Heterogeneity Can Enhance the Adaptivity of Robot Swarms to Dynamic Environments

Raina Zakir[1], Mohammad Salahshour[2], Marco Dorigo[1], and Andreagiovanni Reina[2,3]

[1] IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
{raina.zakir, marco.dorigo}@ulb.de
[2] Department of Collective Behaviour, Max Planck Institute of Animal Behavior, Konstanz, Germany
{msalahshour, areina}@ab.mpg.de
[3] CASCB, Universität Konstanz, Konstanz, Germany

Abstract. We study how robot swarms can collectively adapt to dynamic environments by changing what they collectively select as the best among a set of n possible options. While the robots rely on local communication with one another, follow simple rules, and make estimates of the option's qualities subject to measurement errors, the swarm as a whole can infer the change in the environment to make accurate collective decisions. Most studies focusing on dynamic environments have achieved adaptive behaviour by including random noise or threshold-based approaches to continuously explore alternatives and prevent opinion stagnation once consensus is achieved. In this study, we investigate whether or not swarms of robots with heterogeneous behaviours can be more adaptive than homogeneous swarms. We consider two behaviours from the literature which robots use to update their opinions: the majority rule, where robots gather information from all neighbours, and the voter rule, where robots use information from a single neighbour. In static environments, swarms of majority-rule robots, by using a larger amount of social information, typically make quicker decisions than swarms of voter-rule robots. However, our multiagent and robot simulations show that including voter-rule robots within a swarm of majority-rule robots can increase the group's responsiveness to environmental changes. This result shows the potential benefits of mixing simpler and relatively more advanced robots in the same swarm.

1 Introduction

Collective decision-making is an essential capability to enable autonomy in swarm robotics. Swarm robotics is inspired by the self-organising behaviours observed in biological systems, particularly evident in eusocial insects [13]. Examples of collective decision-making are honeybees' choice of a site for their nest among various alternatives [46], or ants' selection of the shortest path from their nest to a food source [17]. These examples found in nature and characterised by the absence of a global coordinator serve as the foundation for developing collective decision-making algorithms in swarm robotics [39], especially in swarms of minimalistic robots that execute simple behaviours to reach collective decisions in diverse scenarios, such as the selection of an aggregation site [8,48], the coordination of motion in a common direction [29], or the identification of the most abundant environmental feature [58,6]. The use of minimalistic robots, constrained by limited memory, computational power, and communication capabilities, is often mandated by the application scenario. For instance, size is crucial for nano-robots navigating blood vessels, whereas budget constraints may lead to the use of inexpensive, disposable robots in hazardous environments [19], where the risk of robot loss is high. As a result of their limitations, minimalist robots often present a high level of sensory noise, which leads to uncertainty in environmental and social estimations, and ultimately poses a significant challenge to the process of making collective decisions.

An important category of collective decision-making is the "best-of-n problem" where a group needs to choose the best option among n alternatives that can differ in quality [55]. Given the noisy assessment of the quality of the options in minimalistic systems, undertaking best-of-n decisions can be challenging for a robot swarm. This challenge is further amplified in dynamic environments, where conditions change over time [34].

Making collective decisions in the best-of-n problem while requiring minimalism relies on voting algorithms governed by simple rules, typically studied in opinion dynamics [5], where each opinion shared by the robots is treated as a vote. Among the computationally simplest algorithms is the voter model (or voter rule) [7,18], where robots consider only one vote (opinion) from a randomly chosen neighbour. This model has been expanded into the weighted voter model in [53], wherein robots express their votes for a duration proportional to the quality of the communicated option, often resulting in the selection of the best option due to modulation of positive feedback. Another well-known approach to collective decision-making in robot swarms is the use of the local majority rule [16,23,28,11], where each robot chooses the option with the most votes from its neighbours. Compared to the voter rule, the local majority rule has higher computational costs (as robots need to process and aggregate multiple votes instead of randomly selecting one), allowing the robots to pool more accurately neighbour's opinions and, in turn, enabling quicker collective decisions [54,40]. The concept of majority rule, where options favoured by over 50% of neighbouring agents are selected, can be extended by employing various sub- and super-majority quorums [26], such as the k-unanimity rule [44] or the q-voter rule [27]. Besides selecting social information using methods like voter or majority rules, a robot must integrate new social information with its personal opinions. This integration can be done using the direct-switch rule to overwrite the opinion with new social information [53] or by temporarily dropping any personal opinion before adopting a new one, as in the cross-inhibition rule [37].

Besides a few exceptions, e.g., [31,34,49], most research efforts for solving the best-f-n problem in robot swarms have primarily examined scenarios in static environments, where environmental conditions and option qualities remain constant over time. Even more rare are studies on collective decision-making in environments with $n > 2$, i.e., more than two options [14,51,25,3,10]. In real-life applications (e.g. drug-delivering nanobots in human bodies or detection or clearing of chemical spills), robot swarms may need to operate in dynamic environments with several alternatives that change quality over time. Hence, an important aspect to consider in designing robot swarms is their adaptivity, i.e., the swarm's capability to reconsider its opinion in response to environmental changes, such as when the quality of an option diminishes, or when a new, more desirable option appears [22]. Survival in an uncertain environment, even for biological systems, requires the ability to infer and respond to changing environmental conditions [9,12,43]. For adaptivity to occur in robot swarms, it is imperative that the swarm does not become 'locked-in' on an outdated opinion. One way to enable the system to keep reconsidering decisions is through the introduction of noise, which can be either random noise [51], or noise introduced by periodically resetting the robots' personal opinions, in this way allowing the integration of new environmental evidence [31,49]. In general, relying too much on social information can prevent the system from adapting to changes [35,34,42], as indicated by previous work that highlighted the benefits of lower robot connectivity to keep social exchanges low and enable group adaptivity [51,2]. Other work avoided opinion stagnation by including in the swarm a group of robots that never change opinion (called stubborn or zealot robots) [35,34]. These robots occasionally influence the other robots to reconsider opinions that are not shared by the majority; although this solution enables adaptivity, it requires the designer to know in advance all the available options so that stubborn robots can be allocated accordingly.

