obtained with the initial experimental setup proved that self-organising synchronisation can be actually achieved with a minimal complexity at the level of the control and communication strategy. However, the analysis of the scalability results also pointed to some critical aspects of the system that hindered the group from scoring a good performance. We identified the problem in the communication system being global and binary, and to the effects of physical and communicative interferences. To solve this problem, we re-engineered the arbitrarily-chosen communication system exploiting the knowledge acquired by analysing the evolved behaviours. The newly devised continuous signals resulted in better synchronisation behaviours, and in an optimally scaling communication system.

This paper demonstrates that it is possible to engineer some features of a system undergoing artificial evolution on the basis of the outcome of the evolutionary process itself. Contrary to trial and error methods without any guidance, we showed that an attentive analysis of negative results conveys knowledge on how to modify the system for evolving better solutions. Note that this is not in contradiction with respect to the need of little a priori knowledge in the design of the evolutionary experiment, as mentioned in the introduction. In fact, the knowledge we put into engineering the evolutionary system should not be related to the design of the solution, which is left to the evolutionary process, but rather to the preconditions required for obtaining good solutions. This is a subtle but germane difference which should be discussed further. When we put knowledge about the solution into the experimental design—either explicitly in the fitness function or implicitly in the experimental setup—we try to constrain evolution into a specific path. However, it may be very difficult to force the evolutionary dynamics and to obtain the desired solutions. And it may also be the case that the obtained solutions are too specific for the evolutionary conditions that have been defined, missing important features like robustness, flexibility

and adaptivity. Additionally, if a specific behaviour is desired and known, it may be easier to engineer it directly without relying on an automatic design methodology, and possibly use evolutionary methods just for parameter tuning.

On the contrary, we propose an iterative engineering methodology that allows to exploit the knowledge acquired from negative results and previous evolutionary experiments to define the preconditions for obtaining better results. This does not mean forcing the system within specific evolutionary paths, but it rather liberates evolution from constraining attributes of the experimental setup. In this way, the solutions that can be generated should be more efficient, and should better generalise to conditions not directly experienced during evolution.

As a final remark, it is worth discussing the choice of the swarm robotics domain for the proposal of an engineering approach to ER. We actually believe that most of the methodological aspects mentioned in this paper can be easily generalised to other domains. However, within swarm robotics a structured and formal methodology is even more compelling. This is related to the idea of evolving self-organising processes through guidance rather than control: swarm robotics has the additional complexity of dealing with emergent systems, with emergent properties resulting from complex interactions among individuals. In this case, knowing which are the degrees of freedom in setting up an experiment can be of utmost importance, as well as knowing in advance which are the possible effects of certain choices. In future work, theoretical and experimental studies should be conducted to devise a comprehensive methodological framework for the design of ER systems: much like in the software engineering field, engineering practices should be developed to guide the experimenter through the different phases of setup, evolution and analysis of an evolutionary robotic system.

References

- [1] Ampatzis, C., Tuci, E., Trianni, V., Christensen, A. L., and Dorigo, M. (2009). Evolving self-assembly in autonomous homogeneous robots: Experiments with two physical robots. *Artificial Life*, 15(4):465–484.
- [2] Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. (2008). Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behaviour*, 16(1):5–26.
- [3] Baldassarre, G., Trianni, V., Bonani, M., Mondada, F., Dorigo, M., and Nolfi, S. (2007). Self-organised coordinated motion in groups of physically connected robots. *IEEE Transactions on Systems, Man and Cybernetics - Part B: Cybernetics*, 37(1):224–239.
- [4] Bongard, J. C. (2000). The legion system: A novel approach to evolving heterogeneity for collective

