



Deposited via The University of Sheffield.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/179392/>

Version: Accepted Version

Article:

Talamali, M.S., Saha, A., Marshall, J.A.R. et al. (2021) When less is more: Robot swarms adapt better to changes with constrained communication. *Science Robotics*, 6 (56). eabf1416. ISSN: 2470-9476

<https://doi.org/10.1126/scirobotics.abf1416>

This is the author's version of the work. It is posted here by permission of the AAAS for personal use, not for redistribution. The definitive version was published in *Science Robotics* on Vol 6 No 56 28 Jul 2021, DOI: <https://doi.org/10.1126/scirobotics.abf1416>.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

When less is more: robot swarms adapt better to changes with constrained communication

Mohamed S. Talamali^{1,2}, Arindam Saha¹, James A. R. Marshall¹, Andreagiovanni Reina^{1,3*}

¹ Department of Computer Science, University of Sheffield, UK

² Department of Computer Science, University College London (UCL), UK

³ IRIDIA, Université Libre de Bruxelles, Belgium

*andreagiovanni.reina@ulb.be

Short title: When less is more in robot swarm adaptation

Summary: Swarms of minimalistic robots can better respond to environmental changes when communication among robots is reduced

To effectively perform collective monitoring of dynamic environments, a robot swarm needs to adapt to changes by processing new information and discarding outdated beliefs. We show that in a swarm composed of robots relying on local sensing, adaptation is better achieved if the robots have a shorter rather than longer communication range. This result is in contrast with the widespread belief that more communication links always improve the information exchange on a network. We tasked robots with reaching agreement on the best option currently available in their operating environment. We propose a variety of behaviours composed of reactive rules to process environmental and social information. Our study focuses on simple behaviours based on the voter model—a well-known minimal protocol to regulate social interactions—that can be implemented in minimalistic machines. While different from each other, all behaviours confirm the general result: the ability of the swarm to adapt improves when robots have fewer communication links. The average number of links per robot reduces when the individual communication range or the robot density decrease. The analysis of the swarm dynamics via mean-field models suggests that our results generalise to other systems based on the voter model. Model predictions are

confirmed by results of multiagent simulations and experiments with 50 Kilobot robots. Limiting the communication to a local neighbourhood is a cheap decentralised solution to allow robot swarms to adapt to new information that is locally observed by a minority of the robots.

Introduction

Monitoring an environment through a swarm of minimalistic robots can be useful in adverse scenarios that impose constraints on the individual robots' capabilities [1, 2, 3]. Examples are biodegradable devices—simple by design constraints—to monitor remote locations, such as ocean floors or in-body blood vessels [4], or disposable devices—simple by budget constraints—deployed in hazardous search and rescue missions with a high risk of damage [5]. This type of application may not allow for centralised control or human supervision, whereas controlling the robots via minimalistic decentralised behaviours can be a viable solution. Minimal computing provides the advantage of higher transferability to simpler platforms, such as nano and micro-robots [6, 7]. In this paper, we investigate the general scenario in which the environment has n alternative target sites, each with an intrinsic importance (or quality), and the swarm is tasked with reaching a consensus in favour of the most important site, the so-called *best-of- n* problem (see Fig. 1).

Considerable work has been dedicated to the design of decentralised robot behaviours to solve the best-of- n problem [10, 11, 12]. Compared with previous work, we solve a more general variant of the problem using simpler robots in terms of required capabilities. While most studies have investigated solutions of the best-of- n problem in static and binary ($n = 2$) setups, here we provide minimal behaviours to reach consensus decisions in a dynamic environment in which the number of target sites n , as well as their quality, vary over time. Despite previous studies having shown qualitative changes in the system dynamics for $n > 2$ [13], only a few studies have considered the best-of- n problem with more than two options. The proposed solutions typically employ robots with higher requirements than ours in terms of computation, memory, and communication capabilities, and assume prior knowledge of the number and location of the alternatives [14, 15, 16, 17, 18] (see also Text ST1). Exceptions are our previous works [19, 20], which rely on similar minimalistic behaviour and absence of prior knowledge, but are limited to static environments. There are other minimal behaviours, *e.g.* [21], that could potentially be modified to remove limiting assumptions on robot's prior




Figure 1: **Adaptive monitoring of time-varying environments.** (A) The robot swarm may be deployed to monitor a forest fire and collectively select the best site (*e.g.* most urgent) where the action of firefighters on the ground is necessary [8]. Robots explore the environment to acquire individual information and communicate with each other to exchange opinions. (B) We test the behaviour on a swarm of 50 Kilobots, which are simple robots with limited capabilities. The target sites, here superimposed on the image as coloured circles, are perceived by the robots through the Augmented Reality for Kilobot (ARK) system [9], which allows the robots to perceive a simulated time-varying environment. Over time, new sites appear and existing sites change in quality or disappear. The swarm adapts accordingly. Videos of the experiments are available as supplementary electronic material as Movies 1, 2, and 3.

knowledge, however, it has been shown they are unable to adapt to environmental changes [21, 22]. The few studies on best-of- n decisions in time-varying environments are limited to $n = 2$ binary decisions with robots knowing *a priori* the options [23] or their location [24, 22, 25], and only agreeing on the one with the highest quality. Requiring prior knowledge about the number of alternatives n may reduce the applicability of such solutions. As well as collectively selecting the best alternative, a decentralised process of decision making should also include the phase of decentralised discovery of the available alternatives [26, 27]. Behaviours that proved successful in the voting phase may suffer a drastic reduction of their performance, when both the discovery and voting phases are considered [20]. We study collective behaviours to operate in time-varying environments, therefore they include both the mechanisms to discover environmental changes and to spread the new information. Time-varying environments are an

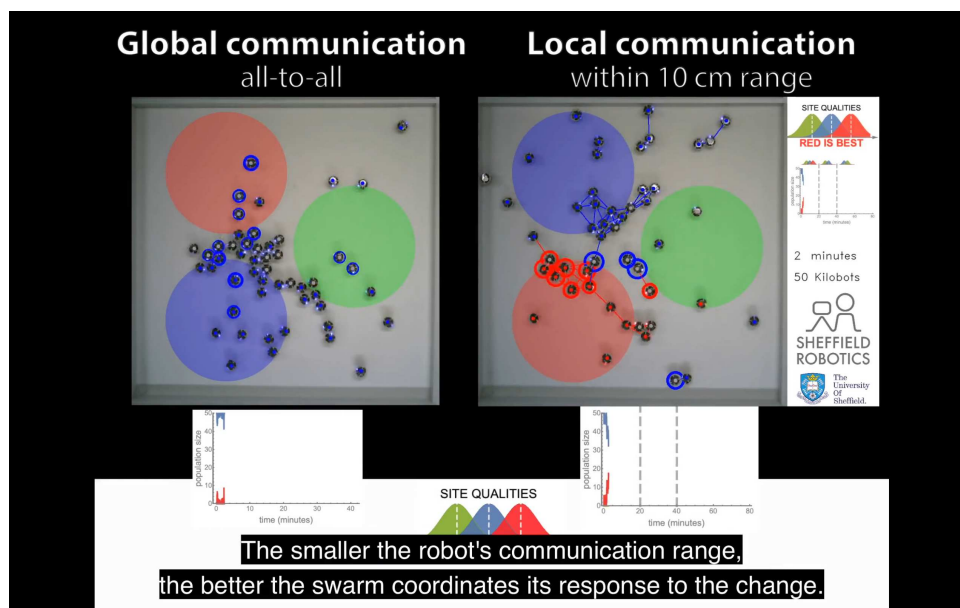
intrinsic characteristic of several application scenarios, and efficient solutions that consider this aspect are therefore necessary for the deployment of robot swarms into the real world.

In this study, the robots have no prior knowledge of the problem. Instead, a robot can only know about a target site either via individual exploration of the environment (it discovers the site) or through social interactions with other robots (it receives the site’s location). We consider robots that have minimal sensory, memory, and communication capabilities. In terms of sensory capabilities, robots in close proximity to a target site are able to make noisy estimates of its quality. Additionally, as the task requires the selection of the best site, robots are able to estimate their own approximate position and move within the environment. In terms of memory capabilities, each robot can only memorise the location and quality of the selected site—that is, the robot has one opinion about the best site. In terms of communication capabilities, a robot can only locally broadcast one single piece of information: the location of the site it considers to be the best.

The robots combine information that they locally acquire in the environment with information that they receive from other swarm members. Our behaviours are based on the classical voter model [28], in which each robot is iteratively influenced by a single random neighbour. Individual and social information is combined through a behaviour based on the *cross-inhibition* pattern [19, 11], by which conflicting information between two communicating robots causes the robots to reset their own opinion and poll other robots’ opinions. Via cross-inhibition a swarm can reach a consensus on the best available option while avoiding decision-deadlocks, as shown in theoretical models [29, 13], honeybee nest-site selection [30], and robotic applications [31, 20]. A widely employed alternative behaviour is based on the *direct-switching* pattern [32]. This, however, has the limitation of only breaking the deadlock of symmetric decisions—when options have the same quality—through noise [33, 34, 35], and therefore can be slow. Through a combination of experiments, simulations, and mathematical analysis, we study when behaviours based on cross-inhibition and direct-switching can adapt to changes in the environment, in particular when the best site appears, disappears, or changes its quality. Through analysis at multiple description levels we measure to what extent these behaviours are scalable to increasing swarm sizes, are sensitive to social information, and are robust to sensorial noise.

To precisely control the swarm behaviour and predict its dynamics in different scenarios, we model the collective dynamics through a system of ordinary differential equations (ODEs).

In swarm robotics, accurate models are necessary but generally hard to obtain [1, 36, 12]. We show that our model accurately predicts the swarm dynamics and highlights a counter-intuitive mechanism: by reducing the range of communication, the swarm can better adapt to changing environments. This result is general across our tested behaviours, and through the model we can understand the cause of this effect. In our experiments, we observe that *less*, in terms of fewer communication links per robot, leads to *more* effective spreading of information within the swarm (Movie 1). This result is in contradiction to the widely accepted belief that more connected networks share information more effectively [37, 38, 21, 39], and is instead congruent with works that document the emergence of the “*slower is faster*” effect [40]. This effect occurs when increasing the performance at the individual level causes a decrease in the collective performance, and has been found in several other contexts, such as ecology [41], voting dynamics [42, 43], and collective animal behaviour [44, 45]. In this study, by reducing the number of communication links, robots sacrifice information spreading speed, which is maximised in highly connected swarms, to facilitate adaptation. Such a solution is simple and highly effective.



Movie 1. Less is more: when simple robots have access to less social information, due to fewer communication links, they can adapt better to environmental changes.

Results

We designed two collective robot behaviours to solve the problem of selecting the best site (best-of-n) in dynamic environments. We opted for minimalistic behaviours which can also run on minimal machines with limited capabilities. Both behaviours extend the classical voter

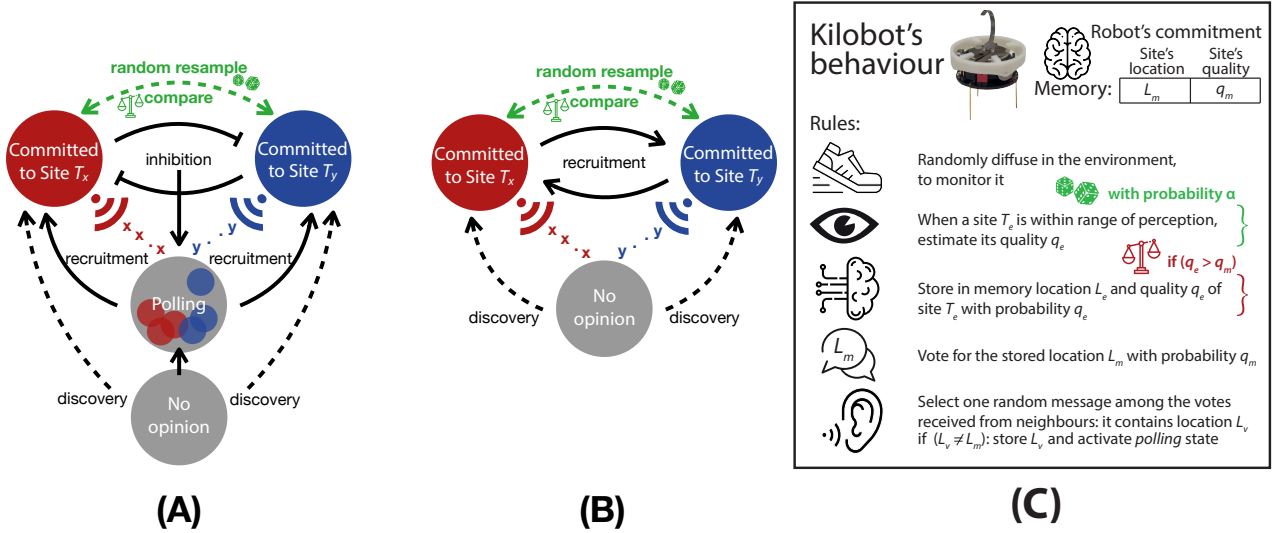


Figure 2: **Individual robot behaviour.** (A)-(B) The robot iteratively updates its opinion based on environmental (dashed transition lines) and social information (solid transition lines). The behaviours can include any of two alternative exploration rules (green dashed transition), *compare* and *resample*, and any of two alternative social interaction patterns (black solid lines), cross-inhibition (A) and direct-switching (B). Through direct-switching, any robot can get recruited and immediately changes its commitment state. Instead, through cross-inhibition only polling robots can be recruited. Therefore, when a committed robot receives a message from a robot committed to a different site, it resets its commitment (inhibition) and activates the polling state. (C) Pseudo-algorithm of the robot behaviour in the committed state (and in the uncommitted state when memory is void). The green and red conditions are mutually exclusive and indicate the effect of the *resample* and *compare* rule, respectively. These minimalistic behaviours can be potentially implemented on extremely simple devices. See more details in Text ST2.

model [28], in which, at each control step, a robot randomly selects one of the messages from its neighbours to update its opinion (see Fig. 2). The message only contains the location of the sender's preferred site i , thus a robot, once informed about a new site, goes to assess its quality q_i . While doing so, the robot remains in a *polling* state in which it listens to incoming messages that it uses to update the location of the target site. Equipped with social information, the polling robot follows a biased random walk until a target site is reached and estimates its quality. During this biased random walk, the robot most often reaches the target site that has more support among its neighbours. Robots, to avoid the quick spreading of erroneous information, do not share the quality of their preferred site but each make an independent noisy estimate, a method that has been shown to improve collective decision accuracy [46, 20]. Non-polling robots instead diffuse in the environment through a random walk to monitor the available target sites and to share their opinion with each other on the best site. When a robot does not have an opinion (*i.e.* it is uncommitted) and encounters a new target site T_x , it makes a noisy estimate of its quality q_x and selects T_x as the best site (commits to T_x) with probability

proportional to q_x . Every time a robot committed to \mathcal{T}_x moves in the proximity of site \mathcal{T}_x , it updates the noisy estimate q_x to keep track of possible quality variations over time. The swarm converges towards the best option because each robot communicates with a frequency linearly proportional to the estimated quality [47]. Such quality-dependent communication was inspired by the house-hunting behaviour of social insects [48, 49] and was successfully implemented in several swarm robotics systems [46, 47, 50, 21, 19].

