

Robot Swarms Break Decision Deadlocks in Collective Perception through Cross-inhibition

Raina Zakir, Marco Dorigo, and Andreagiovanni Reina

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
{raina.zakir, andreagiovanni.reina, marco.dorigo}@ulb.be

Abstract. We study how robot swarms can achieve a consensus on the best among a set of n possible options available in the environment. While the robots rely on local communication with one another, follow simple rules, and make estimates of the option’s qualities subject to measurement errors, the swarm as a whole is able to make accurate collective decisions. We compare the performance of two prominent decision-making algorithms that are based, respectively, on the direct-switching and the cross-inhibition models, both of which are well-suited for simplistic robots. Most studies used these models to let robots achieve consensus by solely relying on social interactions and ignored the aspect of enabling robots to self-source information from the environment. However, in order to select the best option, we deem sampling environmental information crucial for the successful performance of the task. Through robot-swarm simulations, we show that swarms programmed with the direct-switching model are only able to make consensus decisions in asymmetric environments where options have different quality values. Instead, using cross-inhibition, the robot swarm can also break decision deadlocks and reach a consensus in symmetric environments with equal quality options. We investigate the mechanistic causes of such differences and we find that the time the robots spend in a state of indecision is a key parameter to break the symmetry. This research highlights the importance of considering both social and environmental information when studying collective decision-making.

1 Introduction

Swarm robotics is the research field that studies how to apply principles of swarm intelligence [5] to the design of decentralised systems consisting of large numbers of relatively simple robots that collectively perform tasks or solve problems [8]. As the robots within a swarm do not have global knowledge, the swarm’s collective behaviour emerges from the local interactions among the robots and from the interactions the robots have with their surrounding environment. Collective decision-making is a particular type of collective behaviour that is paramount to achieve group coordination—as such it is very often found in group-living animals [6]. For example, honeybees make consensus decisions on the site where to build their nest among several alternative locations [26], ants are able to collectively select the shortest path from their nest to a profitable food source [10], and

flocks of birds on the move select the same direction of motion in a decentralised way [3]. These natural systems have inspired the development of many different types of algorithms to enable robot swarms to make consensus decisions, such as selecting the aggregation site [27], selecting the direction of motion [7, 17, 9], selecting the predominant environmental feature [33], or selecting the shortest path for transporting items efficiently [29]. These algorithms need to be simple—to run on simple robots—and, at the same time, robust to robot malfunctions and flexible to changing environments—to work in real-life applications. A particularly important collective decision-making problem for swarm robotics is the so-called “best-of- n problem” [34], that is, how the swarm can select the best option among a set of n alternatives.

In this study, we consider the best-of-2 problem in which a minimalistic robot swarm is tasked with making a consensus decision on an environmental feature [33]. The environment floor is covered with yellow and blue tiles, and the environmental feature to decide on is which is the predominant colour. Therefore, the two colours represent the two alternative options to decide between, and the abundance of each colour (i.e. the proportion of yellow, or blue, tiles) represents the quality of the option. To achieve consensus on one of the two environmental features, the robots utilise a minimalistic decision-making algorithm. Each robot is committed to the option it considers the best and broadcasts voting messages about this option to its neighbours. Robots apply the decision-making algorithm to update their commitment to an option; the update can be based on either social information (received from neighbours as voting messages) or self-sourced information (obtained through independent exploration). Robots receiving voting messages from their peers update their opinion (i.e. the option to which they are committed) using minimalistic *opinion update models*. These models are minimalistic in nature and are therefore a viable solution for reaching consensus in simplistic robots. Periodically, robots choose to ignore social information and self-source information from the environment by independently switching their commitment to the option locally sensed in the environment. The individual self-sourcing of information through independent exploration of the environment can allow the swarm to achieve better adaptability in dynamic environments [36, 30, 2] where qualities of options may change over time. However, self-sourcing information is a form of asocial behaviour that also increases fluctuations (or noise) in the consensus formation [31, 21] that may result in decision deadlocks in certain decision-making algorithms [16, 21, 12]. Hence, the opinion update models that are used in collective decision-making need to be resilient to decision deadlocks when the amount of noise increases, either due to the exploration of the environment to achieve adaptability or to other sources, such as malfunctioning sensors on the robots that make them asocial (stubborn or zealot) and threaten the resiliency of symmetry-breaking in collective decision-making [14, 15].

