Further investigations on the characteristics of neural network based opinion selection mechanisms for robotic swarms

ALMANSOORI, Ahmed; Mohammed, Muhanad Hayder Mohammed; Tuci, Elio

Publication date:
2023

Document Version
Other version

Further investigations on the characteristics of neural network based opinion selection mechanisms for robotic swarms

Ahmed Almansoori[1], Muhanad Alkilabi[1], and Elio Tuci[1]

[1]Faculty of Computer Science, University of Namur, Namur, Belgium

February 27, 2023

Abstract

Collective decision-making is a process that allows a group of autonomous agents to make a decision in a way that once the decision is made it cannot be attributed to any agent in the group. In the swarm robotics literature, collective decision-making mechanisms have generally been designed using behaviour-based control structures. That is, the individual decision-making mechanisms are integrated into modular control systems, in which each module concerns a specific behavioural response required by the robots to respond to physical and social stimuli. Recently, an alternative solution has been proposed which is based on the use of dynamical neural networks as individual decision-making mechanisms. This alternative solution proved effective in a perceptual discrimination task under various operating conditions and for swarms that differ in size. In this paper, we further investigate the characteristics of this neural model for opinion selection using three different tests. The first test examines the ability of the neural model to underpin consensus among the swarm members in an environment where all available options have the same quality and cost (i.e., a symmetrical environment). The second test evaluates the neural model with respect to a type of environmental variability related to the spatial distribution of the options. The third test examines the extent to which the neural model is tolerant to the failure of individual components. The results of our simulations show that the neural model allows the swarm to reach consensus in a symmetrical environment, and that it makes the swarm relatively resilient to major sensor failure. We also show that the swarm performance drops in accuracy in those cases in which the perceptual cues are patchily distributed.

1 Introduction

Swarm robotics is a particular type of multi-robot system in which each robot has its own controller, perception is local, and communication is based on spatial proximity [10]. Among the most studied mechanisms in swarm robotics, there are those that allow the swarm to make a decision collectively. Collective decision making refers to a situation in which robots collectively make a choice among two or more alternative options in a way that, when the decision is made, it is no longer attributable to any single individual robot [28]. According to [8], collective decision-making is considered one of the key behaviours in swarm robotics and can be classified into: task allocation and consensus achievement. Task allocation refers to a process by which the group's performance is increased by splitting the swarm members into multiple subgroups, each of which is dedicated to solving a particular task. Consensus achievement refers to a process which allows all the swarm members to share the same opinion with respect to alternative options. When the number of available options is finite, the consensus problem, referred to as the "best-of-n" problem [28], requires the swarm to reach a consensus on the best among the available option, when options differ in qualities, or on any option, when options have equal quality or the same utility for the swarm members [19]. In other words, when confronted with the "best-of-n" type of problems, the swarm must avoid separating into two or more distinct subgroups in which robots of a group share an opinion different from the robots of another group.

In the swarm robotics literature, the mechanisms underpinning collective decision-making have been generally designed using behaviour-based type of control structures. That is, the individual decision-making mechanisms are integrated into modular control systems, in which each module takes care of developing a specific behavioural response required by the robots to respond to physical and social stimuli. Control structures developed with these principles have been demonstrated to be effective in supporting the collective decision-making process in a variety of scenarios [27, 26, 23]. However, the capability of these swarms to adapt to different sources of variability tend to be limited to those circumstances that have been clearly predicted by the designer, leaving the robots potentially unprepared to overcome unexpected and unpredictable events that may occur in any complex natural settings. As stated in [17], further research is needed to design opinion selection mechanisms that fit the needs of swarm robotics systems to allow them to mimic natural swarms in terms of robustness, scalability, and flexibility.

In [1], the authors proposed an alternative approach to the classic hand-designed controller, based on the use of artificial neural networks synthesised using evolutionary computation techniques as individual opinion selection mechanisms [25]. They have tested this design method on a type of best-of-n problem originally described in [26]. This problem is characterised by two or more options whose quality concerns the relative proportions with which they are distributed in the environment. The swarm's task is to reach a consensus on which option is the most represented in the environment. Given that every single robot can

only explore a small portion of the environment, only the swarm, by relying on collective intelligence, can correctly evaluate the options' quality and eventually choose the best one. By investigating the behaviour of a swarm engaged in this collective perceptual discrimination task, they have shown that artificial neural networks synthesised using evolutionary computation techniques can be an effective design method to allow the robots to reach a consensus on the best available options in the environment.