The presented behaviour implements the cross-inhibition pattern for the update of social information [29, 19, 13, 30]. The peculiarity of this social update pattern is the inhibition between robots committed to different sites (Fig. 2(A)). Upon inhibition, the robot enters a state of ‘indecision’, the polling state, during which it temporarily suspends active recruitment and polls other robots’ opinions until it gets recruited. The alternative social update pattern is the direct-switching pattern, by which robots committed to different sites directly recruit each other (Fig. 2(B)). A recruited robot directly switches its commitment without activating the polling state. We tested both social information transfer patterns in theory, and the cross-inhibition pattern, which analysis predicts is more robust [29, 13], in physical robot swarms.

Our focus is on the ability of the swarm to collectively select the best site and, more importantly, to adapt to environmental changes when a better site appears, the best site disappears, or a site’s quality changes. After any of these changes, we want the swarm to converge to a stable consensus in favour of the best site, with a supermajority of the population (quorum 80%) having the same opinion. In order to let the swarm adapt, we introduce two alternative rules to allow individual robots to reconsider their opinion when exposed to new environmental evidence; `compare` and `resample`.

Minimal rules are sufficient to let the swarm adapt to dynamic environments

We propose the `compare` and the `resample` exploration rules to extend the base behaviour of Fig. 2. These rules allow committed robots to constantly process new information they locally acquired from the environment, as otherwise the swarm may ‘freeze’ into an absorbing state. The `compare` rule lets the robot compare the quality of its chosen site with the quality of any site found in the environment and probabilistically commit to the new site only if it has a higher quality. In this way, individual robots locally filter environmental information with a response threshold that dynamically changes with the current quality estimate [51, 52] (see

Materials and Methods and Text ST2). The `resample` rule does not require any comparison but committed robots process new environmental information—upon discovery—with a small constant probability α [53]. In this way, swarms that reached a consensus for the best location maintain on average a small proportion of robots reconsidering their opinion.

Comparing the performance of the two exploration rules —`compare` or `resample`—Fig. 3 shows the relationship between increasing individual robot capabilities and faster collective adaptation. On the one hand, when robots individually filter the environmental information (`compare` rule), the swarm shows a faster collective response to changes (Fig. 3(B)). However, filtering through comparison requires slightly more computation and the possibility to store the option’s quality for subsequent comparison; such requirements may not be available at every degree of individual complexity or even necessary [54]. On the other hand, the `resample` exploration rule is a reactive technique that does not require any additional individual computation nor capability, at the cost of a slower collective adaptation to changes (Fig. 3(D)). Additionally, the individual-level simplicity of this rule requires the selection of the parameter α for probabilistic environmental sampling. Very low α will not let the swarm adapt, while values of α that are excessively high can cause the swarm to remain undecided, with a considerable fraction of robots constantly changing opinion for any of the n alternative sites (see Fig. S1). Therefore the value of α needs to be appropriately selected depending on the scenario and individual characteristics (see sensitivity analysis in Text ST3).

Communication range negatively correlates with swarm adaptability

Multiagent simulations show a counter-intuitive result. Figs. 3(A,C) show the probability of the swarm adapting for increasing communication range r_s and swarm size S . The swarm starts from a full consensus in favour of a target site \mathcal{T}_x when a new target site \mathcal{T}_y with better quality ($q_y > q_x$) appears. We observe that through both rules, `resample` and `compare`, the swarm has a lower probability of adapting with an increased communication range. In the extreme case of a fully connected network—attained with the maximum communication range $r_s = 0.5$ —a simulated swarm of $S > 10$ robots is never able to adapt to new better sites. From the point of view of social interactions (not considering physical interactions), an increase in communication range is equivalent to an increase in robot density. In dense conditions, the high number of neighbours per robot undermines the ability to adapt. Figs. 3(B,D) instead show that extremely low values of the communication range can slow down the adaptation. There is

therefore an intermediate value for which adaptation occurs at maximum speed. Qualitatively similar results can be obtained with more sophisticated mechanisms to sample the neighbours' votes. In Text ST4, we show that collective behaviours based on the local majority rule (i.e. selecting the site that has been voted the most by the neighbours) also benefit from a limited number of communication links per robot.

A strongly-opinionated minority encounters competition among voters

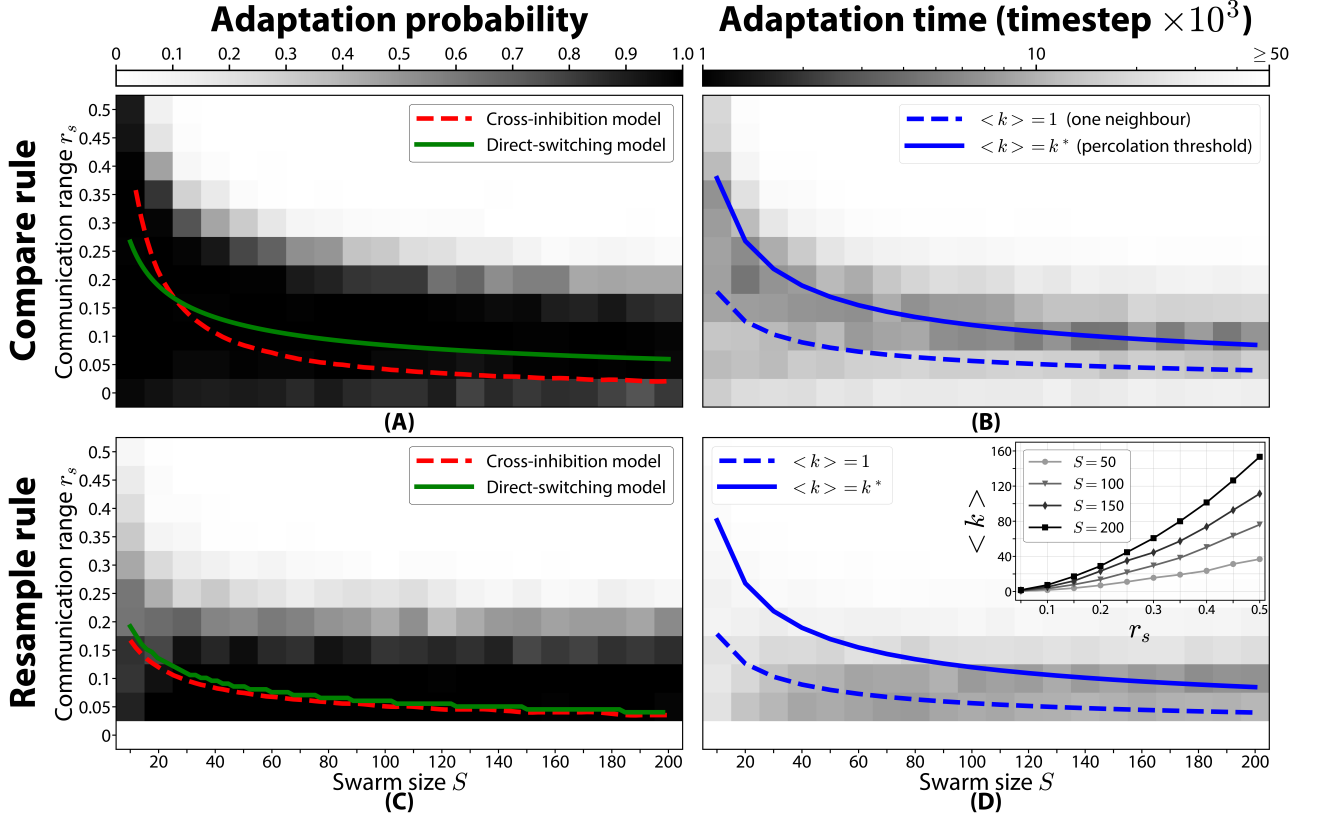


Figure 3: **Large communication ranges and swarm sizes are detrimental to swarm adaptability.** The probability and speed of adaptation to a new better target site (with quality $q_x = 0.8$), when the swarm starts from a full commitment in favour of an inferior site (with quality $q_y = 0.7$), for the four analysed behaviours. Greyscale maps show results for 100 multiagent simulations of $T_{\max} = 6 \cdot 10^4$ time steps; simulated agents only use the cross-inhibition pattern. Superimposed lines are theoretical predictions; theory and simulations show good qualitative agreement. **(A)-(C)** Adaptation probability is the proportion of runs that adapted over the total number of simulations. Adaptation probability decreases for increasing swarm size S and increasing communication ranges; lines show the bifurcation point (see Materials and Methods) for both social interactions patterns. **(B)-(D)** Adaptation speed is high for low communication ranges and swarm sizes; superimposed lines show predicted connectivity transitions: the dashed curve predicts, on average, one neighbour per robot $\langle k \rangle = 1$ and the solid curve $\langle k \rangle = k^* = 4.51$ neighbours per robot (corresponding to the giant-component transition [55]). The best performance in terms of both speed and ability to adapt can be achieved with intermediate values of $\langle k \rangle$. The inset shows that increasing the robots' communication range r_s , or the swarm size S , the average number of communication links per robots per timestep $\langle k \rangle$ increases accordingly.

To understand and predict the swarm behaviour and the effect of the parameters on the performance, we modelled the collective decision making process through a system of ODEs. Each equation describes how subpopulations (groups of robots with the same opinion) change over time as a function of environmental characteristics and robots' capabilities. While we control the robots with individual-level behaviours (Fig. 2), we are interested in understanding and predicting the resulting collective behaviour, which we describe with our models.

A classical approach to model the collective behaviour of a robot swarm is to build a mean-field model describing the average behaviour of an infinite-sized swarm of fully-connected individuals [36]. While this type of model has proved very useful in several scenarios [56, 57, 58], their assumptions make them of little utility to explain the dynamics observed in the scenario we consider here. A model that assumes an infinite-size system cannot describe size-dependent dynamics. We observed above that the investigated swarm has a qualitative change in its environmental response depending on its size (see Fig. 3). Additionally, in swarm robotic systems local communication limits the interaction at each voting iteration to a limited neighbourhood (a small fraction of the entire population), therefore assuming a fully-connected mean-field communication topology may typically be inaccurate. Therefore we developed an ODE model that has explicit dependence on swarm size and robot density, and is able to describe effects determined by a sparse communication topology.

Typically, at the start of every adaptation, the swarm has reached a consensus for the previously best site \mathcal{T}_y , when a new better site \mathcal{T}_x appears. Therefore subpopulations committed to different sites have sizes very different from each other and the *few-vs-many* condition arises—that is, there is a small minority competing against a large majority. Depending on how frequently their members vote, subpopulations can be considered as strongly- or weakly-opinionated. Since communication frequency is proportional to site qualities q_x and q_y , with $q_x > q_y$, the small fraction of robots that spontaneously discover \mathcal{T}_x —the strongly-opinionated minority of size S_x —will vote more frequently than robots committed to \mathcal{T}_y —the weakly-opinionated majority of size S_y . However, in the *few-vs-many* condition, *competition* among voting messages may nullify the bias from quality-dependent communication frequency, and lock the swarm into a consensus for the inferior site \mathcal{T}_y . Competition arises because robots select messages following the voter model approach [28]. Therefore a robot with m neighbours will select (and process the information of) one message among the m received messages (assuming that all m neighbours have sent their message). As a consequence, each neighbour

(voter) of a robot (receiver) has a $1/m$ probability that its message would be read. It is therefore clear that increasing the number of neighbours that each robot has would dilute the impact of each voter. We model such a competition among voters via the Holling type II functional response of Fig. 4(A), which was originally formulated in ecology to describe the interplay between populations of prey and predators [59]. This functional response accounts for the fact that a predator requires time to consume prey. Therefore the biomass of the consumed prey increases sublinearly with the biomass of the prey population. The same functional form has also been used in different fields with different names, for example, the Michaelis-Menten equation in chemical kinetics [60] and the Hill equation in biochemistry [61].

Borrowing the terminology from ecology, we show in Fig. 4(A) the extreme few-vs-many condition of a single ‘predator’ (minority committed to \mathcal{T}_x of size $S_x = 1$, red agent in Fig 4(A)) and $S - 1$ ‘prey’ (susceptible robots, majority committed to \mathcal{T}_y). The ‘predator’ can ‘eat’ (recruit) a number of ‘prey’ that is a function of the majority size $S_y = S - 1$. In small swarms, the recruitment rate per voter is limited by the event of susceptible robots entering the communication range of the single committed robot (which occurs with rate proportional to the robot’s communication area πr_s^2). By increasing the swarm size, the recruitment rate to \mathcal{T}_x increases as the probability of committed-susceptible interaction increases, until it saturates at a maximum rate where the population density is high (Fig. 4(A)). In large swarms, the recruitment rate saturates because in a high density situation each message competes with several others to be read. Note that, in our system, the robots both receive and send messages, therefore increasing the number of susceptible robots (majority) also increases the competition among messages.