Based on the literature, one of the most widespread models for updating a robot’s opinion upon receiving new social information is direct-switching [35], in which a robot switches to a random neighbour’s opinion during the voting phase. Direct-switching has been extensively used to engineer decentralised systems be-

cause of its simplicity and favourable tractability in minimalistic systems. However, theoretical studies on opinion dynamics [16] predict that direct-switching leads to decision deadlocks in presence of noise (e.g. self-sourcing environmental information). Despite being a highly relevant process in making decisions on the best option, there has been limited research focusing on the impact of self-sourcing environmental information on the collective dynamics of swarms using the direct-switching model. An alternative to the direct-switching model is the cross-inhibition model [18, 23, 22], which is inspired by the house-hunting process in honeybees [26]. The cross-inhibition model has comparable simplicity to direct-switching, and theory predicts a higher resilience to the presence of noise. Unlike what happens with the direct-switching model, when a robot using the cross-inhibition model receives a contrasting opinion from one of its neighbours, it gets uncommitted and remains without an opinion—i.e. it becomes undecided. Using robot swarm simulations, we estimate to what extent the decision-making algorithm based on the cross-inhibition model is resilient to increasing noise and show that the time spent by the robots in the uncommitted state is fundamental to the ability of being resilient to noise induced from the self-sourcing information.

The outline of the rest of the paper is as follows. Section 2 defines the best-of- n problem, the collective decision-making algorithms, and the mechanism to self-source information from the environment. In Section 3, we describe the experimental setup and explain the parameters that have been analysed in this study. In Section 4, we present the results, and finally, in Section 5, we conclude and discuss possible directions in which this work could be extended.

2 The Models

We consider the $n = 2$ instance of the best-of- n decision problem, in which the swarm has to converge to the best between two options, A or B. Each option has a quality, q_A and q_B , and the parameter $q = q_A/q_B$ represents the ratio between the two qualities. Without loss of generality, in our study, we assume that $q_A \geq q_B$. Each robot is committed to an option, which corresponds to the robot’s opinion, or uncommitted, that is, without an opinion. The robot behaviour is based on the same finite state machine of [35] characterised by two continuously alternating states: exploration and dissemination shown in Figure 1A. In the exploration state, the robots assess the quality q_i of their current opinion by sampling the environment (with $i = \{A, B\}$). The amount of time a robot stays in exploration is drawn randomly from an exponential distribution with a rate equal to λ^{-1} . In the dissemination state, the robots disseminate their opinion i locally to their neighbours. The amount of time a robot spends disseminating its opinion is drawn from an exponential distribution with a rate $q_i g$, which is directly proportional to the option’s quality q_i and is scaled by the average duration of dissemination g . The parameter g is set based on the requirements of the considered scenario. By scaling the time spent in the dissemination state proportionally to the quality of the options assessed in the exploration state,

the probability of receiving messages from peers committed to the best opinion increases because they disseminate for a longer time. As a result, it will be more likely to observe neighbours that are in favour of the best option than observing neighbours that are supporting the lower quality option. The dissemination state is followed by either a polling state or a self-sourcing state (see Figure 1A). The decision to go to either states is random, based on the noise probability η . With probability η , the robot self-sources a new opinion from the environment, and with probability $(1 - \eta)$ polls other robots’ information. In the self-sourcing state, the robot replaces its opinion with the option (i.e. the colour) found in its current location of the environment. Including the self-sourcing mechanism allows the robots to periodically monitor the environment and reconsider the best option with new environmental evidence. On the other hand, the polling state involves collecting the opinions of the neighbours, choosing one at random and then applying an opinion update mechanism—either direct-switching or cross-inhibition. In this study, to simplify the behaviour for minimalistic robots and minimise memory use, the robots in the polling state only consider the first message they receive from their neighbours. Finally, after either using social information or self-sourcing environmental information, a robot returns to the exploration state to continue the cycle.

Direct-switching When it uses direct-switching as its opinion update model, the robot reads the message of one randomly chosen neighbour (which is disseminating within its communication range) and adopts that neighbour’s opinion regardless of whether it is the same or different from the robot’s own opinion. This mechanism allows accurate consensus formation among neighbours [35]. However, it can also result in unstable group dynamics due to the formation of echo chambers among robots with the same opinion that can prevent consensus formation in the swarm [28].

Cross-inhibition According to the cross-inhibition model, the robot can either be committed to an option or uncommitted. During polling, when a committed robot reads a (randomly chosen) message from a robot committed to a different option (e.g. a robot committed to A reads a message from a robot committed to B), it gets inhibited and becomes uncommitted. When an uncommitted robot receives any opinion (A or B) from one of its neighbours, it gets recruited to the received option.

