

COLLECTIVE PERCEPTION IN A SWARM OF AUTONOMOUS ROBOTS

Giuseppe Morlino, Vito Trianni, Elio Tuci

Institute of Cognitive Sciences and Technology, National Research Council, Via San Martino della Battaglia 44, 00185 Rome, Italy
{giuseppe.morlino, vito.trianni, elio.tuci}@istc.cnr.it

Keywords: Swarm robotics, collective behaviour, perceptual discrimination, evolutionary algorithms, artificial neural networks.

Abstract: We present a study that aims at understanding how perception can be the result of a collective, self-organising process. A group of robots is placed in an environment characterised by black spots painted on the ground. The density of the spots can vary from trial to trial, and robots have to collectively encode such density into a coherent flashing activity. Overall, robots should prove capable of perceiving the global density by exploiting only local information and robot-robot interactions. We show how we can synthesise individual controllers that allow collective perception by exploiting evolutionary robotics techniques. This work is a first attempt to study cognitive abilities such as perception, decision-making, or attention in a synthetic setup as result of a collective, self-organising process.

1 INTRODUCTION

How can a distributed system collectively encode the magnitude of a macroscopic variable? This question holds over multiple domains, and at different scales. First and foremost, in the context of cognitive neuroscience, this question can be reformulated as: what are the neural mechanisms underlying perception? This is a fundamental question, which must be answered first in order to lay the foundations for further investigations on other cognitive processes, such as decision-making, attention or learning. For this reason, the literature abounds of models about neural coding of every sort of stimuli, from the basic ones—e.g., vibro-tactile or visual stimuli (Romo and Salinas, 2003; Loffler, 2008)—to more complex perceptual conditions—e.g., multi-stability, face recognition or numbers (Leopold and Logothetis, 1996; Rubin, 2003; Grill-Spector, 2003; Dehaene, 2003).

The problem of suitably encoding environmental stimuli, however, does not pertain exclusively individual animals, but is of fundamental importance also for collective systems, such as bird flocks and honeybee swarms. Similar systems are often considered as super-organisms, due to their high cohesion and in-

tegrated functioning (Hölldobler and Wilson, 2008; Detrain and Deneubourg, 2006). It is therefore interesting to look at how super-organisms can achieve a coherent perception of macroscopic features of the environment they inhabit. For instance, while searching a new nesting site, honeybees explore the environment thanks to scouts that report their discoveries to the nest. In there, a collective perception and a decision-making process is carried on, which results in the recognition and selection of the best site among the discovered choices (Passino et al., 2008). In this process, no single bee has the full picture. However, the partial information of many bees is aggregated in the nest and through a self-organising process decision-making is successfully performed. Passino et al. (2008) recognise strong similarities between honeybees behaviour and the mechanisms that support perception and decision-making in neural systems. In particular, cross-inhibition within neural populations is functionally similar to negative feedback between bee workers committed to different nesting sites. The parallel between cognitive systems and swarm behaviour goes beyond qualitative considerations. Marshall et al. (2009) compared the nest site selection behaviour in ants and honeybees

and the brain dynamics during decision making in a perceptual choice task. They show that the swarm behaviour can be described by the same model that was proposed for decision making in (Ratcliff and Smith, 2004). As a consequence, the two decision processes can be directly compared, and similarities can be drawn between cognition in the brain and in the swarm.

In this paper, we aim at studying collective perception in a robotic swarm. The goal of this study is understanding which are the self-organising processes underlying the collective perception of a macroscopic environmental feature, which is not accessible to the individual robots due to their limited perceptual abilities and due to the nature of their individual exploration strategies. Therefore, multiple robots need to interact in order to give a collective response that correlates with the macroscopic variable. It is worth noticing that the perceptual discrimination task employed could in principle be solved by a single robot, given an effective exploration strategy and enough time to accomplish it. The reason why we let a group of robots to find a collective solution is because we believe that the study of successful collective discrimination strategies in this particular artificial scenario may shed a light on the mechanisms of collective perception in natural organisms.