Through the functional response of Fig. 4(A), our model introduces recruitment rates that are asymmetric with respect to the number of recruiters and susceptible robots. Fig. 4(B) shows that asymmetry in the interactions can lead to a higher recruitment rate for the inferior site than for the superior one, when $S_x \ll S_y$. Instead, with the standard infinite-size approximation [36], the mean-field model always has symmetric recruitment rates (with the red curve of Fig. 4(B) always higher than the blue). Asymmetric recruitment rates cause a non-monotonic commitment rate to \mathcal{T}_x and thus a bistability of consensus (Figs. 4(C)-(E)). The asymmetry vanishes as robot density is decreased (Figs. 4(C)-(D)), restoring mono-stability for the best site. As a consequence, in low robot densities, there is a small difference between ours and the infinite-size models. This difference increases as either swarm size S or communication range

r_s increases. Stability analysis reveals, for both the `compare` and `resample` rules and both the cross-inhibition and direct-switching patterns, the presence of a bifurcation as the robot density is increased (see Fig. 3). Prior to bifurcation (small S , small r_s), the single stable fixed point corresponds to the entire group adapting to the best option, in agreement with infinite-size approximations. After the bifurcation, a second stable fixed point appears which corresponds to the swarm being unable to adapt to the new better option, remaining instead trapped in the current consensus on \mathcal{T}_y (Fig. S6). Employing classical mean-field approximation, no bifurcation would be present and it would fail to describe the robot swarm dynamics.




Figure 4: Modelling asymmetric recruitment rates explains the interplay between a strongly-opinionated minority and weakly-opinionated majority. (A) A single robot committed to the superior site \mathcal{T}_x (red agent in the insets) recruits susceptible robots (blue agents committed to the inferior site \mathcal{T}_y) at a rate that grows sublinearly with the number of susceptible robots (here $S - 1$). The recruitment rate, based on the Holling type II functional response (red solid curve), increases with S when the number of susceptible robots is low and interactions sporadic. For high numbers, the rate saturates to the frequency of transmission (horizontal dot-dashed line). This is in contrast to the infinite-size approximation, in which the recruitment rate (green dashed line) has linear dependence on the number of susceptible robots. (B) Recruitment rate to a superior (red $q_x = 0.8$) and an inferior (blue $q_y = 0.7$) site is asymmetric in our model (solid lines) and symmetric in infinite-size models (dashed lines). We fix $S = 100$ and vary the number of recruiters to \mathcal{T}_x on the horizontal axis (where the recruiters to \mathcal{T}_y are the complement to x-axis values, $S - x$, for direct-switching). (C)-(E) Rate of change of robots committed to site \mathcal{T}_x (y-axis) against proportion of robots committed to \mathcal{T}_x (x-axis), as from Eq. (1), for varying swarm density (C), swarm size S (D), and for the compare (C)-(D) and the `resample` rules (E).

Fewer communication links make the swarm more flexible

Our results hint at a counter-intuitive solution to the challenge of operating large scale swarms: adaptability of the swarm increases as the robots' communication range decreases (Fig. 3). That is, interacting with fewer robots at a time can improve the ability of the swarm to disseminate new information collected locally. We conducted a set of experiments (Fig. 5) with swarms of 50 Kilobots—small robots for collective intelligence studies [62]. When the robots were able to communicate with any other robot—forming a fully-connected topology—the swarm failed to adapt its decision to new better sites (Fig. 6). Once the swarm reached a consensus in favour of one alternative, robots that discovered new sites, even with a better quality than the previous, were a minority compared with the rest of the swarm. That minority immediately faced a large majority that quickly reverted its mind. Limiting communication, minorities with better opinions could gradually gain traction in the population and eventually steer the swarm towards the correct choice (Fig. 5(C)- 5(F)). Figs. 5(A) and 5(B) show the results from simulated and physical robot experiments.

While a small number of communication links can bring advantages in terms of better adaptability, it is important to note that other types of processes can instead benefit from long communication ranges, or even from no communication whatsoever. Text ST5 and Fig. S3 show that information spreading speed is maximised in fully connected networks and decreases by reducing the number of links. Therefore, in processes where no voting among robots is necessary and information needs to spread quickly, large communication ranges are beneficial. Text ST6 and Fig. S4, instead, show that robots with noiseless sensors can achieve high performances without communicating with each other. However, both analyses indicate that a swarm of robots, that rely on noisy sensors and exchange votes to make a collective decision, maximise the ability to adapt to environmental changes via short range communication. The two proposed behaviours can also scale up to large number of options. In the experiments of Figs. 5 and 6 we tested the swarm in collective decisions among up to $n = 3$ sites. In Text ST7, we run a set of simulations to test the swarm adaptability in an environment with an increasing the number of alternative sites, up to $n = 9$. We show that both behaviours naturally scale to higher number of options.




Figure 5: **The robot swarm can collectively select the best target site in a dynamic environment.** (A) Timeseries of six experiments with swarms of 50 Kilobots monitoring time-varying environments with a collective behaviour based on the `compare` rule and the cross-inhibition pattern (thick lines are the mean, and thin lines are single runs). The swarm successfully adapts in all three phases: appearance and disappearance of the best site, and swap of quality between the best and second best site. (B) Similar results are obtained through physics-based simulations with behaviours based on both the `compare` (top) and the `resample` (bottom) rules (lines are the mean of 100 runs with 95% confidence interval as shades). (C)-(J) Overhead view of one experiment at salient moments; there are three target sites—here superimposed on the images as red, green, and blue circles—that can be locally perceived by robots through augmented reality (ARK [9]). (C) The robots initially have a full consensus in favour of the previous best site (blue $q_B = 0.6$), when a new better site (red $q_R = 0.8$) appears. (D) A minority of the swarm discovers the new site. (E) Through local interactions, the red opinion spreads throughout the robot swarm. (F) The swarm converges to a consensus for red. (G) The red site disappears. (H) The swarm reverts to a consensus for blue. (I) The quality of the blue and green sites swap ($q_B = 0.4$ and $q_G = 0.6$). (J) The swarm once again adapts its decision to the best available site. Full videos are available as supplementary electronic material in Movie 2.

Value-sensitive collective adaptation

We can further extend our analysis to examine adaptation as a function of option values; the model predicts a value-sensitive adaptation as shown in Fig. 7. A swarm committed to site \mathcal{T}_y

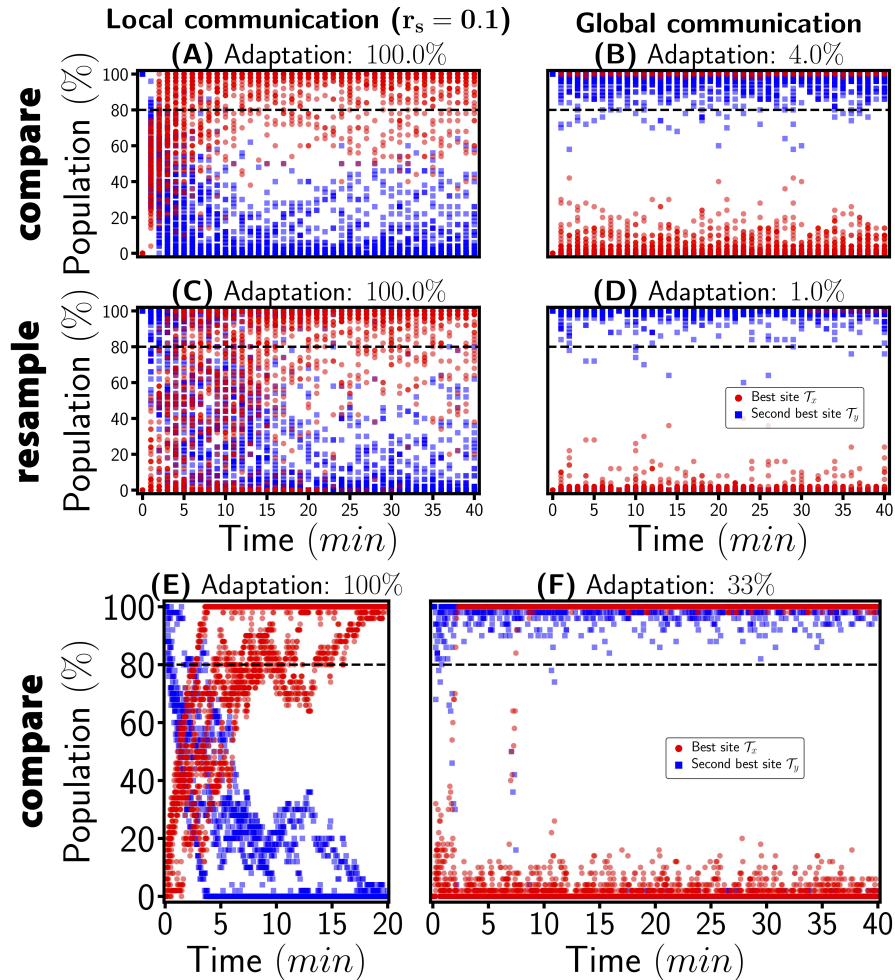


Figure 6: **The Kilobot swarm adapts to environmental changes when robots use short-range communication.** Results from 100 simulations, in (A)-(D), and 6 physical robot experiments, in (E)-(F), with 50 Kilobots with a short (left column) and a long (right column) communication range. The swarm starts with a consensus in favour of the blue site (with $q_B = 0.6$) and is expected to adapt to the red site, which has a higher quality $q_R = 0.8$. Both simulations and physical robot experiments show that the swarm successfully adapts to changes when robots exchange messages locally (a short communication range grants a 100% success rate), but fails to reliably adapt when the communication is global. The proportion of successful adaptations are reported above each scatter plot. A run is considered successful when the average size of the population committed to red in the last 10 minutes is above 80% quorum (horizontal dashed line). Fluctuations are due to a relatively high level of noise $\sigma_q = 0.1$ in robots' estimates of the site qualities. Global communication in the Kilobots is achieved through a virtual transceiver implemented via ARK [9] (Text ST12). Videos of the Kilobot experiments for both conditions (6 repetitions each) are available with the paper in Movies 2 and 3.

with quality q_y will adapt to a new better site \mathcal{T}_x depending both on the quality q_y and the difference $\delta = q_x - q_y$. The minimum quality improvement δ required for adaptation increases with q_y . In other words, the swarm with a consensus in favour of a good location (high q_y) adapts to a better location only if it has a much higher quality (large δ), while adaptation in swarms with low quality opinions (low q_y) also happens for small improvements (small δ). This value-sensitive mechanism is not directly encoded in the individual agent rules, rather it is the

observed emergent behaviour of the collective (see also Text ST8).

Value-sensitivity has been predicted and observed in a variety of natural systems [29, 30, 63, 64] and engineered in robot swarms [31, 65, 11] in a variety of processes, such as decision making or foraging. While most work on decision making focuses on accuracy [15, 22, 21], in which only the best option is rewarded, a value-based metric has a reward dependent on the chosen option’s quality, independently from it being the best [66]; in such scenarios value-sensitive decision dynamics can be beneficial [29]. Our observations of value-sensitive collective adaptation align with previous analysis on costs for switching between options when consensus for one particular option is already established [67].

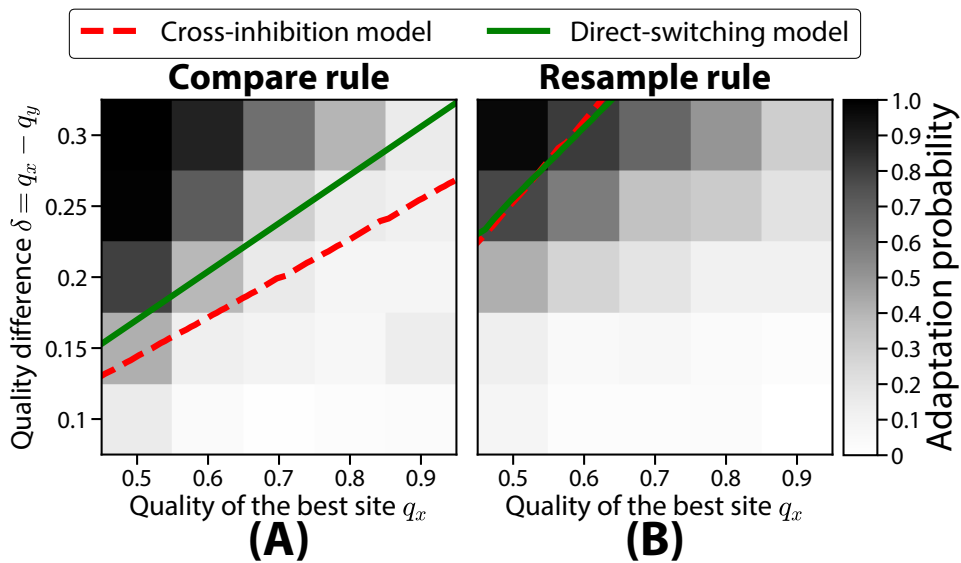


Figure 7: **Value-sensitive adaptation emerges.** The swarm displays a value-sensitive response to environmental changes. This means that a new better site \mathcal{T}_x is selected depending on both the quality of the previous site \mathcal{T}_y and the difference $\delta = q_x - q_y$ between the qualities of \mathcal{T}_x and \mathcal{T}_y . A consensus for a good site (high q_y) is changed only if the quality improvement is high (high δ). Instead, the swarm is less selective (small δ) when the current site’s quality q_y is low. The grayscale maps show the proportion of 100 multiagent simulations that adapted within $T_{\max} = 6 \cdot 10^4$ time steps using the **compare** (A) or the **resample** (B) exploration rules. The $S=100$ agents only implement the social information pattern of cross-inhibition. The superimposed lines are the bifurcation points for the ODE models based on the two alternative social information patterns.

Discussion

We propose two collective behaviours to allow swarms of minimalistic robots to track the best target site in a time-varying environment (the dynamic best-of-n problem). Robotic systems that aim to be deployed in the real world, where real-time changes can be the norm rather than the exception, need to be able to operate in time-varying environments. Our behaviours

enable the robot swarm to successfully adapt to environmental changes, which can be the appearance or disappearance of a site, or a change in sites' qualities. The requirements in terms of individual robot capabilities are minimal, making the implementation possible even in simpler robots than the ones used in our experiments, such as organic nanorobots or disposable devices [4, 5]. Despite the individual simplicity, the swarm is able collectively to track the site with the highest quality and show a value-sensitive response to changes (Movie 1).

