

Collective Robustness of Heterogeneous Decision-Makers Against Stubborn Individuals

Nemanja Antonic

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium;
DEIB, Politecnico di Milano, Milano, Italy;
naXys, Université de Namur, Namur, Belgium
nemanja.antonc@unamur.be

Marco Dorigo

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
mdorigo@ulb.ac.be

Raina Zakir

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
raina.zakir@ulb.be

Andreagiovanni Reina

IRIDIA, Université Libre de Bruxelles, Brussels, Belgium;
CASCB, Max Planck Institute of Animal Behavior,
University of Konstanz, Konstanz, Germany
andreagiovanni.reina@gmail.com

ABSTRACT

Can heterogeneity be a cost-effective solution for swarm robotics? Motivated by what we see in animal groups, especially eusocial insect colonies, that exploit behavioural heterogeneity as the cornerstone of their success, we investigate whether or not swarms of robots with different behaviours can be more cost-effective than homogeneous swarms. We focus on the process of collective decision-making where robots must achieve a consensus on the best alternative between two options with different qualities, the best-of-2 problem. We consider four behaviours from the literature where robots use rules of voter-like models to exchange and update their opinions. We study the swarm's ability to be robust to the presence of zealots, i.e., stubborn robots that do not change their opinions. Our analysis is based on mean-field models that describe the change of the sub-populations holding different opinions. We show that heterogeneous swarms can be more efficient when we include in the analysis the cost of social interactions between robots. Normally, more interactive behaviours (e.g., pooling many neighbours' opinions at each timestep rather than one per timestep) are quicker in making a decision and more robust to zealots. Heterogeneous swarms combine high performance with lower costs, as not the entire group must be highly interactive to maximise collective performance. Our results are useful when seeking a balance between making accurate collective decisions and minimising the cost of social interactions, the objective of artificial and natural swarms.

KEYWORDS

Collective robustness; collective decision making; swarm robotics; group heterogeneity

ACM Reference Format:

Nemanja Antonic, Raina Zakir, Marco Dorigo, and Andreagiovanni Reina. 2024. Collective Robustness of Heterogeneous Decision-Makers Against

Stubborn Individuals. In *Proc. of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024), Auckland, New Zealand, May 6 – 10, 2024*, IFAAMAS, 10 pages.

1 INTRODUCTION

Swarms of autonomous robots, that use local knowledge and local interactions to coordinate their actions and accomplish collaborative tasks, are a technology with tremendous potential [14]. However, finding the rules to design the robot behaviours that generate the desired collective outcome is a hard challenge [19]. This study focuses on the behaviour of robots that exchange opinions in order to make a consensus decision on the option with the highest quality among n alternative options (best-of- n problem [57]); we focus our analysis on binary decisions, $n = 2$. Several of the algorithms developed to solve this problem are inspired by biology and social sciences [5, 42]; e.g., house-hunting social insects [41, 49] and voter-like models in opinion dynamics [7, 44]. Individual insects, or robots, sample the environment to assess the quality of the option they prefer. They hold one opinion at a time. Their opinion can change when they receive messages from their neighbours who express a different opinion. A collective decision is made when, and if, a majority agrees on one site. It has been shown that the collective decision-making process can be undermined by even a small number of misbehaving robots [30, 31, 45], e.g., damaged or hacked robots [52]. Therefore, the design of swarm robotics algorithms also requires an analysis of their robustness against misbehaving robots [21, 54]. In this study, we consider the presence of a minority of stubborn robots, also called zealots [24, 32, 45], that never change their opinion, despite the social information they receive.

We consider different algorithms that control the way the robots filter and update their opinion. The main difference between these algorithms is the amount of social information that each robot needs to process. We construct mathematical models, in the form of ordinary differential equations (ODEs), to describe how the swarm population splits into sub-populations composed of agents sharing the same opinion. First, we establish a baseline understanding of the robustness of homogeneous swarms, where every robot within the swarm uses the same algorithm. Then, we consider heterogeneous swarms composed of robots running different algorithms in order to assess how varying the proportions of robots using a particular



This work is licensed under a Creative Commons Attribution International 4.0 License.

algorithm affects the robustness of the swarm. In particular, we quantify the performance of the heterogeneous swarms in terms of two quantities, the robustness against stubborn robots (zealots) and the amount of social information that the robots need to process to reach an agreement. The former represents how often a swarm is able to reach a consensus on the best option despite the presence of zealots, while the latter counts the total amount of messages exchanged during the decision-making process. Our analysis shows that algorithms processing more social information are, in general, more robust. By changing the composition of the heterogeneous swarms, we show that it is possible to balance the trade-off between social information processing cost and collective robustness.

The novelty of this work resides in building and analysing models of heterogeneous swarms tasked with solving the best-of-2 problem, in the presence of zealots. In the state of the art, mathematical models for homogeneous robot swarms are developed, as in [44, 45, 58, 60], both with the presence of zealots and without. Our study extends these models to heterogeneous swarms and is motivated by the intuition that heterogeneous swarms could provide better, or equal, performance than homogeneous ones requiring however less social-information processing. If true, such a solution can increase the efficiency of our robot swarms and provide useful insights for designing swarms that are robust to malicious agents. While this work focuses on the engineering perspective of maximising the efficiency of artificial swarms, the results of our analysis can also be relevant to collective animal behaviour research where animal groups have been selected on the same criteria of efficient computation which minimises costs. A possible solution—which we explore in this paper—is exploiting group heterogeneity.

2 RELATED WORK

Best-of- n decision problems are relevant in swarm robotics as they can lead to autonomy in selecting the best course of action in several application scenarios [57]. Numerous algorithms have been developed in order to solve practical instances of this problem, such as in [55, 56, 58, 60]. Existing solutions differ in the capabilities the robots are required to have in terms of memory, computation and communication. For example, there are algorithms to let each robot update the probabilities that each option is the correct one based on environmental and social evidence. These algorithms typically require relatively sophisticated computation [2, 11, 17], frequent communication of messages with a relatively rich content [50], or relatively high memory [27]. Our study focuses on simpler algorithms based on the idea of minimal computing [63].

Our robots have minimal requirements in terms of memory, computation and communication as they can only store one opinion that they share with their neighbours. Solutions to the best-of- n problem that require minimalism are based on voting algorithms that leverage simple voting rules, typically studied in opinion dynamics [7], where each opinion shared by the robots is treated as a vote. One of the computationally simplest algorithms is based on the voter model [9, 20] (also referred to as voter rule), where the robots only read one vote (opinion) from a randomly selected neighbour. The voter model has later been extended to the weighted voter model in [59], where robots solve the best-of- n problem by expressing their vote for a time proportional to the quality of the

communicated option; in this way, the best option is often selected. Another popular voting rule used to engineer simple robot swarms is the local majority rule [12, 18, 26, 33], where the robots select the most voted option among the votes of their neighbours. While this rule has a higher computational cost compared to the voter rule (because the robots must process and pool multiple votes rather than just picking one), it has been shown to enable quicker collective decisions [44, 58]. The majority rule—which selects options that have been voted by more than 50% of the neighbours—can be generalised to use different sub- and super-majority quorums [28], e.g., through the k -unanimity rule [48] or the q -voter rule [31]. In addition to investigating what social information the robot selects (e.g., voter rule, majority rule), previous work has studied how new social information is integrated with the previous opinion. Simple models are based either on direct overwriting of the opinion with the new social information (direct-switch rule [59]) or on robots temporarily dropping any opinion before adopting a new one (cross-inhibition rule [40]). Recent work has compared these two rules highlighting a trade-off between the ability of the cross-inhibition rule to make consensus decisions quickly and the high decision accuracy of the direct-switch rule [45, 61].