In this robotic model, we synthesise the robot neural controllers through evolutionary robotic techniques, and we afterwards analyse the obtained results in order to uncover the mechanisms that support the collective perception process.

The usage of evolutionary techniques for collective and swarm robotics has been demonstrated in various recent studies. For instance, Trianni and Nolfi (2009) evolved self-organising synchronisation for a group of robots that presented an individual periodic behaviour, while Sperati et al. (2010) showed how a robotic swarm evolved through an evolutionary process managed to collectively explore the environment and forming a path to navigate between two target areas, which were too distant to be perceived by a single agent at the same time. Hauert et al. (2009) exploited artificial evolution to synthesise Swarming Micro Air Vehicles (SAMVs) able to organise autonomously, relying only on local informations, to establish a wireless communication network between users located on the ground.

Our working hypothesis is that the evolutionary process can produce optimal solutions to the given task. Therefore, by analysing these solutions, we can discover general mechanisms for collective perception, which are adapted to the experimental conditions we have devised. This allows us to discuss the dis-

covered mechanisms with respect to known processes performed by individuals and collectives.

2 EXPERIMENTAL SETUP

As mentioned above, in this paper we study how a swarm of robots can collectively encode a macroscopic feature of the environment. We have set up an experimental arena in which black circular spots are painted on a grey background. The macroscopic feature that must be encoded by the robotic swarm is the density of black spots, which may vary from trial to trial in the range $d \in [0, 1]$. Robots can perceive the colour of the ground only locally, through a noisy infrared sensor placed under their chassis. Robots can emit flashing signals, which can be perceived by all other robots. By combining the locally acquired information through this kind of simple communication, the group should encode the global density through the frequency of the emitted signals: the higher the density, the higher the frequency of the collective flashing signal. In the following, we give the details of the experimental setup and of the evolutionary algorithm we used to synthesise the robot neural controllers.

2.1 The robots and the environment




Figure 1: Two snapshots of the simulated arena are shown. The black disks spots are painted on a grey floor. The spots are positioned on a grid of 40×40 cells. The density (i.e. the probability to find a spot in a cell) varies in the range $[0, 1]$ (left: $d = 0.26$, right: $d = 0.66$).

The experimental arena is square (side $l = 2m$) and surrounded by walls. Circular black spots are painted on the ground in order to probabilistically obtain a desired global density. The spots are homogeneous in colour and size (radius $r = 2.5cm$), and are aligned to a square grid of 40×40 cells (see Fig. 1). The density d represents the probability that each cell of the grid is filled with a black spot. Therefore, when the density is 0, no spot is present and the arena ground is completely grey; when the density is 1, the arena is completely filled with black circular spots. In this

way, we can control the black spot density with a single parameter, and we can create multiple instances for the same macroscopic value.

Ten robots (radius 3.75cm) are randomly placed in the environment. Each robot is equipped with two wheels that provide a differential drive motion (maximum linear speed: $v_{max} = 8.2\text{cm/s}$). Robots can perceive walls and other obstacles by means of eight infrared sensors placed around the turret (see Fig. 2). The infrared sensors can be exploited for obstacle avoidance. The ground colour is perceived through an infrared sensor placed under the chassis of the robot, in the front part (indicated by ‘G’ in Fig. 2). In the absence of noise, the ground sensor returns 0 when is over a black spot, and 0.5 when is over the grey background. Additionally, we make this sensor very unreliable by adding 30% white noise to the absolute sensor reading. Finally, each robot r can emit a flashing signal $S_r(t)$ switching on for a time-step ($\Delta t = 0.1\text{s}$) the LEDs placed around its turret. This signal can be perceived by all the other robots present in the environment in a binary way: $s(t) = 1$ if there is at least one robot r emitting a signal, otherwise $s(t) = 0$. A robot can perceive the flashing signals through the omni-directional camera, including the signals emitted by the robot itself.




Figure 2: The robot and the neural controller. Left: a schema of the simulated robot. Eight proximity sensors (IR_{0-7}) are positioned at 3.25cm from the ground pointing horizontally. The sensors detect obstacles at a maximum distance of $\approx 5\text{cm}$. The floor colour is perceived through an infrared sensor (G) positioned on the robots’ front and pointing the floor. Right: the agents’ controller is a CTRNN with 10 sensory neurons, 4 hidden and three motor units.