The majority of studies investigating collective decision-making have considered behavioural homogeneity, where all the robots in the swarm follow the same rules for selecting and updating opinions, i.e., every robot runs the same algorithm. While behavioural heterogeneity can enhance the functionality of both robotic and natural swarms [30,24,57], models of swarms comprising robots with different behaviours are more difficult to analyse as even minor behavioural differences can trigger significant, often unpredictable changes in the collective dynamics. For instance, even unintentional differences, such as differences in actuation errors, can yield qualitatively different collective responses [36]. Although the literature on this topic remains relatively sparse, prior research has highlighted the huge potential of heterogeneity, showing performances surpassing those of homogeneous swarms [20,21,1].

In this paper, we exploit behavioural heterogeneity to improve the adaptivity of the robot swarm in dynamic environments. We focus on a best-of-5 problem in a collective perception scenario used in several collective decision-making studies [56,14,58,3,47]. In this scenario, the environment floor is covered by coloured tiles (see Fig. 1), where each colour represents an environmental attribute. The environment is dynamic due to the periodic change of colour of the tiles. The goal of the robot swarm is to collectively reach a consensus on the predominant floor colour (i.e., the most frequent tile colour) and to adapt the consensus when the environment changes. Because the robots can only make local noisy estimates of the environmental state, they rely on information exchanges between nearby robots to make accurate decisions. The robot algorithms are based on an extension of the agent-based model proposed in [42]. In Sec. 2, we present our agent-based model and in Sec. 3 we describe how our model can be relevant for and implemented in a swarm robotics application. In Sec. 4, through abstract multiagent simulations and realistic physics-based robot simulations, we show that heterogeneous swarms that comprise robots that use the majority rule and robots that use the voter rule can improve the speed response to adapt collectively to environmental changes, surpassing, under some conditions, the performance of a homogeneous swarm. Sec. 5 concludes the paper by indicating how the results of this study could be extended to characterise when either behavioural homogeneity or behavioural heterogeneity enables a better collective adaptivity to environmental changes.

2 The Model

Our model extends the agent-based model proposed in [42]. In our model, a population of individuals live in an uncertain time-varying environment which can be in one of n possible states. Each individual can access the n representations, corresponding to the environmental states, and can communicate these using n different types of messages to form an opinion about the state of the environment. In a simulation run, the environmental state changes every t time steps, transitioning into one of the other $n-1$ states chosen uniformly at random. In between each environmental change, each agent undergoes an update τ times. Individuals can make noisy observations of the environmental state, which constitutes their *personal information*. We model the observation noise through a single parameter, the error probability $\eta \in [0,1]$. When an individual makes an observation, with probability $1-\eta$ the observation is correct corresponding to the true environmental state and with probability η, the observation is incorrect corresponding to one of the wrong states (chosen uniformly at random). In addition to personal information, individuals can also access their *social information* which corresponds to the opinions of their neighbours. Two individuals are neighbours when they are directly connected in the communication network. Individuals combine social and personal information by weighting their personal information by a factor ω and using a simple decision-making rule.

We consider two decision-making rules: the majority rule and the voter rule. With the majority rule, the individual counts how many neighbours have their opinion in favour of each of the n environmental states. The individual adds the value ω to the count for the option corresponding to its personal information. In other words, the count for option i is $M_i + \omega M \delta_i$, where M is the number of neighbours, M_i is the number of neighbours with opinion i, and δ_i is one when the individual's personal information is equal to i and is zero otherwise. Finally, the individual adopts the opinion with the highest count. If more than one state

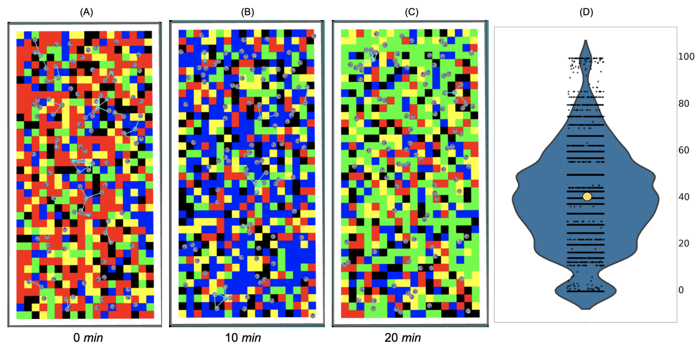

Fig. 1. (A-C) Snapshots of robotics simulations with 100 Kilobots in the tested collective perception scenario. The floor (i.e., the Kilogrid table) is composed of coloured tiles. The robots are tasked with selecting the predominant colour among five alternatives. (A) At the beginning of the experiments, the floor has 40% of red tiles and the remaining 60% are evenly split among blue, green, yellow, and black tiles. (B-C) The environment has changed to a majority for blue and green tiles, respectively. (D) A strip plot overlaid on a violin plot for 5,500 individual environmental observations when the true value is 40% (yellow dot). The simulated robots observe the true environmental state (i.e., the majority colour) on average 71% of the times, thus noise $\eta = 0.29$.