Security is a critical aspect of autonomous robot swarms because they are composed of physical devices that operate in the real world and malfunctioning robot swarms can physically harm people and their environment [21]. While there is a growing interest in the swarm robotics community to investigate the security aspect [8, 15, 39, 53, 54, 62], the number of studies is still limited. In this work, we focus on the presence of robots that do not follow the behaviour with which they have been programmed. In the context of collective decision-making, a number of studies considered solutions to this problem based on outlier detection [47] and tamper-proof protocols based on blockchain smart contracts [52, 54], enabling both the identification and neutralisation of the misbehaving robots. These approaches are not in agreement with the minimal computing paradigm as they require sophisticated algorithms; here, instead, we investigate which of the considered algorithms is intrinsically more robust to the presence of misbehaving robots (without aiming to identify or actively neutralise them). We consider relatively simple misbehaviour, comparable to those of the well-functioning robots. Previous work that considered such type of misbehaviour in best-of- n decision-making looked at the impact of zealots, contrarians, and independents [6]. Zealots, also called stubborn robots, are robots that never update their opinion, thus ignoring the opinion of their neighbours [13, 30, 32, 45]. Contrarians, also called hipsters, pick the opinion that is shared by the minority of the neighbouring robots [25, 29]. Finally, independents, also called wishy-washy, adopt a random opinion each time they partake in the voting process [6, 10].

Behavioural heterogeneity can be a functional feature of robotic and natural swarms [16, 34]. However, studying swarms of robots with different behaviours makes the model analysis complicated as any small behavioural difference can lead to large, often hard-to-predict, changes in the collective dynamics. In fact, even unintentional differences, such as different levels of actuation errors, can cause qualitatively different collective responses [38]. Despite

the number of papers on this subject being relatively limited, previous work showed the great potential of heterogeneity with performances superior to homogeneous swarms, e.g., [22, 23]. Our work focuses on behavioural heterogeneity where robots use different rules to pool social information; to the best of our knowledge, this study is the first looking at this aspect.

3 MODELS

We model robots as agents because we do not consider their movement in, their perception of, or their interactions with a physical environment. We only look at the change of opinions of the robots (agents) over time in response to social information. Each agent receives and sends messages expressing its vote in favour of the option to which it is committed (i.e., the option that it believes to be the best). The agent behaviour is composed of two parts: first, an agent filters the opinions it receives from its neighbours using an opinion filtering rule; then, it updates its opinion with an opinion update rule. These rules are described in detail in the next two subsections. In order to reach a consensus on the best option, agents communicate their opinion with a frequency that is proportional to the quality of the opinion they hold. This method is inspired by the behaviour of social insects [35, 45]. Zealot agents, instead, never change their opinions regardless of their neighbours' opinions; we assume zealots have the malicious intent of disrupting the collective decision-making process, therefore they communicate with the maximum frequency regardless of the opinion they propagate.

We study a binary decision case, i.e., $n = 2$, where the two options are A and B, each with an associated quality q_A and q_B . Without loss of generality, we fix option A's quality to $q_A = 1$ and vary option B's quality in the range $q_B \in [0, 1]$; thus, the difficulty of the decision problem is computed as the quality ratio $q = \frac{q_B}{q_A}$. The variables A and B indicate the proportions of non-zealot agents holding opinion A and opinion B, respectively, which change over time. The proportions of zealots with opinions A and B are two additional parameters: z_a and z_b , representing the percentage of zealots with respect to the total number of collaborative agents. Therefore, in a swarm composed of S susceptible (non-zealot) agents, the numbers of zealots with opinion A and B are $z_a S$ and $z_b S$, respectively. Both z_a and z_b have values in the range $[0, 0.5]$, as we do not consider robustness against a number of zealots exceeding half the swarm size. Furthermore, we consider a well-mixed system, where an agent has an equal probability to interact with any other agent.

3.1 Filtering rules

An agent may receive more than one message from its neighbours expressing their opinions, therefore the agent employs a filtering rule to select one opinion from the received messages. We consider two filtering rules: the voter rule and the majority rule. With the voter rule, the agent selects a random opinion from the received messages. With the majority rule, the agent selects the opinion which is present in more than 50% of the messages it receives.

To mathematically model the filtering rules, we define $n_i^\#$ as the weighted average of the relative number of agents holding opinion i (with $i \in \{A, B\}$). The weights are the qualities of the two options as we recall that the agents vote with a frequency proportional to

the quality. Hence, for options A and B, we respectively have

$$n_A^\# = \frac{A + z_a}{A + z_a + qB + z_b} \quad \text{and} \quad n_B^\# = \frac{qB + z_b}{A + z_a + qB + z_b}. \quad (1)$$

When using the voter rule, the agent filters a message expressing an opinion for option A or B with probabilities v_A and v_B , respectively, which are defined as

$$v_A = n_A^\#, \quad \text{and} \quad v_B = n_B^\#. \quad (2)$$

When using the majority rule, we define the two filtering probabilities as m_A and m_B , corresponding to discrete integrations of a Bernoulli distribution, as previously modelled in [58]. The Bernoulli distribution models the probabilistic events of receiving votes for option A or B. The probabilities that one vote is for A or B correspond to the quantities $n_A^\#$ and $n_B^\#$. We consider each agent to sample, on average, G messages. Therefore, we account for the different combinations of the G votes through the binomial coefficient $\binom{G}{i}$, with i indicating the number of votes in favour of A or B when computing m_A and m_B , respectively. The probabilities are computed as the sum of i ranging from $\lceil \frac{G}{2} \rceil + 1 - G \pmod{2}$ to G , therefore indicating that i is more than half of the number of messages G . The lower bound of the summation takes into account both the cases of even and odd G . More precisely, when G is even, an agent will choose the opinion if there are at least $\frac{G}{2} + 1$ samples sharing that opinion. When G is odd, the minimum number of samples with a given opinion must be $\lceil \frac{G}{2} \rceil$. Hence, for the majority rule, the probability for opinion $j \in \{A, B\}$ is:

$$m_j = \sum_{i=\lceil \frac{G}{2} \rceil + 1 - G \pmod{2}}^G \binom{G}{i} [n_j^\#]^i [n_{-j}^\#]^{G-i}. \quad (3)$$

3.2 Update rules

Update rules allow an agent to update its opinion based on the opinion it selects through the filtering rule. We consider two update rules: direct-switch and cross-inhibition.

An agent that updates its opinion using the direct-switch rule adopts the filtered opinion directly, hence the name. Since an agent can hold either opinion A or B, this rule implies that $A + B = 1$.

On the other hand, according to the cross-inhibition update rule, when the filtered opinion is different from the agent's current opinion, the agent becomes undecided and holds no opinion. Instead, undecided agents directly adopt the filtered opinion. Let U be the proportion of agents with no opinion. Since an agent can either hold opinion A or B or have no opinion, the cross-inhibition update rule implies that $A + B + U = 1$.

