International Journal of Computing and Digital Systems
Int. J. Com. Dig. Sys. 16, No.1 (Sep-24)

http://dx.doi.org/10.12785/ijcds/160189

Enhancing Robustness of Swarm Robotics Systems in a Perceptual Discrimination Task

Rusul Ibrahim[1], Muhanad Alkilabi[1,2,3], Ali Retha Hasoon Khayeat[1] and Elio Tuci[3]

[1]*Department of Computer Science, University of Kerbala, Iraq*
[2]*Department of Medical Instruments Techniques Engineering, AlSafwa University College, Iraq*
[3]*Department of Computer Science, University of Namur, Belgium*

Received 15 Feb. 2024, Revised 17 May 2024, Accepted 26 May 2024, Published 7 Sep. 2024

Abstract: The automation of tasks such as environmental monitoring, toxin detection, and mineral resource identification requires artificial agents with perceptual discrimination capabilities to identify the predominant features in environments much larger than their sensing range. The key challenge is developing collective decision-making methods that allow agents to predict a global perspective of the environment from local observations. Our research explores the leverage of collective decision-making for binary perceptual discrimination tasks, using evolutionary computation techniques to synthesise an artificial neural network controller. We focus on strategies that generalised better to patchy and clustered feature distribution environments. We investigate three communication strategies - *close-neighbour*, *rand-neighbour*, and *far-neighbour*- in which robots exchange opinions about the dominant colour of the environment based on the distance between sender and receiver robots. The results show that the *rand-neighbour* strategy significantly improves performance, particularly in unseen patchy patterns. The extensive analysis of the communication dynamics among the robots indicates that the effectiveness of *rand-neighbour* strategy is attributed to its efficient circulation of opinions among both close and distant robots. Our findings support the hypothesis that primordial communication between one receiver robot and a randomly chosen emitter robot is sufficient to develop an effective collective decision-making strategy for swarm of robots engages in perceptual discrimination tasks.

Keywords: Swarm robotics, Evolutionary robotics, Collective decision-making, Communication strategies

1. Introduction

In environments where the size greatly exceeds the sensing capabilities of individual agents, there is a significant risk of inaccurate evaluations, leading to incorrect decisions. Tasks such as environmental monitoring, toxin detection, and mineral resource identification require artificial agents with perceptual discrimination abilities to identify key features in unknown environments. Employing multiple agents to cover a larger area can make possible to gather more comprehensive information about the quality of environmental features or options.

Nevertheless, effective mechanisms are needed to allow the group of agents to make autonomous collective decisions. Swarm robotics is the research domain that tries to identify the individual mechanisms underpinning collective decision-making as well as other complex collective responses. Generally speaking, swarm robotics systems address tasks that require collaborative efforts of a large number of agents interacting with each other to solve complex and extensive problems that would be otherwise impossible for a single agent to handle. Inspired to the behaviour of social insects, the distinctive characteristics of swarm robotics is self-organisation, distributed control, and local sensing, which endow the swarm with a higher level of fault tolerance, scalability, and adaptability to environmental disturbances [1].

The design methods in swarm robotics require roboticists to identify individual behaviours that generate the swarm desired collective response [2]. However, this is a particularly challenging design problem, since the collective response is a phenomenon that emerges from complex and difficult-to-predict dynamics involving both robot-robot and robot-environment interactions. This design problem can be found in the study of many swarm responses, including those requiring collective decision-making, which is the process of choosing an option between those available by the swarm of robots in a collective way. The characteristics of collective decision-making is that once a consensus is reached, it cannot be attributed to any specific member of the swarm. Rather, it emerges from the complex spatial and

E-mail address: {*rusul.i, muhanad.hayder, ali.r*}*@uokerbala.edu.iq, elio.tuci@unamur.be*

temporal interactions involved in the opinion exchange process among swarm members [3]. In the literature regarding swarm robotics, collective decision-making techniques are generally investigated in two-options scenarios in which the swarm robots have to find a consensus on the best option between the two available. This type of scenarios is generally referred to as best-of-*n* (where the number of available options is two) problem [4], [5]. A specific type of best-of-2 decision-making problem is a perceptual discrimination task where the two options are spread within the environment, with the better quality option associated to the one that appears in a larger quantity than the alternative one. Since the perceptual capabilities of the robots are limited, a consensus on the best-quality option can be achieved only through a collective decision process in which robots explore different areas of the environment, and interact in order to "integrate" their perceptual experiences to converge to a common opinion on the best option.

One effective method to design individual control mechanisms underpinning collective decision-making in perceptual discrimination tasks is the hand-coded approach. In this methodology, designers meticulously craft individual mechanisms that drive the collective response to the problem at hand. The literature has demonstrated the effectiveness of the hand-coded approach, especially when adopting the Voter model principles, where robots modify their opinions based on selecting a random neighbour [6], [7], or the majority rule, where the opinion aligns with the opinion that supported by the majority of a group of neighbours [8], [9]. However, the design of hand-coded controls often relies on strong assumptions made by designers regarding how the problem should be addressed. These assumptions can limit the ability of the swarm to exploit subtle irregularities in physical and social perceptual cues, which could otherwise enhance the collective decision-making process [10]. Recent research has highlighted weaknesses in the hand-coded approach, particularly its adaptability in dynamic environments where the optimal option changes over time [11], [9].

Recently, an alternative design approach based on evolutionary robotics (ER) has been introduced [12]. In this approach, the decision making unit generating the agent's opinion is an artificial neural network synthesised using evolutionary computation techniques [13]. A notable feature of the ER approach is the automation of the design process, which significantly reduces the influence of designer assumptions. Recent research work [14] provides evidence that the ER approach outperformed the hand-coded approach with respect to robustness, scalability, and adaptability of the collective response of the group.