Previous work that investigated simple voting behaviours to reach swarm consensus on the best option [21] has shown an increase in decision performance, in terms of consensus speed, when individuals use more social information (a result also confirmed by our analysis in Texts ST4 and ST5). In particular, they replaced the voter model with the local majority rule (selection of the most voted site by neighbours). Their study also showed that such an increase of social information makes the swarm unable to adapt to changing environmental conditions once a consensus has been reached. Given the importance of adapting to time-varying environments, we focussed our analysis on mechanisms that allow or prevent the swarm from effectively adapting the collective decision. Previous theoretical mean-field models based on the infinite-size assumption, predicted that reducing social information—that is, replacing the majority rule with the voter model—would facilitate adaptation to the best available site [50, 21]. Finite-size simulations, however, conflicted with this prediction and showed that adaptation is not possible without strategies that keep the swarm from reaching full consensus (for instance, by using asocial agents) [22]. In this study, we reconcile theory and application: theoretical models allow us to understand the adaptation dynamics and design minimal behaviours that can adapt to changing environments. Physical robot experiments with a swarm of 50 Kilobots and extensive simulations confirm our findings.

Our analysis shows the counter-intuitive result that reducing the connections between individuals improves the spreading of localised information, and, in turn, allows an informed minority to effectively change the opinion of the entire group. This finding is opposed to the widely accepted and intuitive belief in network science that more connections lead to more effective information exchange [37, 38, 21, 68, 39, 69]. While information spreading speed may indeed increase (see Text ST5 and Fig. S3), we show that adaptation—the ability to modify the group's belief in light of new information—is impaired. Adaptation can be restored by reducing the average number of connections per robot; this can be achieved by reducing either the robot's communication range or the robot density (see Fig. 3). Through transition rates that

depend on both the swarm density and the size of subpopulations committed to different sites, we model a form of ‘competition’ among voters that stems from the voter model. The model is complex but tractable and allows us to study, via bifurcation analysis, the swarm’s ability to adapt when the sizes of committed subpopulations are unbalanced—that is, when there is a large majority and a small minority. When robots have a limited number of communication links, the influence of just a few strongly-opinionated robots (high-quality site) can succeed in recruiting other robots. Instead, when the communication can happen within large groups—due, for instance, to a large communication range—the minority’s opinion is suppressed by the large majority, even when the latter is less opinionated (*i.e.* when the majority is committed to a lower quality site). The minority’s inability to spread better information is exacerbated in largely unbalanced subpopulations (few-vs-many) and vanishes when the two factions have comparable sizes (Fig. 4).

Our theory is in agreement with observations from previous swarm robotics studies that investigated the best-of-n problem in dynamic environments. In particular, Prasetyo et al. [22] showed that adaptation can be obtained by freezing a proportion of the swarm committed to every inferior site (through so-called stubborn robots). Our model explains the cause of their empirical observations, as stubborn robots in fact improve adaptability by reducing the imbalance between committed subpopulations (reduce the few-vs-many ratio, see also Text ST16). Despite the promising results, their solution limits the applicability of the behaviour, because robots—as in other studies [23]—need prior knowledge of the alternative options (*e.g.* site locations). Therefore, these approaches may not scale to scenarios with options that dynamically appear and disappear. Our solution is more general as it includes the possibility and necessity of the spontaneous discovery of the options. Balancing the frequency of spontaneous discovery and social interactions is crucial to achieve coordinated responses to environmental changes in collective systems [39, 70, 71, 72, 20]. Again, our model is now able to explain the mechanisms from previous empirical observations. For instance, we previously documented that relatively frequent social interactions speed up the decision but reduce the ability of the swarm to modify its decision once a consensus for an inferior option is made [20]. The best empirically-found solution comprised a first phase with spontaneous discovery only, and a second phase of social interactions. Retrospectively, we can now understand the mechanism that is at the source of the success of such a collective behaviour; it allows the swarm to split into committed subpopulations of comparable sizes before triggering quick consensus. Additionally, we would like to

reiterate the importance of having an adaptable system, as adaptation can act as a means of correction of earlier mistakes, and have a dramatic impact on accuracy. In summary, existing solutions achieved adaptability by avoiding largely unbalanced distributions of opinions in the population. Our understanding of the model allowed us to also propose alternative strategies, for instance, the communication range reduction, that allow adaptability even in case of extremely unbalanced starting conditions.

Our work has the potential to impact on various disciplines. The investigated problem—that is, how an opinionated minority can spread its opinion throughout a large population that holds a different belief—is relevant in biology [73], social sciences [74], and swarm robotics [75]. The underlying mechanism of our collective behaviour—that is, individuals have social interactions based on the voter model [28]—is also widely employed to model opinion dynamics in humans [76], collective behaviour of animals [77], and natural evolution in ecosystems [78], as well as to design robot swarms [10]. The results are not limited to the voter model, as we also tested collective behaviours with social interactions based on the local majority rule and observed the same dynamics (Text ST4). The number of communication links per individual determines when a more-opinionated minority is able to persuade a less-opinionated majority. We conjecture that any voting system, in which probability of adoption of an opinion by a voter is a sublinearly increasing function of its representation amongst the voter’s neighbours, will exhibit the less is more pattern reported here. Reducing the interaction range at the individual level to collectively adapt to changes is a cheap solution that might be exploited by both natural and artificial swarms. Recent studies have indeed observed that animal groups reduce their interaction network to effectively respond to environmental changes [45, 44]. The behaviours proposed in this study also have similarities with the decision-making behaviour of social insects in terms of quality-dependent communication protocols [48, 49] and individual rules to adapt to environmental changes [79]. While it has been shown that direct comparison of alternatives is not necessary to reach a consensus in favour of the best site [53, 80], it remains unclear in which contexts an experience-dependent filtering—similar to our **compare** rule—is adopted by individual insects when reconsidering their choice [51, 81]. While more demanding at the cognitive level, the **compare** rule shows better performance than the cognitively-simpler **resample** in terms of both adaptation speed and robustness to parameter variations (Fig. 3(B,D) and Texts ST3, ST8). In the same way, when the individuals have more capabilities to process more social information via the local majority rule, the swarm adapts faster (Text ST4).

Our study therefore shows a link between the collective performance and the individual cognitive abilities, in terms both of environmental sampling and social interactions.

The simplicity of our approach is one of its strengths as it reduces the complexity at both the individual level, granting a wider applicability, and at the group level, allowing a better understanding of the emergent dynamics. The performance of our behaviours can be improved by increasing the requirements at the individual level. Robots capable of storing probabilities on multiple opinions and updating them recursively could improve the accuracy and speed of collective decisions [14, 15, 17, 18]. Recent theoretical analysis has also shown that distributed learning by computationally more demanding agents can benefit from limited communication [82]. It will be interesting to test how the same mechanism can be applied to a robot swarm. The exploration of the environment could also be made more efficient by replacing the diffusive random walk with more elaborated collective search strategies [83, 84] that, for example, use Lévy walks and larger individual memory [85], or include a constant probability to return to a site and re-estimate its quality [86]. Furthermore, we can envision that the use of a time-varying communication range—varied by individual robots that increase and decrease it depending on how old or new their environmental information is—could further improve the collective behaviour. Such a solution could exploit the benefits of both a quick consensus by highly connected individuals, and effective adaptation to environmental changes by individuals that reduce their response to social influence when they have recent information. Temporarily exploiting more knowledgeable individuals by modifying the communication network is an effective strategy that provides adaptive benefits in animal groups [87, 88] and could be ‘exported’ to engineered solutions. However, even without such refinements, a simple strategy of less is more allows for sophisticated group adaptation in time-varying environments.

Materials and methods

Formalisation of collective adaptation in a time-varying environment

The robot swarm operates in an environment \mathcal{E} that is a plane with n target sites that can vary over time. A target site \mathcal{T}_i is characterised by its location $L_i \in \mathcal{E}$ and its quality $q_i \in [0, 1]$. When L_i is within the robot’s sensing range r_e , the robot can individually estimate the site’s quality $\hat{q}_i \sim \mathcal{N}(q_i, \sigma_q)$ with noise σ_q , truncated to line in the interval $[0, 1]$. We investigate three types of sudden and instantaneous environmental changes: firstly, a new site appears with a

higher quality than any other site in the environment, secondly, the best site disappears, or thirdly, the best and the second-best sites switch their quality. The robot swarm is tasked to react to these changes and to always converge to a consensus in favour of the currently best target site in the environment. We consider the swarm to have adapted to the best site \mathcal{T}_i when the average size of subpopulation S_i committed to \mathcal{T}_i in the last 5000 temporal steps is above the quorum of 80%. We choose this metric to avoid counting random fluctuations as decisions (see Text ST9), as this could impair the subsequent phase of decentralised measuring of the decision state [89, 90, 91].

Individual behaviour for a collective response

Robots combine environmental exploration with social interactions in order to reach agreement with others on the best site. Each robot uses the information that it obtained through exploration and interactions, to iteratively update its commitment state—every 2 seconds—through the finite state machines of Fig. 2. Robots in any commitment state explore the environment, except for the robots in the polling state which move towards the site they have been recruited to.

Environmental exploration is implemented through random diffusion on the plane \mathcal{E} in order to allow robots to both discover target sites and change their interaction neighbourhood. Agent mobility is important as limited mobility can jeopardise the ability to reach a consensus [31, 92, 93]. In the Kilobots, we implemented the random diffusion via the random waypoint mobility model [94] (see further details in Text ST2).

Social interactions consist of the exchange of messages between neighbouring robots which are within the range of communication r_s . Robots committed to site \mathcal{T}_i send their message every 500 ms with probability equal to the estimated quality \hat{q}_i . The message only contains the location L_i of the site \mathcal{T}_i , but not \hat{q}_i . Therefore, robots that receive a recruitment message and change their commitment state do not know the value of \hat{q}_i . These robots change their state to *polling* during which they do not communicate, as they lack information about the site’s quality. Polling robots move through a biased random walk towards the most frequently voted site. Once the target site is reached, they estimate \hat{q}_i , change their state to committed, and start to periodically broadcast their vote message (Fig. 2(C)). Therefore, both *polling* and *committed* robots have an opinion in favour of one site, however the former do not broadcast their opinion while the latter do. We differentiate these two states in terms of individual behaviour, as shown

in Fig. 2(A), however, we count both populations when measuring the collective opinion and swarm consensus.

Adaptability is obtained by allowing robots to integrate new information from the environment after they have committed to a site. The robot changes its commitment in favour of a new site \mathcal{T}_i that is discovered through exploration with probability $d \propto \hat{q}_i$. The relationship between the probability d and the site's quality \hat{q}_i favours the selection of the best site and is determined by the function $f(\hat{q}_i)$, which in our experiments we set as $f(\hat{q}_i) = \hat{q}_i$, as $\hat{q}_i \in [0, 1]$. Through the **compare** rule, a robot committed to \mathcal{T}_j makes this probabilistic change of commitment only if the new site's quality \hat{q}_i is better than \hat{q}_j , *i.e.*, $\hat{q}_i > \hat{q}_j + \epsilon$. The parameter ϵ sets the minimum difference for which adaptation is worthwhile (see Fig. S1(A)), as changing consensus may have a cost [67]. Therefore, the resulting probability of discovering a new site \mathcal{T}_j is $d = f(\hat{q}_i) H(\hat{q}_i + \epsilon - \hat{q}_j)$, where H is the Heaviside step function. Through the **resample** rule, instead, a committed robot considers new environmental information with a constant probability α . Therefore, when a committed robot encounters site \mathcal{T}_i , the total probability of committing to it is $d = \alpha f(\hat{q}_i)$. The probability term α balances the trade-off between a large stable consensus and the ability to react to changes (see Fig. S1(B)). A small α makes the robots' use of new information from the environment rare, whereas a high α makes the swarm more undecided and the consensus subject to large fluctuations.

Kilobot augmented reality experiments

Kilobots are simple robots widely employed for studies of collective robotics and swarm intelligence [62, 95, 96, 21, 65, 97]. Kilobots move on a plane by modulating the vibration frequency of two motors. The robot moves forward at a speed of about $\nu = 1$ cm/s and rotates at about 40° /s. The robots communicate with one another via infrared messages of 9 bytes. The maximum communication range is about $r_s = 10$ cm. Finally, the Kilobot has an RGB LED light to show us its internal state, in our case its commitment. The Kilobots' capabilities are increased by the Augmented Reality for Kilobot (ARK) system [9] that is described in Text ST10.

In our experiments, the Kilobots are augmented via two virtual sensors: a position and a site sensor. Through the position virtual sensor (see implementation details in Text ST11), polling robots use their location and orientation to effectively move to the target sites they want to estimate. Robots also use the position sensor to perform random waypoint exploration and to avoid collisions with the boundary walls. While we resort to the position sensor, efficient robot

navigation can also be attained with other methods that do not rely on any global positioning system, such as social odometry [98] or self-organised navigation [99]. Through the site virtual sensor, instead, robots estimate the site’s quality $\hat{q}_i \sim \mathcal{N}(q_i, \sigma_q)$ when the site’s location is within the perception range $r_e = 0.2$ m of the robot’s virtual sensor. When the quality estimate is outside the sensing range $[0, 1]$, the estimate \hat{q}_i is set to the nearest boundary value. While in local communication experiments ($r_s = 10$ cm), robots exchange messages via an onboard infrared transreceiver, in the experiments with the global communication range, the Kilobots are also equipped with a virtual transreceiver to exchange messages with the entire swarm as illustrated in detail in Text ST12.

Simulators

We analysed the effect of the various parameters of the system through simulations at two levels of abstraction: self-propelled particles and accurate physics-based models. In the former, robots are modelled as point-size agents that move in a 2D space with periodic boundary conditions. Movement and communication is synchronous and noiseless, rotation in place is instantaneous, and collisions are not taken into consideration. The multiagent simulation is not tailored to the specific robotic platform we use, the Kilobot, but rather describes a generic and simplified agent with capabilities equivalent to our robots. The physics-based simulation, instead, accurately simulates the Kilobot’s sensors and actuators, as well as collision and friction between embodied robots. The Kilobots and the ARK systems are both simulated through a dedicated plug-in for ARGoS [100], which is a fast and accurate simulator for swarm robotics. ARGoS has been configured to simulate noise in motion and communication that quantitatively matches with the noise of the real Kilobots [100]. Additionally, ARGoS uses the identical code that runs on the robot, which improves the fidelity of the simulations and facilitates testing and development. All simulation and robot code is available with the paper [101].