3.3 Homogeneous swarms

The models for homogeneous swarms describe the evolution of sub-populations of agents holding the same opinion when the entire population uses the same filtering rule and the same update rule. We build four models, resulting from the combination of the two filtering rules and the two update rules. Models using the direct-switch update rule have two variables, thus they need two equations: one ordinary differential equation that describes the evolution of the sub-population committed to A and an algebraic equation for mass conservation, i.e., $B = 1 - A$. On the other hand, models using the cross-inhibition update rule have three variables, thus they need

three equations: two ordinary differential equations that describe the evolution of the sub-populations committed to A and B, and the mass conservation equation $U = 1 - A - B$.

For all models, the change in a sub-population is described by two terms. There is a positive term, which indicates an increase in the sub-population as a result of agents inside that sub-population recruiting other agents through their messages. Depending on the update rule, the recruited population are the agents in the other sub-population for the direct-switch rule, or agents with no opinion for the cross-inhibition rule. The second term is negative and indicates a decrease in the sub-population as a result of agents inside that sub-population receiving a message from the other sub-population.

The two models based on the direct-switch update rule have the same form:

$$\frac{dA}{dt} = Bp_A - Ap_B \quad (4)$$

where p_A and p_B refer to the probability of receiving opinion A or B, respectively, which depends on the filtering rule ($p_i \in \{v_i, m_i\}$ and $i \in \{A, B\}$). On the other hand, the change in the sub-populations of a cross-inhibition model has the following structure:

$$\frac{di}{dt} = Up_i - ip_j \quad (5)$$

where $p_i \in \{v_i, m_i\}$, $i, j \in \{A, B\}$ and $j = -i$.

3.4 Heterogeneous swarms

Heterogeneous swarms are composed of robots (modelled as agents) that run different algorithms. We consider heterogeneous swarms comprising agents employing two different filtering rules but using the same update rule. More precisely, in a heterogeneous swarm comprising S susceptible (non-zealot) agents, there are kS agents using the voter rule and the rest (i.e., $(1 - k)S$ agents) using the majority rule. Thus, the parameter $k \in [0, 1]$ represents the proportion of susceptible agents using the voter rule and for values of $k = 0$ and $k = 1$, the heterogeneous swarm models reduce to the homogeneous swarm models of Eqs. (4) and (5). We analyse two heterogeneous models, one for each opinion update rule.

Let A_V (resp. B_V) and A_M (resp. B_M) be the proportion of agents holding opinion A (resp. B) while using the voter rule and the majority rule, respectively. Then, $A = A_M + A_V$ and $B = B_M + B_V$.

For the algorithms using the direct-switch update rule, we model the change in the two sub-populations holding opinion A as

$$\frac{dA_V}{dt} = B_V v_A - A_V v_B \quad \frac{dA_M}{dt} = B_M m_A - A_M m_B, \quad (6)$$

where $A_V + B_V = k$ and $A_M + B_M = 1 - k$.

For algorithms using the cross-inhibition update rule, the models need to keep track of the number of agents in the undecided state for both filtering rules. Let U_V and U_M be the proportions of agents with no opinion while using the voter rule and the majority rule, respectively. Then, $A_V + B_V + U_V = k$ and $A_M + B_M + U_M = 1 - k$. Finally, we model the change in the two sub-populations holding opinions A and B as

$$\begin{aligned} \frac{dA_V}{dt} &= U_V v_A - A_V v_B & \frac{dA_M}{dt} &= U_M m_A - A_M m_B \\ \frac{dB_V}{dt} &= U_V v_B - B_V v_A & \frac{dB_M}{dt} &= U_M m_B - B_M m_A. \end{aligned} \quad (7)$$

4 ANALYSIS

We study the long-term dynamics of the system, i.e., we look at the final state of the ODEs after a large time, $t \rightarrow \infty$. Given the non-linearity of the system, it is difficult to obtain a symbolic equation of the stable states of the ODE systems, especially for the heterogeneous case. Therefore, we resort to numerical integration of the ODEs given an initial condition. We consider the system initialised at symmetry, with the two sub-populations holding opinions A and B having the same size, i.e., $A(0) = B(0) = 0.5$. To avoid initialising the system at an unstable equilibrium where the two populations balance each other, we include a small deviation from the 50-50 perfect split and we initialise the system at $A(0) = 0.5 + \zeta$ and $B(0) = 0.5 - \zeta$, with $\zeta = 10^{-6}$. We numerically integrate the ODE systems for $t_{MAX} = 10^3$ time units. We also tested conditions with no bias and minimal bias towards the inferior option B, the results (not shown in this article) did not show any qualitative change to the results reported in Sec. 5 but only minimal quantitative differences in the cost-performance analysis.

4.1 Metrics of performance

To compute the performance of the swarm in making a collective decision, we measure both the goodness of the decision outcome and the resources invested in making it. Regarding the decision outcome, one needs to consider that in the presence of stubborn agents, reaching a full consensus (100% agreement) can be impossible [30, 45]; therefore, we set a quorum $\theta = 0.7$ to be reached in order to consider the swarm made a decision for either alternative. By looking at the long-term (equilibrium) state of the system, our mean-field analysis allows us to predict if a robot swarm running those rules will eventually select either of the two options, or will remain deadlocked at indecision with no sub-population imposing itself on the other.

For each configuration and each initial condition, we compute the following three measures, that a decision-maker should minimise: **regret** ρ : it is a numerical value that indicates the loss that the collective decision led to. When the population selects the option with the highest quality (option A), the regret is zero as the outcome of the decision could not be better. However, when it selects the option with the inferior quality (option B), the regret is $\rho = q_A - q_B = 1 - q_B$ (as we fix $q_A = 1$). Finally, in case of decision deadlock, $\rho = q_A - 0 = 1$. This loss function is termed regret in agreement with decision theory literature [3].

decision time τ : it is the amount of time that any of the two sub-populations take to reach the decision threshold $\theta = 0.7$.

cognitive cost ϕ : it is a measure of how much social information each individual needs to process during the collective decision process. It is computed as the number of opinions each agent gathers from its neighbours. This quantity is different for the two considered filtering rules, voter rule and majority rule. With the former, the agent reads one message per timestep, while with the latter, it reads G messages per timestep. We multiply the number of these social interactions by the timesteps taken to reach the decision quorum. Thus, for the models using the voter rule, the cognitive cost is $\phi_V = \tau$ and for the models using the majority rule, $\phi_M = \tau G$. For heterogeneous swarms, we compute the cognitive cost as

$$\phi = \tau [k + (1 - k)G], \quad (8)$$

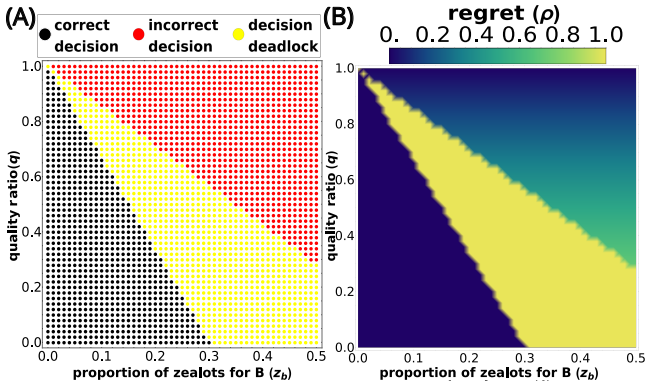


Figure 1: Representative results of the homogeneous swarms composed of agents using the direct-switch and voter rule (system (6) with $k = 1$) under wrong-addressing attack in the parameter space (q, z_b) . In (A), the three colours show the decision outcome for each (q, z_b) : the black points represent convergence of the group opinions to option A ($A(t_{MAX}) \geq \theta$), the red points represent group convergence of the opinions to option B ($B(t_{MAX}) \geq \theta$) and the yellow points represent a decision deadlock ($(A(t_{MAX}) < \theta) \wedge (B(t_{MAX}) < \theta)$). In (B), we use the decision outcome to compute the regret ρ which we represent as a colour map, where decision deadlocks lead to the highest regret (yellow region with $\rho = 1$), and accurate decisions to the zero regret (dark blue region with $\rho = 0$).

recalling that τ is the decision time and G is the (average) number of neighbours sampled by agents using the majority rule.