The study illustrated in [15] highlights that the challenge in perceptual discrimination tasks lies not only in the magnitude of the difference between the quality of the two options, but also in their distribution patterns. This result has been found in a type of perceptual discrimination task in which the options are two colours covering the floor of an arena that the robots explore with a random movement, and the quality refers to the proportion of floor covered by each option. In particular, the authors focused on nine benchmark environments with varying feature distribution patterns, as shown in Figure 1. The results achieved by this study show that, regardless of the difference in quality, groups designed to perform optimally in the Random type of environments (see Figure 1, Random) experience a performance drop when they are post-evaluated in the Off-diagonal and Stripe environment (see Figure 1, Off-diagonal, and Stripe). This observation has been corroborated by other recent studies [12], [16], [17] which report the same type of performance drop in spite of the fact that they employ artificial neural network as robots controllers to improve the robustness of the collective response.

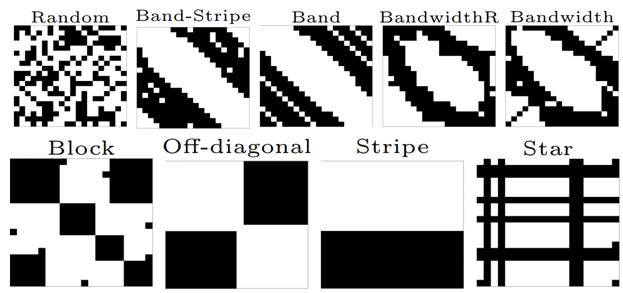

Figure 1. Images of the nine floor patterns used in this experiment. The Random is the floor pattern experienced by the robots during the design phase. The other eight floor patterns, originally introduced in [15], are used as robustness test of the decision-making mechanisms used by the robots to perform this binary perceptual discrimination task.

The main focus of this study is to overcome the limitations illustrated in [15], [12], [16], [17] by developing individual decision-making mechanisms underpinning a collective response that allow a swarm of robots to perform sufficiently well in all the nine types of floor distribution patterns illustrated in Figure 1. In order to achieve this objective, we focus on multiple elements such as the type of individual random walk used by the robots to explore the arena, the structure of the neuro-controller, as well as on the characteristics of the communication strategy used to exchange individuals opinions. We found out that, this latter element is the one that allowed us to achieve an important improvement in terms of robustness of the collective decision with observing specific floor patterns. Particularly, we found out that only when the communication events happens between a robot receiver and a randomly chosen (rather than the closest as in [15], [12], [16], [17]) emitter among those within communication range, no performance drop is observed while moving from Random to all the other nine floor patterns. We show that the superior robustness observed in group in which the communication happens between a robot receiver and a randomly chosen emitter can be attributed to a more effective circulation of opinions among both the spatially close and distant robots, thereby maintaining a high accuracy rate in the decision process throughout the eight floor patterns not experienced by the robots during the design phase. We would like to bring to the reader's attention that this study is an extension of our previous study in [18], where we developed an effective collective decision-making strategy for a group of 20 e-puck robots. However, this research extends [18] by focusing on enhancing communications strategies to improve group performance that generalised better to environments with patchy, clustered feature distribution.

2. Methods

The task the robots are required to perform in this experiment is a binary perceptual discrimination problem. The robots have to collectively choose which colour covers the largest proportion of a closed square arena (200×200 cm), tiled with black and white 10×10 cm tiles. We consider two scenarios: i) a simple scenario (hereafter, referred to as *S-env*), in which the difference in the proportion of black and white tiles is relatively large, since 66% of the arena floor is covered by one colour (the dominant one), while the other colour covers the remaining 34% of the arena floor; ii) the hard scenario (hereafter, referred to as *H-env*), in which the ratio difference of black and white tiles is smaller than in *S-env*, since one colour (the dominant one) covers 55% of the arena floor, while the other colour covers the remaining 45% of the arena floor. For both the *S-env* and the *H-env* scenario, the robots experience both environments: black dominant environment (hereafter, referred to as *BD-env*), and white dominant environment (hereafter, referred to as *WD-env*).

Random positions and orientations are initially chosen to place a swarm of 20 robots in the arena (see Figure 2a). The robots have to explore the arena and reach consensus on the option with the best quality (i.e., choosing the currently dominant colour) over a period of 400 seconds. While exploring the arena, the robots can communicate with spatially proximal neighbours their current opinion. Consensus to the correct option is attained whenever all the 20 robots shared the same correct opinion about which colour is currently dominant for at least 10 s.

Our simulation model the e-puck robot [19], a popular miniature robot commonly utilised in swarm robotics. The simulated robot is equipped with a floor sensor for binary colour detection (0 for black and 1 for white) and eight infrared sensors for obstacle detection. The robots communicate using Range and Bearing sensors, with the communication range limited to 50 cm. To bridge the gap between the simulation and reality, a uniform noise of 10% is appended to all motor outputs, sensor readings, robot positions, and orientations.

The robot's exploration of the environment is based on ballistic motion [20], a variant of random walk used in robotic swarm mapping. This movement pattern involves the robot travelling in a straight line (ballistic trajectory) until encountering an obstacle (other robots or the arena wall), at which point it randomly changes direction. Ballistic motion has proven to be effective in exploring enclosed environments [21]. Figure 2b illustrates the finite state machine controlling the robot's movement.

The robot's decision-making process is controlled by a Continuous-Time Recurrent Neural Network (CTRNN) [22], optimised using artificial evolutionary techniques. The CTRNN comprises 2 sensor neurons, 4 internal neurons, and 1 output neuron representing the robot's opinion. The topology of the CTRNN is depicted in Figure 2c. The network input includes the readings of the floor sensor and the received communication signal from a randomly selected neighbour chosen from those located at a distance less than 50 cm away from the receiver. The network outputs a binary value where 1 corresponds to the opinion that the dominant colour is white, and 0 corresponds to the opinion that the dominant colour is black. This binary value, corresponding to the current robot's opinion, is communicated among spatially proximal robots, as mentioned above. If, for a robot receiver there are no neighbouring robots within communication distance (i.e., < 50 cm), the reading of the receiver's sensor neuron for communication is set to 0.5. During communication, only one neighbour's opinion is selected from the set of available neighbours.

