

RESEARCH OUTPUTS / RÉSULTATS DE RECHERCHE

On the evolution of mechanisms for collective decision making in a swarm of robots

ALMANSOORI, Ahmed; Mohammed, Muhanad Hayder Mohammed; COLIN, Jean-Noel; Tuci, Elio

Published in:

Proc. of the XV International Workshop on Artificial Life and Evolutionary Computation (WIVACE)

Publication date:

2021

[Link to publication](#)

Citation for pulished version (HARVARD):

ALMANSOORI, A, Mohammed, MHM, COLIN, J-N & Tuci, E 2021, On the evolution of mechanisms for collective decision making in a swarm of robots. in *Proc. of the XV International Workshop on Artificial Life and Evolutionary Computation (WIVACE)*.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal ?

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

On the evolution of mechanisms for collective decision making in a swarm of robots^{*}

Ahmed Almansoori^{1,2}[0000-0002-2600-3344], Muhanad Alkilabi^{1,2}[0000-0002-6921-4993], Jean-Noël Colin¹[0000-0003-4754-7671], and Elio Tuci¹[0000-0001-7345-671X]

¹ University of Namur, Namur, Belgium
{ahamed.almansoori, jean-noel.colin, elio.tuci}@unamur.be
<https://www.unamur.be/info/>

² University of Kerbala, Karbala, Iraq
{muhanad.hayder, Ahmed.kamil}@uokerbala.edu.iq

Abstract. Collective decision-making refers to a process of generating a group decision which cannot be attributed to any agent in the group. In swarm robotics, the individual mechanisms for collective decision making are generally hand-designed and limited to a restricted set of solutions based on the voter or the majority model. This study demonstrates that it is possible to take an alternative approach in which the individual mechanisms are implemented using artificial neural network controllers automatically synthesised using evolutionary computation techniques. We qualitatively describe the group dynamics underpinning the collective process leading to consensus. Moreover, this study demonstrates the evolutionary-tailored mechanisms do not follow the principles of the classic hand-coded solutions.

Keywords: Swarm robotics · Collective decision making · automatic design.

1 Introduction

Swarm robotics is a particular type of multi-robot systems in which each robot has its own controller, perception is local and communication is based on spatial proximity [6]. The swarms designer operates at the individual level, by providing each robot the mechanisms that generate its behaviour. The group-level or swarm response emerges, through a self-organisation process, from the interactions between the robots and their social and physical environment. Due to the distributed and random nature of this self-organisation process, it is notoriously difficult for the designer to predict which set of individual actions leads to the emergence of the desired collective response [2]. In this study, we focus on the design of individual mechanisms for collective decision making; that is,

^{*} A.A. thanks UNamur for the financial support. M.A. thanks the research institute naXys for the financial support.

a decision process in which none of the individuals can be accounted for the group choice once a collective decision is made [18]. In nature, similar processes can be observed in social insects, which collectively choose a foraging or a nesting site without any agent knowing the quality of all available options [3]. The swarm robotics literature has shown that in multiple contexts, a consensus can be reached by swarms in which individuals change opinion by following the rules of the voter model, whereby agents change opinion copying a random neighbour, or of the majority model, whereby agents change opinion to the option held by the majority of a group of neighbours [19, 20, 13, 5, 4]. In this study, we take a complementary approach. Instead of making strong assumptions (e.g., voter or majority model) on the nature of the individual mechanisms underpinning collective decision making processes and showing their effectiveness in specific swarm robotics tasks, we select an already investigated scenario and we exploit the evolutionary robotics methodology to synthesise individual neural mechanisms for options selection. The evolutionary approach is based on the use of evolutionary computation techniques to synthesise artificial neural networks as robots' controller [15]. We chose a scenario in which the robots explored a close arena with the floor made of black and white tiles. The swarm has to reach a consensus on which type (black or white) of tiles covers the majority of the arena floor. This scenario has been originally investigated in [17] where the authors tested the effectiveness of three different mechanisms for options selection (a weighted voter model, the majority model, direct comparison) both with simulated and physical robots. In [14], the authors replicated the study described in [17] to test the effectiveness of a blockchain-based smart contract to protect the collective decision making process from agents (i.e., byzantine robots) acting in order to disrupt the decision process. The blockchain-based layer operates on top of the decision making mechanisms based on the majority rules. In [8], the collective perception scenario described in [17] has been transformed into a multi-feature collective decision making problem, in which the simulated and physical robots reach consensus using hand-designed options selection mechanisms.