With extensive comparative experimental work, the authors in [2] have shown that the neural network based opinion selection mechanisms (also referred to, in this paper, as the neural model) is more effective than the classic hand-designed Voter model (see [26]) in a set of environmental conditions generated by varying the level of difficulty of the perceptual discrimination task, by varying the maximum distance for robot-robot communication, and also by dynamically varying the option quality when the swarm has already reached a consensus. The study also showed that the performances of a swarm controlled by the neural model are less touched by variations in the swarm size than those of a swarm controlled by the hand-coded Voter model. The neural model is more effective than alternative hand-designed solutions because it allows each robot to integrate physical and social evidence in a more adaptive and effective way than classic hand-designed approaches in which these different sources of information are treated following some designer "imposed" principles (see [2]).

In this paper, we further explore the characteristics of the neural model with respect to its effectiveness in underpinning opinion selection in a swarm of robots engaged in the above mentioned collective perceptual discrimination task. In particular, we describe the results of three different evaluation tests on the best evolved neural model, designed as illustrated in the previous research work (see [2, 1], for details). The first test is related to the ability of the swarm to reach a consensus in a perfectly symmetrical environment where options have the same quality [28]. The capability to break environmental symmetries is an important feature of opinion selection mechanisms because it allows the swarm to keep on operating as a coherent unit and to overcome the individual limitations with the group responses even when feedback modulation based decision processes are not triggered by environmental structures [3, 4]. This is the reason why the symmetry-breaking process has been studied in a different types of best-of-n scenarios, such as the prey-hunting scenario [29], the double-bridge scenario [16], and the aggregation scenario [9, 15, 12, 24, 14, 13]. The second test evaluates the neural model with respect to a type of environmental variability related to the spatial distribution of the options. In particular, we use the benchmarks proposed in [5] to evaluate the neural model in eight different environments in which the options are more patchily distributed than the environment experienced by the swarm during the control system design phase. These different spatial distributions of the perceptual cues have also been used in [6] to test the robustness of decision making mechanisms for a swarm of robots controlled by a statistically grounded algorithm against spatial correlations in an unknown environment. The third set of tests examines the extent to which the neural model is tolerant to the failure of individual components. In par-

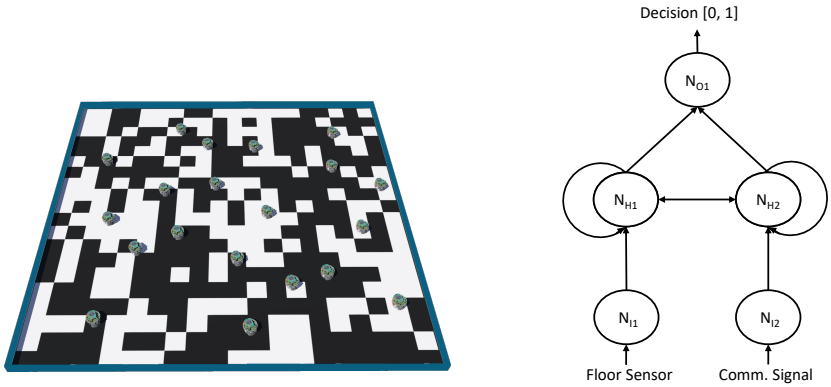

(a) (b)

Figure 1: (a) The simulated arena with the robots engaged on the perceptual discrimination task. (b) The dynamic neural network that underpins the opinion selection in each simulated robot.

ticular, we investigate the robustness of the neural model against failure of the floor sensor in a progressively higher number of robots within the swarm. The floor sensor is of fundamental importance in this task since it allows the robots to evaluate the quality of available options individually.

The results of these tests contribute to generating a more informative estimation of the effectiveness of the neural model within the swarm robotics community. Although our tests are limited to the collective perceptual discrimination task, we show that swarms in which robots are controlled by the neural model can converge to consensus in spite of environmental symmetries. Moreover, they can keep on operating effectively even when more than 40% of the robots suffer from a major sensor failure. As in other opinion selection models, we show that a swarm controlled by the neural model undergoes a significant performance drop in most environments in which the perceptual cues are patchily distributed. The significance of these findings will be further elaborated in the Section 4.