2.2 The controller and the evolutionary algorithm

Each robot is controlled by a continuous time recurrent neural network (CTRNN) (Beer, 1995). The neural network has a multi-layer topology, as shown in Fig. 2: neurons $N_{I,1}$ to $N_{I,10}$ take input from the robot’s sensory apparatus, neurons $N_{O,1}$ to $N_{O,3}$ control the robot’s actuators, and neurons $N_{H,1}$ to $N_{H,4}$ form a fully recurrent continuous time hidden layer. The input neurons are simple relay units, while the output neurons are governed by the following equations:

$$o_j = \sigma(O_j + \beta_j), \quad (1)$$

$$O_j = \sum_{i=1}^4 W_{ij}^O \sigma(H_i + \beta_i), \quad (2)$$

$$\sigma(z) = (1 + e^{-z})^{-1}, \quad (3)$$

where, using terms derived from an analogy with real neurons, O_j and H_i are the cell potentials of respectively output neuron j and hidden neuron i , β_j and β_i are bias terms, W_{ij}^O is the strength of the synaptic connection from hidden neuron i to output neuron j , and o_j and $h_i = \sigma(H_i + \beta_i)$ are the firing rates. The hidden units are governed by the following equation:

$$\tau_j \dot{H}_j = -H_j + \sum_{i=1}^4 W_{ij}^H \sigma(H_i + \beta_i) + \sum_{i=1}^{10} W_{ij}^I I_i, \quad (4)$$

where τ_j is the decay constant, W_{ij}^H is the strength of the synaptic connection from hidden neuron i to hidden neuron j , W_{ij}^I is the strength of the connection from input neuron i to hidden neuron j , and I_i is the intensity of the sensory perturbation on neuron i . The weights of the connection between neurons, the bias terms and the decay constants are genetically encoded parameters. Cell potentials are set to 0 each time a network is initialised or reset. State equations are integrated using the forward Euler method with an integration step-size of 0.1 seconds.

Eight input neurons— $N_{I,1}$ to $N_{I,8}$ —are set from the infrared sensors. Input neuron $N_{I,9}$ is set from the ground sensor. Finally, input neuron $N_{I,10}$ is a binary input set by the perception of the flashing signal $s(t)$. The neurons $N_{O,1}$ and $N_{O,2}$ are used to set the speed of the robot’s wheels. Neuron $N_{O,3}$ is used to switch on the LEDs. In order to emit a flashing signal that lasts a single time-step, the LEDs are switched on only when the neuron activation surpasses the threshold 0.5:

$$S_r(t) = 1 \iff o_3(t) \geq 0.5 \wedge o_3(t-1) < 0.5. \quad (5)$$

This means that in order to flash again, the activation o_3 of neuron $N_{O,3}$ must go below the threshold, and up

again. The minimum period for oscillations is therefore 2 time-steps, that is, $0.2s$.

The free parameters of the robot’s neural controller are encoded in a binary genotype, using 8 bits for each real number. Evolution works on a population of 100 randomly generated genotypes. After evaluation of the fitness, the 20 best genotypes survive in the next generation (elitism), and reproduce by generating four copies of their genes with a 2% mutation probability of flipping each bit. The evolutionary process lasts 5000 generations. During evolution, genotype parameters are constrained to remain within the range $[0, 1]$. They are mapped to produce CTRNN parameters with the following ranges: connection weights $W_{ij} \in [-4, 4]$; biases $\beta \in [-4, 4]$; concerning decay constants, the genetically encoded parameters are firstly mapped onto the range $[-1, 2]$ and then exponentially mapped onto $\tau \in [10^{-1}, 10^2]$. The lower bound of τ corresponds to the integration step size used to update the controller; the upper bound is arbitrarily chosen.