3 Experimental Setup

To analyse the models introduced in Section 2, we implement the collective decision-making behaviour on a swarm of $N = 100$ simulated robots. For this analysis, we use Kilobots [24]—small-sized and low-cost robots that communicate using infrared (IR) transceivers with other robots in a range of 10 cm, move at a speed of 1 cm/s and have a control loop of approximately 32 ms. We simulate

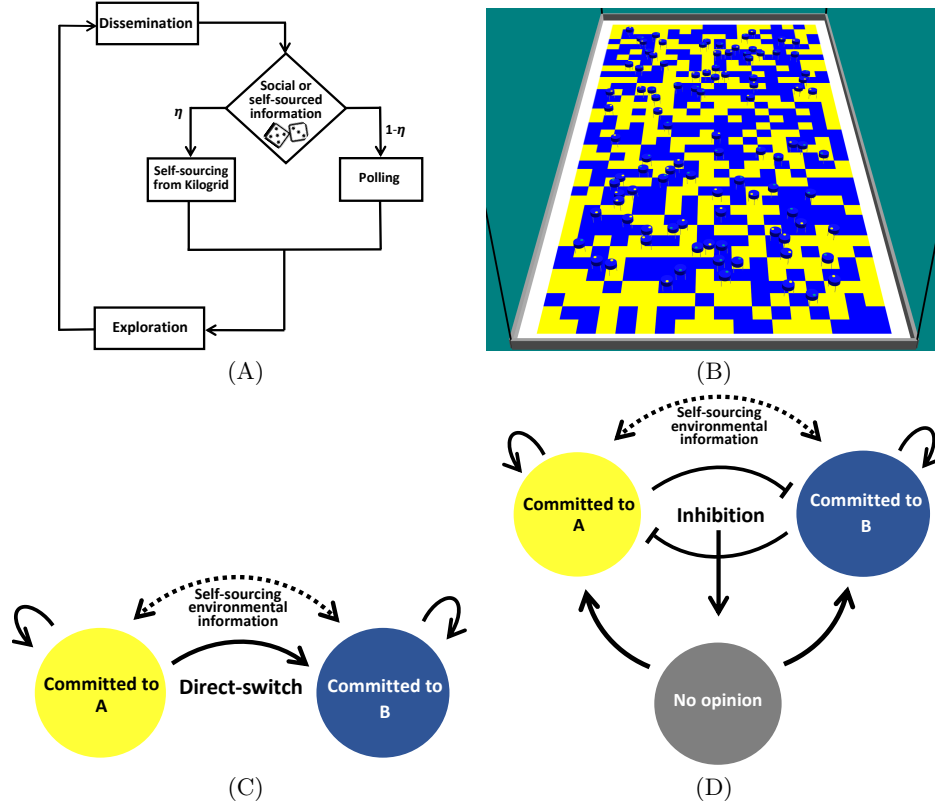


Fig. 1. (A) The finite state machine (FSM) describing the robots' behaviour, based on the FSM of [35] and extended to include the possibility of self-sourcing information. The rectangles represent the four FSM's states and the arrows represent the transitions among them. (B) Snapshot of an experiment showing 50 simulated Kilobots in the ARGoS Kilogrid arena comprising yellow and blue tiles. (C-D) Robot's opinion update model of direct-switching and cross-inhibition, respectively. The robot updates its opinion based on either social information (solid lines) or self-sourced environmental information (dashed lines). In direct-switching (C), the robot that gets recruited changes its commitment immediately. In cross-inhibition (D), when a committed robot receives a message from a robot committed to a different option, it resets its commitment (it gets inhibited).

the robot swarm in ARGoS, a state-of-the-art swarm robotics simulator [20, 19]. To provide robots with a virtual environment from which they can self-source information, we simulate the Kilogrid [1, 32]. The Kilogrid is an electronic table sized $1 \times 2 \text{ m}^2$, composed of 800 cells that interact with the Kilobots through IR and that can be easily simulated in ARGoS [2]. With the exception of the Kilogrid cells at the borders (depicted in white in Figure 1B), all the cells are set to send constantly IR messages signalling their ID and their colour, either the yellow colour associated with option A or the blue colour associated with option B. The proportion of cells allocated to emit messages for each option can be symmetric (50% for A and 50% for B) or asymmetric. In cases of asymmetric environment, as a convention, we keep option A with higher quality, i.e. there are more Kilogrid cells signalling option A than cells signalling option B.

The Kilobots use the IR messages from the Kilogrid’s cell beneath it both in the self-sourcing state to collect new information from the environment, and during exploration to estimate their opinion’s quality (i.e. the proportion of cells of a given colour). As Kilobots are not equipped with any proximity sensors, the Kilogrid cells also send a ‘wall flag’ to signal proximity to a wall (the flag is a binary value that can be either high/low) that the Kilobots use to avoid collisions. The white cells at the borders and the non-white cells adjacent to the white cells send a high wall flag, while all the other internal cells send a low flag. Without such wall flags to detect proximity to the walls, a large number of Kilobots would remain clustered on the arena walls.

