

Miscommunication between robots can improve group accuracy in best-of- n decision-making

Raina Zakir¹, Marco Dorigo¹ and Andreagiovanni Reina^{1,2,3}

Abstract—Making fast and accurate consensus decisions through local communication and decentralised control in a swarm of simple robots can be a very challenging endeavour. In swarms of robots with limited capabilities, consensus decisions can be made using simple voting rules. In our study, the robots use rules based on the cross-inhibition model, which describes a voting mechanism observed in the house-hunting honeybee, that has been shown to efficiently allow consensus achievement in distributed robotic systems. The cross-inhibition mechanism has been shown to lead to a highly stable consensus, preventing the correction of possible group decision errors which can happen, for example, due to high noise in robots’ estimations. In this paper, we investigate the impact of miscommunication on the speed-accuracy trade-off in consensus decision-making in the context of a binary discrimination problem—i.e., choosing collectively the best of two options. We evaluate the accuracy of decision-making theoretically, using continuous and finite-size models, and experimentally in a collective perception scenario, using swarms of 100 simulated robots and 50 real Kilobots. Our study suggests that a certain level of miscommunication (or communication noise) among agents can increase the decision’s accuracy and, thus, can serve an important functional role in making collective decisions in robot swarms.

I. INTRODUCTION

In self-organised decentralised systems, a large number of agents need to coordinate to perform tasks or solve problems exclusively based on the local interactions among the agents themselves and the interactions they have with the surrounding environment [1]. Decentralised decision-making, which plays a crucial role in natural systems [2], particularly among social animals [3], [4], [5], holds equal importance for autonomous systems [6]. Collective decision-making algorithms inspired by natural systems have been implemented in synthetic autonomous systems, such as robot swarms, to make consensus decisions in various situations, such as when selecting an aggregation site [7], [8], a common direction of motion [9] or the dominant environmental feature [10], [11]. Frequently, the application scenario imposes restrictions on the type of robot that can be employed in terms of its computational, communication, and memory capabilities. For

instance, size is the main constraint for a swarm of (nano)-robots navigating in blood vessels. Likewise, the available budget is the main constraint for robot swarms operating in hazardous environments, because during each operation, numerous robots could suffer damage and become unusable, and it is therefore important to use cheap, disposable single-use robots. Such constraints require solutions based on minimal algorithms that can run on minimalistic robots.

An important category of collective decision-making problems is the “best-of- n problem” where a group needs to choose the best option among n alternatives that can differ in quality [12]. Assuming that no individual member of a swarm has the capability to make an accurate assessment of the quality of the options, undertaking best-of- n decisions can be challenging for a robot swarm. Several studies have utilised cross-inhibition as a decision-making algorithm for minimalistic robot swarms, which combine environmental and social information to solve the best-of- n problem [13], [14], [13]. The crucial aspect of this algorithm is inhibitory signalling, a fundamental mechanism widespread in natural systems, where individuals become uncommitted from any option when they encounter another individual with an opinion about what is the best option conflicting with their own. This inhibitory mechanism is present in many different systems, including cellular metabolism [15], neuron activity [16], honeybee house-hunting [3], and human societies [17]. Previous studies have found cross-inhibition not only to be crucial for biological systems seeking to achieve coordinated actions, but also to be important in artificial systems due to its ability to prevent decision deadlocks. For instance, [18] and [14] showed cross-inhibition’s superiority over the direct-switch model in resolving decision deadlocks in the presence of asocial behaviour, even between options of equal or similar quality. However, this ability to break decision deadlocks comes at a cost. In situations where options have different qualities, the goal is to not only reach a consensus but also to agree on the best option. Yet, when options’ qualities are similar, cross-inhibition can often lead to a consensus in favour of the inferior option due to initial random fluctuations caused by erroneous individual estimates. This is because, thanks to the inhibitory mechanism, the dynamics of the group opinion are highly stable [18]. Therefore, once a large majority of agents reach a consensus for an option, the collective agreement is maintained and is difficult to be overturned, even if it happens to be a consensus for the inferior choice. As a result, cross-inhibition trades accuracy for system stability when options are of similar qualities.

Accuracy is an important and widely used metric for

¹IRIDIA, Université Libre de Bruxelles, Brussels, Belgium (email: raina.zakir@ulb.be, mdorigo@ulb.ac.be)

² Centre for the Advanced Study of Collective Behaviour, Universität Konstanz, Konstanz, Germany (email: andreagiovanni.reina@uni-konstanz.de)

³ Department of Collective Behaviour, Max Planck Institute of Animal Behavior, Konstanz, Germany

R. Zakir, M. Dorigo, and A. Reina acknowledge support from the Belgian F.R.S.-FNRS. A. Reina also acknowledges support from the German DFG under Germany’s Excellence Strategy – EXC 2117 – 422037984.

The authors thank Guillermo Legarda Herranz and David Garzon Ramos for the technical support with the robot demonstrations on the Kilogrid.

assessing collective decision-making mechanisms [19], [20], [10], [13]. Measuring the performance of a swarm of autonomous robots exclusively based on accuracy, i.e., their ability to choose the highest quality option, can be useful in mission-critical applications, especially in hazardous or dangerous situations where choosing a lower quality option can have detrimental consequences. In natural systems, the presence of noise and errors can diminish the accuracy of various biological processes, such as system regulation, environmental adaptation, information communication, and signal detection, which disrupts the potential performance of these systems. However, studies have also shown that a moderate amount of noise, e.g., in the form of information error, can contribute to an improvement in the foraging success of honeybees [21]. Similarly, ants exploit noise in decision-making for efficient exploitation of food sources [22]. Humans can also benefit from noise, e.g. caused by anti-conformist behaviour, to collectively rescue a population from states of unproductive conformity [23]. The inevitability of noise in real-world scenarios makes it essential to acknowledge that artificial systems like robot swarms are subject to such disturbances during their operations. It is, therefore, important to study how to design collective robotics systems that, similar to their biological counterpart, can exploit errors and disturbances to enhance their collective performances rather than suffer from them. The property of benefiting from disturbances is also called *antifragility* [24].