To the best of our knowledge, this is the first study demonstrating that the evolutionary approach can be successfully used to synthesise the robots mechanisms in the above described collective perception task in which the group consensus has to be reached with partial individual knowledge and communication limited to the closest neighbour. We provide a preliminary analysis of the evolved collective behaviour showing that if we automate the design of individual mechanisms with a process that evaluates the group performance without imposing specific solutions at the operational level, the evolutionary-tailored mechanisms do not resemble what commonly assumed necessary to reach consensus. In our case, we show that the evolved individual mechanisms for options selection do not follow the rule of the voter model. This result demonstrates that the evolutionary approach generates alternative solutions that may improve the effectiveness of the collective decision making process through individual mechanisms that are task-specific instead of being of task-independent. Moreover, the plasticity, resiliency, and generalisation capabilities of artificial neural networks,

may contribute to improve the robustness and the adaptivity of the swarm with respect to a larger spectrum of sources of environmental variation.

As mentioned above, this is one of the first study using the evolutionary approach to the design of options selection mechanisms for robots of a swarm engaged in a collective decision making scenario. The authors in [12] explored a similar scenario with the evolutionary methods, but with the support of global communication which significantly reduces the complexity of the decision process and its significance for the swarm robotics community. In our scenario, each robot can communicate its opinion only locally to the nearest neighbour and interactions are possible only between agents within 50 cm to each other. To reduce the complexity of the evolutionary task, we take a modular approach in which the movement of the robots is determined by an hand-coded algorithm for pseudo-random walk, while the neural mechanisms for options selection are synthesised using artificial evolution. The modularisation of the control structure does not limit in any way the significance of our results for the swarm robotics community. A detailed analysis of the results of this research work is provided in section 3, while a discussion on the contributions of this study and on the possible developments for future work are illustrated in section 4.

2 Methods



Fig. 1. (a) The simulated collective decision-making scenario. (b) Image of a physical e-puck robot.

We study a collective decision-making scenario that has been described for the first time in [17]. In this scenario, a swarm of 20 robots navigate a close arena of 2x2 m. The arena floor is made of black and white tiles, 10x10 cm each, which are randomly distributed on the arena floor (see Figure 1a). During the design phase, the swarm experiences two types of environment: a) the black-dominant environment in which 52% of the tiles are black and 48% are white; and b) the white-dominant environment in which 52% of the tiles is white and 48% is black. As in [17], the task of the swarm is to reach a consensus on the type of environment on which they are placed. In this study, consensus refers to a state

in which all the robots of the swarm share the same opinion on which tiles' colour covers the largest portion of the arena floor. We have studied two experimental conditions: the comm. experimental condition, which refers to an experimental setup in which the robots can exchange their opinion through communication. To show that consensus can not be reached unless the robots exploit the collective intelligence, we have replicated the study in the no-comm. conditions in which the robots can not communicate their opinion to each other.

During evaluation, the robots move according to a isotropic random walk, with a fixed step length (5 s., at 20 cm/s), and turning angles chosen from a wrapped Cauchy probability distribution characterised by the following PDF:

$$f(\theta, \mu, \rho) = \frac{1}{2\pi} \frac{1 - \rho^2}{1 + \rho^2 - 2\rho \cos(\theta - \mu)}, \quad 0 < \rho < 1, \quad (1)$$

where $\mu = 0$ is the average value of the distribution, and ρ determines the distribution skewness (see [9]). For $\rho = 0$ the distribution becomes uniform and provides no correlation between consecutive movements, while for $\rho = 1$ a Dirac distribution is obtained, corresponding to straight-line motion. In this study $\rho = 0.5$. While moving around, the robots continuously perform an obstacle avoidance behaviour. To perform obstacle avoidance, first a robot detects an obstacle, then stops and keeps on changing its headings of a randomly chosen angle uniformly drawn in $[0, \pi]$ until no obstacles are perceived. Given the robots' pseudo-random walk and the random distribution of black and white tiles on the arena floor, we estimated, by simulating multiple times the task, that each robot explores on average only about 18% of the arena tiles during each evaluation period (i.e., 200 s.). Thus, robots have to develop their opinion based on partial knowledge of the environment and by exploiting the collective intelligence through communication.