2.3 The Fitness Function

A genotype is translated into $N = 10$ identical neural controllers which are downloaded onto N identical robots (i.e., the group is homogeneous). Each group of robots is tested for 20 trials, which last either 1000 or 2000 time-steps (one time-step corresponds to $0.1s$). The density is varied systematically, making the group experience 20 different values, equally distributed in $[0, 1]$. The robots’ neural controllers are not reset from trial to trial, therefore the order in which trials are presented is relevant. At each fitness evaluation, we randomly shuffle the sequence of environments experienced by the same group, in order to remove regularities that could be exploited by spurious behaviours. In order to evaluate the fitness of




Figure 3: Schematic representation of the collective flashing signal, through which the group of robots encodes the black spot density.

a genotype, we measure how well the corresponding robotic group encodes the black spot density d . To do so, we demand that robots as a group display a periodic flashing activity with a frequency that correlates with the black spot density. The group flashing activity is measured on the global signal $s(t)$ that results from the coupled activity of each robot. When

robots flash in subsequent time-steps, their signals are perceived as a single *flashing event* (i.e., a sequence of consecutive flashes is perceived as a square signal, see Fig. 3). We measure the period T_i as the time between the start of two subsequent events. In this way, we obtain a series of inter-flash periods that we use to compute the fitness. First of all, we compute through an exponential moving average the average period \hat{T} and the average difference between two consecutive periods ΔT :

$$\hat{T} = \alpha \hat{T} + (1 - \alpha)T_i, \quad (6)$$

$$\Delta T = \alpha \Delta T + (1 - \alpha)|T_i - T_{i-1}|, \quad (7)$$

where $\alpha = 0.9$ is the time constant of the moving average. At the end of the trial θ , \hat{T} should encode the density. We measure the encoded density by linearly scaling the average period:

$$d_{enc} = \frac{T_M - \hat{T}}{T_M - T_m}, \quad (8)$$

where $T_M = 5s$ and $T_m = 1s$ are respectively the maximum and minimum periods, arbitrarily chosen. Finally, the two fitness components are computed: F_d^θ rewards the group for suitably encoding the black spot density:

$$F_d^\theta = \Phi(1.0 - |d - d_{enc}|), \quad (9)$$

where $\Phi(x)$ is a piecewise linear function that simply constrains the fitness value in the interval $[0, 1]$. This component therefore rewards the group for minimising the difference between the black spot density and the group encoded density. However, it does not assure that the system converges towards a periodic signalling. For this purpose, a second fitness component is computed, that minimises the difference between consecutive periods:

$$F_\Delta^\theta = \Phi(1.0 - \frac{\Delta T}{\Delta T_M}), \quad (10)$$

where $\Delta T_M = 2s$ is the maximum difference allowed. By minimising the difference among consecutive periods, the system is rewarded to produce periodic signals. Finally, the fitness of a genotype is the product of the two fitness components, averaged over multiple trials:

$$F = \sum_{\theta=1}^{20} F_d^\theta \cdot F_\Delta^\theta. \quad (11)$$

A trial is stopped and the fitness is zero when no flashing event is detected within the last $10s$, therefore promoting a sustained flashing activity during the whole trial. Similarly, a trial is stopped if any robot collides with another robot or with a wall, and the fitness is zero for that trial. This indirect selective pressure allows to evolve obstacle avoidance.




Figure 4: Generalisation test for the excitatory strategy. Boxes represent the inter-quartile range of the data, while the horizontal lines inside the boxes mark the median values. 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.

3 RESULTS

We have performed 20 evolutionary runs for 5000 generations. For each evolutionary run, we selected a single genotype to be further analysed. To this aim, we evaluated the performance of the 20 best individuals of the last generation, measuring the fitness over 100 trials for each of the 20 density values (2000 trials in total), and we selected the genotype with the highest mean performance to represent the evolutionary run. Among the selected genotypes, 9 out of 20 resulted in a good collective behaviour while the remaining ones resulted in sub-optimal solutions, in which the group always converged to a fixed signalling frequency, therefore failing to suitably encode the black spot density. The performance of the best genotypes is presented in Table 1. Despite the variability in performance, the behaviours evolved in different evolutionary runs are qualitatively similar: robots mainly rotate on the spot, in some cases slightly moving away from the initial position. While rotating on the spot, the ground sensor positioned on the robot front gives a very local and noisy estimate of the ground colour. The ground information is integrated over time, and modulates an internal oscillator that allows to tune the frequency of a periodic signalling. However, this frequency is related just to the local density perceived by the robot, which may be significantly different from the global density: in fact, an individual robot rotating in one place can perceive only a limited number of different ground patterns, which do not represent well the global density, above all for intermediate density values. Moreover, the 30% white noise of the ground sensor makes it dif-