2 Methods

This study is run in a simulation environment which models the wheeled mobile robot e-puck2, a robotic platform commonly used in swarm robotics experiments [22]. The robot sensory apparatus used in this experiment includes eight infra-red sensors to measure the proximity of obstacles, a floor sensor to perceive the colour of the floor in binary format (i.e., the sensor reads 0 if the robot is on a black tile and 1 if on a white tile), and the range&bearing board for local communication. In particular, each robot emits a binary signal which refers to

its opinion about which colour covers the majority of the arena floor (i.e., 1 for white and 0 for black). The maximum robot-robot communication distance is set to 50 cm. This communication system can be reliably implemented on the physical e-puck2 robot with the range&bearing board. The robot movements are computed with a differential drive kinematic model [11]. 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.

The simulation environment is characterised by a close arena of 2x2 m with the floor covered with black and white tiles, 10x10 cm each, distributed randomly on the floor (see Figure 1a). 20 robots are randomly initialised within the arena. During the evaluation, they move randomly, while avoiding obstacles (i.e., the arena walls and other robots) for 400 s, corresponding to the length of an evaluation trial. As in [26], the task of the 20 robots swarm is to reach a consensus (i.e., all robots sharing the same opinion) about which colour (black or white) covers the largest portion of the arena floor. At each simulation update cycle, the robots sample the arena floor underneath their body and listen to the closest neighbour's opinion. After that, they disseminate their current opinion and update their positions. 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 floor during each evaluation period. Thus, a consensus has to be reached by exploiting collective intelligence through local communication for opinion exchanges.

A hand-coded algorithm makes the robots moving within the arena according to an 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:

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

where $\mu = 0$ is the average value of the distribution, and ρ determines the distribution skewness (see [20]). 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 with its infra-red sensors, then stops and keeps on changing its headings of a randomly chosen angle uniformly drawn in $[0, \pi]$ until no obstacles are perceived from the front four sensors.

The process underpinning the development of the individual opinion is regulated by a continuous time recurrent neural network (CTRNN) [7], synthesised using evolutionary computation techniques. The neural network has a multi-layer topology, as shown in Figure 1b: 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 $\sigma\left(H_i + \beta_i\right)$ 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. 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.

The network parameters, that is, the weights of the connections between neurons, the bias terms, and the decay constants, are genetically encoded parameters, set using simple tournament-based selection evolutionary algorithms as illustrated as illustrated in [2, 1]. The swarm is homogeneous; that is, the neural network in charge of the individual opinion selection process is cloned in each of the 20 robots of the swarm.

3 Results

In this section, we illustrate the performances of a swarm of simulated robots engaged in three different tests related to the robustness of collective decision making strategies in a collective perceptual discrimination task. In all these tests, the robots' individual decisions are underpinned by neural network-based mechanisms synthesised using evolutionary computation techniques as described in Section 2.

In all tests, the swarm is evaluated in condition A (the hard scenario), in which the most represented colour covers 55% of the floor, and in condition B

(the simple scenario), in which the most represented colour covers 66% of the arena floor. In test I, the swarm is also evaluated in condition S (the symmetry case), in which each colour covers exactly 50% the arena floor.

Since the operational principles underpinning the individual opinion selection process are not functionally symmetric with respect to the dominant colour, in each test conditions, the swarm undergoes 50 trials in a black-dominant and 50 trials in a white-dominant environment. A trial starts with the 20 robots randomly positioned in the arena and terminates after 400 seconds. During this time, each robot performs a random walk while avoiding obstacles, and it interacts with the other robots using communication signals, as illustrated in Section 2.

To evaluate the group performance, we employ two metrics. The first metric is the *accuracy* of the decision, which corresponds to the proportion of trials (over 50, for each type of environment) in which the swarm reached consensus on the opinion corresponding to the currently most represented colour on the arena floor. In test I, for condition S (the symmetry case), accuracy refers to the proportion of trials (over 50, for each type of environment) in which the swarm reached consensus on any opinion. Consensus refers to the circumstance in which all robots share the same opinion for at least 10 s within a trial. The second metric is the *time*, within a trial, required to the swarm to converge to a consensus state. This metric is calculated on the successful trials only.