4.2 Measuring the swarm robustness

We measure the swarm robustness in the presence of two types of attacks, wrong addressing and denial of service. These swarm-level attacks borrow names from traditional cyber-security because they generate a similar system dysfunction. Wrong addressing corresponds to leading the swarm to select the inferior option and can be implemented by the zealots by only voting for the option with the lowest quality, in our case option B. Therefore, we study robustness against a wrong-addressing attack by setting $z_a = 0$ and $z_b > 0$. Instead, denial of service corresponds to keeping the swarm undecided, unable to reach any agreement and, therefore, unable to deliver the service—make a collective decision—that it has been programmed for. Previous work [6, 24, 45] has shown that the presence of two sub-populations of zealots, one for each option, can keep the system hung at indecision. Therefore, we study robustness against denial-of-service attacks by setting $z_a = 0.05$ and $z_b > 0$.

For heterogeneous systems, we also vary the heterogeneity parameter k in the range $[0, 1]$ with a step of 0.05. We then compute each metric of swarm performance—regret, decision time, and cognitive cost—as the average value over all tested conditions. To find a single measure that combines both the goodness of the decision outcome and the resources invested in making it, we borrow from cognitive psychology and neuroscience the metric of Bayes risk β , which has been used to study normative models of decision

making [4, 36] and is defined as

$$\beta = c_1 DT + c_2 ER = c_1 \phi + c_2 \rho. \quad (9)$$

The parameters c_1 and c_2 indicate how the decision-maker weighs the two costs of making the decision; more precisely, c_1 weighs the resources invested in making the decision, that in the cognitive psychology literature is measured as the decision time (DT). In our case, we use the cognitive cost ϕ . The parameter c_2 , instead, weighs the cost of making an error (error rate ER). In our case, we use the regret ρ . Thus, in our analysis, we look at systems that can minimise the Bayes risk composed of the weighted sum of cognitive cost and regret. We set $c_1 = 1$ (the cost of each pooled message is 1) and vary the parameter $c_2 \in [0, 10^3]$ weighting the cost of decision mistakes (the cost of the loss in quality for incorrect decision). Our goal is to investigate what is the parameter range of c_2 (if it exists) by which employing a heterogeneous swarm is more convenient (lower β) than using a homogeneous one.

5 RESULTS

We analyse robustness to wrong-addressing and denial-of-service attacks by fixing the proportion of zealots in favour of option A (z_a) to zero and 0.05, respectively, and computing the decision performance for parameters $z_b \in [0, 0.5]$ and $q \in [0, 1]$, denoting the proportion of zealots in favour of option B and the quality ratio. We initially consider that the agents running the majority rule make a sampling effort of reading $G = 8$ neighbour messages. Fig. 1 shows a representative example of the results for homogeneous swarms based on direct-switch and the voter rule ($k = 1$) under a wrong-addressing attack ($z_a = 0$) in terms of the decision outcome in panel (A) and regret ρ in panel (B), in the (q, z_b) parameter space. These plots show the typical trend that can be observed in the large majority of tested cases: for relatively weak attacks (low number of zealots z_b) and for relatively easy decision problems (low quality ratio q), the swarm makes accurate decisions with low regret; for strong attacks (high z_b) and difficult problems (high q), the swarms make incorrect decisions with regret $\rho = 1 - q$; for the intermediate cases, there is a triangular region in the parameter space (q, z_b) where the swarm is unable to make any decision and has maximum regret $\rho = 1$.

When considering heterogeneous swarms, i.e., $0 < k < 1$, the curves separating these three regions of the parameter space (q, z_b) gradually move according to the swarm composition k ; this leads to changes in the collective decision performance. Fig. 2 shows the change of the three regions in (q, z_b) for a system based on the direct-switch, for values of $k \in \{0.8, 0.9, 1\}$. For parameters in the bottom-left and top-right corners, the swarm consistently (for any k) makes accurate and inaccurate decisions, respectively. The yellow region in the bottom-right corner, which denotes decision deadlocks, becomes smaller as k is reduced. In other words, as the proportion of agents using the majority rule increases, the swarm improves its decision-making performance by avoiding decision deadlocks. However, this requires agents to process more social information as the majority rule requires checking the opinions of G neighbours each timestep. We show this cost-performance trade-off in Fig. 3.

Fig. 3 shows the regret ρ and the cognitive cost ϕ for the two heterogeneous swarm models (Eqs. (6) and (7)) for the two considered

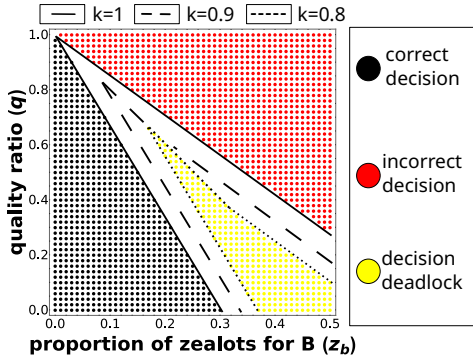


Figure 2: The decision outcome in the parameter space (q, z_b) changes with the swarm composition. We report the results for a system based on the direct-switch rule (with $G = 8$) for three different values of swarm composition $k \in \{0.8, 0.9, 1\}$ under a wrong-addressing attack ($z_a = 0$). The black region represents an agreement on the correct option A, the red region represents an agreement on the inferior option B, and the yellow region represents a decision deadlock. The three types of lines divide the region of convergence for the different values of k . When the system is homogeneous ($k = 1$), the boundary of the yellow region of decision deadlock extends to the solid lines. As heterogeneity in the system increases, the area of deadlock is reduced to the dashed (corresponding to $k = 0.9$) or dotted lines (corresponding to $k = 0.8$).

types of attacks. The results—presented in the cost-performance space (ϕ, ρ) —indicate that, in agreement with previous results on homogeneous models [6], majority-based algorithms ($k = 0$) are more robust than the ones based on the voter rule ($k = 1$). However, this robustness comes at the expense of an increased information processing cost ϕ . Here, we expand previous analyses and find that swarm heterogeneity regulates the cost-performance trade-off for both types of attacks.