As a first step in this direction, in this study, we analyse the influence of noise on the accuracy of the cross-inhibition algorithm in making best-of-2 decisions. We study the algorithm in the same collective perception scenario explored in [14] and [25], whereby a robot swarm is tasked with making a consensus decision between two environmental features (presented in Sec. II). We introduce noise in the system through miscommunication among the robots. We analyse the dynamics of the system with ordinary differential equations (ODEs) in the limits of infinitely-large swarms, and with master equations based on a chemical reaction network for finite-size swarms (described in Sec. III). The Gillespie algorithm [26] is used to numerically simulate this network. Our results, presented in Sec. IV, surprisingly reveal that noise can improve the accuracy of a robot swarm by reducing the probability of selecting the inferior option, at the expense of an increase in decision time. Using ODEs, we find the minimum amount of noise required to shift the consensus from the inferior option to the higher quality option to improve the accuracy. We confirm our theoretical predictions through a set of robot simulations and real robot demonstrations. In Sec. V, we discuss and conclude the study.

II. THE DECISION-MAKING ALGORITHM

In our analysis, we adapt the algorithm proposed in Zakir et al. [14], which is based on the one of Valentini et al. [19], [27], by including a mechanism to regulate and systematically study the communication noise between robots. We focus on the best-of-2 decision problem where a swarm of S robots must converge to the best option between

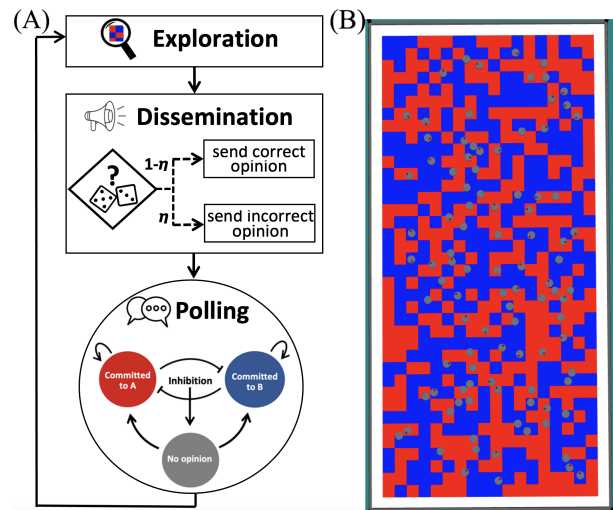


Fig. 1. (A) The finite state machine describing the robots' behaviour consists of the exploration and dissemination states, followed by a polling phase, during which committed robots receiving a message from robots committed to a different option, reset their commitment (they get inhibited). (B) A snapshot of an experiment showing 100 simulated Kilobots (small grey circles) in the ARGoS Kilogrid arena comprising red and blue tiles.

A and B . Both options are associated with a quality value, q_a and q_b , respectively, and the ratio of the two qualities is represented by $q = q_b/q_a$. In our study, we assume, without loss of generality, that $q_a \geq q_b > 0$, and therefore $q \in (0, 1]$.

The robots run an algorithm based on the finite state machine shown in Fig. 1A, that has minimal requirements in terms of computation, communication, and memory capabilities of the robot. Each robot is capable to store and communicate information about one option only, which represents its opinion. The algorithm consists of two alternating states: exploration and dissemination. The robots can be in one of five possible states: A_D (disseminating for opinion A), B_D (disseminating for opinion B), A_E (exploring for opinion A), B_E (exploring for opinion B) and U (uncommitted from any option). At the end of the dissemination state, the robot undergoes a polling phase to update its commitment to an option based on social information it receives from its neighbours, and then switches to the exploration state. In case the robot receives an opinion different from its own during the polling phase, it becomes uncommitted (state U).

Robots in the exploration state (A_E or B_E) evaluate the quality of their current opinion (i.e., option A or B) from the environment. The duration of time a robot stays in the exploration state is determined by a random draw from an exponential distribution with rate $\lambda_e = t_e^{-1}$, which corresponds to the time distribution of a Bernoulli event with probability t_e^{-1} to happen. Therefore, robots leave the exploration state after a mean time t_e and transition to the dissemination state.

Robots in the dissemination state (A_D or B_D) share their opinions (i.e., option A or B) locally with their neighbouring robots. The duration of time a robot disseminates its opinion is determined by a random draw from an exponential distribution with rate $\lambda_d = q_i t_d$; in this way, the average time

spent in the dissemination state is linearly proportional to the estimated option's quality q_i (with $i = \{a, b\}$). Correlating the dissemination duration with option quality induces modulation of positive feedback: the robots that are committed to the best option are more likely to disseminate their opinion for longer times. As a result, a robot is more likely to receive messages from neighbours committed to the higher quality option than from neighbours committed to the inferior option.

Uncommitted robots stay in the uncommitted state U for an average time t_u , during which they do not estimate or disseminate any option.

After the dissemination or uncommitted state, the robots enter the polling phase, which requires them to collect the opinion of one randomly selected neighbour and update their own opinion using the cross-inhibition mechanism. With cross-inhibition (Fig. 1A), when a committed robot (either A_D or B_D) reads a message from a randomly selected robot that is committed to a different option (e.g., a robot committed to option A reads a message from a robot committed to option B), it gets inhibited and becomes uncommitted (U). If an uncommitted robot receives an opinion (either A or B) from one of its disseminating neighbours (A_D or B_D), it will become committed to the corresponding option. After polling, the robot returns to the exploration state to resume the cycle as A_E , B_E , or U depending on the polling outcome.

In our system, we model noise in the form of information miscommunication, i.e., the robots committed to an option have a probability η to flip the opinion they disseminate to the neighbouring robots. That is, a robot with opinion A (resp. B) in the dissemination state A_D (resp. B_D) will occasionally disseminate a message for the opposing option B (resp. A) with probability η .

III. EXPERIMENTAL METHODS

With our algorithm defined, we aim to understand the influence of miscommunication (communication noise η) on the system's accuracy across various q by using tools of nonlinear dynamics and large-scale robotics experiments. In this section, we define the models and the experimental setup, and in the following section, we report the analysis results.