3.1 Test I: the symmetry-breaking test

The first test is related to the ability of the swarm to reach a consensus in a perfectly symmetrical environment. As we mentioned above, symmetry-breaking indicates the ability of the swarm to converge on a single shared opinion (instead of multiple opinions) even in those cases in which the alternative options offered by the environment have equal quality [28]. In our collective perceptual discrimination task, the symmetry case corresponds to environments in which the floor is covered by the same proportion of black and white tiles. The results of the symmetry-breaking tests are illustrated in Figure 2, which depicts the accuracy of the decision (see Figure 2a) and the distribution of times required for the swarm to reach a consensus (see Figure 2b) in condition S (i.e., the symmetrical environment). The graphs in Figure 2 also show the performances of the swarm in condition A and B, as terms of comparison. In Figure 2a, for condition S, the black and the white bars refer to the proportion of trials (over 50 trials) in which the swarm reaches a consensus on the black and on the white opinion, respectively. In Figure 2b, for condition S, the black and white boxes refer to the distribution of times to convergence to consensus to the black opinion and to the white opinion, respectively. For condition A and B, the black bars/boxes refer to performances (i.e., accuracy in Figure 2a, and time to convergence in Figure 2b) in black-dominant environments, while the white bars/boxes to performances in white-dominant environments. These results show that, in condition S, the swarm always reaches a consensus (i.e., the black plus the white bar in Figure 2a, condition S, add to accuracy 1). We also

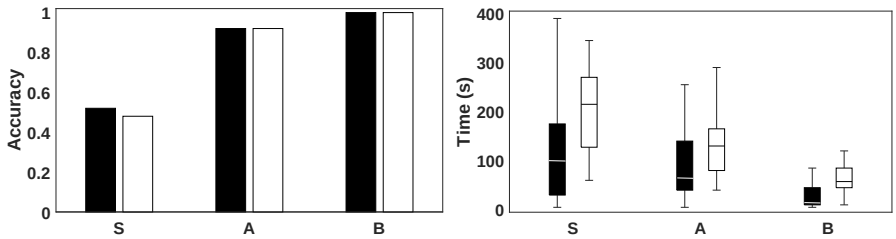

(a) (b)

Figure 2: The symmetry breaking test. In (a), the bars refers to the accuracy of the collective decision, while in (b) the boxes refers to the distribution of times to convergence to consensus. In both graphs, the label S refers to evaluation trials in condition S, that is in symmetrical environments with 50% of black and 50% of white tiles; the A labels refers to evaluation trials in condition A, while the B label refers to trials in condition B. In (a), for condition S, the black and the white bars refer to the proportion of trials (over 50) in which the swarm reach a consensus on the black and on the white opinion, respectively. In (b), for condition S, the black and white boxes refer to the distribution of times to convergence to consensus to the black opinion and to the white opinion, respectively. For condition A and B, the black bars/boxes refer to performances (i.e., accuracy in graph a, and time to convergence in graph b) over 50 trials in black-dominant environments, while the white bars/boxes to performances over 50 trials in white-dominant environments.

note that, due to the stochastic nature of the decision process, the frequency of convergence to opinion black is only marginally higher than the frequency of opinion white. The accuracy in perfectly symmetrical environments is higher than the accuracy in condition A, where the swarm's performance attains about 90% accuracy in both the black-dominant and the white-dominant environment (see Figure 2a, condition A, black and white bars). As already discussed in [2], for progressively simpler perceptual discrimination tasks, the swarm's performance tends to the 100% accuracy (see Figure 2a, condition B, black and white bars). For what concerns the distribution of times to convergence, symmetrical environments are those that require longer time to the swarm to reach a consensus (see Figure 2b, condition S). This finding is in line with previous similar research studies in collective decision making illustrated in [18, 19, 21]. The results of our symmetry-breaking test also show that the swarm requires longer time to reach consensus in the white-dominant than in the black-dominant environment (see Figure 2b, white boxes). As explained in [2], this is due to a genetic bias, by which each robot starts an evaluation trial with opinion white. Paradoxically, this bias delays the convergence to consensus to white. This is due to the fact that the swarm has to go through a series of global states in which the initial genetically-induced consensus to white is progressively lost, and subsequently recovered through a genuine collective decision process triggered