has the highest count, the individual adopts one of them chosen uniformly at random. With the voter rule, the individual chooses the state to adopt as its opinion with a probability proportional to the counts. This rule is equivalent to adopting the personal information with probability $\hat{\omega} = \omega/(1+\omega)$, and adopting the opinion of one randomly-chosen neighbour with probability $1 - \hat{\omega}$.

We form heterogeneous swarms comprising individuals employing one of the two decision-making rules, either voter or majority, where a fraction $k \in [0, 1]$ of the agents use the voter rule and the rest (i.e., fraction $1 - k$) use the majority rule. Thus, for values of $k = 0$ and $k = 1$, the swarm is homogeneous.

We consider asynchronous opinion updates, i.e., at each time step only one individual makes an environmental observation and updates its opinion following either of the two rules. Once it updates its opinion, the individual communicates it to all its neighbours. We consider two types of communication networks: structured and random. In a structured communication network, individuals live on a first-nearest-neighbour square lattice with periodic boundaries and von Neumann connectivity. In a random communication network, each individual is connected to four randomly chosen individuals to whom it transmits its signals. The random network is also dynamic as the neighbours of each individual are drawn randomly each t time steps.

3 Swarm Robotics Simulations

While abstract multiagent simulations offer quick computations of the system dynamics, they may not fully encompass the complexities and real-world constraints of robotic systems. Therefore, to provide a more comprehensive verification of transferability, we test the collective behaviour through physics-based simulations of a swarm of $N = 100$ autonomous robots.

Collective Perception Scenario We test the collective decision-making algorithms in a collective perception scenario where the robots' objective is to collectively infer the environmental state, which comprises randomly distributed floor tiles of five different colours: red, blue, green, yellow, and black, see Fig. 1. This corresponds to an instance of the best-of-5 problem, where the correct environmental state corresponds to the colour that occurs most frequently in the floor tiles. In line with the decision-making literature [38,45,50,15], when scaling the problem to decisions between more than two options, there is one best option (i.e., the predominant colour representing the correct environmental state) and $n-1$, in our case four, remaining options with the same lower quality (i.e., four colours appear in the floor tiles in equal minority proportion).

The environment changes every t time steps, meaning that the tiles' colours change and, after every change, one of the colours that was in minority becomes the predominant one. In the robot simulations, t corresponds to 10 minutes. We run the experiments that last $3t = 30$ minutes, therefore we can test the ability of the robots to adapt to three environmental changes. At the start of each simulation, all robots are initialised at random locations and committed to the blue option while the environment is initialised in a red state (i.e., the red tiles are the most numerous). In this way, we simulate that at the beginning of the simulation, an environmental change has happened. Every 10 minutes, the environment changes, at minute 10, the blue colour becomes predominant and at minute 20, the green colour becomes predominant, as shown in Fig. 1.

Simulation Setup In this analysis, we use Kilobots [41]—small-sized, minimalistic, and cost-effective robots that can broadcast infrared (IR) messages with 9-byte payload in a range of 10 cm, move at $1\,\mathrm{cm/s}$, and have a control loop of approximately 33 ms. Because the Kilobots have limited sensing capabilities, we run our experiments putting the robots in a virtual environment, the Kilogrid [52], that allows them to make environmental readings otherwise impossible. The Kilogrid is an electronic table measuring $1 \times 2\,\mathrm{m}^2$, consisting of 800 cells that can display any RGB colour and interact with the Kilobots via IR messages. All cells, except white border cells (Fig. 1), continuously transmit IR messages with their ID and colour. We simulate the robot swarm using ARGoS [33], a state-of-the-art swarm robotics simulator, which has dedicated plugins to simulate accurately both the Kilobots [32] and the Kilogrid [2].

Robot Behaviour Kilobots perform a random walk (alternating 10 s straight motion with 5 s rotation) to explore the environment and interact with other robots. Since Kilobots lack proximity sensors, Kilogrid cells transmit a binary 'wall flag' (high or low) to indicate proximity to a wall. Border cells and their adjacent non-white cells send a high wall flag, while other internal cells send a low wall flag. Upon receiving a high wall flag, a robot executes a simple obstacle avoidance routine to avoid collisions with the wall.

The robots repeatedly observe the environment, with each observation lasting about 800 robot control loops (rcl). Thus, a robot makes roughly $\tau = 22$ observations between environmental changes (every 10 minutes). During an observation, the robot counts the tiles (Kilogrid cells) of each colour it encounters. The robot can only read the colour of the Kilogrid cell it is on. Through a random walk, during each observation, the robot visits on average 8 tiles. After 800 rcl, the robot sets as its environmental observation the most observed colour (the tile counter is reset every observation cycle of 800 rcl). Because of the limited number of cells visited by the robot, the environmental observation is subject to noise η, i.e., the proportions of tile counts are often different than their true proportions. The value of η varies depending on the true proportions. We test two different proportions (η values) in the robot experiments. When the predominant colour appears in 40% of the tiles and each of the other four colours in 15% of them (as shown in Fig. 1A-C), the average observation error is $\eta = 0.29$ (Fig. 1D). When the predominant colour appears in 30% of the tiles and each of the other four colours in 17.5% of them, the average observation error is $\eta = 0.47$.