A. The ODEs

Let a_d, a_e, b_d, b_e , and u be the proportion of robots in the states A_D, A_E, B_D, B_E , and U , respectively. We study the behaviour of the model when the number of agents tends to infinity ($S \rightarrow \infty$), using the following system of ODEs:

$$\begin{aligned} \frac{da_d}{dt} &= -\frac{a_d}{q_a t_d} + \frac{a_e}{t_e} \\ \frac{db_d}{dt} &= -\frac{b_d}{q_b t_d} + \frac{b_e}{t_e} \\ \frac{da_e}{dt} &= \frac{(a_d - a_d \eta + b_d \eta)}{(a_d + b_d)} \left(\frac{a_d}{q_a t_d} + \frac{u}{t_u} \right) - \frac{a_e}{t_e} \\ \frac{db_e}{dt} &= \frac{(b_d + a_d \eta - b_d \eta)}{(a_d + b_d)} \left(\frac{b_d}{q_b t_d} + \frac{u}{t_u} \right) - \frac{b_e}{t_e} \end{aligned}$$


where $u = 1 - a_d - a_e - b_d - b_e$.

The proportions of robots a_d and b_d (respectively disseminating options A and B) increase when robots complete the exploration of the respective option, thus after an average time t_e they leave states a_e and b_e , respectively. The proportions a_d and b_d diminish when robots complete the dissemination to start exploration, which happens at a rate inversely proportional to the estimated quality, i.e., $1/(q_a t_d)$ and $1/(q_b t_d)$, respectively.

The proportions of exploring robots increase when committed robots terminate dissemination (at the rates indicated before) or uncommitted robots have spent an average time t_u in the uncommitted state (i.e., at rate $1/t_u$). Whether the robot starts the exploration of option A or B (increasing proportion a_e or b_e) depends on the message it reads in the polling phase (happening at the end of dissemination). Note that while uncommitted robots always move to exploration (either for A or B when they read an A or B message, respectively), the committed robots only begin exploration when they read a message supporting their option, otherwise they become uncommitted. The proportions of exploring robots a_e or b_e decrease at rate $1/t_e$, i.e., after an average time t_e .

B. Chemical reaction network

Although distributed systems typically consist of numerous agents, the finite size of these systems ($S < \infty$) can exert a significant influence on their dynamics [28], [29]. To understand the finite-size effects on our model, we use the formalism of chemical reaction networks and numerically approximate the solution of the master equations using Gillespie's stochastic simulation algorithm [26]:



To study the system at convergence in the Gillespie simulations, the initial agent distribution of the simulations is set to $t_d/(t_e+t_d)$ agents in the dissemination state, divided equally between A_D and B_D , and the remaining agents equally allotted between A_E and B_E .

C. Simulated robot experiments

Large-scale robotic experiments serve as a crucial means to validate self-organised decision-making strategies. While stochastic Gillespie simulations allow quick computation of the dynamics of a discrete number of interacting agents, they may not fully capture the intricacies and real-world constraints of a robotic system. Hence, to provide a more convincing verification of transferability, we implemented

our decision-making algorithm on a physics-based simulator called ARGoS [30], and also ran real-robot demos. All the code used for this study, including the mathematical models and the algorithms run by the robots, are open-source and available as supplementary material [31].

Scenario: We test the decision-making algorithm in a collective perception scenario where the robots are tasked with selecting the most frequent element in the environment, in our case, the most frequent colour of the floor which is composed of randomly distributed red and blue tiles (see Fig.1B). Therefore, the colours are the two options— A is red, B is blue—and the abundance of each colour represents its quality.

Simulation setup: We use a swarm of 100 low-cost and small-sized simple robots called Kilobots [32]; they communicate with each other in a range of 10 cm through infrared (IR), and their control loop operates at an approximate interval of 32 ms. We employ the Kilogrid system [33], an interactive and customisable environment specifically designed for the Kilobots. The Kilogrid consists of 800 square cells covering an area of $2 \times 1 \text{ m}^2$; all the cells, except for those at the borders (shown in white in Fig. 1B), are programmed to continuously transmit IR messages containing their ID and colour. Using the IR messages from the Kilogrid cells, the Kilobots collect information from the environment, to estimate their opinion’s quality (i.e. the proportion of cells of a given colour). Both the Kilobots and the Kilogrid can be simulated in ARGoS through dedicated plugins [30], [34].

Initialisation: At the beginning of each simulation run, we initialise the robots in the exploration state with a random initial opinion, with half of the swarm committed to option A (in state A_E), and the other half committed to option B (in state B_E). We deploy the robots on the Kilogrid at uniformly random initial positions. The distribution of Kilogrid tiles is randomly regenerated for each run. Each run lasts $T = 200000$ simulation timesteps, equivalent to 110 minutes.

Robot behaviour: To explore different areas of the environment and interact with different robots, the Kilobots perform a random walk in the environment. As the Kilobots do not have any proximity sensors, the Kilogrid cells also transmit a binary ‘wall flag’ (either high or low) to indicate proximity to a wall. The white cells at the borders and the non-white cells adjacent to them send a high wall flag, while all the other internal cells send a low wall flag. Whenever a robot receives a high wall flag from the Kilogrid, it executes a simple obstacle avoidance routine, irrespective of its current state, to avoid a collision with the wall.

A robot in the exploration state only estimates the quality of the option to which it is committed; to do so it reads Kilogrid messages to keep track of the number T of cells encountered and the count C_i of cells that match its current opinion i . The robot ensures that each cell is counted only once by using the cell’s ID. At the end of the exploration cycle, which lasts on average $\lambda_e^{-1} = t_e = 2800$ control cycles (approximately 93 seconds), the robot estimates the quality of its opinion: $q_i = C_i/T$. The values T and C_i represent the counts obtained during a single exploration cycle and are




Fig. 2. Strip plots overlaid on violin plots illustrating 6000 individual quality estimates made by simulated robots for option A (red) and B (blue) for three distinct q values. The yellow dots on the plot represent the actual proportion of red and blue tiles in the Kilogrid environment.

reset before entering the dissemination state. These quality estimates, made by the robots, are subject to noise, as depicted in Fig. 2, which shows that the estimation noise, caused by a limited number of cell readings, frequently leads the robots to perceive a quality value that is higher or lower than the true value.

