

Université Libre de Bruxelles

*Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*

**Evolution of Neuro-Controllers for the
Alignment of Physical Robots using a
Local Communication System**

Álvaro GUTIÉRREZ MARTÍN, Elio TUCI, Alexandre CAMPO

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2008-008

March 2008

IRIDIA – Technical Report Series

ISSN 1781-3794

Published by:

IRIDIA, *Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*

UNIVERSITÉ LIBRE DE BRUXELLES

Av F. D. Roosevelt 50, CP 194/6

1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2008-008

Revision history:

TR/IRIDIA/2008-008.001 March 2008

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsibility for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

Evolution of Neuro-Controllers for the Alignment of Physical Robots using a Local Communication System

Álvaro GUTIÉRREZ MARTÍN[†] `aguti@etsit.upm.es`

Elio TUCI^{††} `etuci@ulb.ac.be`

Alexandre CAMPO^{††} `acampo@iulb.ac.be`

[†]ETSIT, Universidad Politécnica de Madrid, Madrid, Spain

^{††} IRIDIA, CoDE, Université Libre de Bruxelles, Bruxelles, Belgium

March 2008

Abstract

In this paper, we use artificial evolution to design homogeneous neural network controller for groups of robots required to align. Aligning refers to the process by which the robots managed to head towards a common arbitrary an autonomously chosen direction starting from initial randomly chosen orientations. The cooperative interactions among robots require local communications that are physically implemented using infrared signalling. We study the performance of the evolved controllers, both in simulation and in reality for different group sizes. In addition, we analyze the most successful communication strategy developed using artificial evolution.

1 Introduction

Swarm robotics refers to research work in collective robotics in which autonomous cooperating agents are controlled by distributed, and local rules (Dorigo and Şahin, 2004). Each agent uses individual mechanisms and local perception to decide what action to take. Swarm robotics is of particular interest for roboticists since it promotes desired properties of swarms: (i) the failure of individual components does not significantly hinder the performance of the group (robustness); (ii) cooperative behaviour makes it possible to reduce the complexity of the individuals (simplicity of single units), and (iii) the control mechanisms used are not dependent on the number of agents in the swarm (scalability).

On going research work in swarm robotics is focusing on the development of design methods to obtain effective group level behaviours from the definition of individual mechanisms as well as on the development of effective communication systems that allow the coordination of actions in scenarios that require cooperation among the agents (Ampatzis et al., 2008; Tuci et al., 2008; Kube and Zhang, 1993; Trianni and Dorigo, 2006; Winfield et al., 2005). Following this lines of investigation, we are interested in studying the effectiveness of recent technological advances in the domain of localisation system as means for communication in groups of autonomous cooperating robots.

Committed to the principle of the Occam's razor, we focus on a simple task in which a group of homogeneous physical robots are required to align. Aligning refers to the process by which the robots managed to head towards a common arbitrary an autonomously chosen direction starting from initial randomly chosen orientations. Evolutionary robotics method (hereafter ER) is employed to design control and communication mechanisms to allow the robots to align. ER is a

methodological tool to automate the design of robots’ controllers based on the use of artificial evolution to find sets of parameters for artificial neural networks that guide the robots to the accomplishment of their task (Nolfi and Floreano, 2000). As far as it concerns the subject of this study, ER allow us to develop adaptive communication mechanisms that are grounded on the perceptual experience of the agents and fully integrated with all the other underlying structures that underpin the behavioural repertoire of each robot. This is because, with respect to other design methods, evolutionary robotics doesn’t require the designer to make strong assumptions concerning what behavioural and communication mechanisms are needed by the robots. Robots controllers are designed in simulation and then ported and evaluated on physical robots.

The results of this study are a “proof-of-concept”: they show that dynamic artificial neural networks can be successfully synthesised by artificial evolution to design the neural mechanisms required by the robots to align. In particular, owing to the use of this design method, we take advantage of the properties of the communication system used to allow the agents to obtain the required collective behaviour. Post-evaluation analyses unveil operational aspects of the best evolved behavioural strategy. By exploiting only a part of the potentialities of their communication system, the robots prove to be capable of solving the alignment task with simple and effective coordinated movements.

1.1 Localisation and Communication

The communication system the robots use in the considered task is based on technologies developed for computing relative distance and bearing of infrared signals’ sources. The system used manages transfer and processing of signals in software and does not require any additional hardware. We rely solely on the technologies (sensors and main board) available on the *e-pucks* (see Section *The E-pucks and their task* for a description of the robot).