3.1 Robot Behaviour

The robots start from a uniformly random position in the environment and with a random initial opinion; we initialise half of the swarm committed to option A and the other half to option B. To explore different portions of the environment and exchange messages with different robots, the Kilobots always perform a random walk in the environment, alternating between a rotation phase of approximately 5 s (in a randomly chosen direction—clockwise or counterclockwise) and a straight motion phase of approximately 10 s. The random walk allows the robots to encounter different robots in their neighbourhood during the dissemination phase and allows more accurate estimation of the option qualities from the Kilogrid during the exploration phase.

A robot that receives a high wall flag from the Kilogrid executes—regardless of its state—a simple obstacle avoidance routine. The robot starts a random rotation phase of approximately 4 s followed by a straight motion phase of approximately 7 s. If the wall flag is detected again, the obstacle avoidance routine is reinitialised till the robot receives a low flag.

All robots start the experiment in an exploration state. During the exploration, a robot committed to i reads the Kilogrid messages to keep the count T_i of the number of cells it encountered (it uses the cell’s ID to count each cell only once) and the count C_i of how many of the visited cells have the same colour as its own opinion i . At the end of the exploration cycle, the robot estimates the quality $q_i = \min(1, 2C_i/T_i)$, hence $0 \leq q_i \leq 1$. When $C_i/T_i \geq 0.5$, the quality

is set to its maximum $q_i = 1$ as the goal is to select the predominant colour, and because the robot has found more than half of the readings have colour i , it assigns to i the maximum quality. For $C_i/T_i < 0.5$, the quality scales linearly in $[0, 1]$. The values T_i and C_i correspond to the counts of one exploration cycle only and are reset before entering the dissemination state.

Based on q_i , the robot computes the dissemination time using an exponential distribution with $\lambda_d^{-1} = q_i g_c$ where $g_c = 1\,300$ is the average number of control cycles in dissemination when $q_i = 1$, which corresponds to λ_d^{-1} of about 40 s. In case the robot is uncommitted, the parameter λ_d^{-1} is set to $0.5 g_u$; using the default value for $g_u = 400$, the uncommitted robot spends an average of approximately 6 s in the dissemination state. At the end of the dissemination, the robot decides with probability η whether to perform either an individual environmental observation (enter the self-sourcing state) or a social interaction (enter the polling state). Once the environmental observation or the polling are terminated, the robot computes the exploration time and enters the exploration state again. A committed robot computes the exploration time using $\lambda_e = 0.0003$, resulting in an average exploration time of approximately 100 s. Instead an uncommitted robot uses the same rate used to compute the dissemination time, i.e. $\lambda_e^{-1} = \lambda_d^{-1} = 0.5 g_u$, with $g_u = 400$. The total duration of each simulation run is 110 minutes.

4 Experiments and Results

We run simulations to test the effect of different values of the noise probability $\eta = 0$ (no noise), $\eta = 0.01$ (low), $\eta = 0.05$ (medium) and $\eta = 0.25$ (high), on both opinion update models for different quality ratios q . In the first set of experiments, we test direct-switching and cross-inhibition models in a symmetric environment (50% of the Kilogrid cells signal option A and 50% option B), i.e. $q = q_A/q_B = 1$. The second set of experiments includes the direct-switching and cross-inhibition model in asymmetric environments with three values of quality ratio q : 1.08 (Kilogrid cells: 52%A, 48%B), 1.22 (Kilogrid cells: 55%A, 45%B), and 1.5 (Kilogrid cells: 60%A, 40%B). For each condition, we run 50 simulations that we use to generate the histograms of Figure 2. The histograms show how frequently the swarm distributes between robots supporting option A and B in the last 1 000 timesteps (approximately 30 s) of a run. For each timestep, we subtract the proportion of robots supporting option B from the proportion of robots supporting option A, i.e. (number robots for A - number robots for B) / N , and report the results as histograms in Figure 2.