Equations 1, 2, and 3 illustrate how the sensory, internal, and opinion neurons are updated at every simulation cycle.

$$y_k = gI_k;\; k \in \{1, ..., M\}; \quad \text{with } M = 2; \tag{1}$$

$$\tau_k \dot{y}_k = -y_k + \sum_{j=1}^{j=M+4} \omega_{jk}\sigma(y_k + \beta_j);\; k \in \{M{+}1, .., M{+}4\} \tag{2}$$

$$y_k = \sum_{j=M+1}^{j=M+4} \omega_{jk}\sigma(y_j + \beta_j);\; k \in \{M + 5\}; \tag{3}$$

with $\sigma(x) = (1 + e^{-x})^{-1}$. These equations incorporate terms reminiscent of real neuron functions: cell potential is denoted by y_k, the decay constant is τ_k, g represents the gain factor, and I_k is the activation of the k^{th} sensor neuron where $k = 1, .., M$ (refer to Figure. 2c for mapping between sensor neurons and their corresponding sensors), ω_{kj} the synaptic connection strength from neuron j to neuron k, the bias term β_j, the firing rate $\sigma(y_j + \beta_j)$. The same bias (β_I) for all sensory neurons, and the same holds for opinion neuron (β_O). The genetically specified network parameters include: τ_k and β_k of internal neurons, β_I, β_O, ω_{kj} represents the weights of all the network connections, and g. When the network is initiated or reset, the cell potentials reset to 0. For integrating equation 2, the forward Euler method is employed with an integration time step of $\Delta T = 0.1$

A simple evolutionary algorithm that uses tournament-based selection, as illustrated in [12] is used to optimise the networks' parameters. This population comprises 64 genotypes. New generations emerge from a mixture of elitist selection, recombination, and mutation. Each generation

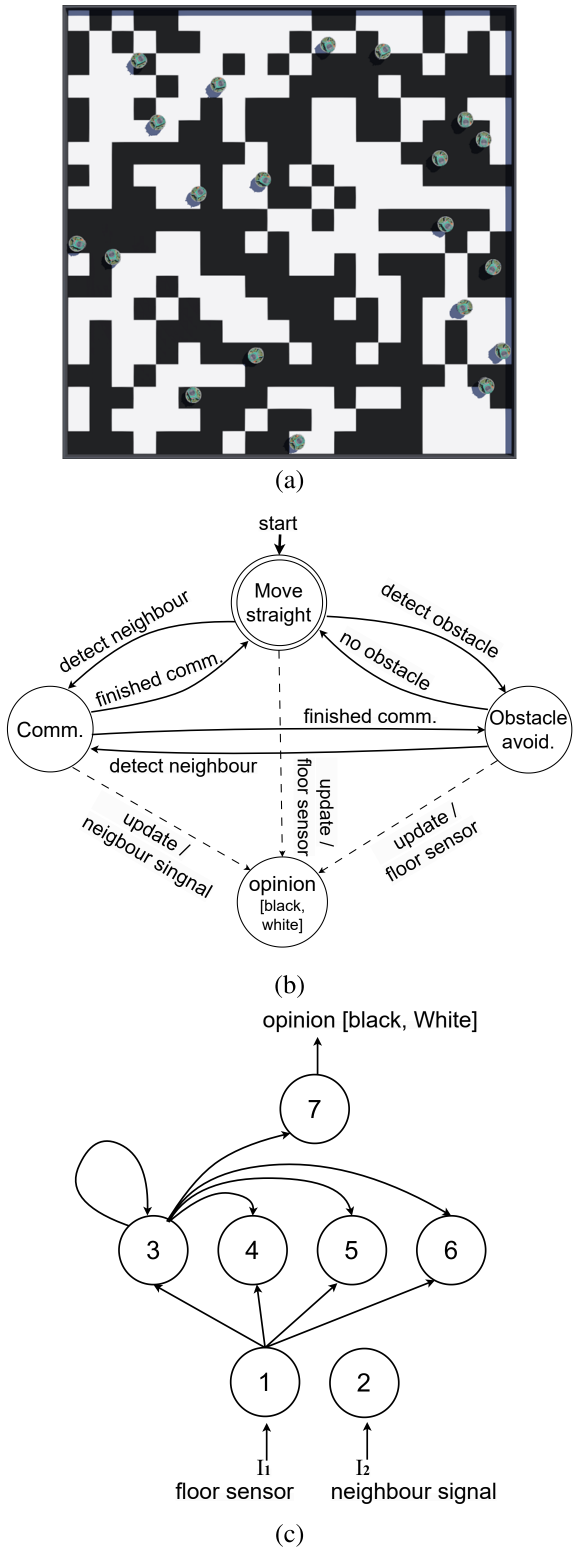

Figure 2. (a) The simulation environment. (b) The finite-state machine controlling the robots' movements. (c) Continuous-Time Recurrent Neural Network (CTRNN) generating the robots' opinion.

preserves the six best performing individuals (i.e., 'elite') from the preceding generation without any change. The rest of the new generation is formulated by proportionally selecting the fittest 40 of the prior generation.

During the evolutionary phase, each group undergoes eight evaluations in the *S-env* condition (four in BD-env and four in WD-env), with each evaluation lasting 400 seconds (equivalent to 4000 simulation steps). In every evaluation, the genotype is decoded into a neuro-controller, which is then replicated in all 20 robots (considering a homogeneous swarm). The robots are placed randomly in the arena, both in terms of position and orientation. After the first 2000 simulation steps, the robot r opinion is evaluated in every simulation step t (i.e., O_t^r). The average opinion of the 20 robots R is calculated and fitness assigned to the group according to Equation 4.

$$F_e = \begin{cases} \frac{T}{2} \sum_{t=\frac{T}{2}}^{T} \sum_{r=1}^{R} O_t^r & \text{in } \textit{WD-env} \\ \frac{T}{2} \sum_{t=\frac{T}{2}}^{T} \sum_{r=1}^{R} (1 - O_t^r) & \text{in } \textit{BD-env} \end{cases} \tag{4}$$

The evaluation of the fitness score in the latter half of the trial time is deliberate to avoid instability of opinion state during the initial exploration phase. In this early stage, robots have not yet accumulated sufficient physical and social experience of the environment.