Throughout the experiment, the robots broadcast a message expressing their opinion (i.e., the environmental state they believe to be true) every 2 s on average. They also display their opinion to a human observer by lighting their LEDs in the same colour. The majority-rule robots process all received messages, grouping them by colour in a set r, whereas the voter-rule robots only store the colour indicated in the last message they receive, overwriting the content at each new message. After each environmental observation (every 800 rcl), the robot uses either the majority or voter rule for decision-making and updates its opinion. In the majority rule process, the robot counts messages received within r as m, calculates $\omega \times m$, and adds it to the count of the most observed tile colour in r. It then applies the majority rule on r to update its opinion. In the voter rule, the robot uses personal observation with probability $\hat{\omega}$ to switch to the most observed tile colour, or with probability $1 - \hat{\omega}$ uses social information to update its opinion. This process repeats throughout the experiment.

4 Results

We measure the system performance Λ as the proportion of robots that infer the correct environmental state on average throughout the experiment. To do so, we compute this proportion for each time step and divide it by the experiment length. We report results for different levels of personal observation noise η and different personal information weights ω (which is used in the decision rule).

4.1 Multiagent Simulation Results

The results of stochastic multiagent simulations with 100 agents are presented in Figs. 2 and 3. We ran one long simulation with 500 environmental switches

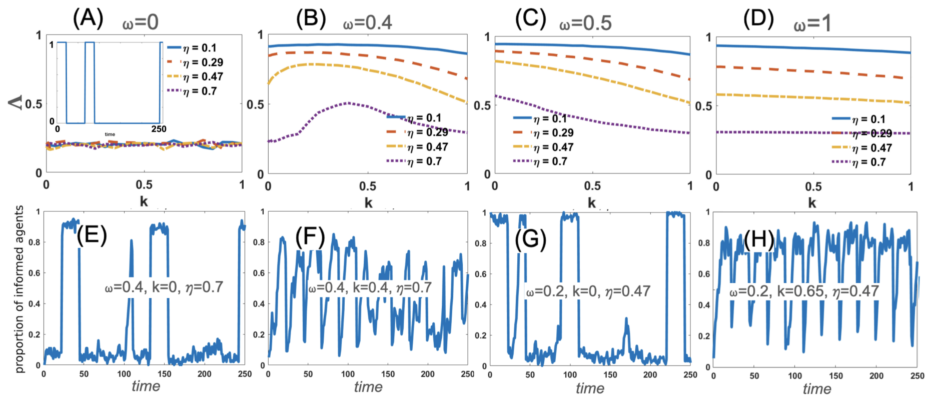

Fig. 2. Multiagent simulation results for swarms of 100 agents interacting on random time-varying communication network. (A-D) Time-averaged fraction of agents inferring the correct environmental state (Λ on the y-axis) for different swarm compositions (k on the x-axis), for various observation noise values $\eta \in \{0.1, 0.29, 0.47, 0.7\}$ and personal information weights $\omega \in \{0, 0.4, 0.5, 1\}$ (average of 11,000 time steps, with environmental changes every $\tau = 22$ time steps). The inset in panel A shows a snippet of the simulation with time on the x-axis and the proportion of correct agents on the y-axis. (E-H) The temporal evolution of some selected cases for a short segment of 250 time steps for $\omega = 0.4$ and $\eta = 0.7$ (in E, F), and $\omega = 0.2$ and $\eta = 0.47$ (in G, H).

(thus a total of $500t = 11,000$ time steps) per configuration. Since the simulation process is ergodic, the results are akin to running multiple simulations.

Figures 2A-D shows results for four values of $\omega \in \{0, 0.4, 0.5, 1\}$ and four personal observation noise $\eta \in \{0.1, 0.29, 0.47, 0.7\}$. Relying prevalently on social or personal information ($\omega = 0$ and $\omega = 1$, respectively) leads in most cases to poor results. For instance, when $\omega = 0$ in Fig. 2A, the agents only use social information, making the population blind to environmental changes and stuck in an immutable consensus for one option (also represented by the inset). Instead, for intermediate values of ω, the system is able to combine social and personal information achieving higher performance Λ. On the horizontal axis of Fig. 2A-D, we vary the swarm composition k indicating the proportion of voter-rule agents (where the other agents, proportion $1-k$, use the majority rule). While Fig. 2C shows that a homogeneous swarm of majority-rule agents ($k = 0$) has the best performance Λ, Fig. 2B also shows that there are conditions when heterogeneous swarms ($0 < k < 1$) are superior to homogeneous ones. In particular, when observation noise is high and the personal information is weighted less than the social information ($\omega < 0.5$), combining the two types of agents can be beneficial.

The temporal evolution of the fraction of informed individuals shown in Figs. 2E-H helps to understand the dynamics of the system in different configurations. With majority-rule agents only (homogeneous swarm, $k = 0$), the swarm is rarely able to adapt to changes (see Figs. 2E and 2G), behaving in a way comparable to an observation-blind swarm (shown in the inset of Fig. 2A).