Based on the estimated value of q_i , the robot calculates its dissemination time using an exponential distribution with $\lambda_d^{-1} = q_i t_d$; in our setup, $t_d = 1800$ control cycles (approximately 60 seconds). The robots influence each other by communicating for a period of time whose length is proportional to the estimates q_i . Despite this modulation of positive feedback relying on inaccurate estimations (Fig. 2), the robots can effectively reach a consensus for the best option (results shown in Sec. IV). Disseminating robots broadcast in a local range (10 cm) IR messages indicating their opinion (these messages represent the robot’s votes for either option). When a robot receives a message, it stores its content (i.e., a vote for option A or B). Because robots have minimal memory, they only store the last received vote by overwriting previous messages. Once dissemination is terminated, the robot starts the polling phase: it reads the last message that it received from its neighbours and uses it to update its opinion. After the polling phase, the robot returns to the exploration state to make a new estimate of the option’s quality. Uncommitted robots move randomly in the environment for an average of $t_u = 800$ control cycles (approximately 26 seconds) without making any quality estimate or sending any message. Then, they move to the polling phase.

The average durations of exploration, dissemination, and uncommitment (t_e , t_d and t_u) have been chosen considering the motion speed of the robots. As Kilobots move forward at 1 cm/s , rotate at $45^\circ/\text{s}$, and communicate in a range of 10 cm, we have chosen time lengths that allow the robots to obtain meaningful estimates, and to encounter a diversity of neighbours while disseminating their opinion (in this way, also approximating a well-mixed system).

IV. THE ANALYSIS

A. Predictions from the mathematical models

We analyse the mathematical models to study the influence of the communication noise η on the decision-making algorithm by identifying when a population can reach an agreement or remains polarised between the two options. To do so, we study the long-term state of the system ($t \rightarrow \infty$) by computing the equilibria of the ODE system and recording

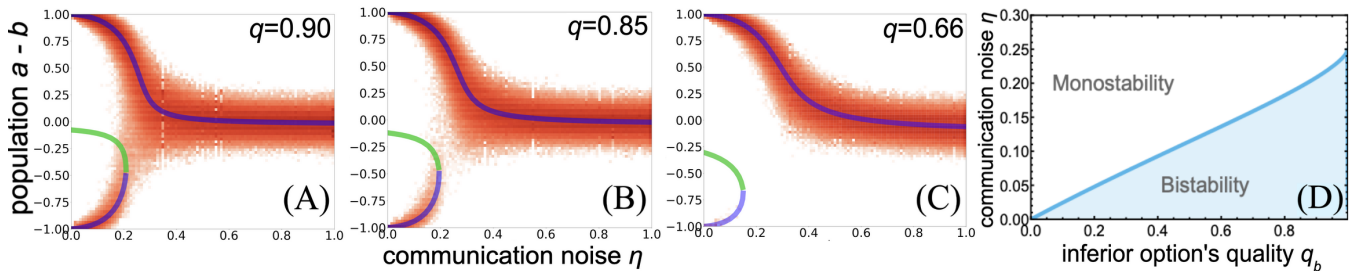


Fig. 3. (A-C) Long-term distribution ($t \rightarrow \infty$) of robots supporting option A or option B (y-axis) for communication noise $\eta \in [0, 1]$ (x-axis) and option quality ratio $q \in \{0.9, 0.85, 0.66\}$. The 2D histograms are computed via the Gillespie simulations by subtracting the proportion b of robots supporting option B from the proportion a of supporters for A (y-axis). The overlying lines show the fixed points of the ODE system (green lines are unstable, blue lines are stable). (D) The blue line shows the bifurcation point, i.e., the value of η for which the lower branches shown in panels A-C disappear, for $q_a = 1$ and $q_b \in [0, 1]$. The bifurcation marks a phase transition from bistability (possible collective mistakes) to monostability (accurate collective decisions) happening at moderate levels of communication noise η . However, as shown in (A-C), too much noise can hamper consensus.

the final state of the Gillespie simulations with $N = 100$ agents. The values of t_d and t_e in both the ODEs and in the Gillespie simulations are fixed to the same value as in the robot simulations (see Sec. III C). We carry out 10^4 Gillespie runs for various η across four different q values. We record the proportion of agents $a = a_d + a_e$ and $b = b_d + b_e$ (supporting options A and B, respectively) at the final step ($T = 200000$) of each run to compute $a - b$ averaged across all the runs, and report the results as 2D histograms in Figs. 3A-C (colour maps). On top of the histograms, we show the bifurcation diagram computed from the ODE system. The agreement between ODEs and Gillespie simulations is good.

Regardless of the quality difference q , we notice the bistability of the system at $\eta = 0$, which represents the situation where the swarm makes a consensus decision for either option, compromising group accuracy. Counterintuitively, as we increase the communication noise η , the lower branch in favour of the inferior option disappears, therefore predicting an increase in collective accuracy. We also notice that the minimum amount of miscommunication η to transition from bistability to monostability, and make the lower branch disappear, increases with the difficulty of the decision problem (i.e., higher quality ratios q). The critical value of η at which the transition from bistability to monostability occurs is a bifurcation point that can be represented as a function of q (considering for simplicity $q_a = 1$). While we could compute the symbolic equation of the bifurcation point (provided in the supplementary material [31]), its complexity makes it impractical for analysis; therefore, we resort to numerical computation to compute the stability diagram of Fig. 3D.