Regarding computational complexity, the time required to complete a single evolutionary run, when executed on a Dell PowerEdge server equipped with 64 cores and 256 GB of main memory, is approximately 10 hours.

3. Results

For designing the controller of the robots, we perform five separate evolutionary runs, each one lasting 2000 generations. We remind the reader that during the evolutionary phase, the robots experience only the Random floor pattern (see first image in Figure 1). In order to select the best group (i.e., the best genotype), the highest-ranked groups from the 1000^{th} to the 2000^{th} generation of each evolutionary run are re-evaluated 50 trials in *BD-env* and 50 trials in *WD-env* environment. The best group out of these re-evaluations is chosen for the demonstration of the ability of the decision-making techniques that based on the neural-network to grant a swarm of simulated robots to attain consensus in the *WD-env* and the *BD-env*. Moreover, we show that the group can adapt to different floor patterns to the one examined through the design phase. In particular, this study show the accuracy of the decision-making process on the best group in eight extra floor patterns shown in Figure 1.

As far as it concerns the performances in the Random floor pattern, Figure 3 shows the decision-making development process by displaying the opinions of all the robots of the best group in both the *S-env* and in the *H-env* conditions, respectively. In both graphs, white boxes indicate the robots' number which have the right opinion in the *WD-env* environment, while black boxes indicate the robots' number which have the right estimation in the *BD-env* environment. Where each box is produced by 50 points (each point is corresponding to a different seeded re-evaluation trial). Figure 3a refers to the robots' opinion in *S-env* (i.e., the dominant colour takes 66% of the tiles that consisting the arena floor), while Figure 3b to the robots' opinion in *H-env* (i.e., the dominant colour takes 55% of the tiles that consisting the arena floor).

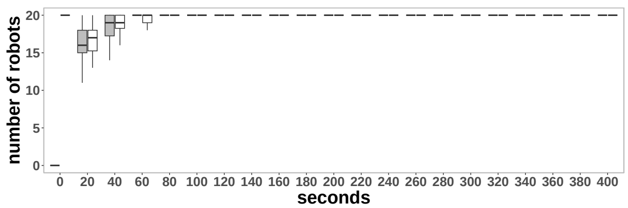

(a)

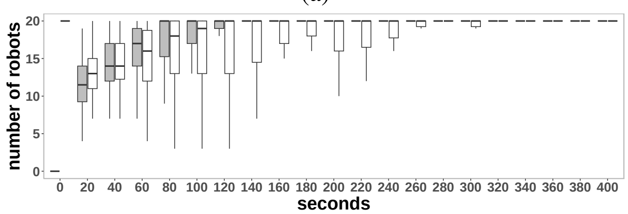

(b)

Figure 3. Boxs plot presenting the robots' number which have the right estimation in the *WD-env* environment (represented by the white boxes) and in the *BD-env* environment (represented by the grey boxes) at constant time durations of 20 s until the trial end (400 s). (a) *S-env* condition; (b) *H-env* condition. Each box is produced by 50 points (where each point is corresponding to a different seeded trial). The inter-quartile range of the data is represented by boxes, while the median value is marked by horizontal bars inside the boxes. The whiskers extend to the most extreme data points within 1.5 times the inter-quartile range from the box.

The graphs indicates that the best group reaches a consensus on the correct option, in both types of environments and in both the *S-env* and in the *H-env*. Moreover, the consensus is reached more quickly in the *S-env* than in the *H-env* condition. The consensus on the correct option in the *H-env* condition is reached in approximately 200 seconds in both types of environments. Note that, the robots' control system has been designed in the *S-env*. Thus, the *H-env* represents a rather novel environmental condition for these robots. It should be noted that in Figures 3a and 3b, the group converges to the white opinion in *WD-env* at time 0. This is due to the genetic basis of the evolved controller which even in the absence of any perceptual evidence—as it happens at the beginning of each trial—it selects opinion *WD-env*. The emergence of a genetic bias in binary collective and individual robot decision scenarios has been documented in previous research (e.g., see [23]), where robots are managed by analogous neural network architectures.

A. Robustness to Different Floor Patterns

To evaluate the robustness of the best group to environments with floor patterns different from those experienced during the design phase, we estimated the accuracy in the decision-making process of this group in eight extra floor patterns shown in Figure 1. The results are shown in Figure 4. The graphs show accuracy, that is the trials number (over 50 trials) where the consensus state to the correct opinion is reached by the group for at least 10 s, in nine distinct floor patterns. That is the Random pattern, already experienced during the evolutionary design phase, and eight extra patterns never experienced before. In this post-evaluation test, each trial lasts 1000 s. Figure 4a shows the results in *S-env* (i.e., the ratio of dominant colour tiles is 66%). The graph demonstrates good performances with a relatively high success ratio in all the floor patterns. In particular, its is worth noticing the accuracy in the Off-diagonal and Stripe, which remains above 80% in both floor patterns and for both the *WD-env* (see Figure 4a, white bars for Off-diagonal and Stripe) and the *BD-env* (see Figure 4a, black bars for Off-diagonal and Stripe). This is a significant performance improvement with respect to previous related works [16], [17], in which the authors report a significant performance drop, in term of accuracy, of robots required to operate in those patchy floor patterns (i.e., the Off-diagonal and the Stripe) without having experienced them during the design phase. In the next Section, we provide further evidence which illustrates the significance of the communication strategy in allowing the best group to extend its good accuracy rate to those eight floor patterns not experienced during the design phase. Figure 4b presents the results in *H-env* (i.e., the ratio of dominant colour tiles is 55%). When the variation in the proportion of floor overlay by the two colour reduced, a slight performance degradation observed across all the environment patterns, particularly in the Stripe environment. This accuracy drop can be, in large parts, accounted for by considering that the criteria for defining success in our experiment setup (i.e., all 20 robots must agree on the correct option for at least 10 s) is very stringent. It is worth mentioning that in many unsuccessful trials in the *H-env*, the majority of the robots (e.g., 18 or 19 robots) shared the same opinion about the correct option. However, in spite of the large convergence of the robots to the right opinion, given our definition of consensus, those trials are not considered successful. Generally speaking, the performance shown in Figure 4a and 4b represent a significant improvement compared to the results reported in recent research works [16], [17], where poor performance was reported even in *S-env* condition, particularly in patchy environments (e.g., Off-diagonal and Stripe).