The swarm controllers are evolved in a simulation environment which models some of the hardware characteristics of the e-puck2 robots (see Figure 1b), small wheeled cylindrical robots, 70 mm diameter, equipped with a variety of sensors, and whose mobility is ensured by a differential drive system (see [11] for details). In this study, the simulated e-pucks are equipped with infra-red sensors, placed all around the robot's body, the floor sensor, placed underneath the robot chassis. The signal of the infrared sensors is a function of the distance between the robot and any perceived obstacle (in this task, an obstacle can be another robot or the arena walls). The floor sensors return 0 when the centre of the robot base is on a black tile, and 1 when is on a white tile. For the communication, we simply assume that whenever two robots are at less than 50 cm from each other, a 1 bit signal sent by one agent can be perceived by the other. Each robots signals its opinion for the entire duration of the evaluation. The signal sender communicates its current opinion which either black-dominant (the signal sent is 0) or white-dominant (the signal sent is 1). This type of communication can be reliably implemented on physical e-puck2 with the range&bearing board. Concerning the function that updates the position of the robots within the environment, we employed the Differential Drive Kinematics equations, as presented in [7].

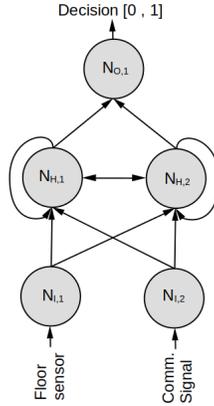


Fig. 2. Architecture of the neural network that control the agents

To compensate for the simulation-reality gap, 10% uniform noise is added to all sensor readings, the motor outputs and the position of the robot [10].

2.1 The robot controller

In the comm. condition, each robot is controlled by a continuous time recurrent neural network (CTRNN) [1]. The neural network has a multi-layer topology, as shown in Figure 2: neurons $N_{I,1}$ and $N_{I,2}$ take input from the robot’s floor sensor and the eventual communication signal (1 for white-dominant, 0 for black-dominant, and 0.5 whenever there is no other robots at less than 50 cm from the receiver), neuron $N_{O,1}$ is used to set the robot opinion, and neurons $N_{H,1}$ and $N_{H,2}$ form a fully recurrent continuous time hidden layer. The input neurons are simple relay units, while the output neuron is governed by the following equations:

$$o = \sigma(O + \beta^O), \tag{2}$$

$$O = \sum_{i=1}^2 W_i^O \sigma(H_i + \beta_i^H), \tag{3}$$

$$\sigma(z) = (1 + e^{-z})^{-1}, \tag{4}$$

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

$$\tau_j \dot{H}_j = -H_j + \sum_{i=1}^2 W_{ij}^H \sigma(H_i + \beta_i^H) + \sum_{i=1}^2 W_{ij}^I I_i, \tag{5}$$

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.

Neuron $N_{O,1}$ is used to set the robot opinion, which corresponds to 1 (i.e., white-dominant) when the neuron firing rate is above the threshold 0.5, and 0 (i.e., black-dominant) otherwise. In the no-comm. condition, $N_{I,2}$ and the connections $\omega_{2,j}^I$ are removed since there is no communication signal.

2.2 The evolutionary algorithm and the fitness function

A simple evolutionary algorithm using linear ranking is employed to set the parameters of the networks. The population contains 100 genotypes. Generations following the first one are produced by a combination of selection with elitism, and mutation. For each new generation, the highest scoring individuals (“the elite”) from the previous generation is retained unchanged. The remainder of the new population is generated by binary tournament selection from the 70 best individuals of the old population. In the comm. condition, each genotype is a vector comprising 15 real values (10 connections, 2 decay constants, 3 bias terms). In the no-comm. condition, each genotype is a vector comprising 13 real values (8 connections, 2 decay constants, 3 bias terms). Initially, a random population of vectors is generated by initialising each component of each genotype to values chosen uniformly random from the range $[0, 1]$. New genotypes, except “the elite”, are produced by applying mutation, which entails a random Gaussian offset applied to each real-valued vector component encoded in the genotype, with a probability of 0.03. 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 $[0, 1]$.