The cost-performance trade-off analysis of Fig. 3 shows that every model can be well described by a piecewise linear function that allows the identification of different trends in how heterogeneous swarms trade regret with cognitive cost. The results show two trends in the cost-performance trade-off of the swarms using the cross-inhibition rule and three trends when using the direct-switch rule. We compute the separation points of the piecewise linear function (Table 1) by testing all values of k , fitting the two (or three lines), and selecting the points that give the minimum cumulative linear fitting error (estimated error variance). The linear fitting shows that swarms composed predominantly of agents using the voter rule ($k \geq 0.5$) have a rapid improvement in performance (lower regret ρ) through the inclusion in the swarm of a minority of agents using the majority rule (red lines in Fig. 3). Instead, when the swarm is predominantly composed of agents using the majority rule ($k \leq 0.5$), adding more majority-rule agents gives a smaller performance improvement, while still incurring an increase in cognitive cost (blue lines). The swarms using the direct-switch rule show a third part of the curve where adding more majority-rule

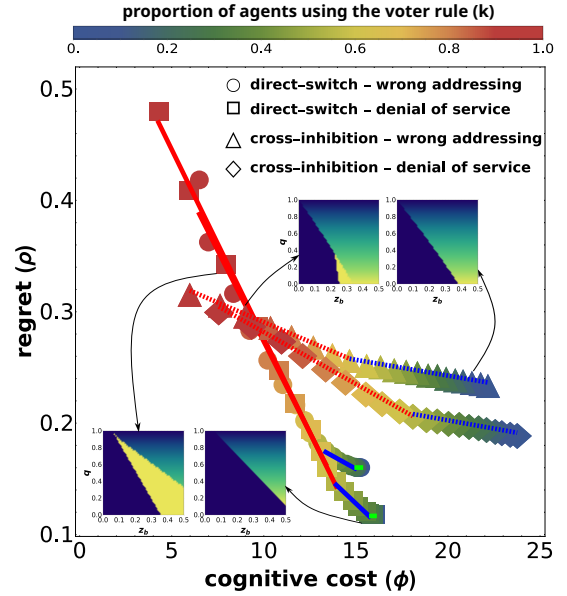


Figure 3: Cost-performance trade-off as the relationship between the regret ρ and the cognitive cost ϕ in making the decision, for both models, direct-switch of Eq. (6) and cross-inhibition of Eq. (7), and for both types of attacks, denial-of-service ($z_a = 0.05$) and wrong-addressing ($z_a = 0$). Each point represents the average regret (y-axis) computed across all tested values of parameters q and z_b ; in other words, we report the average value of the colour maps of Fig. 1B for each $k \in [0, 1]$. Each point’s k value is colour-coded according to the top bar. The insets show the parameter space (q, z_b) computed for some specific k values, which when averaged gives the regret of the plotted point k (the colour maps for all tested conditions are available in the supplementary material [1]). We can observe that the resulting lines in the cost-performance space can be approximated well through piecewise linear functions (lines overlaid onto the data points, see Table 1).

	direct-switch		cross-inhibition	
attack	slope	k range	slope	k range
WA	-0.033	[0.5, 1]	-0.007	[0.65, 1]
	-0.0092	[0.2, 0.5]	-0.0028	[0, 0.65]
	0	[0, 0.2]		
DoS	-0.034	[0.55, 1]	-0.0093	[0.6, 1]
	-0.016	[0.25, 0.55]	-0.0030	[0, 0.6]
	0	[0, 0.25]		

Table 1: Slope of the piecewise linear function and range of k values of the fitted lines of Fig. 3 for wrong addressing (WA) and denial of service (DoS). All reported fitting values are rounded to the second significant digit.

agents does not lead to any improvement in regret (equal values) but only causes an increase in cognitive cost (green lines).

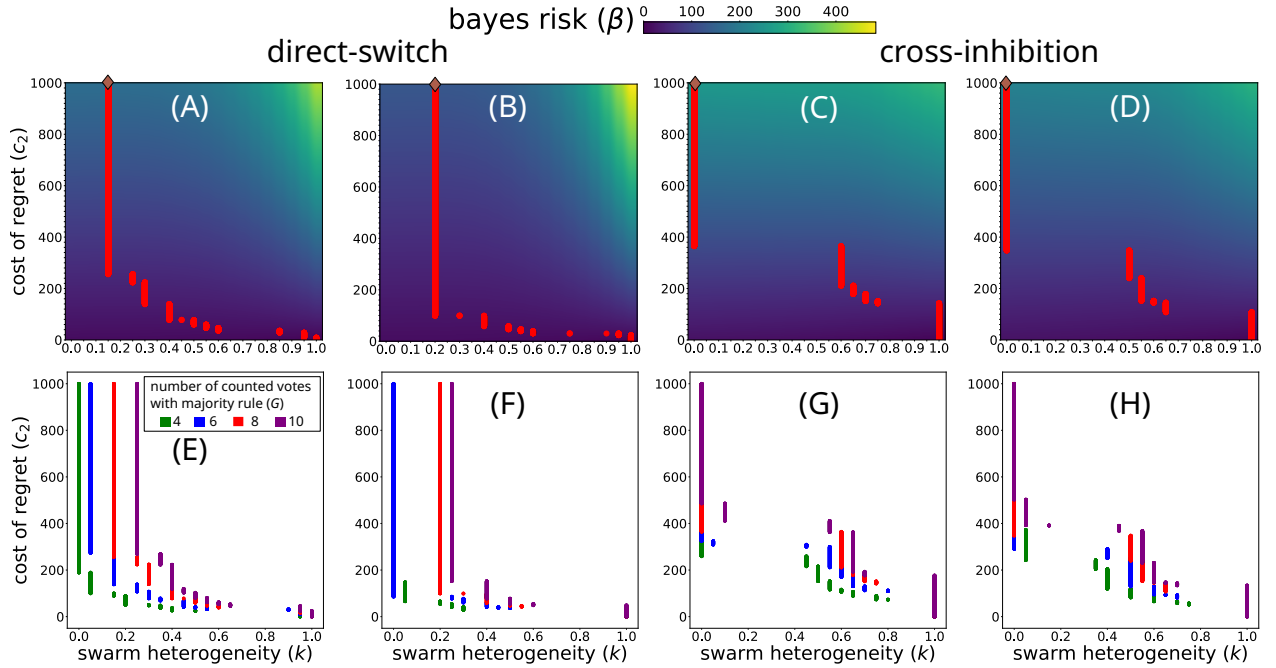


Figure 4: The optimal swarm composition k^* as a function of swarm heterogeneity k (x-axis) and regret weight c_2 (y-axis) for direct-switch algorithms (left panels A,B,E,F) and cross-inhibition (right panels C,D,G,H). In the top row (A-D), we show the Bayes risk of Eq. (9) as a colour map (see top bar), when majority rule agents process $G = 8$ neighbours’ opinions. The red lines show k^* , i.e., where the Bayes risk is minimal and thus the swarm is the most robust on average. In the bottom row (E-H), we report the optimal swarm composition k^* for different values of $G \in \{4, 6, 8, 10\}$ showing a similar trend that shifts with G . The first and third columns (panels A,E,C,G) correspond to denial of service ($z_a = 0.05$) and the second and fourth columns (panels B,F,D,H) to wrong addressing ($z_a = 0$). For all cases, there is a range of regret weight c_2 for which heterogeneous swarms give the best collective performance.

The results of Fig. 3 are in agreement with previous analyses that compared consensus algorithms based on direct-switch and cross-inhibition, showing that in the presence of zealots, direct-switch performs poorly due to frequent decision deadlocks and instead inhibitory signals allow quicker decisions and symmetry-breaking [45, 61]. Fig. 3 also shows that the situation is reversed with majority-rule agents, where direct-switch outperforms cross-inhibition in terms of both metrics, which is in line with previous analysis that has shown absence of decision deadlocks with majority-rule agents [44].