Figure 5 shows the time to converge to consensus which is calculated over successful experiments only (out of 50 post-evaluation trials), in nine distinct floor patterns. The white boxes refer to time to convergence in the *WD-env*, while the black boxes refer to time to convergence in the *BD-env*. Each trial lasts 1000 s. In Figure 5a, the graph shows the results in *S-env* (i.e., the proportion of

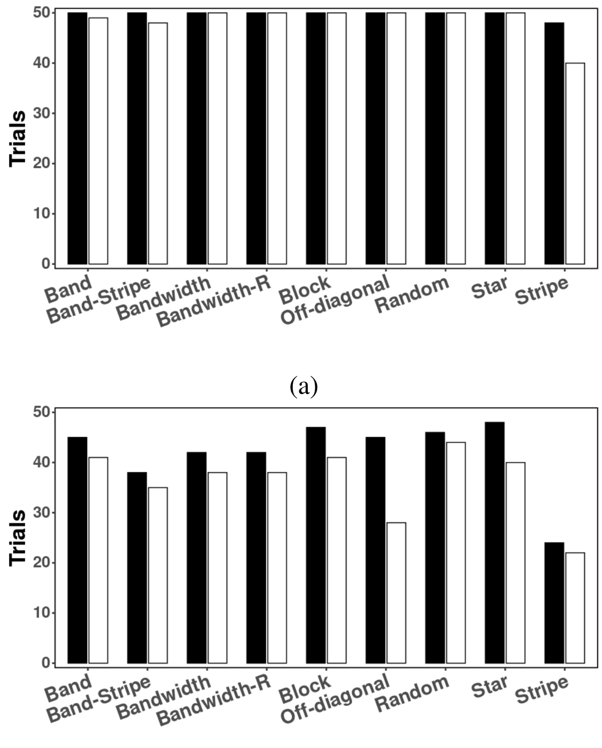

(a)

(b)

Figure 4. Bar plots representing accuracy, which is measured by the trials' number (from 50 experiments) where the swarm robots attained the consensus state to the correct opinion for at least 10 s, in nine distinct floor patterns (see [15] and Section 1). The white bars refer to accuracy in the *WD-env* environment, while the black bars refer to accuracy in the *BD-env* environment. Each trial lasts 1000 s. In (a), the graph shows the results in *S-env* (i.e., the proportion of the dominant colour tiles is %66) while in (b), the graph shows the results in *H-env* (i.e., the proportion of the dominant colour tiles is %55).

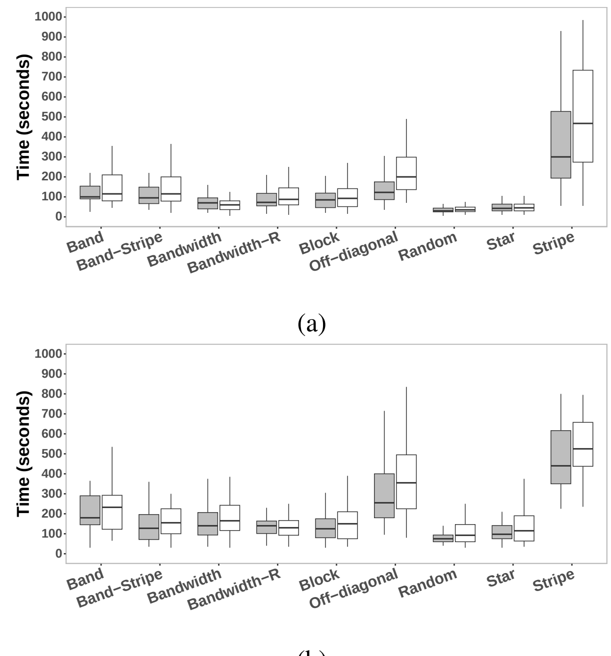

(a)

(b)

Figure 5. Box plots indicating time to converge to consensus which is calculated within successful experiments only (out of 50 post-evaluation trials), in nine distinct floor patterns (see [15] and Section 1). The white boxes refer to time to convergence in the *WD-env* environment, while the black boxes refer to time to convergence in the *BD-env* environment. Each trial lasts 1000 s. In (a), the graph shows the results in *S-env* (i.e., the ratio of the dominant colour tiles is %66) while in (b), the graph shows the results in *H-env* (i.e., the ratio of the dominant colour tiles is %55).

the dominant colour tiles is %66) while in Figure 5b, the graph shows the results in *H-env* (i.e., the proportion of the dominant colour tiles is %55). Similar trends are observed in both graphs, with slightly longer time to converge to consensus in *H-env*. An obvious increase in the time to convergence is observed in the Stripe environment in both graphs. This explains why the trial duration was increased from 400 s during the design phase, to 1000 s for this post-evaluation tests.

B. Further Investigation On the Communication Strategies

In the previous Section, we have shown that our experimental setup allowed us to design decision-making mechanisms using evolutionary-designed neuro-controllers, that allow a group of robots to accurately choose the correct options between two alternatives in a perceptual discrimination task. More importantly, we have shown that our best group manages to generalise its performance to floor patterns not experienced during the design phase. This result is particularly relevant because it represents a step forward compared to the results of previous research works [16], [17], which all reported a large drop in decision accuracy in patchy distributed floor patterns (i.e., the Off-diagonal and the Stripe). In order to achieve the good accuracy rate at the robustness test shown above, we have modified several elements of the original experimental setup as illustrated in [16], [17]. In particular, we have modified the type of random walk with which the robots explore the arena, the structures of the neuro-controller by increasing the number of neurons in the hidden layers, and the communication strategy allowing a receiver robot to receive communication signals from a randomly chosen robot among those in the communication range. These three modifications have been introduced progressively, one after the other with the intent to improve the accuracy at the Robustness test. However, a significant improvement in performance accuracy in the patchy distributed floor patterns has been observed only after having introduced the modification concerning the communication strategy. Thus, this indicates that the new communication strategy has the largest merit in improving the accuracy rate.