Figure 2 shows that when the environment is symmetric ($q = 1$), both models are able to break the symmetry in the absence of noise ($\eta = 0$). However, the performance of direct-switching deteriorates as soon as noise is introduced ($\eta \geq 0.01$), and the swarm cannot reach any agreement but remains in a state of decision deadlock. Direct-switching with noise $\eta > 0$ can only reach convergence to a stable majority towards the best option for a high quality ratio ($q = 1.22$ for low noise and $q = 1.5$ for medium noise); in all other conditions the swarm using

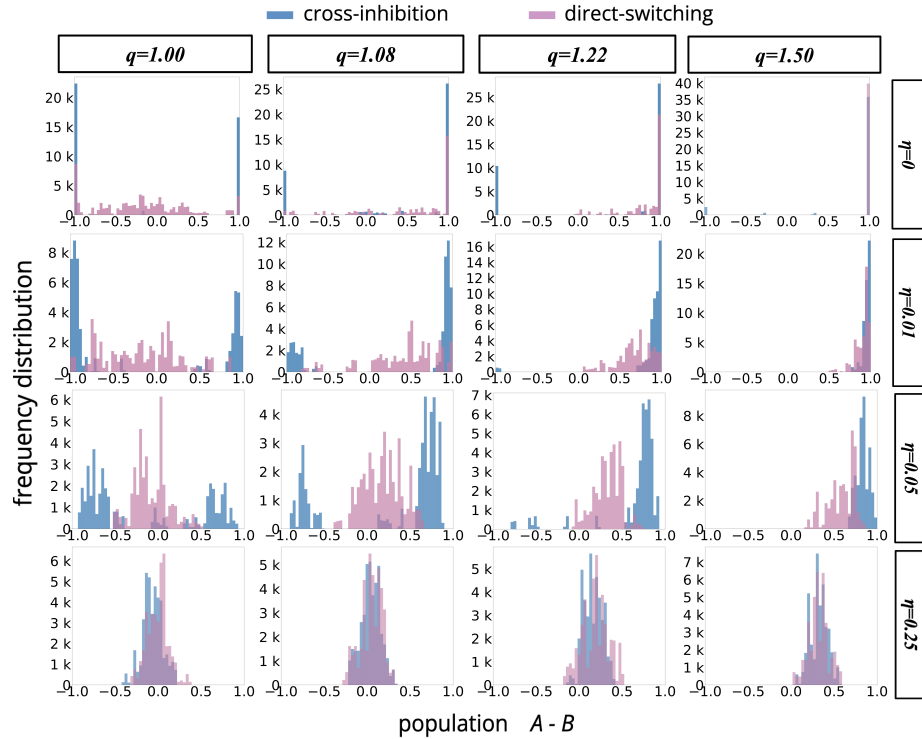


Fig. 2. Histograms for cross-inhibition and direct-switching models when $N = 100$, $g_c = 1300$ and $g_u = 400$ (statistics over 50 runs) showing the effect of increasing the probability of self-sourcing environmental information (η) and the quality ratio ($q = q_A/q_B$) on the collective decision-making process. The histograms are computed as the difference between the proportion of robots supporting A and B for each of the last 1000 timesteps of every run ($(A-B)/N$ on the x-axis).

direct-switching remains in an undecided state. Instead, the cross-inhibition is consistently able to break the symmetry for both low and medium levels of noise ($\eta \leq 0.05$) for any tested value of q . With higher levels of noise, both models fail to break the symmetry, even when the quality ratio increases. In summary, cross-inhibition is always better than direct-switching to break decision deadlocks and make consensus decisions, except for cases in which noise is very high.

To further analyse the mechanism through which the cross-inhibition model is resilient to decision deadlocks, we test the influence of the amount of time a robot spends in the uncommitted state and its ability to break the symmetry. To do so, we vary the average duration of dissemination and exploration of uncommitted agents by varying the parameter g_u from 0 to 2000 (corresponding to an average temporal duration from 0s to approximately 62s). When $g_u = 0$, the voting mechanism becomes equivalent to direct-switching. Increasing g_u corresponds to

increasing the time the robot spends in an uncommitted state, as g_u determines the average dissemination and exploration time of uncommitted robots. The change in dynamics with noise $\eta = 0.05$ and $q = 1$ is shown in Figure 3A. When $g_u = 0$, the result obtained corresponds to dynamics similar to those observed in direct-switching with $\eta = 0.05$ in a symmetric environment (Figure 2, $q = 1$ and $\eta = 0.05$); the swarm remains in a decision deadlock. As g_u increases, the bistability becomes more prominent, as observed in $\eta = 0.05$ in the symmetric environment (Figure 2). The results of Figure 3A show that the amount of time spent in an uncommitted state is the key to converging on a large majority for one of the two equivalent options.

The cross-inhibition model has dynamics that are much more stable than direct-switching [26]; therefore, the swarm reaches and maintains an agreement for either option. However, the high stability of the cross-inhibition model can also occasionally lock the system in a consensus for the inferior option (with lower quality), which may have been reached due to initial random fluctuations. Figure 2 shows that the system is in a bistability state (i.e. selection of both options is possible) when the options have similar qualities ($q \leq 1.08$ in the presence of noise) and is instead able to reliably select the superior alternative for larger quality differences. Differently, the direct-switching model, when it is able to break the symmetry, always selects the option with the highest quality. Interestingly, for $q = 1.22$, cross-inhibition's bistability exists for $\eta = 0$ and vanishes for higher levels of noise, $\eta = 0.01$ and $\eta = 0.05$. In this case, occasionally self-sourcing information helps in correcting initial mistakes.