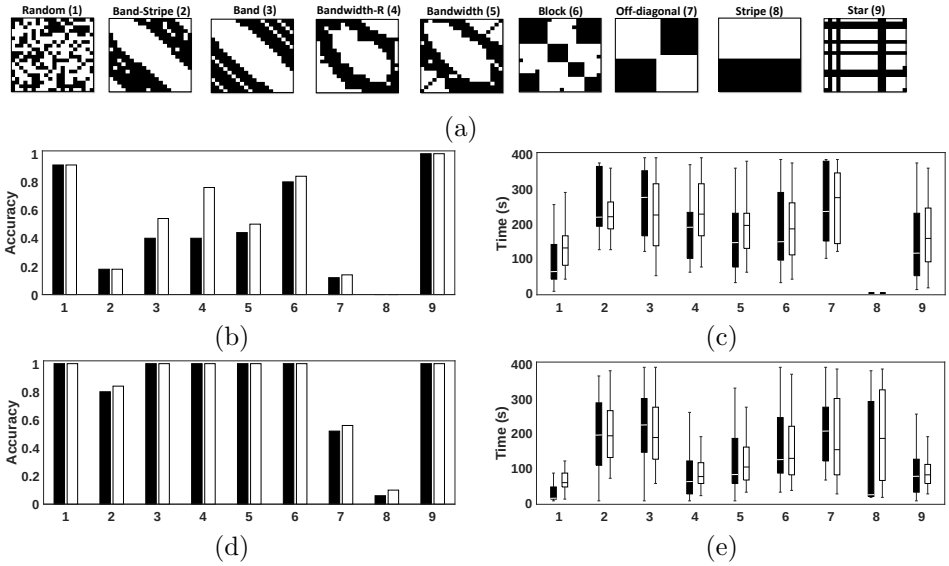

Figure 3: Tiles distribution test. (a) The nine floor patterns. In the first patterns on the left, black and white tiles are randomly distributed. In all other patterns, the tiles are distributed as illustrated in [5]. Graphs in (b) and (d) show the accuracy of the collective decision for each pattern in condition A and in condition B, respectively. Graphs in (c) and (e) show the distribution of times to convergence to consensus for each pattern in condition A and in condition B, respectively. In (b), (c), (d), and (e) the black bars/boxes refers to the performances in black-dominant environment, while the white bar/boxes refers to the performances in white-dominant environment.

by the interactions between the robots (see [2], for details). In summary, the symmetry-breaking test tells us that the neural network based decision making mechanisms are extremely effective in underpinning consensus among the members of the swarm also in a perfectly symmetrical environments. However, the symmetry case requires longer times to the swarm to reach consensus than the non-symmetrical cases. The distribution of times to consensus become progressively shorter in progressively simpler perceptual discrimination tasks (i.e., those in which the proportion of floor taken by the dominant colour is progressively larger).

3.2 Test II: the tiles distribution test

The second test is related to the effectiveness of the swarm to reach consensus in environments in which the perceptual evidence related to the different options is not randomly distributed. Black and white tiles are arranged in order to form specific patterns which may influence the capability of the swarm to reach consensus given the fact that each robot can only explore a limited portion of

the arena floor. The patters we used for this test are those originally illustrated in [5] and shown in Figure 3a.

We remind the reader that pattern 1 (Random) is the one used during the design phase of the neural model. The graphs in Figure 3b, and 3d, show the accuracy of the collective decision for each pattern in condition A and in condition B, respectively. In condition A (see Figure 3b), it turns out that the characteristics of each pattern have a substantial effect on the accuracy of the swarm's collection decision. In general, the swarm performance tend to degrade when the perceptual evidence is spatially arranged in distinctive clusters or patches (see patterns 2, 3, 4, 5, 7, 8 in Figure 3b). The less patchy the distribution of perceptual evidence, the higher the swarm accuracy. This can be accounted for by considering that patchy environments facilitate the emergence of an alignment of opinion among spatially proximal robots randomly wandering within specific clusters. The local alignment on different opinions, supported by the predominant local perceptual evidence hinders the swarm to reach the consensus state. In condition B (see Figure 3d), where the dominant colour covers 66% of the arena floor, the influence of the patchy distribution of the perceptual evidence has smaller impact on the swarm accuracy except for pattern 7 and 8 where a clear performance drop is observed. Moreover, for both conditions, there is a substantial similarity between the swarm performances in the black-dominant and in the white-dominant environment except for pattern 4 in condition A, where the swarm does better in the white-dominant than in the black-dominant environment. Regarding the distribution of times to reach consensus (see Figure 3c, and 3e) it is worth noticing that any pattern requires longer time to the swarm to reach consensus than the times recorded for patterns 1. This phenomenon is more evident in condition A (see Figure 3c) than in condition B (see Figure 3e).