Figs. 3A-C also show that the consensus decision for the best option is only sustained for a limited range of values of η after the bifurcation point. As the value of η increases, the system ultimately has a single stable equilibrium with a population split 50-50 between the two options, indicating a state of indecision where neither population reaches a majority. For harder decision problems (options with similar qualities, i.e., higher values of q), the system enters the indecisive state at lower levels of communication noise than for simpler problems. Therefore, the window between bifurcation point and onset of indecision expands for simpler

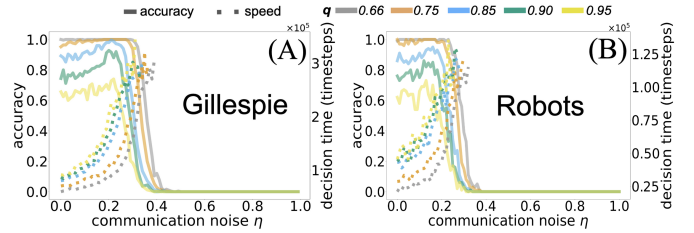


Fig. 4. Results from Gillespie (A) and robot (B) simulations with a swarm of $S = 100$ robots show the relationship between noise η (x-axis) and speed/accuracy (y-axis) for five quality ratios q .

decision problems (low q).

B. Speed vs accuracy

We employ both Gillespie and robot simulations to analyse the relationship between decision accuracy and decision time. We compute accuracy by calculating the proportion of runs, under a given configuration of q and η , that successfully reached the quorum $Q = 0.7$ for the best option, out of the entire set of executed runs. A swarm is considered to have reached Q for option A when 70% of the robots are committed to A ($a_d + a_e \geq 0.7$). We compute decision time as the average number of timesteps taken to reach the quorum Q for the best option A. When, within the designated time limit T , the best option is selected in fewer than 10% of the runs, we omit the time results (insufficiently reliable data).

Fig. 4 shows the speed vs accuracy analysis of 250 Gillespie runs (in panel A) and 50 simulated robot runs (in panel B) for five values of q and $\eta \in [0, 1]$. In general, we can observe a qualitative agreement between the Gillespie and robot simulations. For all the tested values of q , the accuracy increases with increasing miscommunication η , up to a certain value after which performances drop. The improvement in accuracy is more accentuated for difficult problems (high q) than for easier problems (low q) where the collective accuracy remains equally high for a large range of values of η . For instance, in both robot and Gillespie simulations, for $q = 0.66$ the collective accuracy is maximum across all η , and similarly, the improvement in accuracy is minimal for $q = 0.75$. Instead, in situations with high q

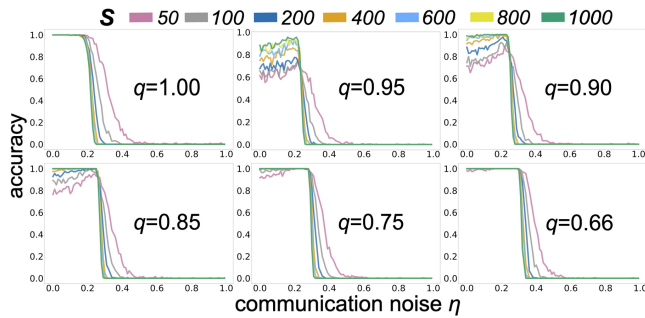


Fig. 5. The accuracy plots, obtained from Gillespie simulations, depict the relationship between noise η (x-axis) and accuracy (y-axis) for different swarm sizes (S) across six distinct quality ratios (q).

values ($q \geq 0.85$), a comparison of accuracy at $\eta = 0$ and $\eta = 0.2$ reveals a notable increase in accuracy, ranging from 5% to 12% improvement. On the other hand, this increase in accuracy comes at the expense of an increase in decision time (see dotted lines in Fig. 4). Therefore, the communication noise η acts as the control parameter to regulate a speed-accuracy trade-off.

C. Our interpretation of the phenomena

Communication noise acts as a perturbation on the collective system, which gets ‘shaken’ and oscillates more around the attraction point. Because the basin of attraction of the equilibrium representing a consensus for the inferior option is smaller than the equilibrium for the correct decision (see Figs. 3A-C), a perturbed system can move from the former to the latter and not vice-versa. Without noise, once the swarm reaches a full (100%) consensus, no robot ever changes its opinion. Moderate levels of system oscillations let the system escape the suboptimal equilibrium, improving accuracy. When perturbations increase, the system goes to an indecision state, as the oscillations are too big to let the system keep any large majority. Other works have looked at other methods to perturb systems composed of robots running algorithms similar to ours, e.g., by including a minority of stubborn robots [35], [18] or letting robots randomly sample the environment [14], [36], [37], showing that these methods are useful to adapt to environmental change, e.g., when the qualities of the two options swap.

D. Scalability

We further extend the analysis and use Gillespie simulations to study how accuracy changes when the swarm size increases. In Fig. 5, we plot the decision accuracy for different swarm sizes (from $S = 50$ to $S = 1000$ robots) as a function of $\eta \in [0, 1]$ for various quality ratios q . Fig. 5 shows that as the system size increases, the probability of choosing the inferior quality option diminishes across all q . When S is small (e.g., $S = 50$), the accuracy has a linear increase till the transition point, showcasing an improvement in accuracy. However, as the swarm size increases, the accuracy plateaus at its maximum value and transforms into a step function, suggesting that larger swarm sizes

allow the system to achieve and maintain optimal accuracy consistently. For instance, for $S = 1000$ and easy decision problems ($q \leq 0.85$), the accuracy remains maximal till the transition point, after which it drops to 0 as the quorum Q is never attained. Nevertheless, even for large S , in difficult decision problems ($q > 0.85$) the accuracy is not maximal at $\eta = 0$, leading to an improvement in accuracy, albeit small, as miscommunication increases. Higher accuracy for larger S is due to the effect of the stochastic fluctuations: as the system size S is larger, the stochastic fluctuations are smaller [38]. When the initial starting conditions of robots committed to options A and B are symmetric, swarms with smaller sizes have larger random fluctuations and, therefore, are more likely to fall in the basin of attraction of the fixed point for consensus for the inferior option (lower branch in Fig. 3). While Gillespie simulations allow us to study the effects of size-dependent random fluctuations (showing increased performance with increasing S), the deployment of a robotic system should also consider the effect of physical interactions and interference among the robots. In fact, physical interference in large swarms can congest the system and limit robot mobility, leading to a decrease, rather than an increase, of the group performance [29], [39]. Fig. 5 also shows that, thanks to cross-inhibition, the swarm is consistently capable of breaking the symmetry and selecting any of the two options when they have the same quality ($q = 1$), in accordance with the existing literature [14], [18], [40], [41].