Through the Bayes risk β of Eq. (9), we compute a single metric as a linear combination of the cost of decision errors (ρ) and the average information processing cost of each agent (ϕ). We fix $c_1 = 1$, meaning that the cost of processing one neighbour’s opinion is one. Then, we increase the regret weight c_2 to scale the cost of decision errors and find which swarm composition k^* leads to the lowest Bayes risk β . In terms of regret under both attacks (Fig.3), the homogeneous swarms using the majority rule always obtained the best performance, however they were not the only ones. When agents use the direct-switch rule, heterogeneous swarms mostly comprising majority rule agents can be mixed with 15% (denial of service) to 20% (wrong addressing) of agents using the voter rule and make decisions with the same average regret ρ . The analysis on β in Figs. 4A-D allows us to combine the performance on

the decision outcome with the cost of processing one or G neighbour’s opinions per timestep. In this light, heterogeneous swarms offer the advantage of providing comparable levels of regret with fewer social exchanges; they are more efficient than homogeneous swarms. For instance, considering a high regret weight c_2 , the Bayes risk for direct-switch (Figs. 4A-B) shows that with the inclusion of 15-20% ($k^* = 0.15, 0.2$) of agents using the voter rule, the average decision regret remains equal to homogeneous swarms with all the agents using the majority rule ($k = 0$), however incurring lower cognitive cost per agent on average. Figs. 4A-D also show that by reducing the decision-error weight c_2 , the best swarm composition k^* shifts towards higher values (i.e., swarms with larger voter-rule sub-populations), for any tested type of attack and decision rule. Finally, we show that when errors are relatively low, $c_2 \lesssim 140 \sim 20$, the homogeneous swarm using only the voter rule, $k = 1$, is the system with less risk. The obtained results seem to be relatively general as we find qualitatively similar trends when measuring the robustness to both denial-of-service attacks and to wrong-addressing attacks, as well as for algorithms based on both the direct-switch and the cross-inhibition rule (compare panels A to D of Fig. 4). The only notable qualitative difference among the models is the fact that when the agents use the cross-inhibition rule, regret is minimised by homogeneous swarms only, instead, when they use the direct-switch rule, regret is also minimised by

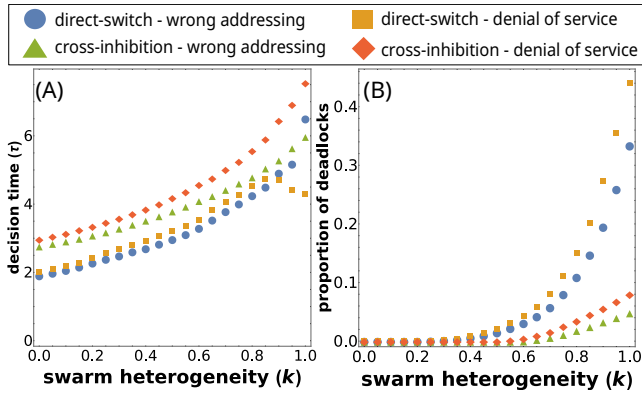


Figure 5: (A) Decision time τ computed as the amount of time to reach the decision threshold $\theta = 0.7$ for either option. (B) The proportion of conditions in the parameter space (q, z_B) where the swarm remains deadlocked at indecision ($(A(t_{MAX}) < \theta) \wedge (B(t_{MAX}) < \theta)$), i.e., the proportion of the yellow area of Fig. 1. In both panels, we report results for both models, both attacks, and various values of the swarm composition $k \in [0, 1]$.

heterogeneous swarms; nevertheless, in both cases, when the ratio c_2/c_1 decreases, the optimal strategy is behavioural heterogeneity.

We also investigate the impact of different implementations of the majority rule, more precisely, how the Bayes risk space changes for different numbers of neighbours’ opinions G . Our analysis in Figs. 4E-H shows that the best swarm composition k^* also changes with the value of G , however qualitatively the trend is the same: for low c_2 , voter-rule agents minimise the Bayes risk, and for increasing c_2 , having more majority-rule agents is better. Figs. 4E-F, for the direct-switch case, show that when G increases, a smaller proportion of agents using the majority rule (i.e., larger k) is enough to reach the minimum Bayes risk. This means that the best swarm composition can comprise a larger number of ‘simpler’ agents using the voter rule.

Before concluding our analysis, we want to touch on another aspect that is especially relevant when considering robustness: the cost of indecision. There are parameter ranges for which, in the long term, swarms subject to the presence of zealots are unable to reach the decision quorum $\theta = 0.7$. In our analysis, the size of this parameter range negatively influences the average regret, however, no temporal cost for indecision is included in the decision time τ , which only includes the time taken to reach the quorum for either option. When we consider that swarms deadlocked in an undecided state imply numerous social exchanges for possibly a long time and we include a temporal cost of indecision higher than the average decision time, swarms with many voter-rule agents are unable to perform well in terms of Bayes risk. Whereas the average time to make a decision increases moderately with k (Fig. 5A), voter-rule algorithms in the presence of zealots go more often into a denial-of-service state (Figs. 2 and 5B). Therefore, including high time cost for indecision penalises more heavily swarms with high k but does not alter the results for low k . As a result, swarms with large proportions of voter-rule agents are optimal for a smaller

range of values of c_2 , however, the results remain the same for large c_2 , marked with a diamond on top of Figs. 4A-D, (see also the supplementary figures in [1]).

6 CONCLUSION

This study shows that heterogeneous robot swarms can outperform homogeneous ones. We consider large swarms of simple robots that use algorithms based on voter-like models of opinion dynamics to process social information and reach a group agreement in favour of the best option between two alternatives. More precisely, we measure the collective robustness in making such decisions when in the swarm there are stubborn robots (zealots) that deviate from conformism rules, possibly leading the swarm to select the inferior option or to a decision deadlock (which relates to the cyber-security attacks of wrong addressing and denial of service, respectively). Our analysis considers both the benefits of making accurate decisions and the cost of processing social information, leading to a cost-performance trade-off. We model the robot swarm as a multiagent system and show that when agents run different rules, depending on the relative weights on cost and performance, heterogeneous swarms can outperform homogeneous ones in managing such a trade-off. In addition to having far-reaching implications for the design of cost-effective and robust robotic systems, our results can help us understand how social animals share cognitive load among individuals exploiting group heterogeneity, as the investigated problem is also relevant for several biological systems.

While this study already considers several conditions, future work should further assess the generality of our results considering a variety of aspects. For instance, our analysis is limited to mean-field analysis with given starting conditions and future work should test whether the results hold in stochastic finite-sized multiagent simulation and for different starting conditions. Further analysis should also investigate how the results are affected when considering more than two options, as previous work has shown that there can be qualitative changes in the group opinion dynamics in best-of- n problems with $n > 2$ [43]. Another relevant aspect of distributed systems is the impact of the network topology on the group dynamics. In our study, we analysed well-mixed models only—which previous work showed can well describe the dynamics of large-scale robotic systems operating in certain scenarios [56, 58]—however, future research should also extend previous studies showing the impact of different network topologies on the group opinion dynamics (e.g., [44, 51]) by considering heterogeneous decision-makers. Finally, we believe that further advantages in using heterogeneous swarms can be highlighted when considering dynamic environments [37, 46], where behavioural heterogeneity may enhance the ability of the group to adapt to changes.

ACKNOWLEDGMENTS

R.Z, M.D, and A.R. acknowledge support from the Belgian F.R.S.-FNRS. A.R. also acknowledges support from DFG under Germany’s Excellence Strategy - EXC 2117 - 422037984. This work was partially supported by SPW Recherche under grant n. 2010235 (ARIAC).