icult to have even a good and stable local perception. For these reasons, robots have to coordinate to better estimate the global density, and to do so, they can exploit the flashing signals. By analysing the communication strategies evolved in the different evolutionary runs, we found that they can be grouped into two classes. In some cases, the flashing signals are *excitatory*, that is, signal reception anticipates or provokes the signal production. This is the case for the behaviour evolved in runs 4, 9, 10 and 15. In the other cases—namely, runs 3, 13, 14, 19 and 20—flashing signals are *inhibitory*, that is, signal reception prevents or delays the signal production. In order to understand the mechanisms behind these two strategies, we analyse the best performing genotype of each class, namely the one obtained in run 4 for the exci-

Table 1: Performance of the genotypes that result in a good collective perception behaviour. Data are sorted in decreasing order and, for each column, the mean and standard deviation are shown. The columns represent the fitness F and the the two components F_d and F_Δ .

run	F	F_d	F_Δ
4	0.87 ± 0.06	0.92 ± 0.06	0.95 ± 0.02
19	0.85 ± 0.08	0.92 ± 0.08	0.93 ± 0.04
14	0.84 ± 0.07	0.89 ± 0.08	0.94 ± 0.03
9	0.83 ± 0.07	0.92 ± 0.06	0.91 ± 0.05
20	0.83 ± 0.08	0.92 ± 0.07	0.90 ± 0.05
10	0.82 ± 0.06	0.91 ± 0.07	0.90 ± 0.03
15	0.81 ± 0.07	0.88 ± 0.08	0.92 ± 0.03
3	0.80 ± 0.10	0.85 ± 0.11	0.94 ± 0.03
13	0.80 ± 0.07	0.89 ± 0.08	0.89 ± 0.04

tatory strategy, and the one obtained in run 19 for the inhibitory one.

3.1 Analysis of the Excitatory Strategy

To better understand the properties of the evolved behaviour, we first perform a generalisation test that aims at revealing how well the system behaves with varying densities. For this purpose, we have recorded the performance of the system over 50 different black spot densities uniformly distributed in the range $[0, 1]$ (200 trials per density value). The results are displayed in Fig. 4. Here, we plot the fitness F and the two components F_d and F_Δ for each density. Moreover, we also plot the encoded density d_{enc} . The figure reveals that the system shows a good behaviour for almost all densities. In average, performance F oscillates in the interval $[0.8, 1.0]$. Moreover, it is possible to observe that the component F_Δ is very high, especially for high densities. This indicates that the group is able to converge to a very regular flashing activity, while for smaller values the period is noisier. However, for small density values the component F_d is higher, revealing that the system performs better in this conditions. The actual abilities of the robotic group can be discussed looking at the encoded density d_{enc} , which is plotted against the ideal case $y = x$. We note that for densities up to $d = 0.6$, the encoded density nicely follows the real one. For larger values, however, a sort of phase transition occurs, in which the robots present a fast signalling behaviour that encodes a density around 0.84.




Figure 5: Excitatory strategy: neural activation and signalling status for two of the ten robots. The bold line indicates the activation of the neuron $N_{O,3}$, which controls the flashing signal. The vertical grey bands indicate the perceived signal $s(t)$. The small dark vertical lines within the grey band indicate the time-step in which the robot itself is signalling. The plot refers to a density $d = 0.2$.