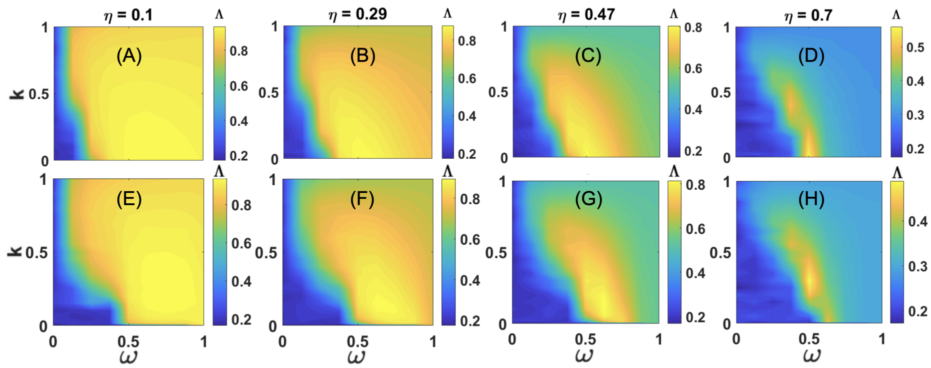

Fig. 3. Results from multiagent simulations with 100 agents on two network types: top row (A-D) shows random time-varying communication, bottom row (E-H) shows structured communication (lattice with von Neuman neighbourhood). The colour map indicates performance Λ, the average fraction of agents inferring the correct environmental state over time. The colour bar scale in each panel is different to better visualise the region of maximal performance in each condition.

In these figures, there are few and high peaks. This means that with $k = 0$ and low $\omega < 0.5$, the agents prevalently rely on social information, obtaining in this way very high levels of group agreement (high peaks), however at the expense of poor adaptability to changes (few peaks). These few peaks mostly happen when, by chance, the environment changes to the state that corresponds to the opinion in which the swarm is locked in. This situation of opinion stagnation can be improved by including in the swarm a proportion of voter-rule agents, as shown in Figs. 2F and 2H for $k = 0.4$ and $k = 0.65$. Here, the system reliably and rapidly adapts to change (periodic peaks at every environmental change), however at the cost of a smaller agreement (lower peaks).

The colour maps of Fig. 3 show a more complete exploration of the parameter space. The two rows show results for the two considered types of networks. The top row shows results for agents interacting on a random time-varying communication network, instead the bottom row shows results for swarms of agents interacting on a static structured communication network. For both types of networks, the trend is similar, however when agents communicate on a structured network, having a heterogeneous swarm can perform better than a homogeneous swarm in a larger range of parameters.

4.2 Swarm Robotics Simulation Results

The swarm robotics results highlight the benefits of having a heterogeneous swarm as combining the two types of robots (majority-rule and voter-rule robots) leads to the best performance in a larger range of parameters. Figure 4 shows the results for two observation noise levels ($\eta = 0.475$ as red points and $\eta = 0.29$ as blue points) and three ω values (different panels). In line with the multiagent

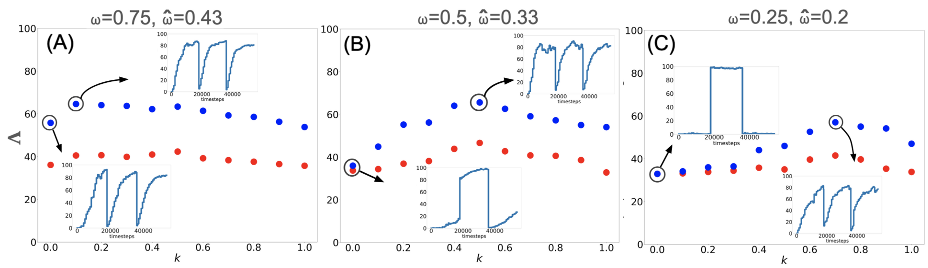

Fig. 4. Swarm robotics simulation results for swarms of 100 Kilobots. Time-averaged fraction of robots inferring the correct environmental state (Λ on the y-axis) for different swarm compositions (k on the x-axis), for various observation noise values (red points for $\eta = 0.475$ and blue points for $\eta = 0.29$) and personal information weights ω. The average is computed on 30 simulation runs per configuration of 30-minute experiments with environmental changes every 10 minutes. Each panel's insets show the temporal evolution of the homogeneous swarm (with majority-rule robots only) and the heterogeneous swarm (with the best performance Λ) for $\eta = 0.29$.

predictions of Sec. 4.1, the robot simulation results show that the benefits of having a heterogeneous swarm are larger when ω is low and the estimation noise is high. The insets in Fig. 4 show the temporal dynamics of some representative cases, confirming that swarms only composed of majority-rule robots ($k = 0$) are unable to rapidly adapt to changes (for high ω, adaptation is slow and for low ω, there is no sign of adaptation). As discussed in the previous section, for $k = 0$ and low ω, the occasional peaks in the temporal evolution are due to the environment transitioning to the state that matches the opinion in which the robots have been initialised. More precisely, this match happens after the first environment change after 10 minutes of simulation, when the predominant colour in the Kilogrid is blue (Fig. 1B). Instead, heterogeneous swarms seem to cope with environmental changes better, to adapt quicker, and have a higher performance Λ than homogeneous swarms in any of the tested conditions. We also conducted a few preliminary real-robot experiments on a swarm of 40 robots, for $\omega = 0.2$ and $\eta = 0.2$, showing promising results with environmental inference improving in heterogeneous swarms (videos in supplementary material [59]).