At the beginning of each evaluation trial, each genotype is decoded into a neuro-controller. Then, the controller is cloned on each of the $R = 20$ robots forming the swarm (i.e., we use homogeneous swarms). The robots are randomly placed in the arena with a randomly chosen orientation. Each trial differs from the others in the initialisation of the random number generator, which influences the robots’ initial position and orientation, and the noise added to motors and sensors. Within a trial, the swarm life-span is $200s$ ($T=2000$ simulation cycles). The fitness of a genotype is its average swarm evaluation score after it has been assessed 2 times, in each type of environment (i.e., the black-dominant, and the white-dominant environment) for a total of 4 trials. In each trial e , the opinion of robot r is evaluated at every time t (i.e., O_t^r) for every robot in the swarm. Finally, the swarm is rewarded by an evaluation function F_e which is computed as follows:

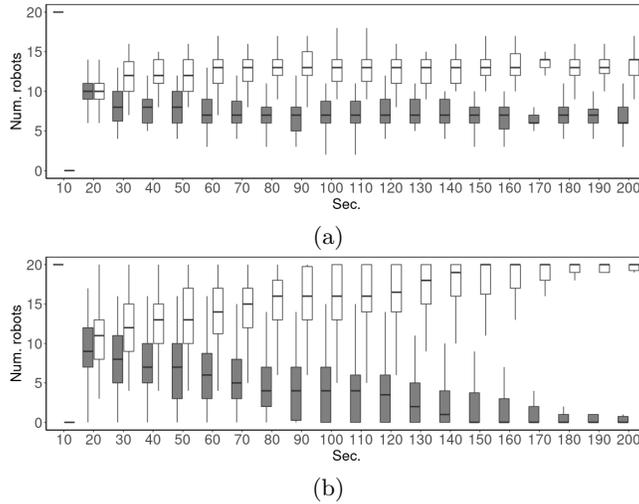


Fig. 3. Box plots showing the number of robots with the opinion white-dominant (white boxes), and those with the black-dominant (grey boxes) every 10 s, during 50 trials in white-dominant environments, for the best group of: (a) the no-comm. condition; and (b) the comm. condition. Each point in the box refers to the group performance in a single trial. Boxes represent the inter-quartile range of the data, while horizontal bars inside the boxes mark the median value. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box.

$$F_e = \begin{cases} \frac{T}{2} \sum_{t=T/2}^T \sum_{r=1}^R o_t^r & \text{if swarm located in a white-dominant env.} \\ \frac{T}{2} \sum_{t=T/2}^T \sum_{r=1}^R (1 - o_t^r) & \text{if swarm located in a black-dominant env.} \end{cases} \quad (6)$$

Note that, within each trial, the fitness score is computed from the simulation cycle $T=1000$ to the end of the trial. This is to exclude from the fitness score the effects of the inevitable fluctuations in the agents’ opinion that are observed at the beginning of each trial when the agents controller is in the “default” initial state.

3 Results

We have performed 20 differently seeded evolutionary simulation runs for each experimental condition. We remind the reader that we have two experimental conditions, the comm. condition where each robot can communicate with the closest robot if its distance is less than 50 cm; and the no-comm. condition where there is no communication between the robots. Each run of each condition lasts 500 generations. In each run, each genotype is evaluated for 4 trials. A trial starts

when the robots are randomly placed within the arena and lasts for 200 s. (i.e., 2000 simulation time-steps). In half of the trials, the robots experience a black-dominant, and in the other half the white-dominant environment. Each trail is an independent event, since at its beginning the robots' controller is reset (see section 2.1 for details). At the end of the evolutionary phase, for each condition, we have re-evaluated the best genotype of each of the last 100 generations of each evolutionary run. In this re-evaluation phase, the performance of each best genotype is estimated on 100 trials (i.e., 50 in a black-dominant, and 50 in a white-dominant environment), using eq 6 as evaluation metrics. In the no-comm. condition, the performances of these group never rose above 80%. This indicates that none of these groups managed to successfully solve this collective decision-making task in totally. In the comm. condition, several genotypes of different evolutionary runs performed very well, with a fitness score very close to the maximum value. This indicates that every run managed to generate groups that can successfully accomplish this task. In the remaining of this section, we show the results of these and of further post-evaluation tests on the very best genotype of each condition. These tests are meant to illustrate some of the characteristics of the best group performance. Although due to shortage of space, we discuss only a single group performance, several successful groups underwent these post-evaluation tests with very similar results. Thus, the results discussed below are representative of the behaviour of several best-evolved groups in each condition.