Several works focus on the development of relative positioning system for autonomous robots. However, many of these works are emulations of relative positioning systems, and in some cases the systems are developed and tested only in simulation. Roumeliotis and Bekey (2002) proposed the use of a Kalman filter to combine dead reckoning and information from an emulated relative positioning system to allow a group of mobile robots to solve the issue of localisation. Ludwig and Gini (2006) used a wireless local area network for a robotic swarm dispersion to cover an unknown area. In this study, one robot was required to be stationary to determine distance and bearing information. The use of infrared signals for the development of relative positioning system has already been studied by Spears et al. (2004), Kelly and Martinoli (2004) and more recently by Pugh and Martinoli (2006).

All the above mentioned works shared with our study an interest in the design of relative positioning systems for autonomous robots that can only use local and decentralised control. However, while these works are mainly focused on the development of relative localisation systems, our goal is to merge these technologies with the ER to allow autonomous agents to use both situated and abstract communication. Situated communication refer to social interactions in which the physical instantiation of the message contributes to define its semantics (Clancey, 1997, for more details). In this work, receivers extract the range of the infrared signal source by considering the intensity of the modulated signal, and the bearing by considering the relative position on their body of the sensor impinged by the received signal. For example, a signal which is detected by a sensor positioned on the back of a robot receiver, indicated that the signal source is more or less at 180° with respect to the receiver heading. Since the localisation of signal source is obtained by exploiting the physical properties of the signals and the morphology of the robots, it follows that our robots are endowed with situated communication capabilities.

Abstract communication refers to communication protocols in which only the content of the message has meaning and the physical signal (the medium) that transports the message does not have any semantic properties (Støy, 2001). In our study, the robots are allowed to associate a binary message to each infrared signal. Although the message is anonymous and simply broadcast into the environment and not directly sent to specific robots, the receivers can identify the physical location of the signal and exploit the message. Since the semantics of the binary message have

nothing to do with the physical properties of the medium that transport it, it follows that our robots are endowed with abstract communication capabilities as well.

Note that, abstract communication tends to be used in multi-agents systems in which a wireless network provides the required structures to transmit messages from a specific sender to a specific receiver/s. In swarm robotics, the decentralised and local approach makes the use of wireless network difficult and abstract communication is not very common. Our swarm robotics system bypasses this limitation owe to the use of the innovative communication technologies described above. Moreover, ER makes possible to ground the semantics of abstract communication into the sensory-motor experience of the agent. That is, the meaning of binary messages is not a priori determined by the designer, but it emerges during the evolutionary process due to its value in the social interactions that characterise the alignment process.

2 The *E-pucks* and their task

E-pucks are modular, robust and non-expensive robots designed by Francesco Mondada and Michael Bonani from *Ecole Polytechnique Fédérale de Lausanne* (EPFL) for research and educational purposes. They are small wheeled cylindrical robots, 7cm of diameter, equipped with a variety of sensors, and whose mobility is ensured by a differential drive system¹ (see Figure 1a). The controllers are evolved in a simulation environment which models some of the hardware characteristics of the *e-pucks*. In particular, we have developed a simulation software based on ODE (<http://www.ode.org>), an open source, high performance library for simulating rigid 3D body dynamics that provides primitives for the implementation of detailed and realistic physics-based simulations. In our simulation, an *e-puck* is modelled as a cylindrical body of 3.5cm radius that holds 8 infrared sensors distributed around the body. A differential drive system composed of two wheels fixed to the body of the simulated robot.

We decided to test the effectiveness of the evolutionary robotics method to the design of adaptive communication mechanisms on a simple but fundamental task of alignment. In this task, robots should manage to adopt similar directions by exchanging relevant communications. Alignment is a fundamental behaviour to a number of tasks in robotics such as cooperative transport or flocking (see Campo et al., 2006; Hayes and Dorminiani-Tabatabaei, 2002; Kelly and Keating, 1996). Artificial evolution coupled to the local communication system may produce new and unexpected strategies for solving these tasks efficiently. The experimental setup consists of a group of homogeneous *e-pucks* that are positioned in a boundless arena at a distance of 10cm from each other, with randomly generated initial orientations, as depicted in Figure 1b. Each agent can only change its orientation through rotational movements. The robots can not move away or approach each other, they only turn on spot. The robots should tackle the task by exploiting the properties of a communication system based on the use of infrared signals from which the range and bearing of the signal source and a 3 bits message can be obtained.

¹Further details on the robot platform can be found at www.e-puck.org.

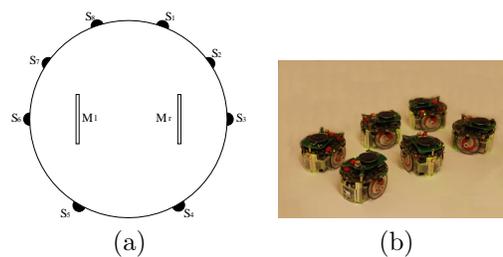


Figure 1: (a) Depiction of an *e-puck*. S_i $i \in [1, 8]$ refer to the infrared sensors which are used in this study for local communications. M_l and M_r are the left and right motor respectively. (b) Six *e-pucks* at the beginning of a trial.