REFERENCES

- [1] Nemanja Antonic, Raina Zakir, Marco Dorigo, and Andreagiovanni Reina. 2024. *Supplementary material for paper 'Collective robustness of heterogeneous decision-makers against stubborn individuals'*. <https://iridia.ulb.ac.be/supp/IridiaSupp2024-001/>
- [2] Palina Bartashevich and Sanaz Mostaghim. 2021. Multi-featured collective perception with Evidence Theory: tackling spatial correlations. *Swarm Intelligence* 15, 1 (2021), 83–110.
- [3] David E. Bell. 1982. Regret in Decision Making under Uncertainty. *Operations Research* 30, 5 (1982), 961–981.
- [4] Rafal Bogacz, Eric Brown, Jeff Moehlis, Philip Holmes, and Jonathan D. Cohen. 2006. The Physics of optimal decision making: A formal analysis of models of performance in two-alternative forced choice tasks. *Psychological Review* 4, 4 (2006), 700–765.
- [5] Thomas Bose, Andreagiovanni Reina, and James A. R. Marshall. 2017. Collective Decision-Making. *Current Opinion in Behavioral Sciences* 6 (2017), 30–34.
- [6] Francesco Canciani, Mohamed S. Talamali, James A. R. Marshall, Thomas Bose, and Andreagiovanni Reina. 2019. 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)*. IEEE Press, NJ.
- [7] Claudio Castellano, Santo Fortunato, and Vittorio Loreto. 2009. Statistical physics of social dynamics. *Reviews of Modern Physics* 81, 2 (2009), 591–646.
- [8] Eduardo Castelló Ferrer, Thomas Hardjono, Alex Pentland, and Marco Dorigo. 2021. Secure and secret cooperation in robot swarms. *Science Robotics* 6, 56 (2021), abf1538.
- [9] Peter Clifford and Aidan Sudberry. 1973. A model for spatial conflict. *Biometrika* 60, 3 (1973), 581–588.
- [10] Nuno Crokidakis and Paulo Murilo Castro de Oliveira. 2015. Inflexibility and independence: Phase transitions in the majority-rule model. *Physical Review E* 92 (2015), 062122.
- [11] Michael Crosscombe and Jonathan Lawry. 2016. A model of multi-agent consensus for vague and uncertain beliefs. *Adaptive Behavior* 24, 4 (2016), 249–260.
- [12] Emilio Cruciani, Hlalo Alfie Mimun, Matteo Quattropiani, and Sara Rizzo. 2023. Phase transition of the k-majority dynamics in biased communication models. *Distributed Computing* 36, 2 (2023), 107–135.
- [13] Giulia De Masi, Judhi Prasetyo, Raina Zakir, Nikita Mankovskii, Eliseo Ferrante, and Elio Tuci. 2021. Robot swarm democracy: the importance of informed individuals against zealots. *Swarm Intelligence* 15, 4 (2021), 315–338.
- [14] Marco Dorigo, Guy Theraulaz, and Vito Trianni. 2020. Reflections on the future of swarm robotics. *Science Robotics* 5 (2020), eabe4385.
- [15] Marco Dorigo, Guy Theraulaz, and Vito Trianni. 2021. Swarm Robotics: Past, Present and Future. *Proceedings of the IEEE* 109, 7 (2021), 1152–1165.
- [16] Marco Dorigo et al. 2013. Swarmanoid: A Novel Concept for the Study of Heterogeneous Robotic Swarms. *IEEE Robotics & Automation Magazine* 20, 4 (2013), 60–71.
- [17] Julia T. Ebert, Melvin Gauci, Frederik Mallmann-Trenn, and Radhika Nagpal. 2020. Bayes Bots: Collective Bayesian Decision-Making in Decentralized Robot Swarms. In *2020 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE Press, NJ, 7186–7192.
- [18] Serge Galam. 1986. Majority rule, hierarchical structures, and democratic totalitarianism: A statistical approach. *Journal of Mathematical Psychology* 30, 4 (1986), 426–434.
- [19] Heiko Hamann. 2018. *Swarm Robotics: A Formal Approach*. Springer, Cham.
- [20] Richard A. Holley and Thomas M. Liggett. 1975. Ergodic Theorems for Weakly Interacting Infinite Systems and the Voter Model. *The Annals of Probability* 3, 4 (1975), 643–663.
- [21] Edmund R. Hunt and Sabine Hauert. 2020. A checklist for safe robot swarms. *Nature Machine Intelligence* 2, 8 (2020), 420–422.
- [22] Bhargav Karamched, Megan Stickler, William Ott, Benjamin Lindner, Zachary P. Kilpatrick, and Krešimir Josić. 2020. Heterogeneity Improves Speed and Accuracy in Social Networks. *Physical Review Letters* 125, 21 (2020), 218302.
- [23] Daniela Kengyel, Heiko Hamann, Payam Zahadat, Gerald Radspieler, Franz Wotawa, and Thomas Schmickl. 2015. Potential of heterogeneity in collective behaviors: A case study on heterogeneous swarms. In *PRIMA 2015: Principles and Practice of Multi-Agent Systems*. LNCS, Vol. 9387. Springer, Cham, 201–217.
- [24] Nagi Khalil, Maxi San Miguel, and Raúl Toral. 2018. Zealots in the mean-field noisy voter model. *Physical Review E* 97 (2018), 012310.
- [25] Nagi Khalil and Raúl Toral. 2019. The noisy voter model under the influence of contrarians. *Physica A: Statistical Mechanics and its Applications* 515 (2019), 81–92.
- [26] Pavel L. Krapivsky and Sidney Redner. 2003. Dynamics of Majority Rule in Two-State Interacting Spin Systems. *Physical Review Letters* 90 (2003), 238701.
- [27] Chanelle Lee, Jonathan Lawry, and Alan F. T. Winfield. 2021. Negative updating applied to the best-of-n problem with noisy qualities. *Swarm Intelligence* 15, 1 (2021), 111–143.
- [28] James A. R. Marshall, Ralf H. J. M. Kurvers, Jens Krause, and Max Wolf. 2019. Quorums enable optimal pooling of independent judgements in biological systems. *eLife* 8 (2019), e40368.
- [29] Naoki Masuda. 2013. Voter models with contrarian agents. *Physical Review E* 88, 5 (2013), 052803.
- [30] Mauro Mobilia. 2003. Does a Single Zealot Affect an Infinite Group of Voters? *Physical Review Letters* 91, 2 (2003), 028701.
- [31] Mauro Mobilia. 2015. Nonlinear q-voter model with inflexible zealots. *Physical Review E* 92, 1 (2015), 012803.
- [32] Mauro Mobilia, Alexander Petersen, and Sidney Redner. 2007. On the role of zealotry in the voter model. *Journal of Statistical Mechanics: Theory and Experiment* 2007, 08 (2007), P08029.
- [33] Marco A. Montes de Oca, Eliseo Ferrante, Alexander Scheidler, Carlo Pinciroli, Mauro Birattari, and Marco Dorigo. 2011. Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence* 5, 3 (2011), 305–327.
- [34] Thomas A O'Shea-Wheller, Edmund R Hunt, and Takao Sasaki. 2020. Functional Heterogeneity in Superorganisms: Emerging Trends and Concepts. *Annals of the Entomological Society of America* 114, 5 (2020), 562–574.
- [35] Chris A. C. Parker and Hong Zhang. 2009. Cooperative Decision-Making in Decentralized Multiple-Robot Systems: The Best-of-N Problem. *IEEE/ASME Transactions on Mechatronics* 14, 2 (2009), 240–251.
- [36] Angelo Pirrone, Tom Stafford, and James A. R. Marshall. 2014. When natural selection should optimize speed-accuracy trade-offs. *Frontiers in Neuroscience* 8, 73 (2014).
- [37] Judhi Prasetyo, Giulia De Masi, and Eliseo Ferrante. 2019. Collective decision making in dynamic environments. *Swarm Intelligence* 13, 3-4 (2019), 217–243.
- [38] Mohsen Raoufi, Pawel Romanczuk, and Heiko Hamann. 2023. Individuality in Swarm Robots with the Case Study of Kilobots: Noise, Bug, or Feature?. In *ALIFE 2023: Ghost in the Machine: Proceedings of the 2023 Artificial Life Conference*. MIT Press, Cambridge, MA, 35.
- [39] Andreagiovanni Reina. 2020. Robot teams stay safe with blockchains. *Nature Machine Intelligence* 2, 5 (2020), 240–241.
- [40] Andreagiovanni Reina, Thomas Bose, Vito Trianni, and James A. R. Marshall. 2018. Effects of Spatiality on Value-Sensitive Decisions Made by Robot Swarms. In *Distributed Autonomous Robotic Systems (DARS 2016): The 13th International Symposium*. SPAR, Vol. 6. Springer, Cham, 461–473.
- [41] Andreagiovanni Reina, Marco Dorigo, and Vito Trianni. 2014. Collective Decision Making in Distributed Systems Inspired by Honeybees Behaviour. In *Proceedings of 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*. IFAAMAS, Richland, SC, 1421–1422.
- [42] Andreagiovanni Reina, Eliseo Ferrante, and Gabriele Valentini. 2021. Collective decision-making in living and artificial systems: editorial. *Swarm Intelligence* 15, 1-2 (2021), 1–6.
- [43] Andreagiovanni Reina, James A. R. Marshall, Vito Trianni, and Thomas Bose. 2017. Model of the best-of-N nest-site selection process in honeybees. *Physical Review E* 95, 5 (2017), 052411.
- [44] Andreagiovanni Reina, Thierry Njougouo, Elio Tuci, and Timoteo Carletti. 2023. Studying speed-accuracy trade-offs in best-of-n collective decision-making through heterogeneous mean-field modelling. *arXiv* 2310.13694 (2023).
- [45] Andreagiovanni Reina, Raina Zakir, Giulia De Masi, and Eliseo Ferrante. 2023. Cross-inhibition leads to group consensus despite the presence of strongly opinionated minorities and asocial behaviour. *Communications Physics* 6 (2023), 236.
- [46] Mohammad Salahshour. 2019. Phase Diagram and Optimal Information Use in a Collective Sensing System. *Physical Review Letters* 123, 6 (2019), 068101.
- [47] David Saldana, Amanda Prorok, Shreyas Sundaram, Mario F. M. Campos, and Vijay Kumar. 2017. Resilient consensus for time-varying networks of dynamic agents. In *2017 American Control Conference (ACC)*. IEEE Press, NJ, 252–258.
- [48] Alexander Scheidler, Arne Brutschy, Eliseo Ferrante, and Marco Dorigo. 2016. The k-Unanimity Rule for Self-Organized Decision-Making in Swarms of Robots. *IEEE Transactions on Cybernetics* 46, 5 (2016), 1175–1188.
- [49] Thomas D. Seeley, P. Kirk Visscher, Thomas Schlegel, Patrick M. Hogan, Nigel R. Franks, and James A. R. Marshall. 2012. Stop Signals Provide Cross Inhibition in Collective Decision-Making by Honeybee Swarms. *Science* 335, 6064 (2012), 108–111.
- [50] Qihao Shan and Sanaz Mostaghim. 2022. Benchmarking Performances of Collective Decision-Making Strategies with Respect to Communication Bandwidths in Discrete Collective Estimation. In *Swarm Intelligence (ANTS 2022)*. LNCS, Vol. 13491. Springer, Cham, 54–65.
- [51] V. Sood, Tibor Antal, and Sidney Redner. 2008. Voter models on heterogeneous networks. *Physical Review E* 77, 4 (2008), 041121.
- [52] Volker Strobel, Eduardo Castelló Ferrer, and Marco Dorigo. 2018. Managing Byzantine Robots via Blockchain Technology in a Swarm Robotics Collective Decision Making Scenario. In *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2018)*. IFAAMAS, Richland, SC, 541–549.