How can the robots produce such behaviour? We answer this question by analysing the behavioural and communication strategy. In this case, robots rotate on the spot without changing position. By observing an

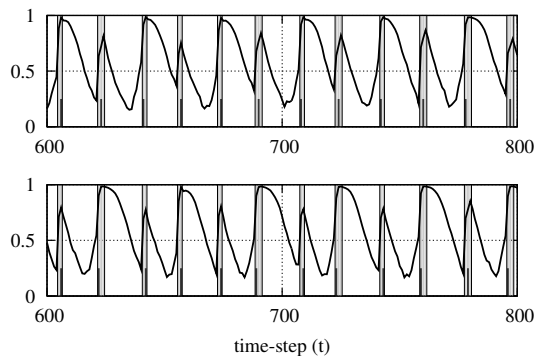


Figure 6: Excitatory strategy: neural activation for $d = 0.8$.

isolated robot, we noticed that the flashing activity is regulated by the locally perceived density: the higher the density, the higher the flashing frequency. However, the individual robot always underestimates the real density, in average. Therefore a collective mechanism must be in place. As mentioned above, in the excitatory strategy the reception of a signal provokes the emission of a signal. The dynamics of the oscillation for a low density $d = 0.2$ can be observed looking at Fig. 5. In correspondence of a perceived signal, the activation o_3 of the neuron controlling the signal output increases until it goes beyond the 0.5 threshold, making the robot itself signal. Upon the sustained perception of a signal, the activation o_3 remains high, therefore delaying the following flash. For instance, in the first signalling event in Fig. 5, the top robot flashes the earliest, and the persistence of the signal afterwards delays the following flash. Instead, if more than one flash is required for o_3 to overcome the signalling threshold, the following flash is anticipated: the bottom plot reveals that in correspondence of a very delayed flash—during the second signalling event—the activation o_3 is just over the threshold and goes immediately down afterwards, allowing the robot to anticipate the following flash. The sequence of perceived flashes functions both as a positive and negative feedback mechanism: robots compete in emitting the first flash, and consequently mutually accelerate their rhythm. This acceleration is however limited by the presence of multiple signals that slow the flashing frequency down. The same mechanism is in place for larger frequencies (see Fig. 6). However, in this case the system converges into a different dynamic regime, in which robots differentiate in two groups that alternately signal. This is evident in the dynamics of the activation o_3 shown in Fig. 6: the asymmetric oscillations indicate that robots engage in a sort of turn-taking, achieving the maximum flashing frequency. This also justifies the phase transition we observed in Fig. 4: for high densities the probability of converging into this fast flashing regime is higher.

In order to further test the hypothesis that robots compete to emit the first signal, we run a series of experiments varying the number of robots in the arena. The results plotted in Fig. 7 show that the average encoded density increases with the number of robots, thus suggesting that robots are able to collectively accelerate their flashing rhythm.




Figure 7: Encoded density for varying group size. Each line represents the average of 100 trials, performed for 100 density values in $[0, 1]$. Data for 1, 2, 3, 4 and 10 robots are shown.

3.2 Analysis of the Inhibitory Strategy

In the case of the inhibitory strategy, we performed similar analyses. The results of the generalisation test are plotted in Fig. 8. It is possible to notice that the system has a very similar performance with respect to the excitatory case: the performance F_d is very high for each density, while F_Δ slightly increases for large d . Therefore, also in this case the group converges towards very regular and precise flashing activity, especially for high densities. Looking at d_{enc} , it is possible to notice that the system presents a phase transition similar to the one discussed for the excitatory strategy.

All these similarities, however, result from radically different mechanisms. As we already mentioned, in this case signals are inhibitory: when a robot perceives a flash, the neural activity o_3 that controls the flashing signal is reset, whatever its value is. This means that there is normally only one robot flashing at any time, that is, the one that reaches the signalling threshold the earliest. This behaviour is evident looking at Fig. 9, in which the dynamics of the neural activity of two different robots are plotted for a density $d = 0.2$: the bottom plot reveals that the corresponding robot flashes the earliest in the first four signalling events, preventing other robots to flash themselves. The situation is similar in the case of $d = 0.8$, shown in Fig. 10, in which we observe that robots compete in order to flash the earliest, similarly to what happens for the excitatory strategy. However, in this case the inhibitory signal does not allow a negative feedback mechanism. In fact, if a robot flashes with