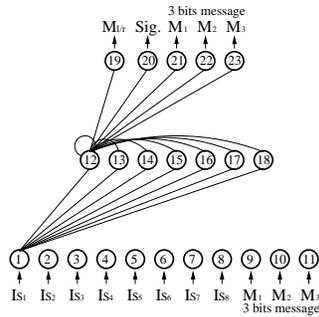


Figure 2: Neural network architecture: only the efferent connections of the first node of each layer are drawn. See the text for the meaning of the labels.

A signal generated by a robot sender through its infrared sensors can be picked up by the infrared sensors of robots receivers that are at a distance less than 25cm from the sender. The intensity of the signal decreases with the inverse of the square of the distance. It can be used by the receiver to infer distance of the sender. Given the directional propagation of infrared signals, the infrared sensor that picks up the signal can be used by the receiver to infer the relative direction of the sender with respect to its own heading. Given the notation used in Figure 1a, a signal emitted by a sender positioned 90° left of a receiver, is picked up by infrared sensor S_6 of the receiver. Each signal takes 25ms to be emitted, that is $1/8$ of the length of the control loop, which is 200ms. In the remaining 175ms the robots read from the sensors, update their control structure, and set the status of their left and right motor. Note that robots can not pick up signals while they send signals. Each signal can carry a binary encoded message, which are 3 bit long in this study. In principle, the 8 possible states of a message give the robots the possibility to refer to each sensor. However, the reader should note that the semantics of the messages is entirely left to the evolution, which shapes the mechanisms underpinning the robots behaviour. Whether or not the messages play a functional role in the accomplishment of the task, and which roles they play, will be ascertained through post-evaluation analysis by looking at behavioural mechanisms implemented by the evolved successful strategies.

In simulation, signals are modelled as an additive field of single frequency and fixed intensity which decrease with the square of the distance from the source. Simulated sensors return 1 bit to indicate if they were impinged by a signal and 3 other bits containing the message, if any. Noise value chosen uniformly random in $\pm 45^\circ$ is added to the relative orientation of the sender with respect to the receiver's heading, and noise chosen uniformly random in ± 2.5 cm is added to the sender-receiver distance. The model takes into account the fact that a receiver can not perceive a message if it is sending a message or if two messages arrive at the same time. A third robot placed on the direct path in between a signal source and the receiver's sensor interrupts the transmission. Moreover, each emitted message can be lost with a probability which varies from 1% when the sender-receiver distance is less than 1cm, to 50% when the two robots are at 25cm from each other. Messages are not altered in simulation: if a message is captured by a robot, it corresponds exactly to what was broadcast by the sender. Additionally, 10% uniform noise is added to the *e-pucks'* motor outputs.

3 The Robots' Controller

The agent controller is composed of a continuous time recurrent neural network (CTRNN) of 7 inter-neurons and an arrangement of 11 sensor neurons and 5 output neurons (see Beer and Gallagher, 1992, for further details on CTRNNs). At each time cycle, the network receives in input a 11 bits vector, which is composed of 8 bits corresponding to the status of all sensors I_{S_i} and the 3 bits corresponding to the message M_i of the sensor that is impinged by a signal, if any (see Figure 2). The inter-neuron network is fully connected. Additionally, each inter-

neuron receives one incoming synapse from each sensory neuron. Each output neuron receives one incoming synapse from each inter-neuron. There are no direct connections between sensory and output neurons. The states of the output neurons are used to set the speed of the robot wheels, and to control the signalling system. The network neurons are governed by the following equation:

$$\tau_i \dot{y}_i = \begin{cases} -y_i + gI_{S_i}; & i \in [1, 8] \\ -y_i + gM_j; & i \in [9, 11], j \in [1, 3] \\ -y_i + \sum_{j=1}^k \omega_{ji} \sigma(y_j + \beta_j); & i \in [12, 23] \end{cases} \quad (1)$$

with $\sigma(x) = \frac{1}{1 + e^{-x}}$

where, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant, g is a gain factor, $I_{S_i} = \{0, 1\}$ the perturbation on sensory neuron i with $i \in [1, 8]$ corresponding to the status of sensor S_i , $M_j = \{0, 1\}$ the perturbation on the 9th, 10th, and 11th sensory neuron corresponding to the 3 bits message, ω_{ji} the strength of the synaptic connection from neuron j to neuron i , β_j the bias term, $\sigma(y_j + \beta_j)$ the firing rate. The cell potentials y_i of the 19th neuron, mapped into $[0.0, 1.0]$ by a sigmoid function σ and then linearly scaled into $[-1, 1]$, set the angular speed of the robot right motor. The speed of the left motor is the inverse of the right motor. The output neurons 20, 21, 22, and 23 are used to control the signalling system. At each time cycle, a 3 bits message is broadcast if the cell potential of output neuron 20 is $y_{20} > 0.5$. The message M_j , $j \in [1, 3]$ is determined by the cell potentials y_i of the 21st, 22nd, and 23rd neuron. M_j is set to 1 if the cell potentials of the corresponding output neuron i is $y_i > 0.5$, 0 otherwise. The first time the network is updated $\forall j, M_j = 0$. If at time cycle t no sensors are impinged by a signal, that is $\forall i, I_{S_i} = 0, i \in [1, 8]$, then $\forall j, M_j$ keeps the value at time cycle $t - 1$. Cell potentials are set to 0 at the beginning of each experiment, and circuits are integrated using the forward Euler method with an integration step-size of $dt = 0.2s$.