- problem solving. In Poli, R., Banzhaf, W., Langdon, W., Miller, K., Nordin, P., and Fogarty, T., editors, *Genetic Programming*. Berlin: Springer Verlag.
- [5] Buck, J. (1988). Synchronous rhythmic flashing of fireflies. II. *The Quarterly Review of Biology*, 63(3):256–289.
- [6] Christensen, A. L., O’Grady, R., and Dorigo, M. (2009). From fireflies to fault tolerant swarms of robots. *IEEE Transactions on Evolutionary Computation*, 13(4):754–766.
- [7] Der, R., Hesse, F., and Martius, G. (2006). Rocking stamper and jumping snake from a dynamical system approach to artificial life. *Adaptive Behavior*, 14(2):105–115.
- [8] Der, R., Steinmetz, U., and Pasemann, F. (1999). Homeokinesis - a new principle to back up evolution with learning. In *Computational Intelligence for Modelling, Control, and Automation*, volume 55 of *Concurrent Systems Engineering Series*, pages 43–47. Amsterdam: IOS Press.
- [9] Floreano, D., Dürr, P., and Mattiussi, C. (2008). Neuroevolution: from architectures to learning. *Evolutionary Intelligence*, 1(1):47–62.
- [10] Floreano, D., Mitri, S., Magnenat, S., and Keller, L. (2007). Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17(9):514–519.
- [11] Floreano, D. and Urzelai, J. (2000). Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13(4–5):431–443.
- [12] Funes, P., Orme, B., and Bonabeau, E. (2003). Evolving emergent group behaviors for simple humans agents. In Banzhaf, W., Christaller, T., Dittrich, P., Kim, J. T., and Ziegler, J., editors, *Advances in Artificial Life. Proceedings of the 7th European Conference on Artificial Life (ECAL 2003)*, volume 2801 of *Lecture Notes in Artificial Intelligence*, pages 76–89. Berlin: Springer Verlag.
- [13] Harvey, I., Di Paolo, E. A., Wood, R., Quinn, M., and Tuci, E. (2005). Evolutionary robotics: A new scientific tool for studying cognition. *Artificial Life*, 11(1–2):79–98.
- [14] Helbing, D. and Lämmer, S. (2008). Managing complexity: An introduction. In Helbing, D., editor, *Managing Complexity: Insights, Concepts, Applications*, pages 1–16. Berlin: Springer Verlag.
- [15] Klyubin, A., Polani, D., and Nehaniv, C. (2005a). All else being equal be empowered. In Capcarrere, M., Freitas, A. A., Bentley, P. J., Johnson, C. G., and Timmis, J., editors, *Advances in Artificial Life. Proceedings of the 8th European Conference on Artificial Life (ECAL 2005)*, volume 3630 of *Lecture Notes in Artificial Intelligence*, pages 744–753. Berlin: Springer Verlag.

- [16] Klyubin, A., Polani, D., and Nehaniv, C. (2005b). Empowerment: A universal agent-centric measure of control. In *Proceedings of the 2005 IEEE Congress on Evolutionary Computation*, pages 128–135. Piscataway, NJ: IEEE Computer Society Press.
- [17] Martius, G. and Herrmann, J. M. (2010). Taming the beast: Guided self-organization of behavior in autonomous robots. In Doncieux, S., Girard, B., Guillot, A., Hallam, J., Meyer, J.-A., and Mouret, J.-B., editors, *From Animals to Animats 11: 11th International Conference on the Simulation of Adaptive Behavior (SAB 2010)*, volume 6226 of *Lecture Notes in Computer Science*, pages 50–61. Berlin: Springer Verlag.
- [18] Martius, G., Herrmann, J. M., and Der, R. (2007). Guided self-organisation for autonomous robot development. In et al., F. A., editor, *Advances in Artificial Life. Proceedings of the 9th European Conference on Artificial Life (ECAL 2007)*, volume 4648 of *Lecture Notes in Artificial Intelligence*, pages 766–775. Berlin: Springer Verlag.
- [19] Meyer, J.-A., Husbands, P., and Harvey, I. (1998). Evolutionary robotics: A survey of applications and problems. In Husbands, P. and Meyer, J.-A., editors, *Proceedings of the First European Workshop on Evolutionary Robotics (EvoRobot98)*, volume 1468 of *Lecture Notes in Computer Science*, pages 1–21.
- [20] Mirolli, M. and Parisi, D. (2008). How producer biases can favor the evolution of communication: An analysis of evolutionary dynamics. *Adaptive Behavior*, 16(1):27–52.
- [21] Mirolo, R. E. and Strogatz, S. H. (1990). Synchronization of pulse-coupled biological oscillators. *SIAM Journal on Applied Mathematics*, 50(6):1645–1662.
- [22] Mitri, S., Floreano, D., and Keller, L. (2009). The evolution of information suppression in communicating robots with conflicting interests. *Proceedings of the National Academy of Sciences*, 106(37):15786–15790.
- [23] Mitri, S., Floreano, D., and Keller, L. (2011). Relatedness influences signal reliability in evolving robots. *Proceedings of the Royal Society B*, 278:378–383.
- [24] Mondada, F., Pettinaro, G. C., Guignard, A., Kwee, I. V., Floreano, D., Deneubourg, J.-L., Nolfi, S., Gambardella, L. M., and Dorigo, M. (2004). SWARM-BOT: A new distributed robotic concept. *Autonomous Robots*, 17(2–3):193–221.
- [25] Nelson, A. L., Barlow, G. J., and Doitsidis, L. (2009). Fitness functions in evolutionary robotics: A survey and analysis. *Robotics and Autonomous Systems*, 57(4):345–370.
- [26] Nolfi, S. and Floreano, D. (2000). *Evolutionary Robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.