an individual frequency higher than the other robots (e.g., the robot locally perceives a higher density), it would impose its frequency to the group by inhibiting all other robots. If this is the case, the group systematically overestimates the black spot density due to those robots that locally perceive a high value. Therefore there must exist another mechanism that serves as negative feedback to control the frequency of the group. By looking at the behaviour of the robots, we notice that at the beginning of the trial robots slightly move from their initial position while rotating on the spot. This allows robots to explore the neighbourhood for a short time. In order to understand the role of these movements, we tested the robotic system fixing the motor outputs to constant values ($o_1 = 1$ and $o_2 = 0$), forcing the robots to turn on the spot without changing position. The results are shown in

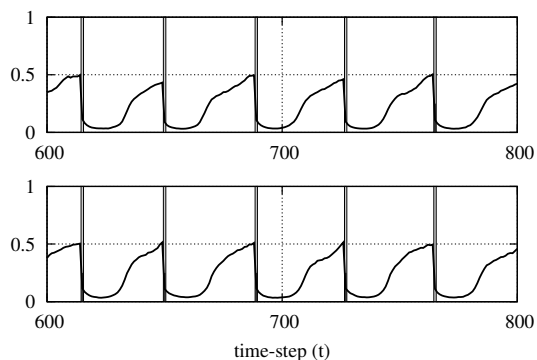


Figure 9: Inhibitory strategy: neural activation for $d = 0.2$.

Fig. 11, for varying density and varying number of robots. From these tests, we infer that the slight motion of the robots is an adaptive mechanism, given that the system without motion performs worse. As predicted, when robots cannot search the neighbourhood of their initial position, they slightly overestimate the density. We also observe that robots exploit the information coming from other robots. In fact, there is no difference between the performance of a single robot when it can and when it cannot move. This means that the motion alone does not allow a single robot to better estimate the global density. We therefore believe that the initial motion of the robots is performed when there are discrepancies between the locally perceived density and the global flashing activity. In other words, a robot moves in search of a local density that corresponds to the globally encoded density. This is brought forth only for a short time at the beginning of the trial. After this short time, the robot stops in place, whatever the local density is. With this mechanism, robots can average out the global density.




Figure 8: Generalisation test for the inhibitory strategy.




Figure 10: Inhibitory strategy: neural activation for $d = 0.8$.




Figure 11: Encoded density for varying group size and for blocked. Each line represents the average of 100 trials, performed for 100 density values in $[0, 1]$. Data for 1, 2, 3, 4 and 10 robots are shown.

4 DISCUSSIONS AND CONCLUSION

In this paper, we have analysed how a group of robots can collectively encode a macroscopic variable that is not accessible to the single individuals in the group. By evolving the collective perception behaviour, we have found two possible strategies that use the communication channel in a opposite way: signals are

either excitatory or inhibitory. In both cases, robots compete in flashing the earliest. In doing so, they share the information gathered locally, allowing to collectively encode an average value close to the actual density.

It is important to remark the fact that, besides the excitatory or inhibitory communicative interaction, a second mechanism is necessary to regulate the activities of the group. On the one hand, this mechanism has been found in the length of the signalling event, which limits robots in producing the first flash for many times consecutively. On the other hand, we observed that robots move from their initial position, therefore exploring the neighbourhood. In both cases, these regulatory mechanisms allow multiple robots to participate in the collective perception in order to have a better estimate of the macroscopic variable. It is therefore possible to identify a general strategy that supports the collective perception in our system: individually, robots encode the local density in a flashing frequency and compete in producing the first flash, which is globally perceived and influences the whole group. At the same time, robots try to hand on the leader role and to listen to the other robots. The balancing of these two tendencies leads to the correct encoding of the global density.