4 The Evolutionary Algorithm

A simple generational genetic algorithm is employed to set the parameters of the networks (Goldberg, 1989). The population contains 80 genotypes. Genotypes of the first generation are generated randomly. Generations following the first one are produced by a combination of selection with elitism, recombination and mutation. For each new generation, the five highest scoring individuals (“the elite”) from the previous generation are retained unchanged. The remainder of the new population is generated by fitness-proportional selection (also known as roulette wheel selection) from the individuals of the old population. Each genotype is a vector comprising 182 real values (i.e., 161 connection weights, 12 decay constants, 8 bias terms, and one gain factor). Initially, random genotypes are generated by initialising each of their components to values chosen uniformly random from the range $[0, 1]$. New genotypes, except “the elite”, are produced by applying recombination with a probability of 0.3 and mutation. Mutation entails that a random Gaussian offset is applied to each real-valued vector component encoded in the genotype, with a probability of 0.15. The mean of the Gaussian is 0, and its standard deviation is 0.1. During evolution, all vector component values are constrained to remain within the range $[0, 1]$. Genotype parameters are linearly mapped to produce network parameters with the following ranges: biases $\beta_i \in [-4, -2]$ with $i \in [12, 23]$, weights $\omega_{ij} \in [-6, 6]$ with $i \in [12, 23]$ and $j \in [12, 18]$; gain factor $g \in [1, 10]$. The outputs neurons share the same bias. Decay constants are firstly linearly mapped into the range $[-1.0, 1.2]$ and then exponentially mapped into $\tau_i \in [2 * 10^{-1.0}, 10^{1.2}]$. The lower bound of τ_i corresponds to the integration step-size (i.e., $dt = 0.2s$) used to update the controller; the upper bound, arbitrarily chosen, corresponds to about 1/3 of the maximum length of a trial (i.e., 300s).

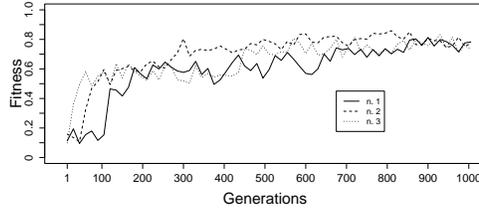


Figure 3: Average fitness of the best groups at each generation of the three best performing evolutionary runs.

5 The Fitness Function

During evolution, each genotype is translated into a robot controller, and cloned onto each simulated *e-puck*. Each group is evaluated in 8 trials. Each trial (e) differs from the others in the initialisation of the random number generator, which determines the amount of noise injected into the system and the robots initial orientation. Within a trial, the robots life-span is 300 simulated seconds (1500 simulation cycles). In each trial e , each group is rewarded by an evaluation function F_e which seeks to assess the ability of the *e-pucks* to perform alignment. F_e is computed as follows:

$$F_e = \frac{2}{T} \sum_{t=\frac{T}{2}}^T \left(1 - \frac{\pi - H_t}{\pi} \right); \quad (2)$$

$$H_t = \min(|\theta_t^{rc} - \theta_t^{cr}|), \text{ with } \theta_t^{rc}, \theta_t^{cr} \in [0^\circ, 360^\circ];$$

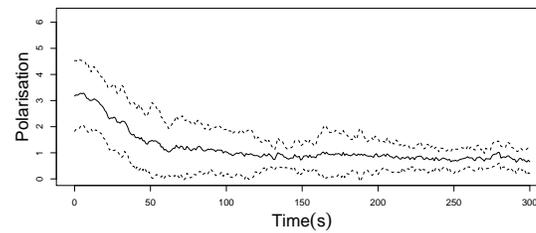
with $T = 1500$ corresponding to the length of a trial in simulation cycles; H_t corresponds the absolute value of the difference between θ_t^{rc} and θ_t^{cr} with $\{r, c \in [1, N] \mid r \neq c, N = 6\}$. At time t , θ_t^{rc} corresponds to the relative orientation of *e-puck* c with respect to the heading of *e-puck* r , and θ_t^{cr} corresponds to the relative orientation of *e-puck* r with respect to the heading of *e-puck* c . Headings are computed with respect to a general polar coordinate system. The fitness assigned to each genotype after evaluation of the *e-pucks* behaviour is given by