- [53] Volker Strobel, Eduardo Castelló Ferrer, and Marco Dorigo. 2020. Blockchain Technology Secures Robot Swarms: A Comparison of Consensus Protocols and Their Resilience to Byzantine Robots. *Frontiers in Robotics and AI* 7 (2020).
- [54] Volker Strobel, Alexandre Pacheco, and Marco Dorigo. 2023. Robot swarms neutralize harmful Byzantine robots using a blockchain-based token economy. *Science Robotics* 8, 79 (2023), eabm4636.
- [55] Mohamed S. Talamali, James A. R. Marshall, Thomas Bose, and Andreagiovanni Reina. 2019. Improving collective decision accuracy via time-varying cross-inhibition. In *Proceedings of the 2019 IEEE International Conference on Robotics and Automation (ICRA 2019)*. IEEE Press, NJ, 9652–9659.
- [56] Mohamed S. Talamali, Arindam Saha, James A. R. Marshall, and Andreagiovanni Reina. 2021. When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics* 6, 56 (2021), eabf1416.
- [57] Gabriele Valentini, Eliseo Ferrante, and Marco Dorigo. 2017. The Best-of-n Problem in Robot Swarms: Formalization, State of the Art, and Novel Perspectives. *Frontiers in Robotics and AI* 4 (2017).
- [58] Gabriele Valentini, Eliseo Ferrante, Heiko Hamann, and Marco Dorigo. 2016. Collective decision with 100 Kilobots: speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems* 30, 3 (2016), 553–580.
- [59] Gabriele Valentini, Heiko Hamann, and Marco Dorigo. 2014. Self-Organized Collective Decision Making: The Weighted Voter Model. In *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014)*. IFAAMAS, Richland, SC, 45–52.
- [60] Gabriele Valentini, Heiko Hamann, and Marco Dorigo. 2015. 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 (AAMAS 2015)*. IFAAMAS, Richland, SC, 1305–1314.
- [61] Raina Zakir, Marco Dorigo, and Andreagiovanni Reina. 2022. Robot Swarms Break Decision Deadlocks in Collective Perception through Cross-inhibition. In *Swarm Intelligence (ANTS 2022)*. LNCS, Vol. 13491. Springer, Cham, 209–221.
- [62] Hanqing Zhao, Alexandre Pacheco, Volker Strobel, Andreagiovanni Reina, Xue Liu, Gregory Dudek, and Marco Dorigo. 2023. A Generic Framework for Byzantine-tolerant Consensus Achievement in Robot Swarms. In *Proceedings of the 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS 2023)*. IEEE Press, NJ, 8839–8846.
- [63] Anıl Özdemir, Melvin Gauci, Salomé Bonnet, and Roderich Groß. 2018. Finding Consensus Without Computation. *IEEE Robotics and Automation Letters* 3, 3 (2018), 1346–1353.