Implementing the same behaviour at two or more levels of abstraction is best practice for the analysis of collective systems. In fact, collective systems are typically difficult to model and fully understand. The effect of certain parameters on swarm dynamics can be counter-intuitive (as this study shows) and modelling assumptions may hide emerging patterns (as for infinite-size approximations, for example). Therefore, the implementation of the same behaviour at various levels of complexity can help in the understanding of the system and the generality of the obtained results.

Experimental setup

We conducted a series of experiments to understand the robot swarm behaviour and validate the modelling results. Experiments in simulation used parameters that agree with the real counterpart tested with the Kilobots (*e.g.*, motion speed $\nu = 1 \text{ cm/s}$, environment size $\mathcal{E} = 1 \times 1 \text{ m}^2$, sensing range $r_e = 20 \text{ cm}$, communication range $r_s = 10 \text{ cm}$). All parameters are indicated and discussed in Text ST13; here for completeness we only briefly report an overview of the parameters and their values. The environment includes $n = 3$ target sites with quality $\{q_x, q_y, q_z\} = \{0.8, 0.7, 0.1\}$ for multiagent simulations and $\{q_x, q_y, q_z\} = \{0.8, 0.6, 0.4\}$ for robot experiments, if not indicated otherwise. In both multiagent simulations and robot experiments, noise in individual estimates is relatively high: $\sigma_q = 0.1$ (see inset of Fig. 1(B)). Exploration rules' parameters are $\epsilon = 0.05$ and $\alpha = 0.01$ for the `compare` and `resample` rule, respectively. In the multiagent simulations, the three events of appearance, disappearance, and quality-exchange have been studied in isolation with dedicated experiments. Instead, the robot experiments are long demonstrations (80 minutes, Fig. 5) comprising three phases, in each of which an environmental change occurs (see Text ST13).

Dynamical systems analysis

The collective behaviours that we investigate in this study comprise of one exploration rule among `compare` and `resample`, and one social interaction pattern among direct-switching and the cross-inhibition. The combination of exploration rules and social interaction patterns lead to four distinct, yet related, systems of ODEs which describe the behaviour of the robots. The ODE systems describe the macroscopic dynamics of swarm subpopulations committed to the different sites. In Text ST14, we formulate a system of ODEs for each of the four investigated behaviours. These models significantly simplify when describe the collective adaptation process. The simpler models allow us to compute the bifurcation point as a function of the system's parameters, as detailed in Text ST15. Here, we report the simplified models that describe the adaptation process.

Let S be the swarm size. Let x and y be the fraction of robots committed to the new best and the previously best target sites, with quality q_x and q_y , respectively. Also, let z be the fraction of polling robots. Our system has finite size (that is, a constant number of robots), therefore we have $x + y + z = 1$ for the cross-inhibition patterns, and $x + y = 1$ for the direct-switching pattern (as $z = 0$). This implies, that in case of adaptation, the dynamics reduces

to one dimension in case of direct-switching and to two dimensions for cross-inhibition. In particular,

- **For compare with direct-switching**

$$\dot{x} = k \pi r_e^2 q_x y + \frac{k \pi r_s^2 S y}{1 + k \pi r_s^2 S y} q_x x - \frac{k \pi r_s^2 S x}{1 + k \pi r_s^2 S x} q_y y. \quad (1)$$

- **For compare with cross-inhibition**

$$\begin{aligned} \dot{x} &= k \pi r_e^2 q_x y + \gamma \frac{k \pi r_s^2 S z}{1 + k \pi r_s^2 S z} q_x x - \frac{k \pi r_s^2 S x}{1 + k \pi r_s^2 S x} q_y y \\ \dot{y} &= -k \pi r_e^2 q_x y + \gamma \frac{k \pi r_s^2 S z}{1 + k \pi r_s^2 S z} q_y y - \frac{k \pi r_s^2 S y}{1 + k \pi r_s^2 S y} q_x x. \end{aligned} \quad (2)$$

- **For resample with direct-switching**

$$\dot{x} = k \pi r_e^2 \alpha q_x y - k \pi r_e^2 \alpha q_y x + \frac{k \pi r_s^2 S y}{1 + k \pi r_s^2 S y} q_x x - \frac{k \pi r_s^2 S x}{1 + k \pi r_s^2 S x} q_y y. \quad (3)$$

- **For resample with cross-inhibition**

$$\begin{aligned} \dot{x} &= k \pi r_e^2 \alpha q_x y - k \pi r_e^2 \alpha q_y x + \gamma \frac{k \pi r_s^2 S z}{1 + k \pi r_s^2 S z} q_x x - \frac{k \pi r_s^2 S x}{1 + k \pi r_s^2 S x} q_y y \\ \dot{y} &= k \pi r_e^2 \alpha q_y x - k \pi r_e^2 \alpha q_x y + \gamma \frac{k \pi r_s^2 S z}{1 + k \pi r_s^2 S z} q_y y - \frac{k \pi r_s^2 S y}{1 + k \pi r_s^2 S y} q_x x. \end{aligned} \quad (4)$$

In these equations, α is a proportionality constant representing the rate at which committed robots resample; γ is the proportion of polling robots which get committed to a target site; and k is a proportionality constant for the probability per unit time of a robot encountering a target site or being in communication range with another robot. These probabilities can be expressed as $P_e = k\pi r_e^2$ and $P_m = k\pi r_s^2$, respectively. In our model, the proportionality constant k depends on factors such as the speed of motion of the robots and their movement patterns, and is indicative of the speed of the collective dynamics. Full details on the derivation and analysis of the models are available in Texts ST14, ST15, and ST16; we also include with the paper a Jupyter notebook to reproduce our analytical results [101].

Supplementary Material

- Movie 2. Six experiments with 50 Kilobots that exchange messages in a local range
- Movie 3. Six experiments with 50 Kilobots that exchange messages through global communication

- Text ST1. Comparison with other behaviours
- Text ST2. Individual robots' algorithm
- Text ST3. Effects of the exploration rule parameters on the collective behaviour
- Text ST5. Information spreading trade-off
- Text ST6. The impact of noise in the individual estimations
- Text ST7. Scalability to large numbers of options
- Text ST8. The cost-performance trade-off of two exploration rules
- Text ST9. Detecting a stable consensus can be useful for decentralised quorum sensing
- Text ST10. The Augmented Reality for Kilobot (ARK)
- Text ST11. Virtualisation of the global positioning system
- Text ST12. Virtualisation of a global-range transreceiver
- Text ST13. Parameters of the experimental setup
- Text ST14. Derivation of the mathematical models
- Text ST15. Stability and bifurcation analysis of the ODEs
- Text ST16. More localised information diffuses less
- Algorithm S1. Individual robot's algorithm
- Fig. S1. Effect of the exploration rule parameters on the collective behaviour
- Fig. S3. Information spreading trade-off
- Fig. S4. The impact of noise in the individual estimations
- Fig. S5. Adaptation in the best-of-9 decision problem
- Fig. S6. Bifurcation plots for the four behaviours

References

- [1] H. Hamann, *Swarm Robotics: A Formal Approach*. Cham: Springer International Publishing, 2018.
- [2] G.-Z. Yang, J. Bellingham, P. E. Dupont, P. Fischer, L. Floridi, R. Full, N. Jacobstein, V. Kumar, M. McNutt, R. Merrifield, B. J. Nelson, B. Scassellati, M. Taddeo, R. Taylor, M. Veloso, Z. L. Wang, and R. Wood, “The grand challenges of science robotics,” *Science Robotics*, vol. 3, no. 14, 2018.
- [3] C. Torney, Z. Neufeld, and I. D. Couzin, “Context-dependent interaction leads to emergent search behavior in social aggregates,” *Proceedings of the National Academy of Sciences*, vol. 106, no. 52, pp. 22055–22060, 2009.
- [4] I. C. Yasa, H. Ceylan, U. Bozuyuk, A.-M. Wild, and M. Sitti, “Elucidating the interaction dynamics between microswimmer body and immune system for medical microrobots,” *Science Robotics*, vol. 5, no. 43, 2020.
- [5] N. T. Jafferis, E. F. Helbling, M. Karpelson, and R. J. Wood, “Untethered flight of an insect-sized flapping-wing microscale aerial vehicle,” *Nature*, vol. 570, pp. 491–495, jun 2019.
- [6] M. Gauci, J. Chen, W. Li, T. J. Dodd, and R. Groß, “Self-organized aggregation without computation,” *The International Journal of Robotics Research*, vol. 33, no. 8, pp. 1145–1161, 2014.
- [7] A. Özdemir, M. Gauci, S. Bonnet, and R. Groß, “Finding consensus without computation,” *IEEE Robotics and Automation Letters*, vol. 3, no. 3, pp. 1346–1353, 2018.
- [8] D. Carrillo-Zapata, E. Milner, J. Hird, G. Tzoumas, P. J. Vardanega, M. Sooriyabandara, M. Giuliani, A. F. T. Winfield, and S. Hauert, “Mutual shaping in swarm robotics: User studies in fire and rescue, storage organization, and bridge inspection,” *Frontiers in Robotics and AI*, vol. 7, p. 53, 2020.
- [9] A. Reina, A. J. Cope, E. Nikolaidis, J. A. R. Marshall, and C. Sabo, “ARK: Augmented Reality for Kilobots,” *IEEE Robotics and Automation Letters*, vol. 2, no. 3, pp. 1755–1761, 2017.

- [10] G. Valentini, E. Ferrante, and M. Dorigo, “The best-of-n problem in robot swarms: Formalization, state of the art, and novel perspectives,” *Frontiers in Robotics and AI*, vol. 4, p. 9, 2017.
- [11] R. Gray, A. Franci, V. Srivastava, and N. E. Leonard, “Multiagent decision-making dynamics inspired by honeybees,” *IEEE Transactions on Control of Network Systems*, vol. 5, no. 2, pp. 793–806, 2018.
- [12] M. Brambilla, E. Ferrante, M. Birattari, and M. Dorigo, “Swarm robotics: a review from the swarm engineering perspective,” *Swarm Intelligence*, vol. 7, no. 1, pp. 1–41, 2013.
- [13] A. Reina, J. A. R. Marshall, V. Trianni, and T. Bose, “Model of the best-of-N nest-site selection process in honeybees,” *Physical Review E*, vol. 95, no. 5, p. 052411, 2017.
- [14] C. Lee, J. Lawry, and A. Winfield, “Combining opinion pooling and evidential updating for multi-agent consensus,” in *IJCAI International Joint Conference on Artificial Intelligence*, vol. 1, pp. 347–353, 2018.
- [15] C. Lee, J. Lawry, and A. Winfield, “Negative updating combined with opinion pooling in the best-of-n problem in swarm robotics,” in *Swarm Intelligence, 11th International Conference, ANTS 2018*, vol. 11172 of *LNCS*, pp. 97–108, Springer, 2018.
- [16] A. Scheidler, A. Brutschy, E. Ferrante, and M. Dorigo, “The k-unanimity rule for self-organized decision-making in swarms of robots,” *IEEE Transactions on Cybernetics*, vol. 46, no. 5, pp. 1175–1188, 2016.
- [17] J. Lawry, M. Crosscombe, and D. Harvey, “Epistemic sets applied to best-of-n problems,” in *Symbolic and Quantitative Approaches to Reasoning with Uncertainty* (G. Kern-Isberner and Z. Ognjanović, eds.), (Cham), pp. 301–312, Springer International Publishing, 2019.
- [18] M. Crosscombe, J. Lawry, and P. Bartashevich, “Evidence propagation and consensus formation in noisy environments,” in *Scalable Uncertainty Management* (N. Ben Amor, B. Quost, and M. Theobald, eds.), (Cham), pp. 310–323, Springer International Publishing, 2019.
- [19] A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni, “A design pattern for decentralised decision making,” *PLoS ONE*, vol. 10, no. 10, p. e0140950, 2015.

- [20] M. S. Talamali, J. A. R. Marshall, T. Bose, and A. Reina, “Improving collective decision accuracy via time-varying cross-inhibition,” in *Proceedings of the 2019 IEEE International Conference on Robotics and Automation (ICRA 2019)*, pp. 9652–9659, IEEE, 2019.
- [21] G. Valentini, E. Ferrante, H. Hamann, and M. Dorigo, “Collective decision with 100 kilobots: speed versus accuracy in binary discrimination problems,” *Autonomous Agents and Multi-Agent Systems*, vol. 30, no. 3, pp. 553–580, 2016.
- [22] J. Prasetyo, G. De Masi, and E. Ferrante, “Collective decision making in dynamic environments,” *Swarm Intelligence*, vol. 13, no. 3, pp. 217–243, 2019.
- [23] M. D. Soorati, M. Krome, M. Mora-Mendoza, J. Ghofrani, and H. Hamann, “Plasticity in collective decision-making for robots: Creating global reference frames, detecting dynamic environments, and preventing lock-ins,” in *2019 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 4100–4105, 2019.
- [24] J. Prasetyo, G. De Masi, P. Ranjan, and E. Ferrante, “The best-of-n problem with dynamic site qualities: Achieving adaptability with stubborn individuals,” in *Swarm Intelligence (ANTS 2018)*, vol. 11172 of *LNCS*, pp. 239–251, Cham: Springer, 2018.
- [25] G. De Masi and E. Ferrante, “Quality-dependent adaptation in a swarm of drones for environmental monitoring,” in *2020 Advances in Science and Engineering Technology International Conferences (ASET)*, pp. 1–6, 2020.
- [26] Y. Khaluf, P. Simoens, and H. Hamann, “The neglected pieces of designing collective decision-making processes,” *Frontiers in Robotics and AI*, vol. 6, p. 16, 2019.
- [27] A. M. Hein and B. T. Martin, “Information limitation and the dynamics of coupled ecological systems,” *Nature Ecology & Evolution*, vol. 4, no. 1, pp. 82–90, 2020.
- [28] R. A. Holley and T. M. Liggett, “Ergodic theorems for weakly interacting infinite systems and the voter model,” *Ann. Probab.*, vol. 3, pp. 643–663, 08 1975.
- [29] D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. R. Marshall, “A mechanism for value-sensitive decision-making,” *PLoS ONE*, vol. 8, no. 9, p. e73216, 2013.