In this section, we show the results of further post-evaluation tests which aim to provide elements to explain why the new communication strategy proven more effective than the previous strategy in making the collective decision process robust enough to deal with all different floor patterns shown in Figure 1.

To understand how effective communication contributes to collective performance, we studied communication strategies focusing on the distance between signal sender and receiver. In particular, we investigated three types of communication strategies: i) a strategy called *close-neighbour* in which the communication events are possible only between a robot receiver and a robot emitter located at the shortest distance to the receiver among those within the receiver communication range (i.e., 50 cm); ii) a strategy called *rand-neighbour* in which communication events are possible between a robot receiver and a randomly chosen robot emitter among those within the receiver communication range; iii) a strategy called *far-neighbour* in which communication events are possible only between a robot receiver and a robot emitter located at the longest distance to the receiver among those within the receiver communication range.

In this post-evaluation tests, we run 50 trials in which we recorded, for each type of communication strategy employed by the robots (i.e., the *close-neighbour*, the *rand-neighbour*, and the *far-neighbour*), the number of communication events falling in each of the following distance category: i) <= 20 cm, ii) (20 cm, 30 cm], iii) (30 cm, 40 cm], and iv) (40 cm, 50 cm]. This post-evaluation test is meant to provide better insights into how opinions are communicated within the group.

Figure 6 shows heatmaps of the communication frequency between the robots during 50 trials, with the communication frequency sampled every 10 s over a trial duration of 400 s. The darker the cell colour in the map, the higher the frequency of communication. Figures 6a, 6b, and 6c represent the frequency of communication in the *close-neighbour*, *rand-neighbour* and *far-neighbour* strategies, respectively. As expected, when the robots employ the *close-neighbour* strategy (Figure 6a), the most frequent interactions are those falling into the category < 20 cm. On the contrary, when the robots employ the *far-neighbour* strategy, the most frequent interactions are those falling into the category (40 cm, 50 cm]. This demonstrates that in *close-neighbour* and *far-neighbour* strategies, opinion exchanging is spatially restricted to robots within a specific range of distances. That is, communication tends to concern either spatially close robots (when the group employs the *close-neighbour* strategy) or spatially distant robots (when the group employs the *far-neighbour* strategy). This bias affects the way in which opinions flow within the group, with a clear negative consequence on the accuracy in the patchy distributed floor patterns. Note that, in those works that reported a significant performance drop in the patchy floor patterns (i.e., [16], [17]), the robots employ the *close-neighbour* strategy. When the robots employ the *rand-neighbour* strategy, we notice a frequency distribution similar to the *far-neighbour* strategy but definitely less biased towards the category (40 cm, 50 cm] (see Figure 6b). This indicates that, when the robots employ the *rand-neighbour* strategy, as in our experimental setup, opinions circulate more frequently than in the *close-neighbour* strategy among distant robots, and also more frequently than in the *far-neighbour* strategy among the nearest robots. This is an

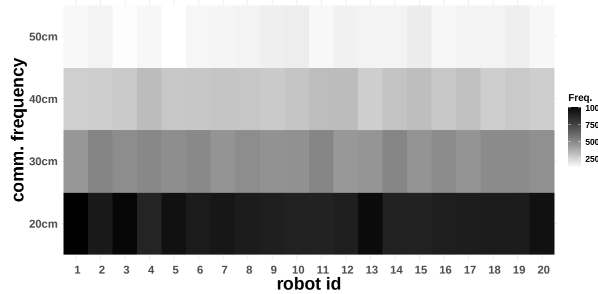

(a)

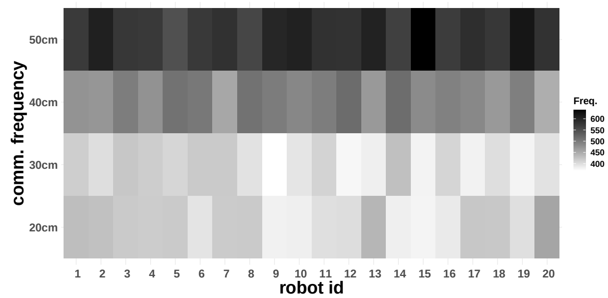

(b)

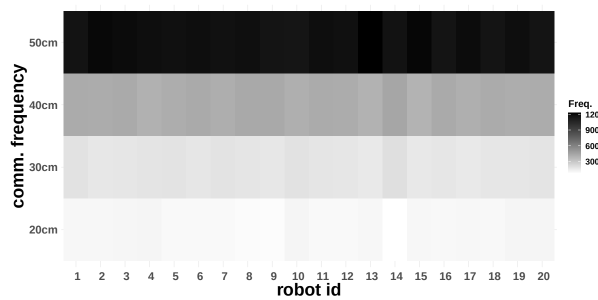

(c)

Figure 6. Heatmaps showing the frequency of communication events, over 50 trials, between two robots located at progressively longer distances. In particular, communication events are categorised into four categories (i.e., <= 20 cm, (20 cm, 30 cm], (30 cm, 40 cm], (40 cm, 50 cm]) based on the distance between the robot emitter and the robot receiver. In (a), the graph refers to the condition in which communication events are possible only between a robot receiver and a robot emitter located at the shortest distance to the receiver among those within the receiver communication range. In (b), the graph refers to the condition in which communication events are possible between a robot receiver and a randomly chosen robot emitter among those within the receiver communication range. In (c), the graph refers to the condition in which communication events are possible only between a robot receiver and a robot emitter located at the longest distance to the receiver among those within the receiver communication range.

element that favours an opinion exchange process that turns out to generate decision-making strategies capable of dealing with the patchy floor distributions without a significant accuracy drop in the group performance (see Figure 4).