Surprisingly, the swarm robotics results more closely align with those of a multiagent system on a fixed structured network, where the benefits of heterogeneity are clearer, rather than on a random time-varying network, where heterogeneous systems are optimal in fewer conditions. In fact, the interactions among robots should be better described by a network that changes over time through random encounters. While we do not have a definitive explanation yet, we believe that one of the causes could be the neighbourhood correlation (i.e., there are cliques where several neighbours of my neighbours are also my neighbours), but further investigation is needed for confirmation.

5 Discussion and Conclusion

In this paper, we have studied swarm robotics algorithms for the best-of-n problem in dynamic environments. The robots continuously attempt to infer the correct environmental state which changes over time. To alleviate the effect of potentially large errors in making personal observations of the environment, robots exchange information on what they believe to be the true environmental state: thanks to these social exchanges, they make collective decisions that are more accurate than what they could do if operating alone. Both our and prior analyses suggest that heavy reliance on social information can lead to group opinion stagnation, hindering the swarm from acquiring new environmental information and locking the collective into former beliefs [31,34,49,51,42]. We differentiate from previous work that investigated adaptability to environmental changes in collective best-of-n decision-making of robot swarms by considering a behaviourally heterogeneous swarm. In our study, some of the robots use the majority rule while others use the voter rule (these are two simple voting rules frequently studied in the opinion dynamics and swarm robotics literature [18,16,56,11,4]).

Even though the majority rule is more sophisticated and allows the robots to quantify better the option of the rest of the group (using larger neighbourhood sampling), the introduction of a proportion of robots using the simpler voter rule, which is based on one random social sample, can lead to collective benefits. Our analysis shows that there are indeed conditions where having a heterogeneous swarm composed of two groups of robots, one using the majority rule and the other using the voter rule, can lead to a quicker response to environmental changes. Relevant conditions where we find swarm heterogeneity useful are when there are high levels of errors in personal environmental observations (high noise parameter η) and when the robots have low confidence in their observations and give more importance to social information rather than personal information (low weight parameter ω). In addition, our analysis shows that heterogeneous groups can also be quicker to respond to an environmental change (e.g., see insets of Fig. 4). Future research may investigate whether the idea of having a mix of simpler and more sophisticated robots could also be beneficial in other cases, such as when robots are heterogeneous due to manufacturing differences (e.g., robots with different sensor noise levels) [36].

Our analysis based on multiagent and swarm robotics simulations shows the potential benefits of heterogeneity, expanding the theoretical analysis conducted in [42]. Further analysis should better characterise in which conditions homogeneity or heterogeneity are the best strategy. Our intuition is that information noise and social network correlations may be critical factors. We believe a fruitful research direction is considering heterogeneity for cost-efficiency, as shown in our previous work [1], where combining simple and complex robot behaviours can reduce the average cost of running the swarm algorithm.

Acknowledgements. R.Z and M.D. acknowledge support from the Belgian F.R.S.-FNRS. A.R. acknowledges support from DFG under Germany's Excellence Strategy - EXC 2117 - 422037984.

References

1. Antonic, N., Zakir, R., Dorigo, M., Reina, A.: Collective robustness of heterogeneous decision-makers against stubborn individuals. In: Proceedings of the 23rd International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2024). pp. 68—77. IFAAMAS, Richland, SC (2024)
2. Aust, T., Talamali, M., Dorigo, M., Hamann, H., Reina, A.: The hidden benefits of limited communication and slow sensing in collective monitoring of dynamic environments. In: Swarm Intelligence (ANTS 2022). LNCS, vol. 13491, pp. 234–247. Springer, Cham, CH (2022)
3. Bartashevich, P., Mostaghim, S.: Multi-featured collective perception with Evidence Theory: tackling spatial correlations. Swarm Intelligence **15**(1-2), 83–110 (2021)
4. Canciani, F., Talamali, M.S., Marshall, J.A.R., Bose, T., Reina, A.: 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). p. 4 (2019)
5. Castellano, C., Fortunato, S., Loreto, V.: Statistical physics of social dynamics. Reviews of Modern Physics **81**(2), 591–646 (2009)
6. Chin, K.Y., Khaluf, Y., Pinciroli, C.: Minimalistic collective perception with imperfect sensors. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 8862–8868. Piscataway, NJ (2023)
7. Clifford, P., Sudberry, A.: A model for spatial conflict. Biometrika **60**(3), 581–588 (1973)
8. Correll, N., Martinoli, A.: Modeling and designing self-organized aggregation in a swarm of miniature robots. The International Journal of Robotics Research **30**(5), 615–626 (2011)
9. Couzin, I., Krause, J., Franks, N., Levin, S.: Effective leadership and decision making in animal groups on the move. Nature **433**, 513–516 (2005)
10. Crosscombe, M., Lawry, J.: Collective preference learning in the best-of-n problem. Swarm Intelligence **15**, 1–26 (2021)
11. Cruciani, E., Mimun, H.A., Quattropani, M., Rizzo, S.: Phase transition of the k-majority dynamics in biased communication models. Distributed Computing **36**(2), 107–135 (2023)
12. Dall, S.R., Giraldeau, L.A., Olsson, O., McNamara, J.M., Stephens, D.W.: Information and its use by animals in evolutionary ecology. Trends in Ecology & Evolution **20**(4), 187–193 (2005)
13. Dorigo, M., Theraulaz, G., Trianni, V.: Swarm robotics: Past, present, and future [point of view]. Proceedings of the IEEE **109**(7), 1152–1165 (2021)
14. Ebert, J.T., Gauci, M., Mallmann-Trenn, F., Nagpal, R.: Bayes bots: Collective bayesian decision-making in decentralized robot swarms. In: 2020 IEEE International Conference on Robotics and Automation (ICRA). pp. 7186–7192. IEEE, Piscataway, NJ (2020)
15. Franks, N.R., Dornhaus, A., Best, C.S., Jones, E.L.: Decision making by small and large house-hunting ant colonies: one size fits all. Animal Behaviour **72**(3), 611–616 (2006)
16. Galam, S.: Majority rule, hierarchical structures, and democratic totalitarianism: A statistical approach. Journal of Mathematical Psychology **30**(4), 426–434 (1986)
17. Goss, S., Aron, S., Deneubourg, J.L., Pasteels, J.: Self-organized shortcuts in the argentine ant. Naturwissenschaften **76**, 579–581 (1989)