During the 100 trials of the first post-evaluation test, we observed that all the robots of both best groups of the comm. and of the no-comm. conditions have a genetic predisposition to chose the black-dominant opinion. Thus, whenever these groups experience a black-dominant environment, the robots tend to believe in the correct opinion for the entire duration of the trial, with minimal and sporadic fluctuations away from consensus, characterised by one or maximum two robots that for very short periods change to a white-dominant opinion. The evolution of a genetic bias in binary collective and single-robot decision problems has been observed in other different studies in which the robots are controlled by similar type of neural networks (see [16], for example). When these groups experience a white-dominant environment, the group dynamics becomes more interesting. These dynamics are illustrated in Figure 3 which shows the number of robots with the white-dominant (see white boxes in Figure 3) and with the black-dominant (see grey boxes in Figure 3) opinion every 10 s, during 50 trials in white-dominant environments, for the best group of the no-comm. condition (Figure 3a) and the comm. condition (Figure 3b). In the no-comm. condition (Figure 3a), the robots start the trial with the wrong opinion, due to the genetic bias mentioned above, but just after 20 s. of trial, half of the robots have changed opinion. After 40 s. the group reach a rather stable state in which only 12/13 robots believe in the correct white-dominant opinion. The robots keep on changing opinion during the rest of the trials, but since there are frequent opinion changes in both ways (i.e., from white-dominant to black-dominant and vice-versa) the median remains rather away from consensus. Based on this evidence, we claim that without communication this collective decision-making task

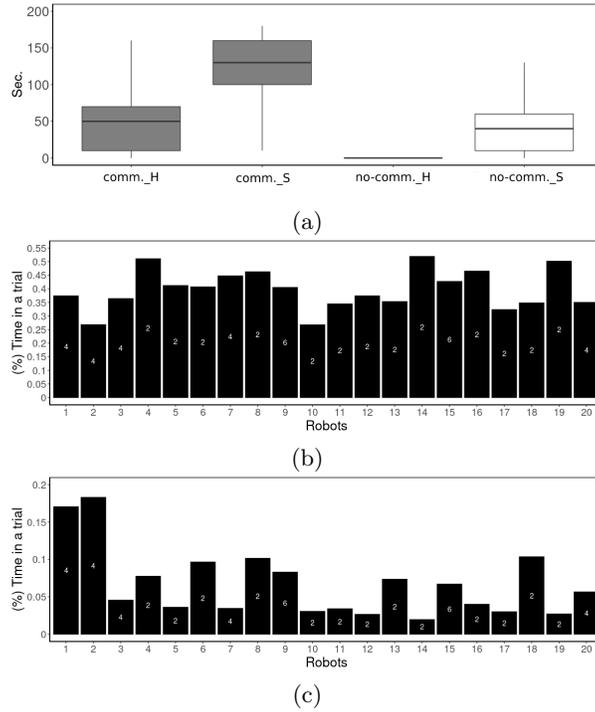


Fig. 4. (a) Box plot showing the length of the longest periods on consensus over 50 trials in white-dominant environments. Grey boxes refer to the best evolved group in the comm. condition, while white boxes refer to the best evolved group in the no-comm. condition. The boxes indicated as S refer to simple environments (i.e., 66% of white and 34% of black tiles); those indicated as H refer to the hard environment (i.e., 52% of white and 48% of black tiles). In (b) and (c), graphs showing for each robot of the best evolved group of the comm. condition, during a trial in a white-dominant environment, the number of times the floor sensor in (b) and the communication sensor in (c) returns a value which contradicts the current robot opinion, during periods in which the robot’s opinion does not change. In both graphs, the numbers inside the bars indicate the number of times the corresponding robot changes opinion within the trial.