Fig. 3. (A) 2D-histogram for increasing g_u for $\eta = 0.05$ and $q = 1$ showcasing the shift from indecision to symmetry-breaking. (B) 2D-histogram for $\eta=0$ and $q = 1.22$ showcasing the switch from consensus on the best option when $g_u = 0$ to bistability as g_u increases. The plots show how the distribution of robots supporting A and B (y-axis) change as g_u (x-axis) increases. We consider the last 1000 timesteps (e.g.6s) of each of the 50 runs per g_u and subtract the proportion of robots supporting option B from the proportion of supporters for A (i.e. (number robots for A - number robots for B) / N) to plot the 2D-histograms.

To understand the accuracy of the two opinion update models in selecting the best option in the presence of similar quality options, we vary the time g_u when $q = 1.22$ and $\eta = 0$ (Figure 3B). For $g_u = 0$, the swarm breaks the symmetry in favour of the option of highest quality (A) as the model is equivalent to direct-switching. When g_u increases, the system gradually moves towards a state of bistability. As observed in Figure 2 ($q = 1.22$ and $\eta = 0$), cross-inhibition can occasionally select the inferior option due to its highly stable dynamics that lock the system into consensus for either option when qualities are similar. Figure 3B shows that bistability becomes more and more pronounced as the time spent in the uncommitted state increases, or, in other words, the probability of selecting the inferior option increases with increasing g_u . As noted earlier, the selection of the inferior option becomes less probable when the cross-inhibition model is subject to moderate levels of noise (compare Figure 2 $\eta = 0$ and $\eta = 0.01$, for $q = 1.22$) as it increases the exploratory behaviour of the robots and enables their ability to correct their collective decision.

Therefore, our results show a trade-off between the ability to make consensus decisions in the presence of noise (when the robots spend long times in the uncommitted state) and the ability to avoid inaccurate decisions (for short times in the uncommitted state). In scenarios where choosing the option with the highest quality is an utmost requirement and noise is a factor not applicable, direct-switching is a better choice for collective decision-making. However, random fluctuations can be inevitable in systems operating in the real world, and we have shown how they can dramatically hamper the performance of direct-switching. Therefore, our study highlights the importance of using cross-inhibition to make collective decisions in realistic application scenarios.

5 Discussion and Conclusions

In this study, we investigated two prominent collective decision-making algorithms for the best-of- n problem in the presence of both social interactions and environmental information. The results of our simulations show that robot swarms running algorithms based on the direct-switching model fail to reach a consensus on the best option when robots use both social information and self-sourced information acquired through individual exploration of the environment. Self-sourcing information from the environment can also be modelled as noise, which is very likely to be present in most real-world scenarios, for example in the form of asocial robots or sensor failures [11]. Therefore, even if direct-switching has the desirable property of being very simple, to deploy systems in the real world the robot algorithms must be resilient to noise. We show that the cross-inhibition model serves as an ideal alternative to direct-switching. By letting robots inhibit each other and become uncommitted for some time, the cross-inhibition model enables stability and symmetry-breaking dynamics that prevent decision deadlocks. This work is limited to simulation; we plan as future work to conduct mathematical analyses based on ODEs and chemical reaction network models in order to understand better the role of the time spent in the

uncommitted state for obtaining high stability and breaking symmetry. We also plan to validate our results through real-robot experiments on the real Kilogrid. Moreover, most of the collective decision-making research in swarm robotics is concentrated on binary best-of- n problems, with only a few studies exploring $n > 2$ [25, 13, 4, 30]. Therefore, as future work, we also aim to expand our analyses and experiments to $n > 2$ scenarios and investigate if the robustness of the cross-inhibition model extends to non-binary environments, as theory predicts [22].

Acknowledgements The authors thank Till Aust and Jonas Kuckling for the technical support on simulating the Kilogrid and running simulations on the HPC. This work was supported by Service Public de Wallonie Recherche under grant n° 2010235 - ARIAC by DigitalWallonia4.AI. M. Dorigo and A. Reina acknowledge support from the Belgian F.R.S.-FNRS, of which they are Research Director and Chargé de Recherches, respectively.