- [30] T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall, “Stop signals provide cross inhibition in collective decision-making by honeybee swarms,” *Science*, vol. 335, no. 6064, pp. 108–111, 2012.
- [31] A. Reina, T. Bose, V. Trianni, and J. A. R. Marshall, “Effects of spatiality on value-sensitive decisions made by robot swarms,” in *Proceedings of the 13th International Symposium on Distributed Autonomous Robotic Systems (DARS2016)*, vol. 6 of *STAR*, pp. 461–473, Springer, 2018.
- [32] J. A. R. Marshall, R. Bogacz, A. Dornhaus, R. Planqué, T. Kovacs, and N. R. Franks, “On optimal decision-making in brains and social insect colonies,” *Journal of The Royal Society Interface*, vol. 6, no. 40, pp. 1065–1074, 2009.
- [33] T. Biancalani, L. Dyson, and A. J. McKane, “Noise-induced bistable states and their mean switching time in foraging colonies,” *Physical Review Letters*, vol. 112, no. 3, p. 038101, 2014.
- [34] Y. Khaluf, C. Pinciroli, G. Valentini, and H. Hamann, “The impact of agent density on scalability in collective systems: noise-induced versus majority-based bistability,” *Swarm Intelligence*, vol. 11, no. 2, pp. 155–179, 2017.
- [35] Y. Khaluf, A. Reina, T. Bose, and J. A. R. Marshall, “Agent density in collective decisions.” https://github.com/Di0DeProject/MuMoT/blob/master/DemoNotebooks/Agent_density.ipynb, 2019. [Online; accessed 7-October-2020].
- [36] K. Elamvazhuthi and S. Berman, “Mean-field models in swarm robotics: a survey,” *Bioinspiration & Biomimetics*, vol. 15, no. 1, p. 015001, 2019.
- [37] V. Sood, T. Antal, and S. Redner, “Voter models on heterogeneous networks,” *Physical Review E*, vol. 77, no. 4, p. 041121, 2008.
- [38] Y. Shang and R. Bouffanais, “Influence of the number of topologically interacting neighbors on swarm dynamics,” *Scientific Reports*, vol. 4, no. 1, p. 4184, 2015.
- [39] I. Rausch, A. Reina, P. Simoens, and Y. Khaluf, “Coherent collective behaviour emerging from decentralised balancing of social feedback and noise,” *Swarm Intelligence*, vol. 13, no. 3–4, pp. 321–345, 2019.

- [40] C. Gershenson and D. Helbing, “When slower is faster,” *Complexity*, vol. 21, no. 2, pp. 9–15, 2015.
- [41] L. B. Slobodkin, *Growth and Regulation of Animal Populations*. New York: Holt, Rinehart and Winston, 1961.
- [42] H.-u. Stark, C. J. Tessone, and F. Schweitzer, “Decelerating microdynamics can accelerate macrodynamics in the voter model,” *Physical Review Letters*, vol. 101, p. 018701, jun 2008.
- [43] H.-U. Stark, C. J. Tessone, and F. Schweitzer, “Slower is faster: fostering consensus formation by heterogeneous inertia,” *Advances in Complex Systems*, vol. 11, no. 04, pp. 551–563, 2008.
- [44] P. Rahmani, F. Peruani, and P. Romanczuk, “Flocking in complex environments—attention trade-offs in collective information processing,” *PLoS Computational Biology*, vol. 16, no. 4, pp. 1–18, 2020.
- [45] M. M. G. Sosna, C. R. Twomey, J. Bak-Coleman, W. Poel, B. C. Daniels, P. Romanczuk, and I. D. Couzin, “Individual and collective encoding of risk in animal groups,” *Proceedings of the National Academy of Sciences*, vol. 116, no. 41, pp. 20556–20561, 2019.
- [46] C. A. C. Parker and H. Zhang, “Biologically inspired decision making for collective robotic systems,” in *2004 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS) (IEEE Cat. No.04CH37566)*, vol. 1, pp. 375–380, IEEE, 2004.
- [47] C. A. C. Parker and H. Zhang, “Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem,” *IEEE/ASME Transactions on Mechatronics*, vol. 14, no. 2, pp. 240–251, 2009.
- [48] N. R. Franks, S. C. Pratt, E. B. Mallon, N. F. Britton, and D. J. T. Sumpter, “Information flow, opinion polling and collective intelligence in house-hunting social insects,” *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences*, vol. 357, no. 1427, pp. 1567–1583, 2002.
- [49] R. Jeanson, A. Dussutour, and V. Fourcassié, “Key factors for the emergence of collective decision in invertebrates,” *Frontiers in Neuroscience*, vol. 6, p. 121, 2012.

- [50] G. Valentini, H. Hamann, and M. Dorigo, “Self-organized collective decision making: The weighted voter model,” in *Proceedings of the 2014 International Conference on Autonomous Agents and Multi-Agent Systems, AAMAS '14*, pp. 45–52, International Foundation for Autonomous Agents and Multiagent Systems, 2014.
- [51] N. Stroeymeyt, E. J. H. Robinson, P. M. Hogan, J. A. R. Marshall, M. Giurfa, and N. R. Franks, “Experience-dependent flexibility in collective decision making by house-hunting ants,” *Behavioral Ecology*, vol. 22, no. 3, pp. 535–542, 2011.
- [52] W. Liu, A. F. T. Winfield, J. Sa, J. Chen, and L. Dou, “Towards energy optimization: Emergent task allocation in a swarm of foraging robots,” *Adaptive Behavior*, vol. 15, no. 3, pp. 289–305, 2007.
- [53] E. J. H. Robinson, N. R. Franks, S. Ellis, S. Okuda, and J. A. R. Marshall, “A simple threshold rule is sufficient to explain sophisticated collective decision-making,” *PLoS ONE*, vol. 6, no. 5, p. e19981, 2011.
- [54] E. J. H. Robinson, O. Feinerman, and N. R. Franks, “How collective comparisons emerge without individual comparisons of the options,” *Proceedings of the Royal Society B: Biological Sciences*, vol. 281, no. 1787, p. 20140737, 2014.
- [55] V. Trianni, D. De Simone, A. Reina, and A. Baronchelli, “Emergence of consensus in a multi-robot network: from abstract models to empirical validation,” *IEEE Robotics and Automation Letters*, vol. 1, no. 1, pp. 348–353, 2016.
- [56] K. Lerman, A. Martinoli, and A. Galstyan, “A review of probabilistic macroscopic models for swarm robotic systems,” in *Lecture Notes in Computer Science*, vol. 3342, pp. 143–152, 2005.
- [57] M. A. Hsieh, Á. Halász, S. Berman, and V. Kumar, “Biologically inspired redistribution of a swarm of robots among multiple sites,” *Swarm Intelligence*, vol. 2, no. 2-4, pp. 121–141, 2008.
- [58] A. Reina, R. Miletitch, M. Dorigo, and V. Trianni, “A quantitative micro-macro link for collective decisions: The shortest path discovery/selection example,” *Swarm Intelligence*, vol. 9, no. 2–3, pp. 75–102, 2015.

- [59] L. A. Real, “The kinetics of functional response,” *The American Naturalist*, vol. 111, no. 978, pp. 289–300, 1977.
- [60] A. Cornish-Bowden, “One hundred years of Michaelis–Menten kinetics,” *Perspectives in Science*, vol. 4, pp. 3–9, mar 2015.
- [61] R. Gesztelyi, J. Zsuga, A. Kemeny-Beke, B. Varga, B. Juhasz, and A. Tosaki, “The Hill equation and the origin of quantitative pharmacology,” *Archive for History of Exact Sciences*, vol. 66, pp. 427–438, jul 2012.
- [62] M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal, “Kilobot: A low cost robot with scalable operations designed for collective behaviors,” *Robotics and Autonomous Systems*, vol. 62, no. 7, pp. 966–975, 2014.
- [63] A. Pirrone, H. Azab, B. Y. Hayden, T. Stafford, and J. A. R. Marshall, “Evidence for the speed–value trade-off: Human and monkey decision making is magnitude sensitive,” *Decision*, vol. 5, no. 2, pp. 129–142, 2018.
- [64] A. Dussutour, Q. Ma, and D. J. T. Sumpter, “Phenotypic variability predicts decision accuracy in unicellular organisms,” *Proceedings of the Royal Society B: Biological Sciences*, vol. 286, no. 1896, p. 20182825, 2019.
- [65] M. S. Talamali, T. Bose, M. Haire, X. Xu, J. A. R. Marshall, and A. Reina, “Sophisticated collective foraging with minimalist agents: A swarm robotics test,” *Swarm Intelligence*, vol. 14, no. 1, pp. 25–56, 2020.
- [66] A. Pirrone, T. Stafford, and J. A. R. Marshall, “When natural selection should optimize speed-accuracy trade-offs,” *Frontiers in Neuroscience*, vol. 8, no. 73, pp. 1–5, 2014.
- [67] S. C. Nicolis, N. Zabzina, T. Latty, and D. J. T. Sumpter, “Collective irrationality and positive feedback,” *PLoS ONE*, vol. 6, no. 4, p. e18901, 2011.
- [68] A. E. Turgut, C. Huepe, H. Çelikkanat, F. Gökçe, and E. Şahin, “Modeling phase transition in self-organized mobile robot flocks,” in *Ant Colony Optimization and Swarm Intelligence* (M. Dorigo, M. Birattari, C. Blum, M. Clerc, T. Stützle, and A. F. T. Winfield, eds.), (Berlin, Heidelberg), pp. 108–119, Springer Berlin Heidelberg, 2008.

- [69] I. Rausch, P. Simoens, and Y. Khaluf, “Adaptive foraging in dynamic environments using scale-free interaction networks,” *Frontiers in Robotics and AI*, vol. 7, no. July, pp. 1–14, 2020.
- [70] T. Latty and M. Beekman, “Keeping track of changes: the performance of ant colonies in dynamic environments,” *Animal Behaviour*, vol. 85, no. 3, pp. 637–643, 2013.
- [71] T. J. Czaczkes, B. Czaczkes, C. Iglhaut, and J. Heinze, “Composite collective decision-making,” *Proceedings of the Royal Society B: Biological Sciences*, vol. 282, no. 1809, p. 20142723, 2015.
- [72] F. Arvin, A. E. Turgut, F. Bazyari, K. B. Arikan, N. Bellotto, and S. Yue, “Cue-based aggregation with a mobile robot swarm: a novel fuzzy-based method,” *Adaptive Behavior*, vol. 22, no. 3, pp. 189–206, 2014.
- [73] I. D. Couzin, J. Krause, N. R. Franks, and S. A. Levin, “Effective leadership and decision-making in animal groups on the move,” *Nature*, vol. 433, no. 7025, pp. 513–516, 2005.
- [74] D. Centola, J. Becker, D. Brackbill, and A. Baronchelli, “Experimental evidence for tipping points in social convention,” *Science*, vol. 360, no. 6393, pp. 1116–1119, 2018.
- [75] E. Ferrante, A. E. Turgut, C. Huepe, A. Stranieri, C. Pinciroli, and M. Dorigo, “Self-organized flocking with a mobile robot swarm: a novel motion control method,” *Adaptive Behavior*, vol. 20, pp. 460–477, dec 2012.
- [76] J. Fernández-Gracia, K. Suchecki, J. J. Ramasco, M. San Miguel, and V. M. Eguíluz, “Is the voter model a model for voters?,” *Physical Review Letters*, vol. 112, no. 15, p. 158701, 2014.
- [77] J. Jhawar, R. G. Morris, U. R. Amith-Kumar, M. Danny Raj, T. Rogers, H. Rajendran, and V. Guttal, “Noise-induced schooling of fish,” *Nature Physics*, vol. 16, no. 4, pp. 488–493, 2020.
- [78] T. Zillio, I. Volkov, J. R. Banavar, S. P. Hubbell, and A. Maritan, “Spatial scaling in model plant communities,” *Physical Review Letters*, vol. 95, no. 9, p. 098101, 2005.

- [79] A. Dornhaus, N. R. Franks, R. M. Hawkins, and H. N. S. Shere, “Ants move to improve: colonies of *Leptothorax albigenis* emigrate whenever they find a superior nest site,” *Animal Behaviour*, vol. 67, no. 5, pp. 959–963, 2004.
- [80] T. Sasaki and S. C. Pratt, “Emergence of group rationality from irrational individuals,” *Behavioral Ecology*, vol. 22, no. 2, pp. 276–281, 2011.
- [81] A. L. Cronin, “Fussy groups thwart the collective burden of choice: a theoretical study of house-hunting ants,” *Journal of Theoretical Biology*, p. 110000, 2019.
- [82] U. Madhushani and N. E. Leonard, “Distributed learning: Sequential decision making in resource-constrained environments,” in *ICLR — 2021 - The Ninth International Conference on Learning Representations*, pp. 1–8, 2021.
- [83] R. K. Ramachandran, Z. Kakish, and S. Berman, “Information correlated lévy walk exploration and distributed mapping using a swarm of robots,” *IEEE Transactions on Robotics*, pp. 1–20, 2020.
- [84] A. Shirsat, K. Elamvazhuthi, and S. Berman, “Multi-robot target search using probabilistic consensus on discrete markov chains,” sep 2020.
- [85] J. Nauta, Y. Khaluf, and P. Simoens, “Hybrid foraging in patchy environments using spatial memory,” *Journal of The Royal Society Interface*, vol. 17, no. 166, p. 20200026, 2020.
- [86] B. Granovskiy, T. Latty, M. Duncan, D. J. T. Sumpter, and M. Beekman, “How dancing honey bees keep track of changes: the role of inspector bees,” *Behavioral Ecology*, vol. 23, no. 3, pp. 588–596, 2012.
- [87] N. Stroeymeyt, N. R. Franks, and M. Giurfa, “Knowledgeable individuals lead collective decisions in ants,” *Journal of Experimental Biology*, vol. 214, no. 18, pp. 3046–3054, 2011.
- [88] N. Pinter-Wollman, R. Wollman, A. Guetz, S. Holmes, and D. M. Gordon, “The effect of individual variation on the structure and function of interaction networks in harvester ants,” *Journal of The Royal Society Interface*, vol. 8, no. 64, pp. 1562–1573, 2011.