E. Real robot demonstrations

To validate our analysis, we ran a set of real-robot demonstrations with a swarm of 50 Kilobots on a square $1 \times 1 \text{ m}^2$ Kilogrid arena (Fig. 6A), using the same setup described in Sec. III.C. Each demonstration started with the swarm opinion equally split between the two options, i.e., 25 robots committed to red and 25 committed to blue. The demonstration is terminated when the quorum $Q = 0.7$ for either option is maintained for one minute or after 30 minutes. To track the robots’ commitment over time, each Kilobot sends every 2 seconds a message with its state to the Kilogrid which logs all data. The robots also showed their opinion to us via their coloured LEDs (option A as red, option B as blue, and no opinion as green). All the demonstration videos are available as supplementary material [31]. We ran a total of five demonstrations with $q = 0.85$ and $\eta = 0.15$. In all runs, the swarm successfully selected the higher quality option A (Fig. 6B). By implementing the algorithm on a swarm of physical robots, we confirm the possibility of using the proposed algorithm for robotics systems even when communication noise is high (in our demos, 15% of the exchanged messages convey the wrong information).

V. DISCUSSION AND CONCLUSION

Information noise, which masks and alters the original information content, is intrinsic to any collective process of group-living organisms that exchange information for coordination, from humans to animals to cells. Indeed, the




Fig. 6. (A) Robot demonstration with 50 Kilobots for $q = 0.85$ and $\eta = 0.15$. (B) All five robot demos led to a consensus for the best option (red). The videos are provided as supplementary material [31].

impact of noise is the subject of numerous studies in opinion dynamics [42], [43], [44], [45], social sciences [23], [46], behavioural ecology [47], [48], [49], and cognitive neuroscience [50], [16]. Noise is intrinsic in robotics too and we are seeing a growing number of swarm robotics studies taking noise as the central component of their analysis [51], [52], [53]. In this study, we consider a type of communication noise where robots send a message with incorrect information with a given probability. We quantify the effect of such noise on the accuracy and speed of collective best-of-2 decisions. Our analysis reveals that group accuracy in selecting the best option is higher when the robots make a modest number of communication errors than when they do not make any errors (Fig. 4). However, as expected, too much miscommunication hampers group agreement. Therefore, this study shows that communication noise, as a form of social noise between the individuals of a population, can help the group in selecting the best available alternative.

In order to make a collective decision, our robots ran a weighted voting algorithm based on the cross-inhibition model [54], [18]. Despite the decision-making algorithm being based on minimal robot requirements (limited computation, memory, communication), and the individual robots making highly erroneous estimates of the option's qualities (Fig. 2), the swarm could, most of the times, collectively select the best option. Having such simple algorithms operating on minimalistic devices can be useful for application scenarios where the robots are extremely small (e.g., nano-robots) or extremely cheap (e.g., disposable robots). In addition to inaccurate cheap sensors, the cause of highly noisy measurements can also be faulty components or malicious intrusions. Indeed, the robustness to individual estimation errors and the ability to collectively filter such noise are among the main factors that led to the first implementation of these types of bio-inspired voting algorithms for controlling robot swarms [55], [56]. The inspiration comes from house-hunting social insects, such as honeybees or ants, which also face the same challenge of noise-prone individual estimates [57]. These biological systems equally face the problem of noisy signalling and communication errors [58]. Interestingly, our results show that moderate levels of communication error can be beneficial, rather than detrimental, to the collective performance. Our study also shows that this performance improvement (higher accuracy) comes at the cost of slower decisions (Fig. 4). Such speed-accuracy trade-off, commonly found in decision-making [19], [20], is, in our case, regulated

by the noise level η .

Our previous studies have shown that the cross-inhibition algorithm is an effective method for making value-based decisions [18] and quickly reaching a consensus (quicker than other similar methods [59]) even when options have similar qualities [13], asocial minorities hindering group consensus are present [18], [59], or individuals deviate from social norms [14]. This paper shows that its performance can increase through moderate levels of miscommunication. Future work may investigate whether our results scale to larger dimensions of the decision problem (best-of- n with $n > 2$) [40], [52].