References

1. Antoun, A., Valentini, G., Hocquard, E., Wiandt, B., Trianni, V., Dorigo, M.: Kilogrid: A modular virtualization environment for the Kilobot robot. In: 2016 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 3809–3814 (2016). <https://doi.org/10.1109/IROS.2016.7759560>
2. Aust, T., Talamali, M., Dorigo, M., Hamann, H., Reina, A.: The hidden benefits of limited communication and slow sensing in collective monitoring of dynamic environments. In: Dorigo, M., et al. (eds.) *Swarm Intelligence (ANTS 2022)*. LNCS, vol. 13491. Springer (2022)
3. Ballerini, M., Cabibbo, N., Candelier, R., Cavagna, A., Cisbani, E., Giardina, I., Orlandi, A., Parisi, G., Procaccini, A., Viale, M., Zdravkovic, V.: Empirical investigation of starling flocks: a benchmark study in collective animal behaviour. *Animal Behaviour* **76**(1), 201–215 (2008). <https://doi.org/10.1016/j.anbehav.2008.02.004>
4. Bartashevich, P., Mostaghim, S.: Multi-featured collective perception with Evidence Theory: tackling spatial correlations. *Swarm Intelligence* **15**(1-2), 83–110 (2021). <https://doi.org/10.1007/s11721-021-00192-8>
5. Bonabeau, E., Dorigo, M., Theraulaz, G.: *Swarm Intelligence: From Natural to Artificial Systems*. Oxford University Press, New York, NY (1999)
6. Bose, T., Reina, A., Marshall, J.A.R.: Collective decision-making. *Current Opinion in Behavioral Sciences* **6**, 30–34 (2017). <https://doi.org/10.1016/j.cobeha.2017.03.004>
7. Campo, A., Nouyan, S., Birattari, M., Groß, R., Dorigo, M.: Negotiation of goal direction for cooperative transport. In: Dorigo, M., Gambardella, L.M., Birattari, M., Martinoli, A., Poli, R., Stützle, T. (eds.) *Ant Colony Optimization and Swarm Intelligence – Proceedings of ANTS 2006 – Fifth International Workshop*. Lecture Notes in Computer Science, vol. 4150, pp. 191–202. Springer, Berlin, Germany (2006)
8. Dorigo, M., Birattari, M., Brambilla, M.: Swarm robotics. *Scholarpedia* **9**(1), 1463 (2014)

9. Ferrante, E., Turgut, A.E., Huepe, C., Stranieri, A., Pinciroli, C., Dorigo, M.: Self-organized flocking with a mobile robot swarm: a novel motion control method. *Adaptive Behavior* **20**(6), 460–477 (2012). <https://doi.org/10.1177/1059712312462248>
10. Goss, S., Aron, S., Deneubourg, J.L., Pasteels, J.: Self-organized shortcuts in the argentine ant. *Naturwissenschaften* **76**, 579 (1989). <https://doi.org/10.1007/BF00462870>
11. Khalil, N., Miguel, M.S., Toral, R.: Zealots in the mean-field noisy voter model. *Physical Review E* **97**(1) (2018). <https://doi.org/10.1103/physreve.97.012310>
12. Khaluf, Y., Pinciroli, C., Valentini, G., Hamann, H.: The impact of agent density on scalability in collective systems: noise-induced versus majority-based bistability. *Swarm Intelligence* **11**(2), 155–179 (2017). <https://doi.org/10.1007/s11721-017-0137-6>
13. Lee, C., Lawry, J., Winfield, A.F.T.: Negative updating applied to the best-of-n problem with noisy qualities. *Swarm Intelligence* **15**(1-2), 111–143 (jun 2021). <https://doi.org/10.1007/s11721-021-00188-4>
14. Mobilia, M., Petersen, A., Redner, S.: On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment* **2007**(08), P08029 (2007). <https://doi.org/10.1088/1742-5468/2007/08/p08029>
15. Mobilia, M.: Does a single zealot affect an infinite group of voters? *Physical Review Letters* **91**(2) (2003). <https://doi.org/10.1103/physrevlett.91.028701>
16. Mobilia, M.: Nonlinear q-voter model with inflexible zealots. *Physical Review E* **92**(1) (2015). <https://doi.org/10.1103/physreve.92.012803>
17. Nouyan, S., Campo, A., Dorigo, M.: Path formation in a robot swarm: Self-organized strategies to find your way home. *Swarm Intelligence* **2**(1), 1–23 (2008). <https://doi.org/10.1007/s11721-007-0009-6>
18. Pais, D., Hogan, P.M., Schlegel, T., Franks, N.R., Leonard, N.E., Marshall, J.A.R.: A mechanism for value-sensitive decision-making. *PLoS ONE* **8**(9), 1–9 (2013). <https://doi.org/10.1371/journal.pone.0073216>
19. Pinciroli, C., Talamali, M.S., Reina, A., Marshall, J.A.R., Trianni, V.: Simulating Kilobots within ARGoS: models and experimental validation. In: M. Dorigo et al. (ed.) *Swarm Intelligence (ANTS 2018)*. LNCS, vol. 11172, pp. 176–187. Springer, Cham (2018). https://doi.org/https://doi.org/10.1007/978-3-030-00533-7_14
20. Pinciroli, C., Trianni, V., O’Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., Birattari, M., Gambardella, L.M., Dorigo, M.: ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence* **6**(4), 271–295 (2012). <https://doi.org/10.1007/s11721-012-0072-5>
21. Rausch, I., Reina, A., Simoens, P., Khaluf, Y.: Coherent collective behaviour emerging from decentralised balancing of social feedback and noise. *Swarm Intelligence* **13**(3–4), 321–345 (2019). <https://doi.org/10.1007/s11721-019-00173-y>
22. Reina, A., Marshall, J.A.R., Trianni, V., Bose, T.: Model of the best-of-N nest-site selection process in honeybees. *Physical Review E* **95**(5), 052411 (2017). <https://doi.org/10.1103/PhysRevE.95.052411>
23. Reina, A., Valentini, G., Fernández-Oto, C., Dorigo, M., Trianni, V.: A design pattern for decentralised decision making. *PLoS ONE* **10**(10), e0140950 (2015). <https://doi.org/10.1371/journal.pone.0140950>
24. Rubenstein, M., Ahler, C., Nagpal, R.: Kilobot: A low cost scalable robot system for collective behaviors. In: *2012 IEEE International Conference on Robotics and Automation*. IEEE Press, Piscataway, NJ (2012). <https://doi.org/10.1109/ICRA.2012.6224638>