$$FF = \frac{1}{E} \sum_{e=1}^E F_e \text{ with } E = 8; \quad (3)$$

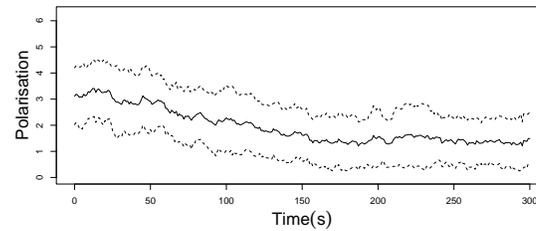
6 Results in simulation

Ten evolutionary simulations, each using a different random initialisation, were run for 1000 generations. These evolutionary processes aim at designing control mechanisms for groups of 6 simulated *e-pucks* required to negotiate a common heading. During evolution, the highest fitness score that a group can reach ($FF = 1.0$) corresponds to the circumstance in which the simulated robots are perfectly aligned. None of the evolutionary runs produced groups that managed to get the highest score. As shown in Figure 3, the fitness of the best evolutionary runs, after 700 generations, reached a plateau around the value of 0.7. If we consider that the fitness is computed by referring to the two less aligned robots of each group (see equation 2), and that the maximum score can hardly be produced by the robots due to the issues mentioned in the introduction, these fitness curves are quite promising.

However, it is generally the case that the fitness of the best groups in evolution is overestimated. Moreover, there is no guarantee that, once ported on the physical robots, the best evolved controllers are as effective as they are in simulation given that simulations approximate reality and that critical aspects of reality may have been neglected. For these reasons, we firstly run a series of post-evaluation tests in which each of the best evolved controllers at each generation of



(a)



(b)

Figure 4: Mean polarisation (\pm standard error) in function of time for 30 repeated experiments. (a) 4 physical *e-pucks*. (b) 6 physical *e-pucks*.

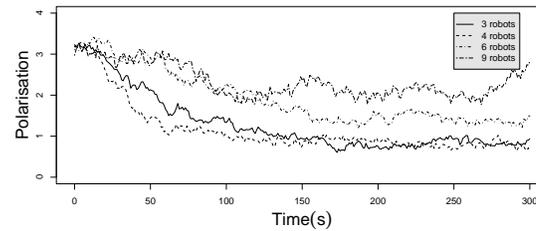


Figure 5: Mean polarisation (\pm standard error) in function of time for 30 repeated experiments for all group sizes tested (3, 4, 6, and 9 *e-pucks*) in superimposition, standard errors not shown for the sake of clarity.

each evolutionary run are evaluated for 1000 trials. In each post-evaluation trial the behaviour of robots is scored using the fitness function previously defined.

The best performing genotype resulting from this set of post-evaluations is decoded into an artificial neural network which is then cloned and ported onto physical robots. In what follows we provide the results of post-evaluation tests aimed at evaluating the success rate of physical *e-pucks* at the task. Post-evaluation was decoded into an artificial neural network which was then cloned and ported onto physical robots. In what follows, first we provide the results of post-evaluation tests aimed at evaluating the success rate of physical *e-pucks* at the task. Subsequently, we analyse the behaviour of simulated *e-pucks* to unveil operational aspects underlying the best evolved behavioural strategy.

Supplementary information is available at <http://iridia.ulb.ac.be/supp/IridiaSupp2008-005/index.html>

7 Results with physical robots

Real experiments are recorded using a digital camera. A tracking software is used to automatically extract the heading of each robot at each second. We use a specific measure of polarisation to calculate the degree of alignment of all the robots. While the fitness function is a measure tuned to facilitate the evolution of good controllers, the polarisation measure is devised to report faithfully observations on real robots.

Polarisation $P(G)$ of a group of robots G is defined using the angular nearest neighbour. For a robot r , the corresponding angular nearest neighbour c is defined such that θ^{rc} , the relative orientation of c with respect to r is the smallest possible : $\theta^{rc} < \theta^{ri}, \forall i \in G \setminus \{c\}$. We denote $\theta_{ann}(r)$ the relative orientation of the angular nearest neighbour of the robot r . The formal definition of polarisation is as follows :

$$P(G) = \sum_{i \in G} \theta_{ann}(i). \quad (4)$$

If all robots are aligned, then $P(G) = 0$. Conversely, if headings are evenly distributed, $P(G) = 2\pi$. Lastly, if headings are random, *i.e.* drawn from a uniform distribution, then $P(G) = \pi$ in average. It is worth mentioning that meaningful comparisons among different group sizes are possible since the average value of $P(G)$ is not affected by the number of robots in G .