REFERENCES

- [1] H. Hamann, *Swarm Robotics: A Formal Approach*. Cham: Springer International Publishing, 2018.
- [2] T. Bose, A. Reina, and J. A. R. Marshall, "Collective decision-making," *Current Opinion in Behavioral Sciences*, vol. 6, pp. 30–34, 2017.
- [3] T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall, "Stop signals provide cross inhibition in collective decision-making by honeybee swarms," *Science*, vol. 335, no. 6064, pp. 108–111, 2012.
- [4] V. H. Sridhar, L. Li, D. Gorbonos, M. Nagy, B. R. Schell, T. Sorochkin, N. S. Gov, and I. D. Couzin, "The geometry of decision-making in individuals and collectives," *Proceedings of the National Academy of Sciences*, vol. 118, no. 50, p. e2102157118, 2021.
- [5] L. Conradt and C. List, "Group decisions in humans and animals: A survey," *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 364, no. 1518, pp. 719–742, 2009.
- [6] A. Reina, E. Ferrante, and G. Valentini, "Collective decision-making in living and artificial systems: editorial," *Swarm Intelligence*, vol. 15, p. 1–6, 2021.
- [7] N. Correll and A. Martinoli, "Modeling and designing self-organized aggregation in a swarm of miniature robots," *The International Journal of Robotics Research*, vol. 30, no. 5, pp. 615–626, 2011.
- [8] A. Sion, A. Reina, M. Birattari, and E. Tuci, "Controlling robot swarm aggregation through a minority of informed robots," in *Swarm Intelligence (ANTS 2022)*, ser. LNCS, vol. 13491. Cham: Springer, 2022, pp. 91–103.
- [9] G. Öntür, A. E. Turgut, and E. Şahin, "Predictive search model of flocking for quadcopter swarm in the presence of static and dynamic obstacles," *Swarm Intelligence*, vol. 18, pp. 187–213, 2024.
- [10] J. T. Ebert, M. Gauci, F. Mallmann-Trenn, and R. Nagpal, "Bayes bots: Collective Bayesian decision-making in decentralized robot swarms," in *2020 IEEE International Conference on Robotics and Automation (ICRA)*. Piscataway, NJ: IEEE, 2020, pp. 7186–7192.
- [11] Q. Shan and S. Mostaghim, "Discrete collective estimation in swarm robotics with distributed bayesian belief sharing," *Swarm Intelligence*, vol. 15, no. 4, pp. 377–402, 2021.
- [12] 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.
- [13] M. S. Talamali, J. A. R. Marshall, T. Bose, and A. Reina, "Improving collective decision accuracy via time-varying cross-inhibition," in *2019 International conference on robotics and automation (ICRA)*. Piscataway, NJ: IEEE, 2019, pp. 9652–9659.
- [14] R. Zakir, M. Dorigo, and A. Reina, "Robot swarms break decision deadlocks in collective perception through cross-inhibition," in *Swarm Intelligence (ANTS 2022)*, ser. LNCS, vol. 13491. Cham: Springer, 2022, pp. 209–221.
- [15] N. R. Nené, J. Garca-Ojalvo, and A. Zaikin, "Speed-dependent cellular decision making in nonequilibrium genetic circuits," *PLoS ONE*, vol. 7, no. 3, p. e32779, 2012.
- [16] L. Oscar, L. Li, D. Gorbonos, I. D. Couzin, and N. S. Gov, "A simple cognitive model explains movement decisions in zebrafish while following leaders," *Physical Biology*, vol. 20, no. 4, p. 045002, 2023.
- [17] A. Franci, A. Bizyaeva, S. Park, and N. E. Leonard, "Analysis and control of agreement and disagreement opinion cascades," *Swarm Intelligence*, vol. 15, p. 47–82, 2021.