- [89] C. A. C. Parker and H. Zhang, “Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem,” *Swarm Intelligence*, vol. 4, pp. 199–220, sep 2010.
- [90] S. C. Pratt, E. B. Mallon, D. J. T. Sumpter, and N. R. Franks, “Quorum sensing, recruitment, and collective decision-making during colony emigration by the ant *leptothorax albipennis*,” *Behavioral Ecology and Sociobiology*, vol. 52, pp. 117–127, jul 2002.
- [91] J. R. Cody, K. A. Roundtree, and J. A. Adams, “Human-collective collaborative site selection,” 2020.
- [92] M. Starnini, M. Frasca, and A. Baronchelli, “Emergence of metapopulations and echo chambers in mobile agents,” *Scientific Reports*, vol. 6, p. 31834, oct 2016.
- [93] W. F. Vining, F. Esponda, M. E. Moses, and S. Forrest, “How does mobility help distributed systems compute?,” *Philosophical Transactions of the Royal Society B: Biological Sciences*, vol. 374, p. 20180375, jun 2019.
- [94] C. Bettstetter, H. Hartenstein, and X. Pérez-Costa, “Stochastic properties of the random waypoint mobility model,” *Wireless Networks*, vol. 10, pp. 555–567, sep 2004.
- [95] M. Rubenstein, A. Cornejo, and R. Nagpal, “Programmable self-assembly in a thousand-robot swarm,” *Science*, vol. 345, pp. 795–799, aug 2014.
- [96] I. Slavkov, D. Carrillo-Zapata, N. Carranza, X. Diego, F. Jansson, J. Kaandorp, S. Hauert, and J. Sharpe, “Morphogenesis in robot swarms,” *Science Robotics*, vol. 3, p. eaau9178, dec 2018.
- [97] J. T. Ebert, M. Gauci, and R. Nagpal, “Multi-feature collective decision making in robot swarms robotics track,” in *Proceedings of the 17th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS 2018)*, pp. 1711–1719.
- [98] Á. Gutiérrez, A. Campo, F. C. Santos, F. Monasterio-Huelin, and M. Dorigo, “Social odometry: Imitation based odometry in collective robotics,” *International Journal of Advanced Robotic Systems*, vol. 6, no. 2, p. 11, 2009.
- [99] F. Ducatelle, G. A. Di Caro, C. Pinciroli, F. Mondada, and L. M. Gambardella, “Communication assisted navigation in robotic swarms: Self-organization and cooperation,”

in *2011 IEEE/RSJ International Conference on Intelligent Robots and Systems*, (Los Alamitos, CA), pp. 4981–4988, IEEE, 2011.

- [100] C. Pinciroli, M. S. Talamali, A. Reina, J. A. R. Marshall, and V. Trianni, “Simulating Kilobots within ARGoS: models and experimental validation,” in *Swarm Intelligence (ANTS 2018)* (M. Dorigo et al., ed.), vol. 11172 of *LNCS*, pp. 176–187, Cham: Springer, 2018.
- [101] M. S. Talamali, A. Saha, J. A. R. Marshall, and A. Reina, “Supplementary code for robot, simulation, and analysis.” <https://github.com/Di0DeProject/AdaptationStudy>, 2020.
- [102] B. Ermentrout, “XPPAUT,” *Scholarpedia*, vol. 2, no. 1, p. 1399, 2007. revision #136177.
- [103] A. Meurer, C. P. Smith, M. Paprocki, O. Čertík, S. B. Kirpichev, M. Rocklin, A. Kumar, S. Ivanov, J. K. Moore, S. Singh, T. Rathnayake, S. Vig, B. E. Granger, R. P. Muller, F. Bonazzi, H. Gupta, S. Vats, F. Johansson, F. Pedregosa, M. J. Curry, A. R. Terrel, v. Roučka, A. Saboo, I. Fernando, S. Kulal, R. Cimrman, and A. Scopatz, “SymPy: symbolic computing in python,” *PeerJ Computer Science*, vol. 3, p. e103, Jan. 2017.
- [104] P. L. Krapivsky and S. Redner, “Dynamics of majority rule in two-state interacting spin systems,” *Physical Review Letters*, vol. 90, no. 23, p. 238701, 2003.

Acknowledgments

The authors thank Michael Port whose technical support has been vital for the success of this project.

Funding

This project was funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement number 647704) and by the Office of Naval Research Global (ONRG) under grant number 12547352. A.R. also acknowledges funding by the F.R.S.-FNRS on the *HECOROS* project.

Author contributions

M.S.T. and A.R. conceived the original idea. M.S.T. and A.R. implemented the multiagent simulator, the robot control code, and conducted the robot experiments. M.S.T. conducted the simulation experiments. A.S. performed the dynamical system modelling and analysis. A.R. directed the project. All authors interpreted results and wrote the manuscript.

Competing interests

The authors declare that they have no competing interests.

Data and materials availability

All the code to generate the data reported and discussed is available online on GitHub at [101].

The twelve videos of the robot experiments are available as Movies 2 and 3.

Supplementary Material

ST1. Comparison with other behaviours

In the literature, the problem of collective decision making to reach an agreement in robot swarms has been extensively investigated (*e.g.* see review in [10]). Our study advances the state of the art by proposing and analysing behaviours of minimalistic robots in a time-varying environment. The two main aspects that allow us to make a clear distinction between the investigated behaviours from previous works are (i) the individual robot’s requirement and (ii) the possibility to operate in unknown time-varying environments. Here, we expand the discussion of the main text to detail the unique aspects of our study, and compare our behaviours with the state of the art.

ST1.i Comparison in terms of individual robot’s requirement

When comparing robots’ algorithms, it is also important to consider the requirements and assumptions on which they are built. There are works that obtain high swarm performance by relying on skilled individual robots that apply sophisticated algorithms to filter and aggregate environmental information [14, 15, 16, 17, 18]. However, our focus is on algorithms for extremely simple machines that are highly limited in memory, communication, and computation. Algorithms that can work with such constraints are based on probabilistic rules that give sub-optimal performance measured in statistical terms. The few works [21, 20] that presented behaviours with individual-level requirements that match our constraints are the basis of the behaviour that we presented. Those behaviours in their original form are not able (as indicated in the literature [22]) to adapt to changes in the environment. Therefore, we modified the behaviours by adding two alternative exploration rules—**compare** and **resample**—that make the swarm capable of adaptation. We provide extensive analysis, numerous results, and detailed discussion on the comparison of the performance of two exploration rules in different conditions.

Collective behaviours that require additional individual capabilities, yet are relatively simplistic, are behaviours based on the local majority rule [21]. Previous work has shown that through this rule the swarm can improve its performance in terms of consensus speed. In order to test the generality of our results on adaptability, in Text ST4 we modify our behaviours to include the local majority rule. As discussed in Text ST4, we obtain qualitative similar results that confirm the general scope of our findings.

ST1.ii Comparison in terms of prior knowledge assumptions

Previous algorithms, that can adapt to environmental changes using low individual requirements [23, 24, 22, 25], cannot cope with an unknown number of options, that appear and disappear, as in our study. In fact, in the literature, the only type of environmental change that has been investigated is the change in option qualities. The state of the art algorithms are limited to operate for binary decisions and require prior knowledge of the environment. The robots are programmed with information about the options [23] or their location [24, 22, 25], and are only tasked with agreeing on the option with the highest quality. Therefore, these algorithms cannot work in an environment with a time-varying number of options that can unexpectedly appear or disappear. An integral component of our algorithm is the environmental exploration and individual discovery of unforeseen changes. As previous algorithms have not been designed to operate in the type of scenarios that we investigate in this paper, we focus our comparison on the algorithms proposed in this study. We conduct a performance comparison of two alternative exploration rules—**compare** and **resample**—, as well as we include the analysis of two alternative interaction patterns—direct switching and cross-inhibition—to show that our results generalise to a larger class of behaviours.

ST2. Individual robots' algorithm

Each robot moves in the environment and updates its commitment state by executing the Algorithm SA1 (which is also reported in Fig. 2(C) in the main text). Both committed and uncommitted robots move randomly to explore the environment. The mobility model for random exploration that we implemented is the random waypoint mobility model [94]. Through this mobility model, the robot selects a random position in the environment, which it sets as its target destination L_t . Once the robot reaches L_t , it selects a new random target destination. The robot avoid collision with the boundary walls by selecting random destinations that are at least two robot-body lengths (approximately 7 cm) far from the walls. The robots do not have any proximity sensor, therefore they do not implement any obstacle avoidance manoeuvre to prevent collisions with each other. To avoid robots remaining stuck for a long time in traffic jams caused by groups of robots moving in opposite directions, or robots not moving due to malfunctioning motors, the robot selects a new random target destination if L_t is not reached within 2 minutes (*i.e.* about the double of the time necessary to traverse the entire environment in a collision-free situation).

Robots in the polling state do not perform random exploration but move towards the last site location L_v received from the neighbours. Every two seconds the information about L_v is updated with the latest vote message, and the target destination is set with the new value: $L_t \leftarrow L_v$. In this way, the polling robot follows a biased random walk that ends when the site is reached. At the site \mathcal{T}_v , the robot estimates the site's quality q_v and becomes committed to site \mathcal{T}_v , as both site's location L_v and quality q_v are stored in the robot's memory L_m and q_m , respectively. Committed robots resume random exploration.

Due to the limited capabilities of the Kilobot, some of the sensors are virtualised through the system of Augmented Reality for Kilobots (ARK [9]). The virtual sensor readings are updated every 5 ~ 6 seconds. See further details in the Texts ST10, ST11, and ST12. All simulation and robot code is available with the paper [101].

SA1 Algorithm: Individual robot's algorithm, repeated every control cycle (~ 30 Hz)

Opinion variables: site \mathcal{T}_m 's location L_m ; site \mathcal{T}_m 's quality q_m

Navigation variables: target destination L_t

Sensors' parameters: sensing range r_e ; communication range r_s

resample rule parameters: resampling probability α

compare rule parameters: minimum difference ϵ

Sensors' readings

| $\hat{q}_e \leftarrow$ site \mathcal{T}_e 's quality estimate $\sim \mathcal{N}(q_e, \sigma_q) \in [0, 1]$ every time \mathcal{T}_e is within range r_e

| $L_v \leftarrow$ incoming vote message from neighbour in range r_s

| $(L_r, \theta_r) \leftarrow$ current robot's location L_r and orientation θ_r , updated every 5 \sim 6 s

PROCESS ENVIRONMENTAL INFORMATION

if site \mathcal{T}_e 's location L_e is within robot's sensing range r_e **then**

| **if** $L_e \neq L_m$ **then** # encountered new site

| | **With probability:** α # condition tested only for **resample** rule

| | **if** $\hat{q}_e > q_m + \epsilon$ **then** # condition tested only for **compare** rule

| | | **With probability** \hat{q}_e

| | | | $L_m \leftarrow L_e$; $q_m \leftarrow \hat{q}_e$ # Store environmental information about site \mathcal{T}_e

| | | **end**

| | **end**

| **else**

| | $q_m \leftarrow \hat{q}_e$ # resampling site \mathcal{T}_m 's quality

| **end**

end

PROCESS SOCIAL INFORMATION

Every 2 s

| **if** $L_v \neq \emptyset$ **and** $L_v \neq L_m$ **then**

| | $L_m \leftarrow L_v$ # Store social information about site \mathcal{T}_v

| | $q_m \leftarrow 0$; $L_t \leftarrow L_v$ # Go to resample quality q_v

| **end**

end

VOTING

Every 500 ms

| **With probability** q_m

| | broadcast a vote for \mathcal{T}_m to neighbour robots within communication range r_s

| **end**

end

NAVIGATION

use (L_r, θ_r) to move towards target destination L_t

if $L_r == L_t$ **then** # Target position reached

| $L_t \leftarrow$ new random destination # Random waypoint mobility model

end

ST3. Effects of the exploration rule parameters on the collective behaviour

We investigate the sensitivity of exploration rules **compare** and **resample** to the parameters ϵ and α , respectively. As desired, ϵ sets the minimum difference that the swarm should consider for adapting to a new site. Figure S1(A) shows the swarm response for three values of $\epsilon \in \{0.05, 0.15, 0.25\}$.

A similar effect can be obtained through modulation of the parameter α in behaviours that use the **resample** exploration rule (Fig. S1(B)). However, a miss-parameterisation of the behaviour may lead to undesired dynamics. Excessively high values of α can leave the system

below quorum and unable to attain a sustained agreement in favour of any site, when the sites have similar qualities. See for instance $\alpha = 0.4$ in Fig. S1(B). Instead, for values of α that are too low, robots resample option in extremely rare occasions, therefore adaptation becomes very slow. For instance, in Fig. S1(B), behaviours that used $\alpha = 10^{-4}$ rarely adapted within $T_{\max} = 6 \cdot 10^4$ time steps.




Figure S1: Through multiagent simulations (100 runs per condition) we study the sensitivity to parameters ϵ (in (A)) and α (in (B)) for the exploration rule `compare` and `resample`, respectively. We report the proportion of multiagent simulations (with $S = 100$ noiseless agents, $\sigma_q = 0$, initialised with commitment to site \mathcal{T}_y with $q_y = 0.5$) that adapted to site \mathcal{T}_x with quality $q_x = q_y + \delta$. Difference δ is varied on the x-axis.