18. Holley, R.A., Liggett, T.M.: Ergodic theorems for weakly interacting infinite systems and the voter model. The Annals of Probability **3**(4), 643–663 (1975)
19. Hunt, E.R., Cullen, C.B., Hauert, S.: Value at Risk strategies for robot swarms in hazardous environments. In: Unmanned Systems Technology XXIII. p. 117580M. SPIE (2021)
20. Karamched, B., Stickler, M., Ott, W., Lindner, B., Kilpatrick, Z.P., Josić, K.: Heterogeneity improves speed and accuracy in social networks. Physical Review Letters **125**(21), 218302 (2020)
21. Kengyel, D., Hamann, H., Zahadat, P., Radspieler, G., Wotawa, F., Schmickl, T.: 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, pp. 201–217. Springer, Cham, CH (2015)
22. Khaluf, Y., Simoens, P., Hamann, H.: The neglected pieces of designing collective decision-making processes. Frontiers in Robotics and AI **6**, 16 (2019)
23. Krapivsky, P.L., Redner, S.: Dynamics of majority rule in two-state interacting spin systems. Physical Review Letters **90**, 238701 (2003)
24. Krause, J., Ruxton, G., Krause, S.: Swarm intelligence in animals and humans. Trends in ecology & evolution **25**, 28–34 (10 2009)
25. Lee, C., Lawry, J., Winfield, A.F.T.: Negative updating applied to the best-of-n problem with noisy qualities. Swarm Intelligence **15**(1-2), 111–143 (2021)
26. Marshall, J.A.R., Kurvers, R.H.J.M., Krause, J., Wolf, M.: Quorums enable optimal pooling of independent judgements in biological systems. eLlife **8**, e40368 (2019)
27. Mobilia, M.: Nonlinear q-voter model with inflexible zealots. Physical Review E **92**(1) (2015)
28. Montes de Oca, M.A., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., Dorigo, M.: Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. Swarm Intelligence **5**, 305–327 (2011)
29. Önür, G., Turgut, A.E., Şahin, E.: Predictive search model of flocking for quadcopter swarm in the presence of static and dynamic obstacles. Swarm Intelligence pp. 1935–3820 (2024)
30. O'Shea-Wheller, T.A., Hunt, E.R., Sasaki, T.: Functional heterogeneity in superorganisms: Emerging trends and concepts. Annals of the Entomological Society of America **114**(5), 562–574 (2020)
31. Pfister, K., Hamann, H.: Collective decision-making and change detection with Bayesian robots in dynamic environments. In: 2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS). pp. 8814–8819. IEEE, Piscataway, NJ (2023)
32. Pinciroli, C., Talamali, M.S., Reina, A., Marshall, J.A.R., Trianni, V.: Simulating Kilobots within ARGoS: models and experimental validation. In: Swarm Intelligence (ANTS 2018). LNCS, vol. 11172, pp. 176–187. Springer, Cham (2018)
33. Pinciroli, C., Trianni, V., O'Grady, R., Pini, G., Brutschy, A., Brambilla, M., Mathews, N., Ferrante, E., Di Caro, G., Ducatelle, F., Birattari, M., Gambardella, L.M., Dorigo, M.: ARGoS: a modular, parallel, multi-engine simulator for multi-robot systems. Swarm Intelligence **6**(4), 271–295 (2012)
34. Prasetyo, J., De Masi, G., Ferrante, E.: Collective decision making in dynamic environments. Swarm Intelligence **13**(3), 217–243 (2019)
35. Prasetyo, J., De Masi, G., Ranjan, P., Ferrante, E.: The best-of-n problem with dynamic site qualities: Achieving adaptability with stubborn individuals. In: Swarm