A final series of post-evaluation tests is run to further investigate how the opinions flow within the group for the three different communication strategies. In particular, we run 50 trials in which, for each type of communication strategy, we recorded the number of communication events between a receiver and senders, within communication range, ordered from closest to farthest. For example, the closest sender is considered the first 1, the second closest to second 2, and so forth. The primary focus of this test is to correlate the communication strategies with the number of available robots within communication range, aiming to understand how this number affects the performance of opinions exchanged between robots. Figure 7 shows the number of communication events between a receiver and senders, within communication range, ordered from closest to farthest. These events are computed over 50 trials in post-evaluation tests in which the robots employ the *close-neighbour* strategy (see Figure 7a), the *rand-neighbour* strategy (see Figure 7b), and the *far-neighbour* strategy (see Figure 7c). The number of communication events is sampled every 10 s over a 400 s trial. It is worth noticing that, as for the previous test, the *rand-neighbour* strategy generates distributions of events more similar to the *far-neighbour* strategy, while recording the highest number of communication events for the first robot.

Generally speaking, the *rand-neighbour* strategy seems to generate a circulation of opinions between both closest and farthest robots, while the *close-neighbour* strategy allows only communication between the closest robots among those within communication range, and the *far-neighbour* strategy only between the farthest robots among those within communication range. Thanks to this property, the *rand-neighbour* strategy, contrary to the other two, allows the group to maintain a high accuracy rate even in the patchy floor patterns.

This study describes a series of experiments designed to develop effective and robust swarm robotics control mechanisms to allow the robots to perform a binary collective perceptual discrimination task, in which we vary not only the options' quality but also the way of the perceptual cues distribution within the environment. Our primary objective was to overcomes certain limitations observed in similar recent studies [15], [12], [16], [17], concerning the robustness of the collective decision-making process. In particular, we focus on a task where the decision making mechanisms for individuals are first optimised to let a swarm of robots to obtain a high accuracy rate in the collective decision in a type of environment in which cues are distributed randomly, and subsequently tested for their robustness in eight different environments where cues are distributed differently.

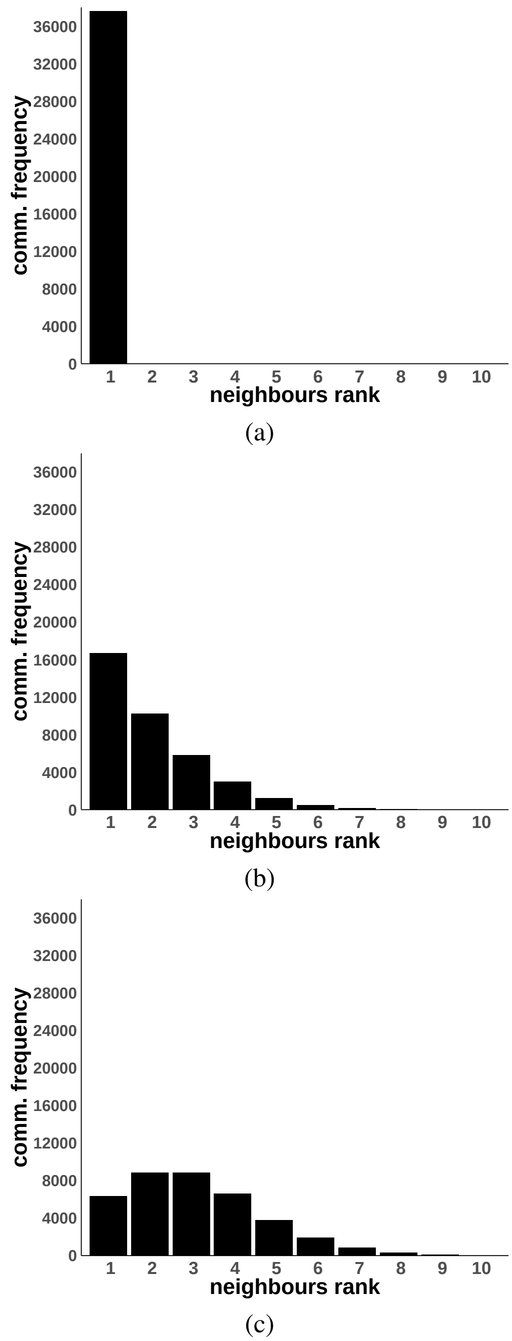

Figure 7. Bar plots showing the number of communication events between a receiver and senders, within communication range, ordered from closest to farthest. The x-axis refers to the ordinal number of the senders, while the y-axis refers to the number of communication events. These events are computed over 50 trials in post-evaluation tests in which the robots employ (a) the *close-neighbour* strategy, (b) the *rand-neighbour* strategy, (c) the *far-neighbour* strategy.

In order to improve the robustness, we modified three elements compared to [12], [16], [17]: the way in which the robots explore the arena, the structure of the neuro-controller, and the communication strategy among the robots. In this paper, we illustrate the important improvements, in terms of robustness of the collective response with respect to [15], [12], [16], [17] emerged thanks to the introduction of an alternative communication strategy, in which the opinions are exchanged between a robot receiver and a randomly chosen (instead from the closest as in [15], [12], [16], [17]) robot sender among those within communication range.

4. Conclusions and Future Work

Our findings indicate that the control mechanism following the new communication strategy significantly enhances performance, particularly in unseen patchy patterns of option distribution in the environment. The superiority of the random strategy over the previously used solutions is due to its more efficient circulation of opinions among both the spatially close and distant robots, thus ensuring high accuracy of opinion even in environments with patchy distributed features. However, we believe further investigation is needed to understand whether this improvement is consistent across other types of environmental patterns and tasks such as site selection.

In the future, we plan to investigate the impact of communication strategies on the performance of larger swarm size. We also intend to transfer the developed controller to physical e-puck robots to validate our findings in a physical system.

References

[1] H. Hamann, *Swarm robotics: A formal approach*. Springer, 2018, vol. 221.

[2] V. Trianni and S. Nolfi, "Engineering the evolution of self-organizing behaviours in swarm robotics: A case study," *Artificial life*, vol. 17, no. 3, pp. 183–202, 2011.

[3] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, "Swarm robotics: a review from the swarm engineering perspective," *Swarm Intelligence*, vol. 7, pp. 1–41, 2013.