25. 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 (2016). <https://doi.org/10.1109/TCYB.2015.2429118>
26. Seeley, T.D., Visscher, P.K., Schlegel, T., Hogan, P.M., Franks, N.R., Marshall, J.A.R.: Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science* **335**(6064), 108–111 (2012). <https://doi.org/10.1126/science.1210361>
27. Sion, A., Reina, A., Birattari, M., Tuci, E.: Controlling robot swarm aggregation through a minority of informed robots. In: Dorigo, M., et al. (eds.) *Swarm Intelligence (ANTS 2022)*. LNCS, vol. 13491. Springer (2022). <https://doi.org/10.48550/arXiv.2205.03192>
28. Starnini, M., Frasca, M., Baronchelli, A.: Emergence of metapopulations and echo chambers in mobile agents. *Scientific Reports* **6** (2016). <https://doi.org/10.1038/srep31834>
29. Talamali, M.S., Bose, T., Haire, M., Xu, X., Marshall, J.A.R., Reina, A.: Sophisticated collective foraging with minimalist agents: A swarm robotics test. *Swarm Intelligence* **14**(1), 25–56 (2020). <https://doi.org/10.1007/s11721-019-00176-9>
30. Talamali, M.S., Saha, A., Marshall, J.A.R., Reina, A.: When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics* **6**(56), eabf1416 (2021). <https://doi.org/10.1126/scirobotics.abf1416>
31. Tsimring, L.S.: Noise in biology. *Reports on progress in physics. Physical Society* **77**(2), 026601 (2014). <https://doi.org/10.1088/0034-4885/77/2/026601>
32. Valentini, G., Antoun, A., Trabattoni, M., Wiandt, B., Tamura, Y., Hocquard, E., Trianni, V., Dorigo, M.: Kilogrid: A novel experimental environment for the kilobot robot. *Swarm Intelligence* **12**(3), 245–266 (2018)
33. Valentini, G., Brambilla, D., Hamann, H., Dorigo, M.: Collective perception of environmental features in a robot swarm. In: Dorigo, M., et al. (eds.) *Swarm Intelligence (ANTS 2016)*. LNCS, vol. 9882, pp. 65–76. Springer (2016). https://doi.org/10.1007/978-3-319-44427-7_6
34. 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>
35. Valentini, G., Hamann, H., Dorigo, M.: Self-organized collective decision making: The weighted voter model. In: *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems*. pp. 45–52. AAMAS '14, IFAAMAS, Richland, SC (2014), <https://dl.acm.org/doi/10.5555/2615731.2615742>
36. Wahby, M., Petzold, J., Eschke, C., Schmickl, T., Hamann, H.: Collective change detection: Adaptivity to dynamic swarm densities and light conditions in robot swarms. In: *Artificial Life Conference Proceedings*. pp. 642–649. MIT Press, Cambridge, MA (2019). https://doi.org/10.1162/isal_a_00233