3.3 Test III: robots' floor sensor readings failure

The third test focuses the robustness of the collective decision-making process under conditions in which a progressively higher number of robots within the swarm suffers from failures of the floor sensor, which instead of correctly reading the colour of the floor, it returns randomly generated binary values. The graphs in Figure 4 show the swarm performances (i.e., accuracy of the group decision, see Figure 4a, and 4c, and time to convergence to consensus, see Figure 4b, and 4d) for different numbers of robots (from 0 to 10) suffering from failures of the floor sensor. Figure 4a, and 4b refer to condition A, while Figure 4c, and 4d refer to condition B. We recall the reader that the swarm size is 20. The results of these tests unequivocally indicate that the collective decision-making progress is relatively robust to this type of disruption. In condition B, the swarm attains 100% accuracy in both types of environment even when 50% of the swarm suffers from failure of the floor sensor (see Figure 4c). In condition A, only when more than 35% of the swarm suffers from floor sensor failure (see Figure 4a, for 7 robots), the accuracy starts to progressively drop only for the white-dominant environment. For the black-dominant environment, the swarm

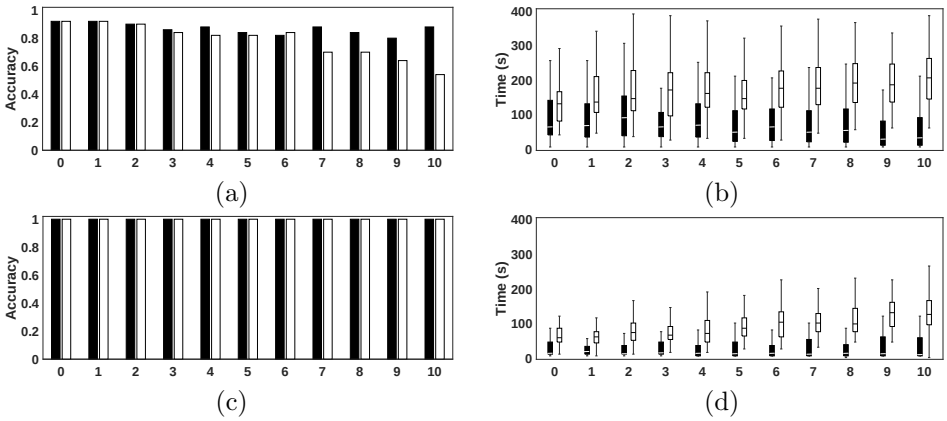

Figure 4: Floor sensor readings failure test. Graphs in (a) and (c) show the accuracy of the collective decision for each pattern in condition A and in condition B, respectively. Graphs in (b) and (d) show the distribution of times to convergence to consensus for each pattern in condition A and in condition B, respectively. In all graphs, black bars/boxes refers to performances in black-dominant environments, while white bars/boxes to performances in white-dominant environments; the labels on the x-axes indicate the number of robots affected by the floor sensor failure.

accuracy tend to remain above 80% even with 10 robots suffering from the floor sensor failure (see Figure 4a. Regarding time to convergence to consensus, the results of the test clearly indicate that the higher the number of robots with floor sensor failure the longer the time to reach consensus in white-dominant environments. This trend is observed both in condition A and in condition B (see Figure 4b, and 4c).