cannot be accomplished, since the group can only reach consensus in one type of environment due to the genetic bias. In the comm. condition (Figure 3b), the group dynamics are quite different, since after the wrong start, the group slowly converge to the correct opinion, in spite of the minimal fluctuations from consensus determined by one or two robots that sporadically and only for a short period of time change opinion. In Figure 4, we will further inspect the behaviour of these successful groups in white-dominant environment. This will help us to get a better idea of the groups behaviour in this type of environment. Figure 4a tells us more about the consensus periods. In particular, it shows the longest periods on consensus over 50 trials in white-dominant environments. The grey

boxes refer to the best evolved group in the comm. condition, while white boxes refer to the best evolved group in the no-comm. condition. Each group underwent two sets of 50 trials, the first set in a hard environment (i.e., 52% of white tiles and 48% of black tiles), and the second set in a simple environment (i.e., 66% of white tiles and 34% of black tiles). In Figure 4a, the results referring to hard environment are indicated with H, while those in the simple environment are indicated with S. For the group with communication (see Figure 4a grey boxes), we notice that, for about half of the trials, the length of the longest period on consensus is more than 50 s. in the hard environment (H grey bar) and more than 130 s. in the simple one (S grey bar). This suggests that the group can benefit from easier environmental conditions to remain for a longer time on consensus. For the group without communication (see Figure 4a white boxes), we notice that, there is no consensus at all in the hard environment (H white bar). Nevertheless, the group manages to reach consensus in the simple environment, where the perceptual evidence generated by the floor sensor is sufficient to drive the robot to a common opinion. In Figure 4b, we show data that rules out the hypothesis that each robot of the group evolved in the comm. condition follows the principle of the voter model. In swarm robotics, the voter model corresponds to the behaviour of a robot that changes its opinion anytime a randomly chosen neighbour among those in communication range, disseminates a different opinion. Figure 4b shows, for each robot of the best evolved group of the comm. condition, during a trial in a white-dominant environment, the number of times the communication sensor, returns a value which contradicts the current robot opinion, during periods in which the robot’s opinion does not change. Since this happens multiple time for each robots, we exclude that the robots are following the rules of the voter model. Data in Figure 4c let us exclude also the hypothesis that the robots are simply choosing their opinion based on the perception of the floor colour. Figure 4c shows for each robot of the best evolved group of the comm. condition, during a trial in a white-dominant environment, the number of times the floor sensor returns a value which contradicts the current robot opinion, during periods in which the robot’s opinion does not change. Again, since this happens multiple times for each robots, we can exclude that the robot opinion is simply determined by the colour of the floor. The robots opinion seems to be generated by mechanisms which follow principles slightly more complex than models in which the robots simply react to either the perception of the floor or to the communication signals. The nature of these mechanisms, which successfully drive the group to a correct consensus in both the black-dominant and the white-dominant environment will be investigate in a future study.

4 Conclusions

The first contribution of this study is to provide a proof-of-concept demonstration of the possibility to synthesise, using evolutionary computation techniques, neuro-controllers for a swarm of robots engaged in a collective decision making tasks requiring consensus of opinions among the robots. We showed that a

relatively small dynamic neural network, made of only 5 neurons, is sufficient to generate the mechanisms required by the robots of a homogeneous swarm to collectively decide which colour (between two) is predominant on the arena floor, despite the fact that each robot can only explore a small portion of the floor, and communication is limited both in terms of numbers of different signals received at one time (just one) and in terms of distance (i.e., 50 cm) between sender and receiver. In other words, the neuro-controller supports the development of a collective process which overcomes the individual limitations. This evidence paves the way to an alternative research line which challenges some of the preconceptions of large part of the previous state-of-the-art in swarm robotics by which mechanisms for collective decision making are generally implemented with the voter and/or the majority model. The evolutionary swarm robotics approach, used in this study makes possible to exempt the design from making *a priori* assumptions on the nature of the individual mechanisms required for reaching a consensus in collective decision-making tasks. Indeed, our analysis demonstrated that, in the comm. condition, none of the robots changes its opinion based on the rules of the voter model. We have kept the controllers as small as possible to facilitate the analysis of the operational principles that underpins the individual opinion-making process in this collective task. This will be the subject of future work. Moreover, the evolutionary swarm robotics approach offers the possibility to synthesise mechanisms that are evolutionary tailored to the task requirements, with the possibility to improve the effectiveness and adaptability of the swarm behaviour. This work can be extended in multiple different ways. First, we intend to integrate into the neuro-controller both the mechanisms for the individual opinion-making process and those to generate the robots movement, thus dropping the assumption that robots has to make a specific pseudo-random walk during the exploration of the arena. This will allow us to investigate which individual movements support the collective decision-making process. Second, we would like to study scenarios with more than two options, to study how the task scenario influences the individual behaviour.