We tested the best genotype obtained by artificial evolution in groups of 3, 4, 6, and 9 physical robots. Figure 4 reports the average polarisation (\pm standard error) of the robots across 30 repeated experiments. We have observed very good abilities to align for groups of 3 and 4 robots. For these group sizes, the polarisation measure is not statistically different. Groups of 6 robots are less performant and need more time to achieve alignment and groups of 9 robots display a polarisation close to random headings (close to π). Figure 5 clearly shows that groups of 3, 4 and 6 robots converge and maintain their alignment till the end of the experiments. The collective behaviour consists in a first phase of negotiation in which robots try to adopt similar headings for a period of time that varies from $t = 100s$ (3 and 4 robots) to $t = 150s$ (6 robots). With 9 robots the phase of negotiation may last the whole experiments as there is always at least one robot not properly aligned. In a second phase of maintenance, groups of robots stay aligned and achieve a stable degree of polarisation by regularly applying slight corrections to the headings. These corrections are needed as unavoidable errors in communications cancel the polarisation of the groups. During the negotiation phase, robots communicate in a pairwise manner that explains why this phase takes more time as group size increases. With respect to simulated experiments shown in Figure 7a and b, physical robots did not achieve such an extreme degree of polarisation. This is due to two main factors, namely the errors of the communication system and the automatic tracking system that has only 10° accuracy in the detection of robots headings, due to computational limitations.

To summarise, results on real robots show that the evolved genotype has been successfully transposed in real robots with no specific tuning, and while evolution was carried out with only groups of 6 robots, the neural network is able to cope with different group sizes and exhibit graceful degradation of performance as the task becomes more difficult.

8 Communication strategies

In order to find out the nature of the mechanisms underlying the communication strategies the robots are using to solve this coordination task, we proceed by carrying out further post-evaluation tests on simulated robots controlled by networks built from the best evolved genotype which is the same used on the real robots. In particular, we look at a simplest possible scenario—*i.e.*, two simulated *e-pucks* randomly oriented—in which the evolved communication mechanisms can be observed. In a single trial test, we record the status of the two *e-pucks*' sensors, the 3 bits messages transmitted, and H_t which is a measure of the convergence of the *e-pucks*' headings (see equation 2). Note that, $H_t = 1$ when the simulated robots are phasing the opposite direction,

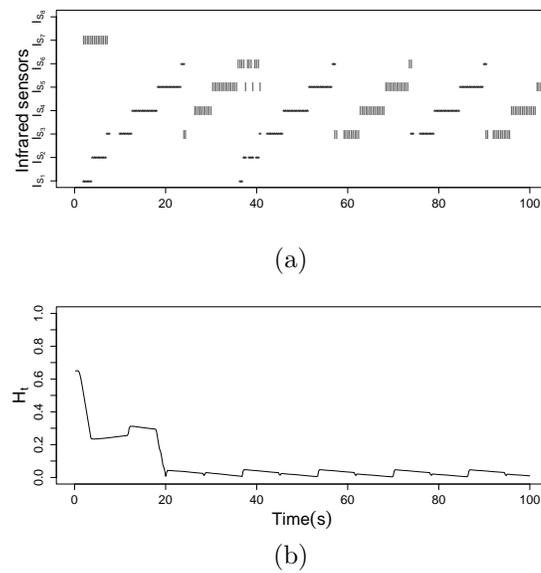


Figure 6: (a) Time during a trial in which the sensors of *e-puck* (|) and *e-puck* (*) are active—that is $I_{S_i} = 1$, with $i \in [1, 8]$. (b) Value of H_t during a trial.

and 0 when both robots head in the same direction. Since the robots can only turn on spot, the range of the signal sources is not relevant to the accomplishment of the task. The results of our post-evaluation test reveal that the 3 bits message is not used by the robots to solve the task given that they are always emitting the same message. Therefore, the behaviour we described in Section *Results with physical robots* seems to result from the exploitation of only the bearing of signal sources. Figure 6a provide information to understand how two robots manage to converge to a state of alignment.

When the robots share the same heading (i.e., from $t = 20$ s till the end of the trial, see Figure 6b), they rotate counterclockwise with the same angular speed, they emit signals for half of a turn period and they alternate signalling. This produces the regular pattern observed in Figure 6a, which shows the time during a trial in which the sensors of *e-puck* (|) and *e-puck* (*) are active. In particular, we notice that when the *e-puck* (|) is emitting signals *e-puck* (*) is receiving and vice versa. Both *e-pucks* switch from receiving to emitting as soon as they receive a signal from sensor S_6 , and from emitting to receiving as soon as they perceive a signal from sensor S_3 . Since the robots are moving with the same angular speed, the switch of roles—sender versus receiver—happens at a regular frequency. That is, the sender continuously emits while doing a 180° turn. This corresponds to the time it takes to the receiver to reach an orientation in which sensor S_6 is impinged by the signal. At this point, both robots are signalling and the one signalling for a longer time is receiving from sensor S_3 . Therefore, it stops emitting. This process is repeated till the end of the trial.