4 Conclusion

In this paper, we further investigated the characteristics of the neural model described in [1, 2] with respect to its effectiveness in underpinning consensus in a swarm of robots engaged in a collective perceptual discrimination task. In particular, we run three tests. The first test focused on the swarm's ability to establish a consensus in a symmetrical environment where all alternatives are of equal quality. The findings of this test indicate that the decision-making processes based on neural networks are exceptionally successful at supporting consensus among the swarm members, even in situations with perfect symmetry. As expected, it takes longer for the swarm to reach consensus in the symmetrical than in the non-symmetrical environments. The third set of tests examined the extent to which the neural model is tolerant to the failure of individual components. More precisely, we investigated the neural model's ro-

bustness against floor sensor's failure in a progressively higher number of robots within the swarm. The results of this test unequivocally indicated that the collective decision-making process is relatively robust to this type of disruption. In condition B, the swarm attains 100% accuracy even when 50% of the swarm suffers from a failure of the floor sensor. In the condition A, where the dominant colour covers 55% of the arena floor, the accuracy remains over 80% even when more than 35% of the swarm suffered from a floor sensor failure.

The second test evaluated the neural model with respect to a type of environmental variability related to the spatial distribution of the options. We employed the benchmarks proposed in [5] to evaluate the neural model in eight different environments in which the options are more patchily distributed than the environment experienced by the swarm during the design phase. This test showed that the swarm's performance tends to degrade when the perceptual evidence is spatially arranged in distinctive clusters or patches. The less patchy the distribution of perceptual evidence, the higher the swarm accuracy. This can be accounted for by considering that patchy patterns facilitate the emergence of an alignment of opinion among spatially proximal robots randomly wandering within specific clusters. The local alignment on different opinions, supported by the predominant local perceptual evidence, hinders the swarm from reaching a consensus state. In the condition B, where the dominant colour covers 66% of the arena floor, the influence of the patchy distribution of the perceptual evidence has a smaller impact on the swarm's accuracy in most studied patterns. There are several elements on which we plan to act in the future, to try to overcome this limitation. One idea is to develop exploration strategies that adapt to the characteristics of the environment. For example, instead of moving with the same pseudo-random walk as in this study, the robots could adaptively mix Lévy flight type of random walk with Brownian motion to sample more distant portion of the environment. Alternatively, we are planning to integrate into the neural model the mechanisms underpinning the robots movements. This would allow us to exploit the adaptivity of the neural model also with respect to the exploration strategies. More effective swarm decision strategies in patchy environments can also come from an increased flexibility of the individual mechanisms in charge of mixing perceptual cues and social evidence.

References

[1] A. Almansoori, M. Alkilabi, J-N. Colin, and E. Tuci. 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*, Communications in Computer and Information Science (CCIS). Springer Verlag, 2021. In Press. Preprints available at: https://drive.google.com/file/d/168s-I8d4f8x26TnyxwYeLuFSpazy8wx/view?usp=sharing.

[2] Ahmed Almansoori, Muhanad Alkilabi, and Elio Tuci. A comparative study

on decision making mechanisms in a simulated swarm of robots. In *2022 IEEE Congress on Evolutionary Computation (CEC)*, pages 1–8. IEEE, 2022.

[3] J-L. Amé, J. Halloy, C. Rivault, C. Detrain, and J-L. Deneubourg. Collegial decision making based on social amplification leads to optimal group formation. *PNAS*, 2006.

[4] J-L. Amé, C. Rivault, and J-L. Deneubourg. Cockroach aggregation based on strain odour recognition. *Animal Behaviour*, 2004.

[5] Palina Bartashevich and Sanaz Mostaghim. Benchmarking collective perception: New task difficulty metrics for collective decision-making. In *EPIA Conference on Artificial Intelligence*, pages 699–711. Springer, 2019.

[6] Palina Bartashevich and Sanaz Mostaghim. Multi-featured collective perception with evidence theory: tackling spatial correlations. *Swarm Intelligence*, 15(1):83–110, 2021.

[7] R. D. Beer. A dynamical systems perspective on agent-environment interaction. *Art. Intell.*, 72:173–215, 1995.

[8] Manuele Brambilla, Eliseo Ferrante, Mauro Birattari, and Marco Dorigo. Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence*, 7(1):1–41, 2013.

[9] A. Campo, S. Garnier, O. Dédriche, M. Zekkri, and M. Dorigo. Self-organized discrimination of resources. *PLoS ONE*, 6(5):e19888, 2010.

[10] M. Dorigo and E. Şahin. Guest editorial. special issue: Swarm robotics. *Autonomous Robots*, 17(2-3):111–113, 2004.