References

1. Beer, R.D.: A dynamical systems perspective on agent-environment interaction. *Art. Intell.* **72**, 173–215 (1995)
2. Brambilla, M., Ferrante, E., Birattari, M., Dorigo, M.: Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence* **7**(1), 1–41 (2013)
3. Britton, N., Franks, N., Pratt, S., Seeley, T.: Deciding on a new home: how do honeybees agree? *Proceedings. Biological sciences* **269** **1498**, 1383–1388 (2002)
4. De Masi, G., Prasetyo, J., Tuci, E., E.Ferrante: Zealots attack and the revenge of the commons: quality vs quantity in the best-of-n. In: *Proc of the 12th Int. Conf. on Swarm Intelligence (ANTS)*. pp. 256–268 (2020)
5. De Masi, G., Prasetyo, J., Zakir, R., Mankovskii, N., Tuci, E.E.: Robot swarm democracy: the importance of informed individuals against zealots. *Swarm Intelligence Journal* **In Press** (2021)
6. Dorigo, M., Şahin, E.: Guest editorial. special issue: Swarm robotics. *Autonomous Robots* **17**(2-3), 111–113 (2004)

7. Dudek, G., Jenkin, M.: *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge, UK (2000)
8. Ebert, J., Gauci, M., Nagpal, R.: Multi-feature collective decision making in robot swarms. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. pp. 1711–1719 (2018)
9. Kato, S., Jones, M.: An extended family of circular distributions related to wrapped cauchy distributions via brownian motion. *Bernoulli* **19**(1), 154–171 (2013)
10. Ligot, A., Birattari, M.: Simulation-only experiments to mimic the effects of the reality gap in the automatic design of robot swarms. *Swarm Intelligence* **14**(1), 1–24 (2020)
11. Mondada, F., et al.: The e-puck, a robot designed for education in engineering. In: *Proc. of the 9th Int. Conf. on Autonomous Robot Systems and Competitions*. vol. 1, pp. 59–65 (2009)
12. Morlino, G., Trianni, V., Tuci, E.: Evolution of Collective Perception in a Group of Autonomous Robots, vol. 399, pp. 67–80. Springer (2012)
13. Scheidler, A., Brutschy, A., Ferrante, E., Dorigo, M.: The k -unanimity rule for self-organized decision-making in swarms of robots. *IEEE Transactions on Cybernetics* **46**, 1175–1188 (2016)
14. Strobel, V., Ferrer, E., Dorigo, M.: Managing byzantine robots via blockchain technology in a swarm robotics collective decision making scenario. In: *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS)*. p. 541549. International Foundation for Autonomous Agents and Multiagent Systems (2018)
15. Trianni, V., Ampatzis, E.T.C., Dorigo, M.: *Evolutionary Swarm Robotics: a Theoretical and Methodological Itinerary from Individual Neuro-Controllers to Collective Behaviour*, pp. 153–178. MIT Press (2014)
16. Tuci, E., Quinn, M., Harvey, I.: An evolutionary ecological approach to the study of learning behaviour using a robot-based model. *Adaptive Behavior* **10**(3-4), 201–221 (2002)
17. Valentini, G., Brambilla, M., Hamann, H., Dorigo, M.: Collective perception of environmental features in a robot swarm. In: *International Conference on Swarm Intelligence*. pp. 65–76. Springer (2016)
18. Valentini, G., Ferrante, E., Dorigo, M.: The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI* **4**, 9 (2017). <https://doi.org/10.3389/frobt.2017.00009>
19. Valentini, G., Hamann, H., Dorigo, M.: Self-organized collective decision making: The weighted voter model. In: *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*. p. 4552. International Foundation for Autonomous Agents and Multiagent Systems (2014)
20. Valentini, G., Hamann, H., Dorigo, M.: Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off. In: *Proceedings of the 2015 International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*. p. 13051314. International Foundation for Autonomous Agents and Multiagent Systems (2015)