The presence of two counteracting mechanisms that regulate the activity of the group is common to systems as diverse as brains and swarms. The positive feedback loop allows to amplify small perturbations and quickly spread information in a system, while the negative feedback loop controls the competition between different options and modulate the information spreading. The relevance of the negative feedback is recognised in neural systems—in which it is

provided by specialised inhibitory inter-neurons and mediated by glycine and gamma-aminobutyric acid (GABA) transmitters (Jonas and Buzsaki, 2007)—and in super-organisms—in which it may result from specific stop signals issued by some individuals (Nieh, 2010). In our system, we have not provided a specific interaction modality different from the flashing signals. Despite this limitation, evolution could synthesise other mechanisms that resulted in regulatory processes.

In future work, we plan to continue the study of cognitive abilities displayed by collective systems (Trianni and Tuci, 2009). The experimental scenario we have presented here has been conceived to investigate the collective perception and decision making. We plan to study whether groups of robots can select the most dense environment among two or more possibilities presented sequentially or segregated in space. By comparing the results obtained in different artificial setups, we aim at discovering general principles about collective perception and decision making that could be generalised also to the biological reality.

ACKNOWLEDGEMENTS

The authors thank the Institute of Cognitive Sciences and Technology of the Italian National Research Council for having funded the research work presented in this paper. The authors also thank the members of LARAL group for the constructive comments during the early preparation of this research work.

REFERENCES

- Beer, R. D. (1995). A dynamical systems perspective on agent-environment interaction. *Art. Intell.*, 72:173–215.
- Dehaene, S. (2003). The neural basis of the weber-fechner law: a logarithmic mental number line. *Trends in Cognitive Sciences*, 7(4):145–147.
- Detrain, C. and Deneubourg, J.-L. (2006). Self-organized structures in a superorganism: do ants “behave” like molecules? *Physics of Life Reviews*, 3:162–187.
- Grill-Spector, K. (2003). The neural basis of object perception. *Current Opinion in Neurobiology*, 13(2):159–166.
- Hauert, S., Zufferey, J.-C., and Floreano, D. (2009). Evolved swarming without positioning information: an application in aerial communication relay. *Autonomous Robots*, 26(1):21–32.
- Hölldobler, B. and Wilson, E. O. (2008). *The Superorganism: The Beauty, Elegance, and Strangeness of Insect Societies*. W. W. Norton & Company, New York, NY.
- Jonas, P. and Buzsaki, G. (2007). Neural inhibition. *Scholarpedia*, 2(9):3286.
- Leopold, D. A. and Logothetis, N. K. (1996). Activity changes in early visual cortex reflect monkeys’ percept during binocular rivalry. *Nature*, 379:549–553.
- Loffler, G. (2008). Perception of contours and shapes: Low and intermediate stage mechanisms. *Vision Research*, 48:2106–2127.
- Marshall, J. A. R., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., and Franks, N. R. (2009). On optimal decision-making in brains and social insect colonies. *Journal of the Royal Society Interface*, 6(40):1065–1074.
- Nieh, J. C. (2010). Negative feedback signal that is triggered by peril curbs honey bee recruitment. *Current Biology*, 20:310–315.
- Passino, K., Seeley, T., and Visscher, P. (2008). Swarm cognition in honey bees. *Behavioral Ecology and Sociobiology*, 62:401–414.
- Ratcliff, R. and Smith, P. L. (2004). A comparison of sequential sampling models for two-choice reaction time. *Psychological Review*, 111:333–367.
- Romo, R. and Salinas, E. (2003). Flutter discrimination: neural codes, perception, memory and decision making. *Nature Reviews Neuroscience*, 4:203–218.
- Rubin, N. (2003). Binocular rivalry and perceptual multistability. *Trends in Neurosciences*, 26(6):289–291.
- Sperati, V., Trianni, V., and Nolfi, S. (2010). Evolution of self-organised path formation in a swarm of robots. In *Proceedings of the 7th International Conference on Swarm Intelligence (ANTS 2010)*. In press.
- Trianni, V. and Nolfi, S. (2009). Self-organising sync in a robotic swarm. a dynamical system view. *IEEE Transactions on Evolutionary Computation, Special Issue on Swarm Intelligence*, 13(4):722–741.
- Trianni, V. and Tuci, E. (2009). Swarm cognition and artificial life. In *Advances in Artificial Life. Proceedings of the 10th European Conference on Artificial Life (ECAL 2009)*.