When robots are misaligned, they act in order to bring forth the cyclic phases described above. At the beginning of a trial, both robots emit signals and turn counterclockwise. When *e-puck* (*) switch state from receiving to emitting due to the perception of a signal form sensor S_6 , *e-puck* (|) does not stop emitting since it is not currently receiving from sensor S_3 but from sensor S_7 . *E-puck* (*) reacts to this contingency by stopping the emission of signals and changing its behaviour. That is, it interrupts its smooth rotational movement and starts oscillating, changing its direction of rotation from clockwise to counterclockwise anytime sensor S_6 is impinged by a signal, and from counterclockwise to clockwise anytime sensor S_5 is impinged by a signal. During the oscillations, *e-puck* (*) begins again to emit signals. While *e-puck* (*) is busy in performing this dance, *e-puck*

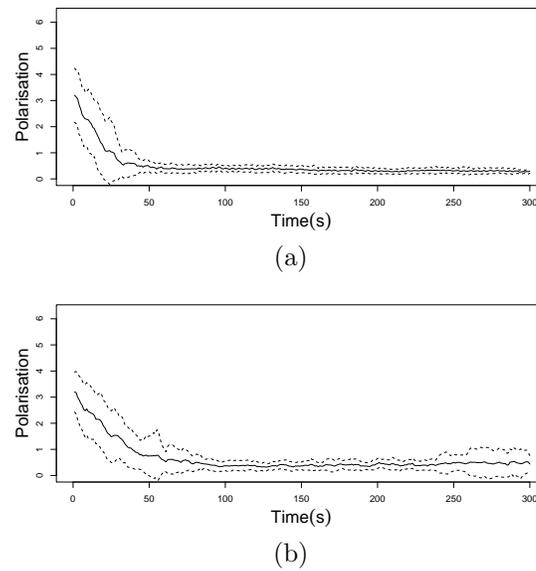


Figure 7: (a) Polarisation results of a team of 4 simulated *e-pucks* for a set of 30 experiments.(b) Polarisation results of a team of 6 simulated *e-pucks* for a set of 30 experiments.

(|) keeps on smoothly moving counterclockwise. As soon as it reaches an orientation in which signals impinged its sensor S_3 , *e-puck* (|) stops emitting. For *e-puck* (*) this is the condition to interrupt its oscillating behaviour and start the smooth rotational movement. At this point ($t = 20s$) the robots have more or less the same heading—see $H_t < 0.1$ in Figure 6b—and the cyclic phases guarantee the maintenance of the coordinated movements.

Note also that, in this trial, the simulated robots are initially oriented in such a way that they are nearly facing opposite directions—see $H_t = 0.85$ in Figure 6b. At the very beginning of the trial, both robots try to wait for the other changing rotation sense, so they both enter into the oscillatory phase mentioned above. Errors in communications and noise on motors allow the robots to get out of the deadlock and successfully accomplish their objective.

We replicate test with physical robots in simulation, to do a comparison between the real and simulated behaviour. Figure 7a shows polarisation measures for a group of 4 robots, and Figure 7b for a group of 6. In simulation robots are behaving slightly better than in reality. They manage to negotiate their common heading before $t = 70s$. We studied these differences by keeping track of the inputs and outputs values of the neural network in both simulated and real experiments. We observe that as the size of the group of robots increases, more communication signals are simultaneously broadcast in the environment and message collisions are therefore more likely to arise. This degradation of the communication signals has not been implemented on the simulated robots due to the difficulty of simulating the shared medium of the communication system. This is the main explanation to the negotiation time difference between real and simulated results.

9 Conclusions

In this paper, we describe a neural network controller designed using artificial evolution that allows robots to implement efficient communication strategies to align. Aligning refers to the process by which the robots managed to head towards a common arbitrary an autonomously chosen direction starting from initial randomly chosen orientations. The communication among the robots is base on a system which provides a receiver information concerning range, bearing of the signal source and a binary message broadcast into the environment by the sender. The results show that

artificial evolution synthesised successful neural network controllers that provide the robots all the behavioural and communication mechanisms required by the task.

The best evolved neural network has been successfully ported on the physical robots. Despite the differences between simulations and reality, the evolved neural network has proved being robust to errors on the communication and kinematics of the real robots. Results are a proof-of-concept that artificial evolution helps designing controllers faster and easily. We also observed that changing the cardinality of the group does not disrupt the effectiveness of the best evolved strategies. However, the robots exhibits graceful degradation of their performance as the group size increases and the task becomes more difficult. Further investigation on the operational principles of the best evolved strategies reveal how the communication system is being used. Robots do not rely on semantic communication as it was expected. Instead, artificial evolution has found a very simple way to deal with the alignment task by simple relying on the physics of the signal and on a signalling behaviour in which sender and receiver alternates period of signalling and period of silence. The simplicity of the synthesised controller is likely to produce robust behaviours, and the lack of semantic in messages makes the controller less prone to defect of alignment due to errors in transmission.