Intelligence (ANTS 2018), LNCS, vol. 11172, pp. 239–251. Springer, Cham, CH (2018)
36. Raoufi, M., Romanczuk, P., Hamann, H.: Individuality in swarm robots with the case study of kilobots: Noise, bug, or feature? In: ALIFE 2023: Proceedings of the 2023 Artificial Life Conference. pp. 35–44. MIT Press, Cambridge, MA (2023)
37. Reina, A., Bose, T., Trianni, V., Marshall, J.A.R.: Effects of spatiality on value-sensitive decisions made by robot swarms. In: Distributed Autonomous Robotic Systems (DARS 2016): The 13th International Symposium, pp. 461–473. SPAR, Springer, Cham, CH (2018)
38. Reina, A., Bose, T., Trianni, V., Marshall, J.A.R.: Psychophysical laws and the superorganism. Scientific Reports **8**(4387) (2018)
39. Reina, A., Ferrante, E., Valentini, G.: Collective decision-making in living and artificial systems: editorial. Swarm Intelligence **15**(1-2), 1–6 (2021)
40. Reina, A., Njougouo, T., Tuci, E., Carletti, T.: Speed-accuracy trade-offs in best-of-n collective decision making through heterogeneous mean-field modeling. Physical Review E **109**, 054307 (2024)
41. Rubenstein, M., Ahler, C., Nagpal, R.: Kilobot: A low cost scalable robot system for collective behaviors. In: 2012 IEEE International Conference on Robotics and Automation (ICRA). IEEE, Piscataway, NJ (2012)
42. Salahshour, M.: Phase diagram and optimal information use in a collective sensing system. Physical Review Letters **123**(6), 068101 (2019)
43. Salahshour, M., Rouhani, S., Roudi, Y.: Phase transitions and asymmetry between signal comprehension and production in biological communication. Scientific reports **9**(1), 3428 (2019)
44. Scheidler, A., Brutschy, A., Ferrante, E., Dorigo, M.: The k-unanimity rule for self-organized decision making in swarms of robots. IEEE Transactions on Cybernetics **46**, 1175 (2016)
45. Seeley, T.D., Buhrman, S.C.: Nest-site selection in honey bees: how well do swarms implement the "best-of-n" decision rule? Behavioral Ecology and Sociobiology **49**(5), 416–427 (2001)
46. Seeley, T.D., Visscher, P.K., Schlegel, T., Hogan, P.M., Franks, N.R., Marshall, J.A.R.: Stop signals provide cross inhibition in collective decision-making by honeybee swarms. Science **335**(6064), 108–111 (2012)
47. Shan, Q., Mostaghim, S.: Discrete collective estimation in swarm robotics with distributed bayesian belief sharing. Swarm Intelligence **15**(4), 377–402 (2021)
48. Sion, A., Reina, A., Birattari, M., Tuci, E.: Controlling robot swarm aggregation through a minority of informed robots. In: Swarm Intelligence (ANTS 2022). LNCS, vol. 13491, pp. 91–103. Springer, Cham, CH (2022)
49. Soorati, M.D., Krome, M., Mora-Mendoza, M., Ghofrani, J., Hamann, H.: 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). pp. 4100–4105. IEEE, Piscataway, NJ (2019)
50. Talamali, M.S., Marshall, J.A., Bose, T., Reina, A.: Improving collective decision accuracy via time-varying cross-inhibition. In: 2019 International conference on robotics and automation (ICRA). pp. 9652–9659. IEEE, Piscataway, NJ (2019)
51. Talamali, M.S., Saha, A., Marshall, J.A.R., Reina, A.: When less is more: Robot swarms adapt better to changes with constrained communication. Science Robotics **6**(56), eabf1416 (2021)

52. Valentini, G., Antoun, A., Trabattoni, M., Wiandt, B., Tamura, Y., Hocquard, E., Trianni, V., Dorigo, M.: Kilogrid: A novel experimental environment for the kilobot robot. Swarm Intelligence **12**(3), 245–266 (2018)
53. Valentini, G.: Self-organized collective decision-making in swarms of autonomous robots. In: Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2014). pp. 1703–1704. IFAAMAS, Richland, SC (2014)
54. Valentini, G., Brambilla, D., Hamann, H., Dorigo, M.: Collective perception of environmental features in a robot swarm. In: Swarm Intelligence (ANTS 2016). LNCS, vol. 9882, pp. 65–76. Springer, Cham, CH (2016)
55. Valentini, G., Ferrante, E., Dorigo, M.: The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives. Frontiers in Robotics and AI **4**, 9 (2017)
56. Valentini, G., Ferrante, E., Hamann, H., Dorigo, M.: Collective decision with 100 Kilobots: Speed versus accuracy in binary discrimination problems. Autonomous Agents and Multi-Agent Systems **30**(3), 553–580 (2016)
57. York, C., Madin, Z.R., O'Dowd, P., Hunt, E.R.: Shaping multi-robot patrol performance with heterogeneity in individual learning behavior. arXiv preprint arXiv:2403.01181 (2024)
58. Zakir, R., Dorigo, M., Reina, A.: Robot swarms break decision deadlocks in collective perception through cross-inhibition. In: Swarm Intelligence (ANTS 2022). LNCS, vol. 13491, pp. 209–221. Springer, Cham, CH (2022)
59. Zakir, R., Salahshour, M., Dorigo, M., Reina, A.: Supplementary material for "Heterogeneity can Enahnce the Adaptivity of Robot Swarms to Dynamic Environments" (2024), `https://iridia.ulb.ac.be/supp/IridiaSupp2024-003/`