[11] G. Dudek and M. Jenkin. *Computational Principles of Mobile Robotics*. Cambridge University Press, Cambridge, UK, 2000.

[12] Z. Firat, E. Ferrante, Y. Gillet, and E. Tuci. On self-organised aggregation dynamics in swarms of robots with informed robots. *Neural Computing and Applications*, 32(17):13825–13841, 2020.

[13] Gianpiero Francesca, Manuele Brambilla, Arne Brutschy, Vito Trianni, and Mauro Birattari. Automode: A novel approach to the automatic design of control software for robot swarms. *Swarm Intelligence*, 8(2):89–112, 2014.

[14] Gianpiero Francesca, Manuele Brambilla, Vito Trianni, Marco Dorigo, and Mauro Birattari. Analysing an evolved robotic behaviour using a biological model of collegial decision making. In *International Conference on Simulation of Adaptive Behavior*, pages 381–390. Springer, 2012.

[15] Simon Garnier, Jacques Gautrais, Masoud Asadpour, Christian Jost, and Guy Theraulaz. Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior*, 17(2):109–133, 2009.

[16] Simon Garnier, Faben Tache, Maud Combe, Anne Grimal, and Guy Theraulaz. Alice in pheromone land: An experimental setup for the study of ant-like robots. In *2007 IEEE swarm intelligence symposium*, pages 37–44. IEEE, 2007.

[17] H. Hamann. *Swarm Robotics: A Formal Approach*. Springer Cham, 2018.

[18] Heiko Hamann, Bernd Meyer, Thomas Schmickl, and Karl Crailsheim. A model of symmetry breaking in collective decision-making. In *International Conference on Simulation of Adaptive Behavior*, pages 639–648. Springer, 2010.

[19] Heiko Hamann, Thomas Schmickl, Heinz Wörn, and Karl Crailsheim. Analysis of emergent symmetry breaking in collective decision making. *Neural Computing and Applications*, 21(2):207–218, 2012.

[20] S. Kato and M. Jones. An extended family of circular distributions related to wrapped cauchy distributions via brownian motion. *Bernoulli*, 19(1):154–171, 2013.

[21] Yara Khaluf. The emergence of collective response to decisions in a group of physical agents. In *ALIFE 2021: The 2021 Conference on Artificial Life*. MIT Press, 2021.

[22] F. Mondada and 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*, volume 1, pages 59–65, 2009.

[23] A. Scheidler, A. Brutschy, E. Ferrante, and M. Dorigo. The k-unanimity rule for self-organized decision-making in swarms of robots. *IEEE Transactions on Cybernetics*, 46:1175–1188, 2016.

[24] A. Sion, A. Reina, M. Birattari, and E. Tuci. Controlling robot swarm aggregation through aminority ofinformed robots. In M. Dorigo, H. Hamann, M. López-Ibáñez, J. García-Nieto, A. Engelbrecht, C. Pinciroli, V. Strobel, and C. Camacho-Villalón, editors, *Swarm Intelligence*, pages 91–103, Cham, 2022. Springer International Publishing.

[25] V. Trianni, E. Tuci, C. Ampatzis, and M. Dorigo. Evolutionary swarm robotics: a theoretical and methodological itinerary from individual neurocontrollers to collective behaviours. In P. A. Vargas, E. Di Paolo, I. Harvey, and P. Husbands, editors, *The Horizons of Evolutionary Robotics*, pages 153–178. MIT Press, Cambridge, MA, 2014.

[26] G. Valentini, M. Brambilla, H. Hamann, and M. Dorigo. Collective perception of environmental features in a robot swarm. In *International Conference on Swarm Intelligence*, pages 65–76. Springer, 2016.

[27] G. Valentini, H. Hamann, and M. Dorigo. 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)*, pages 1305–1314. International Foundation for Autonomous Agents and Multiagent Systems, 2015.

[28] Gabriele Valentini, Eliseo Ferrante, and Marco Dorigo. The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. *Frontiers in Robotics and AI*, 4:9, 2017.

[29] Jan Wessnitzer and Chris Melhuish. Collective decision-making and behaviour transitions in distributed ad hoc wireless networks of mobile robots: Target-hunting. In *European Conference on Artificial Life*, pages 893–902. Springer, 2003.