In the future, we intend to develop the approach described in this paper into more complex scenarios. Since alignment is a fundamental behaviour for a number of swarm robotics tasks, we will investigate more challenging tasks that demand potentially more elaborated communications among the robots. In particular, we are thinking about setups such as flocking, with the expectation to discover new strategies of communication that could be as surprising as the ones observed in this work.

10 Acknowledgements

Authors acknowledge the ASL, LIS and SWIS laboratories of the Ecole Polytechnique Fédérale de Lausanne for the design and creation of the *e-puck* robot, and Valentin Longchamp for his contribution to the implementation of infrared communication on the *e-pucks*. E. Tuci acknowledges European Commission support via the *ECAgents* project, funded by the Future and Emerging Technologies programme (grant IST-1940). A. Gutiérrez acknowledges Consejo Social of the Universidad Politécnica de Madrid support via Fifth PhD Formation programme and Dr. Marco Dorigo for a fruitful visiting period at IRIDIA. A. Campo acknowledges support from the fund for scientific research F.R.S.-FNRS of Belgium’s French Community, of which he is a Research Fellow.

References

- Ampatzis, C., Tuci, E., Trianni, V., and Dorigo, M. (2008). Evolution of signaling in a multi-robot system: Categorization and communication. *Adaptive Behavior*, 16(1):5–26.
- Beer, R. D. and Gallagher, J. C. (1992). Evolving dynamic neural networks for adaptive behavior. *Adaptive Behavior*, 1(1):91–122.
- Campo, A., Nouyan, S., Birattari, M., Groß, R., and Dorigo, M. (2006). Negotiation of goal direction for cooperative transport. In *Ant Colony Optimization and Swarm Intelligence 2006*, volume 4150 of *LNCS*, pages 191–202. Springer Verlag.
- Clancey, W. J. (1997). *Situated Cognition: On Human Knowledge and Computer Representations*. Cambridge University Press, Cambridge, UK.
- Dorigo, M. and Şahin, E. (2004). Swarm robotics – special issue editorial. *Autonomous Robots*, 17(2–3):111–113.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization and machine learning*. Reading, MA: Addison-Wesley.
- Hayes, A. and Dorminiani-Tabatabaei, P. (2002). Self-organizaed flocking with agent failure: Off-line optimization and demostration with real robots. In *IEEE Int. Conf. on Robotics and Automation*, pages 3900–3905.

- Kelly, I. and Keating, D. (1996). Flocking by the fusion of sonar and active infrared sensors on physical autonomous mobile robots. In *Third Int. Conf. on Mechatronics and Machine Vision in Practice*, pages 1–4.
- Kelly, I. and Martinoli, A. (2004). A scalable, on-board localisation and communication system for indoor multi-robot experiments. *Sensor Review*, 24(2):167–180.
- Kube, C. R. and Zhang, H. (1993). Collective robotics: From social insects to robots. *Adaptive Behaviour*, 2(2):189–218.
- Ludwig, L. and Gini, M. (2006). Robotic swarm dispersion using wireless intensity signals. In *Proc. of the Int. Symposium on Distributed Autonomous Robotic Systems*, pages 135–144.
- Nolfi, S. and Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. Cambridge, MA: MIT Press.
- Pugh, J. and Martinoli, A. (2006). Relative localization and communication module for small-scale multi-robot systems. In *IEEE Int. Conf. on Robotics and Automation*, pages 188–193.
- Roumeliotis, S. and Bekey, G. (2002). Distributed multirobot localization. *IEEE Transaction on Robotics and Automation*, 18(5):781–795.
- Spears, W., Heil, R., Spears, D., and Zarzhitsky, D. (2004). Physicomimetics for mobile robot formations. In *Autonomous Agents and Multi-Agent Systems*.
- Støy, K. (2001). Using situated communication in distributed autonomous mobile robots. In *Proc. of the 7th Scandinavian Conf. on artificial intelligence*, pages 44–52. Amsterdam, NL: IOS Press.
- Trianni, V. and Dorigo, M. (2006). Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, 95:213–231.
- Tuci, E., Ampatzis, C., Vicentini, F., and Dorigo, M. (2008). Evolving homogeneous neuro-controllers for a group of heterogeneous robots: Coordinated motion, cooperation, and acoustic communication. *Artificial Life*, 14(2). Forthcoming.
- Winfield, A. F. T., Sa, J., Gago, M. C., Dixon, C., and Fisher, M. (2005). On formal specification of emergent behaviours in swarm robotic systems. *International journal of advanced robotic systems*, 2(3):363–370.