- [18] A. Reina, R. Zakir, G. De Masi, and E. Ferrante, "Cross-inhibition leads to group consensus despite the presence of strongly opinionated minorities and asocial behaviour," *Communications Physics*, vol. 6, p. 236, 2023.
- [19] G. Valentini, E. Ferrante, H. Hamann, and M. Dorigo, "Collective decision with 100 Kilobots: Speed versus accuracy in binary discrimination problems," *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 3, pp. 553–580, 2016.
- [20] M. Raoufi, H. Hamann, and P. Romanczuk, "Speed-vs-accuracy trade-off in collective estimation : An adaptive exploration-exploitation case," in *2021 International Symposium on Multi-Robot and Multi-Agent Systems (MRS)*. Piscataway, NJ: IEEE, 2022, pp. 47–55.
- [21] R. Okada, H. Ikeno, T. Kimura, M. Ohashi, H. Aonuma, and E. Ito, "Error in the honeybee waggle dance improves foraging flexibility," *Scientific reports*, vol. 4, p. 4175, 2014.
- [22] A. Dussutour, M. Beekman, S. Nicolis, and B. Meyer, "Noise improves collective decision-making by ants in dynamic environments," *Proceedings of the Royal Society of London series B*, vol. 276, pp. 4353–4361, 2009.
- [23] W. Toyokawa and W. Gaissmaier, "Conformist social learning leads to self-organised prevention against adverse bias in risky decision making," *eLife*, vol. 11, p. e75308, 2022.
- [24] C. Axenie, O. López-Corona, M. Makridis, M. Akbarzadeh, M. Savariano, A. Stancu, and J. West, "Antifragility in complex dynamical systems," *npj Complexity*, vol. 1, 2024.
- [25] G. Valentini, D. Brambilla, H. Hamann, and M. Dorigo, "Collective perception of environmental features in a robot swarm," in *Swarm Intelligence (ANTS 2016)*, ser. LNCS, vol. 9882. Springer, 2016, pp. 65–76.
- [26] D. T. Gillespie, "Exact stochastic simulation of coupled chemical reactions," *The Journal of Physical Chemistry*, vol. 81, no. 25, pp. 2340–2361, 1977.
- [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 14th International Conference on Autonomous Agents and Multiagent Systems*, ser. AAMAS '15. Richland, SC: IFAAMAS, 2015, pp. 1305–1314.
- [28] Y. Khaluf, C. Pinciroli, G. Valentini, and H. Hamann, "The impact of agent density on scalability in collective systems: noise-induced versus majority-based bistability," *Swarm Intelligence*, vol. 11, no. 2, pp. 155–179, 2017.
- [29] J. Kuckling, R. Luckey, V. Avrutin, A. Vardy, A. Reina, and H. Hamann, "Do we run large-scale multi-robot systems on the edge? More evidence for two-phase performance in system size scaling," in *IEEE International Conference on Robotics and Automation (ICRA 2024)*. Piscataway, NJ: IEEE, 2024, pp. 4562–4568.
- [30] C. Pinciroli, M. S. Talamali, A. Reina, J. A. R. Marshall, and V. Trianni, "Simulating Kilobots within ARGoS: models and experimental validation," in *Swarm Intelligence (ANTS 2018)*, ser. LNCS, vol. 11172. Cham: Springer, 2018, pp. 176–187.
- [31] R. Zakir. (2023, Jul.) Supplementary material for paper "Miscommunication between robots can improve group accuracy in best-of-n decision-making". [Online]. Available: Webpage: <https://iridia.ulb.ac.be/supp/IridiaSupp2023-004/index.html>, Repository: <https://github.com/rainazakir/accuracy-ci>, Video: <https://youtu.be/9JXpeT5rKh0>
- [32] M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal, "Kilobot: A low cost robot with scalable operations designed for collective behaviors," *Robotics and Autonomous Systems*, vol. 62, no. 7, pp. 966–975, 2014.
- [33] G. Valentini, A. Antoun, M. Trabattoni, B. Wiandt, Y. Tamura, E. Hocquard, V. Trianni, and M. Dorigo, "Kilogrid: A novel experimental environment for the Kilobot robot," *Swarm Intelligence*, vol. 12, no. 3, pp. 245–266, 2018.
- [34] T. Aust, M. S. Talamali, M. Dorigo, H. Hamann, and A. Reina, "The hidden benefits of limited communication and slow sensing in collective monitoring of dynamic environments," in *Swarm Intelligence (ANTS 2022)*, ser. LNCS. Cham: Springer, 2022, vol. 13491, pp. 234–247.
- [35] J. Prasetyo, G. De Masi, and E. Ferrante, "Collective decision making in dynamic environments," *Swarm intelligence*, vol. 13, no. 3-4, pp. 217–243, 2019.
- [36] M. D. Soorati, M. Krome, M. Mora-Mendoza, J. Ghofrani, and H. Hamann, "Plasticity in collective decision-making for robots: Creating global reference frames, detecting dynamic environments, and preventing lock-ins," in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. Piscataway, NJ: IEEE, 2019, pp. 4100–4105.
- [37] M. S. Talamali, A. Saha, J. A. R. 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.
- [38] N. G. van Kampen, *Stochastic processes in physics and chemistry*. Amsterdam, NL: Elsevier, 1992.
- [39] H. Hamann and A. Reina, "Scalability in computing and robotics," *IEEE Transactions on Computers*, vol. 71, no. 6, pp. 1453–1465, 2022.
- [40] A. Reina, J. A. R. Marshall, V. Trianni, and T. Bose, "Model of the best-of-N nest-site selection process in honeybees," *Physical Review E*, vol. 95, no. 5, p. 052411, 2017.
- [41] D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. R. Marshall, "A mechanism for value-sensitive decision-making," *PLoS ONE*, vol. 8, no. 9, p. e73216, 2013.
- [42] B. L. Granovsky and N. Madras, "The noisy voter model," *Stochastic Processes and their applications*, vol. 55, no. 1, pp. 23–43, 1995.
- [43] F. Herrerías-Azcué and T. Galla, "Consensus and diversity in multistate noisy voter models," *Physical Review E*, vol. 100, no. 2, p. 022304, 2019.
- [44] Z. Liu, M. Crosscombe, and J. Lawry, "Imprecise fusion operators for collective learning," in *ALIFE 2021: The 2021 Conference on Artificial Life*. Cambridge, MA: MIT Press, 2021, p. 77.
- [45] M. Salahshour, "Phase diagram and optimal information use in a collective sensing system," *Physical Review Letters*, vol. 123, no. 6, p. 068101, 2019.
- [46] H. Shirado and N. A. Christakis, "Locally noisy autonomous agents improve global human coordination in network experiments," *Nature*, vol. 545, no. 7654, pp. 370–374, 2017.
- [47] L. S. Tsimring, "Noise in biology," *Reports on progress in physics. Physical Society*, vol. 77, no. 2, p. 026601, 2014.
- [48] J. A. R. Marshall, A. Dornhaus, N. Franks, and K. Tim, "Noise, cost and speed-accuracy trade-offs: Decision-making in a decentralized system," *Journal of the Royal Society, Interface*, vol. 3, pp. 243–54, 2005.
- [49] J. Jhavar, R. G. Morris, U. R. Amith-Kumar, M. Danny Raj, T. Rogers, H. Rajendran, and V. Guttal, "Noise-induced schooling of fish," *Nature Physics*, vol. 16, no. 4, pp. 488–493, 2020.
- [50] A. Pirrone, A. Reina, and F. Gobet, "Input-dependent noise can explain magnitude-sensitivity in optimal value-based decision-making," *Judgment and Decision Making*, vol. 16, no. 5, pp. 1221–1333, 2021.
- [51] I. Rausch, A. Reina, P. Simoens, and Y. Khaluf, "Coherent collective behaviour emerging from decentralised balancing of social feedback and noise," *Swarm Intelligence*, vol. 13, no. 3–4, pp. 321–345, 2019.
- [52] C. Lee, J. Lawry, and A. F. T. Winfield, "Negative updating applied to the best-of-n problem with noisy qualities," *Swarm Intelligence*, vol. 15, no. 1-2, pp. 111–143, 2021.
- [53] M. Raoufi, P. Romanczuk, and H. Hamann, "Individuality in swarm robots with the case study of Kilobots: Noise, bug, or feature?" in *ALIFE 2023: Proceedings of the 2023 Artificial Life Conference*. Cambridge, MA: MIT Press, 2023, pp. 35–44.
- [54] A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni, "A design pattern for decentralised decision making," *PLoS ONE*, vol. 10, no. 10, p. e0140950, 2015.
- [55] C. A. C. Parker and H. Zhang, "Active versus passive expression of preference in the control of multiple-robot decision-making," in *IEEE International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2005, pp. 3706–3711.
- [56] —, "Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem," *IEEE/ASME Transactions on Mechatronics*, vol. 14, no. 2, pp. 240–251, 2009.
- [57] T. D. Seeley, P. K. Visscher, and K. M. Passino, "Group decision making in honey bee swarms: When 10,000 bees go house hunting, how do they cooperatively choose their new nesting site?" *American Scientist*, vol. 94, no. 3, pp. 220–229, 2006.
- [58] K. Preece and M. Beekman, "Honeybee waggle dance error: adaption or constraint? Unravelling the complex dance language of honeybees," *Animal Behaviour*, vol. 94, pp. 19–26, 2014.
- [59] F. Canciani, M. S. Talamali, J. A. R. Marshall, T. Bose, and A. Reina, "Keep calm and vote on: Swarm resiliency in collective decision making," in *Proceedings of Workshop Resilient Robot Teams of the 2019 IEEE International Conference on Robotics and Automation (ICRA 2019)*, 2019.