- [27] Pfeifer, R., Lungarella, M., and Iida, F. (2007). Self-organization, embodiment, and biologically inspired robotics. *Science*, 318(5853):1088–1093.
- [28] Prokopenko, M. (2007). Design vs. self-organization. In Prokopenko, M., editor, *Advances in Applied Self-organizing Systems*, pages 3–17. Berlin: Springer Verlag.
- [29] Prokopenko, M. (2009). Guided self-organization. *HFSP Journal*, 3(5):287–289.
- [30] Prokopenko, M., Gerasimov, V., and Tanev, I. (2006). Evolving spatiotemporal coordination in a modular robotic system. In Nolfi, S., Baldassarre, G., Calabretta, R., Hallam, J., Marocco, D., Meyer, J., Miglino, O., and Parisi, D., editors, *From Animals to Animats 9: 9th International Conference on the Simulation of Adaptive Behavior (SAB 2006)*, pages 558–569. Berlin: Springer Verlag.
- [31] Quinn, M., Smith, L., Mayley, G., and Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London, Series A: Mathematical, Physical and Engineering Sciences*, 361(2321-2344).
- [32] Siciliano, B. and Khatib, O., editors (2008). *Springer Handbook of Robotics*. Berlin: Springer Verlag.
- [33] Sperati, V., Trianni, V., and Nolfi, S. (2008). Evolving coordinated group behaviours through maximization of mean mutual information. *Swarm Intelligence*, 2(2-4):73–95.
- [34] Strogatz, S. H. (2003). *Sync: The Emerging Science of Spontaneous Order*. New York, NY: Hyperion Press.
- [35] Strogatz, S. H. and Stewart, I. (1993). Coupled oscillators and biological synchronization. *Scientific American*, 269(6):102–109.
- [36] Trianni, V. (2008). *Evolutionary Swarm Robotics. Evolving Self-Organising Behaviours in Groups of Autonomous Robots*, volume 108 of *Studies in Computational Intelligence*. Berlin: Springer Verlag.
- [37] Trianni, V. and Dorigo, M. (2006). Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, 95:213–231.
- [38] Trianni, V. and Nolfi, S. (2009). Self-organising sync in a robotic swarm. A dynamical system view. *IEEE Transactions on Evolutionary Computation*, 13(4):722–741.
- [39] Trianni, V., Nolfi, S., and Dorigo, M. (2008). Evolution, self-organisation and swarm robotics. In Blum, C. and Merkle, D., editors, *Swarm Intelligence. Introduction and Applications*, Natural Computing Series, pages 163–192. Berlin: Springer Verlag.
- [40] Waibel, M., Keller, L., and Floreano, D. (2009). Genetic team composition and level of selection in the evolution of cooperation. *IEEE Transactions on Evolutionary Computation*, 13(3):648–660.

- [41] Wischmann, S., Huelse, M., Knabe, J. F., and Pasemann, F. (2006). Synchronization of internal neural rhythms in multi-robotic systems. *Adaptive Behavior*, 14(2):117–127.

sensory-motor system	<ul style="list-style-type: none"> - selected sensors and actuators - communication channels - pre-processing of raw data
genotype-to-phenotype mapping	<ul style="list-style-type: none"> - direct vs. indirect encoding - genetic relatedness of the group (homogeneous vs. heterogeneous)
fitness function	<ul style="list-style-type: none"> - functional vs. behavioural - implicit vs. explicit - external vs. internal
ecology	<ul style="list-style-type: none"> - sampling of ecological conditions - symmetry breaking

Table 1: The table summarises the relevant choices to be made in setting up an evolutionary experiment within the swarm robotics domain. Refer to Sections 2.1 to 2.4 for more details.




Figure 1: (a,b) The *s-bot*, the robot used in the experiments. (c) Snapshot of a simulation showing three robots in the experimental arena. The dashed lines indicate the reference frame used in the experiments.




Figure 2: Post-evaluation results of the best evolved controllers c_i in each replication $i = 1, \dots, 20$ of the evolutionary experiment. The performance is represented on the horizontal axis, and the controller number on the vertical axis. The boxplot displays the whole dataset: each box represents the inter-quartile range of the data, while the black vertical line inside the box marks the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box. The empty circles mark the outliers. Data from different controllers are sorted according to the median value. Moreover, statistical similarities are represented as vertical bars spanning over the controller numbers (see text for detail).




Figure 3: Scalability analysis. The boxplot shows, for each evolved controller, the performance obtained in tests with 3, 6, 12, 24, 48, and 96 robots. Each box represents the inter-quartile range of the data, while the black horizontal line inside the box marks the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box. Outliers are not shown. Top: scalability of the evolved controllers under normal conditions. Bottom: scalability of the synchronisation mechanism.




Figure 4: Scalability analysis for the continuous communication system. Top: scalability of the evolved controllers under normal conditions. Bottom: scalability of the synchronisation mechanism.