[4] V. Trianni, E. Tuci, C. Ampatzis, and M. Dorigo, "Evolutionary swarm robotics: A theoretical and methodological itinerary from individual neuro-controllers to collective behaviours," *The horizons of evolutionary robotics*, vol. 153, 2014.

[5] G. Valentini, E. Ferrante, and M. Dorigo, "The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives," *Frontiers in Robotics and AI*, vol. 4, p. 9, 2017.

[6] 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*, vol. 46, no. 5, pp. 1175–1188, 2015.

[7] G. Valentini, D. Brambilla, H. Hamann, and M. Dorigo, "Collective perception of environmental features in a robot swarm," in *Swarm Intelligence: 10th International Conference, ANTS 2016, Brussels, Belgium, September 7-9, 2016, Proceedings 10*. Springer, 2016, pp. 65–76.

[8] G. De Masi, J. Prasetyo, E. Tuci, and E. Ferrante, "Zealots attack and the revenge of the commons: Quality vs quantity in the best-of-n," in *International Conference on Swarm Intelligence*. Springer, 2020, pp. 256–268.

[9] M. S. Talamali, A. Saha, J. A. Marshall, and A. Reina, "When less is more: Robot swarms adapt better to changes with constrained communication," *Science Robotics*, vol. 6, no. 56, p. eabf1416, 2021.

[10] P. A. Vargas, E. A. Di Paolo, I. Harvey, and P. Husbands, *The horizons of evolutionary robotics*. MIT press, 2014.

[11] G. De Masi and E. Ferrante, "Quality-dependent adaptation in a swarm of drones for environmental monitoring," in *2020 Advances in Science and Engineering Technology International Conferences (ASET)*. IEEE, 2020, pp. 1–6.

[12] 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 *Italian Workshop on Artificial Life and Evolutionary Computation*. Springer, 2021, pp. 109–120.

[13] S. Nolfi and D. Floreano, *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT press, 2000.

[14] A. Almansoori, M. Alkilabi, and E. Tuci, "A comparative study on decision making mechanisms in a simulated swarm of robots," in *2022 IEEE Congress on Evolutionary Computation (CEC)*, 2022, pp. 1–8.

[15] P. Bartashevich and S. Mostaghim, "Benchmarking collective perception: New task difficulty metrics for collective decision-making," in *Progress in Artificial Intelligence*, P. Moura Oliveira, P. Novais, and L. P. Reis, Eds. Cham: Springer International Publishing, 2019, pp. 699–711.

[16] A. Almansoori, M. Alkilabi, and E. Tuci, "Further investigations on the characteristics of neural network based opinion selection mechanisms for robotics swarms," in *International Conference on the Applications of Evolutionary Computation (Part of EvoStar)*. Springer, 2023, pp. 737–750.

[17] D. Trendafilov, A. Almansoori, T. Carletti, and E. Tuci, "The role of the environment in collective perception: A generic complexity measure," in *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*. MIT Press, 2023.

[18] A. Almansoori, M. Alkilabi, and E. Tuci, "On the evolution of adaptable and scalable mechanisms for collective decision-making in a swarm of robots," *Swarm Intelligence*, pp. 1–21, 2024.

[19] F. Mondada, M. Bonani, X. Raemy, J. Pugh, C. Cianci, A. Klaptocz, S. Magnenat, J.-C. Zufferey, D. Floreano, and A. Martinoli, "The e-puck, a robot designed for education in engineering," in *Proceedings of the 9th conference on autonomous robot systems and competitions*, vol. 1, no. CONF. IPCB: Instituto Politécnico de Castelo Branco, 2009, pp. 59–65.

[20] G. Francesca, M. Brambilla, A. Brutschy, V. Trianni, and M. Birattari, "Automode: A novel approach to the automatic design of control software for robot swarms," *Swarm Intelligence*, vol. 8, pp. 89–112, 2014.

[21] M. Kegeleirs, D. Garzón Ramos, and M. Birattari, "Random walk exploration for swarm mapping," in *Annual conference towards autonomous robotic systems*. Springer, 2019, pp. 211–222.

[22] R. D. Beer, "A dynamical systems perspective on agent-environment interaction," *Artificial intelligence*, vol. 72, no. 1-2, pp. 173–215, 1995.

[23] E. Tuci, M. Quinn, and I. Harvey, "An evolutionary ecological approach to the study of learning behaviour using a robot-based model," *Adaptive Behavior*, vol. 10, no. 3-4, pp. 201–221, 2002.

Rusul Ibrahim received the bachelor's degree in computer science from the University of Kerbala, Iraq, in 2019. Currently she is pursuing a master's degree in computer science as a graduate student at the research stage at Kerbala University, Iraq.

Muhanad Alkilabi obtained his electical engineering bachelors degree from Al-Mustansiriyah University, Baghdad in 2003 and a Masters degree from Pune University, India in 2007. He recevied a PhD in swarm robotics from Aberystwyth University, UK, in 2017. His research interest is in bio-inspired robotics and computational intelligence.

Ali Retha Hasoon Khayeat has a Ph.D. in Computer science / AI (Computer Vision and Machine Learning) from Cardiff University/UK. He is working as an assistant professor in the Computer Science Department/ College of Computer Science and IT/University of Kerbala/IRAQ. He is a professor of Digital Image Processing and Computer Vision for undergraduate and postgraduate in Kerbala CSD. His research interests in pattern recognition, Object detection and recognition, medical imaging, and digital forensics.

Elio Tuci is a professor at the Faculty of Computer Science, University of Namur since 2018. Before, he has lectured at Middlesex University London, and Aberystwyth University in UK. He is currently leading a lab of about ten researchers, whose activity is structured around two main axes: i) the design of conrol mechanisms for robotic systems required to independently perform specific tasks in relatively unstructured environements; ii) the development of mathematical and computer-based models to investigate issues related to the nature of the mechanisms underlying complex individual and social behaviour in natural systems. Further information on past and current research projects of Elio Tuci research group can be found at https://www.naxys.be/robotics/