ST4. Generalisation of the results to other mechanisms of sampling neighbours' opinions

Here, we test if the reported results can generalise to other simple decentralised mechanisms to sample neighbours' opinions that are different from the voter model, which we investigated in the main text. Through the voter model [28], the robot selects one random message among the messages received from its neighbours (in our implementation, it selects the last message received). We replace the voter model mechanism with the local majority rule [104]. Through this rule, the robot combines all the votes from its neighbours, including its own vote, and change opinion using the vote that is more frequent (in case of a tie, it makes a random choice). The local majority rule requires more capabilities at the individual level than the voter model, as the robot needs to store the received messages, count the received votes for each site, and select the highest count. While more demanding, it has been shown that the local majority




Figure S2: We replaced the voter model with the local majority rule and we obtain qualitatively similar results to Fig. 3. We report results the adaptation probability (*i.e.* the proportion of 100 runs that adapted within $T_{\max} = 6 \cdot 10^4$ time steps) and adaptation time. We use the same experimental setup of Fig. 3. For both exploration rules (**compare** and **resample**), the swarm has an improved performance in terms of collective adaptation for a relatively low number of communication links $\langle k \rangle$. The communication links decrease by decreasing the robots' communication range r_s or the swarm size S (as shown in the inset).

rule can also give better performances than the voter model in terms of consensus speed [21], as it will be also confirmed by our results. The cause of such a performance improvement is the processing of more social information (by applying the local majority rule).

In our study, we replace test the local majority rule to assess the generality of the main results of our paper, that is, 'less is more'. Fig. S2 shows that also in this case a lower number of communication links (*less*) corresponds to an increase (*more*) of the collective performance in terms of adaptation ability. The investigated scenario is the same of Fig. 3: the environment has three sites \mathcal{T}_x , \mathcal{T}_y , and \mathcal{T}_z with qualities $\{q_x, q_y, q_z\} = \{0.8, 0.7, 0.1\}$ and the swarm starts committed to the site \mathcal{T}_y (see simulation details in Fig. 3 and in Text ST13). Both the probability that the swarm will adapt to \mathcal{T}_y and the adaptation speed increased for relatively low communication ranges r_s , or in scenarios with low robot density (small swarm sizes S).

Note that, compared with the voter model (Fig. 3), the local majority rule relies on a higher amount of social information and thus has quicker dynamics in terms of speed.

ST5. Information spreading trade-off

Our study shows the benefits of limited communication in terms of better adaptability to environmental changes. While a small number of communication links can bring the advantages of better adaptability as illustrated in the main paper, it is well documented and common knowledge that information can spread quicker in highly connected networks. Therefore, we conduct a set of tests to measure information spreading in our experimental scenario.

We ran multiagent simulations in the same scenario we investigated in Fig. 3 for values of communication range $r_e \in \{0, 0.05, \dots, 0.5\}$ and level of noise $\sigma_q = 0.1$. In this scenario, the robot swarm is fully committed to site \mathcal{T}_y with quality $q_y = 0.7$ and at time $t = 0$ a new better site \mathcal{T}_x with $q_x = 0.8$ appears. Fig. 3 shows the adaptation probability and adaptation time. These two metrics measure with what probability and how quickly the swarm adapted. We consider a swarm has adapted when the average size of the subpopulation committed to \mathcal{T}_x in the last 5000 temporal steps is above the quorum of 80%. The adaptation probability is the proportion of runs that adapted over the total number of simulations, while the adaptation time is the time from $t = 0$ in order reach adaptation. Fig. S3 shows a different metric from Fig. 3. Fig. S3 shows the probability that a new piece of information spreads throughout the swarm, as well as the spreading speed. The piece of information that spreads is the individual discovery of the new site \mathcal{T}_x . This information is shared among robots and can either spread throughout the majority of the population or be suppressed by the subpopulation committed to \mathcal{T}_y . Every time the subpopulation S_x committed to \mathcal{T}_x vanishes (*i.e.* it has size zero), we consider it as an information spreading failure. Instead, when S_x reaches the quorum of 80% of the total swarm S , we consider it as an information spreading success. The information spreading probability reported in Fig. S3 is computed as the proportion of successes over the number of attempts (successes+failures). The information spreading time is computed as the average time from the individual discovery to the reaching of the quorum.

Fig. S3 shows that there exists a trade-off between information spreading speed and information spreading probability. On the one hand, large communication ranges form highly connected networks that can quickly spread information but only rarely they are successful. On the other hand, information spreading in sparse networks formed by robots with low communication ranges is slow but information eventually spreads throughout the population with a low failure rate.




Figure S3: Through multiagent simulations (10^3 runs per condition) we study how quickly and how effectively information spread throughout the population for different communication range values. $S = 50$ robots start committed to the inferior site \mathcal{T}_y (with $q_y = 0.7$) and new better site \mathcal{T}_x with $q_x = 0.8$ appears. The individual discovery of the new site has higher probabilities to spread when robots communicate in a local range but it happens slower. Instead, when the communication range is large, information can spread quickly but most of the times the spreading is unsuccessful. See discussion in Text ST5.

ST6. The impact of noise in the individual estimations

Physical sensor-based estimation noise is unavoidable. When a robot makes its individual estimate \hat{q}_i of a site \mathcal{T}_i , we model \hat{q}_i as a stochastic variable with a normal probability distribution centred on the true quality value q_i and with standard deviation σ_q . That is, $\hat{q}_i \sim \mathcal{N}(q_i, \sigma_q)$. We test our behaviours for different levels of noise $\sigma_q \in \{0, 0.05, \dots, 0.25\}$. Fig. S3 shows important results that are useful to understand the benefit of the proposed strategies. The results indicate, in fact, that when the noise is absent and all robots make correct quality estimates, then the robots can effectively, and efficiently, solve the best-of-n problem individually and communication is not necessary. Without the use of communication (for $r_s = 0$), the swarm can always adapt (100% adaptation) with the quickest adaptation time. However, when robots make noisy individual estimates, the non-communicating robots have a drastic decrease in the collective performance and the swarm is unable to reach a consensus in favour of the best site. Instead, through local communication, the swarm can make accurate decisions regardless of the noise level. Fig. S3 shows that, through intra-robot communication, the swarm collectively filters the noise, and accurate collective decisions are made. Note that for all tested values of

noise, large communication ranges impede adaptation; a result in agreement with Fig. 3 in the main text and, in general, with the main message of the paper: less is more. Noise in physical sensors is unavoidable, especially in simple machines for which we designed our behaviours. Our collective behaviours, based on low-range communication, are effective solutions to achieve swarm adaptation.




Figure S4: In the same scenario of Fig. 3, through multiagent simulations (100 runs per condition) we study adaptation probability and average speed for different levels of noise σ_q . When individual robots make noisy estimates of the site's quality they benefit from communicating with each other.

ST7. Scalability to large numbers of options

Differently from most works in the literature that limit their investigation to binary decision problems ($n = 2$), the behaviours that we propose and analyse in this paper can scale to large numbers of alternatives ($n \geq 2$). In the main text, we report extensive analysis of experiments with $n = 3$ alternative sites. Here, we include additional results of multiagent simulations with

up to $n = 9$ sites. The experiments have been designed as follows. A swarm of 50 uncommitted robots starts operating in an environment with two sites— \mathcal{T}_1 and \mathcal{T}_2 —with qualities $q_1 = 0.1$ and $q_2 = 0.2$, respectively. With regular frequency, every $\tau = 5 \times 10^4$ timesteps, a new site \mathcal{T}_i appears with quality q_i higher by 0.1 of the previous highest quality. We add sites to reach up to 9 sites in the environment with qualities $\{q_1, q_2, \dots, q_9\} = \{0.1, 0.2, \dots, 0.9\}$. Fig. S5 shows that both behaviours—based on the `compare` and `resample` rules—are able to monitor the environment and adapt to all the 7 changes with probability close to 100%.




Figure S5: Change over time of the subpopulations committed to the $n = 9$ sites (mean subpopulation sizes for 100 multiagent simulations with 95% confidence interval as shades). At the beginning of the experiment, the environment has two sites. Every $\tau = 5 \times 10^4$, a new site with better quality appears. On top of each of the eight phases (separated by vertical dashed lines), there are the sites' qualities represented as Gaussian distributions with standard deviation equal to sensing noise $\sigma_q = 0.1$. The numbers superimposed on the plot indicate the proportion of runs in which the subpopulation committed to the best site was above the 80% quorum (horizontal dashed line) τ timesteps after the best site appeared. Almost always the swarm adapts to the best site for both exploration rules.

ST8. The cost-performance trade-off of two exploration rules

Depending on the application scenario and its requirements, the designer can make an informed choice: if adaptation speed is a critical factor, `compare` should be preferred, otherwise if individual simplicity is a key requirement, the `resample` rule is better. Note, however, that the dynamics of a swarm that uses the `resample` rule are influenced by the resampling probability α which therefore needs values in an appropriate range to successfully operate in a given sce-

nario (see Fig. S1(B)). The `compare` rule instead is more effective and its parameter sets the minimum quality difference to trigger a response (see Fig. S1(A)). It is interesting to note the following point: while individual robots have a response to new sites that is insensitive to its site’s quality value q_y and only depends on the quality difference $\delta = q_x - q_y$, instead, at the swarm level, value-sensitivity emerges, with adaptation that depends on both the value and the difference, as shown in Fig. 7.

ST9. Detecting a stable consensus can be useful for decentralised quorum sensing

Rather than casting a decision as soon as the quorum threshold is reached, we instead decided to measure when sustained agreement is maintained over a long time period (that is, average size of committed subpopulation in the last 5000 temporal steps is above 80% quorum). In this way, we avoid miscounting random fluctuations in our measure of collective adaptability. The consensus must be stable in order not to impair the computation of the collective decision outcome in real applications. The swarm decision is in fact distributed in the form of individual commitment among the robots, and, in order to reliably ‘aggregate’ the collective decision outcome, the consensus must be stable during the aggregation process. The individual commitment can be aggregated by the robot themselves, for instance, through algorithms for decentralised quorum sensing [89, 90], or by an external user, for instance, a person coordinating search-and-rescue operations [91] (Fig. 1(A) in the main text).

ST10. The Augmented Reality for Kilobot (ARK)

The Kilobots’ capabilities are increased by the Augmented Reality for Kilobots (ARK) system [9]. A matrix of 2×2 cameras track the position, orientation, and LED colour of the 50 Kilobots in real time (7 fps). The virtual environment composed of target sites is simulated in the PC acting as the base control station of ARK. The control station updates the state of the virtual environment and the values of the virtual sensors of each Kilobot. The sensor readings are communicated to the robot every 5 seconds. The data are transferred to the robots via overhead controllers that broadcast infrared messages to the robots. ARK updates the state of all the robots in about 1 second. This transmission can interfere with the intra-robot communication which is also based on the same infrared channel. Therefore, during the update of the virtual sensors, the communication among robots is temporarily disabled. ARK

also allows us to program and set up the robots, as well as to record the video and log the timeseries of the robots' states (see Movie 2 and 3). In addition, the motors of the Kilobots have been periodically re-calibrated, between repetitions, through the ARK's automated calibration routine.

ST11. Virtualisation of the global positioning system

The robots know their location and orientation through a position virtual sensor. The sensor readings are updated every 5 ~ 6 seconds, which is enough to allow the robots to effectively navigate the environment. ARK computes the readings of each Kilobot's position virtual sensor through real-time camera tracking. The robot's location is computed as the centre of the tracked robot's body. The robot's orientation is computed comparing the robot's current location with its location 1 s earlier. The difference vector indicates the average direction of motion of the robot in the last second. To reduce tracking noise, ARK stores the direction values computed in the latest five camera frames and computes the robot's orientation as the weighted average of these five directions¹. Therefore, both location and orientation are not perfect information but are subject to noise from the camera tracking.

Additional noise in the position virtual sensor is caused by ARK's communication protocol which uses a limited number of bits per message. In all our experiments with Kilobots, the sensing resolution of the position virtual sensor is $1/16 \text{ m} = 6.25 \text{ cm}$ for location and 12° for orientation. In other words, ARK sends to the robots their locations as coordinates in a 16×16 grid and orientations as one of the 30 orientation slices, each of which is 12° wide. The virtual sensor resolution has been tuned to give a good compromise between ARK's communication load (ARK can transmit 2 bytes per robot at 60 Hz) and positioning noise. Despite the low resolution, the robots can effectively move in the environment (*e.g.* see Movie 2 and 3 with videos of the robot experiments).

ST12. Virtualisation of a global-range transreceiver

In the experiments with the global communication range (Fig. 6 in the main text), the Kilobots are equipped with a virtual transreceiver that allows each robot to exchange messages with the entire swarm. ARK virtualises the global communication among all Kilobots by sending a message every 2 s containing the number of robots committed to each site (population counts),

¹ARK's code is open source and available at <https://github.com/Di0DeProject/KilobotArena>

the current location, and quality of the sites. The counts do not include polling robots, as they do not communicate. ARK can compute these counts in real time by tracking the robots' colour and positions. Kilobots show their commitment state through the colour of their LED. Robots are labelled as being in the polling state when they change commitment state (change their LED colour) and have not yet sampled the site's quality (the site is outside the Kilobot's sensing range r_e). The Kilobot uses the population counts to simulate the reception of the message from one random robot of the entire swarm. To do so, the robot reduces by one the size of the population to which it belongs, and multiplies each population count by the respective site's quality (as communication frequency is a linear function of the site's quality). In this way, the Kilobot computes the expected number of messages that it would receive by the rest of the population if communication were to happen globally among the entire swarm. Finally, the Kilobot picks one message at random via a probability weighted by the number of messages from each subpopulation.

ST13. Parameters of the experimental setup

All experiments have been conducted in a square environment $\mathcal{E} = 1 \times 1 \text{ m}^2$ and including $n = 3$ target sites with quality values that varied among setups, as indicated per case. In the multiagent simulations, the site's location was randomly chosen in the environment with a uniform distribution. In the robot experiments, the sites were placed in random order on the vertices of a central equilateral triangle (sized 29 cm), to avoid bias of pathological cases in a limited number of repetitions. Where not differently specified, both multiagent simulations and robot experiments study collective behaviour in noisy conditions with $\sigma_q = 0.1$. The exploration rule's parameters are $\epsilon = 0.05$ and $\alpha = 0.01$ in the `compare` and `resample` rule, respectively. Simulation used parameters for motion speed $\nu = 1 \text{ cm/s}$ and sensing range $r_e = 20 \text{ cm}$ that agree with the real counterpart (the Kilobots).

We investigated the ability of the swarm to adapt to three types of change events: appearance of a new best site, disappearance of the best site, and exchange of the qualities between the best and the second best site. In the multiagent simulations, these three events have been studied in isolation with dedicated experiments. Instead, the robot experiments are long demonstrations comprising three phases, in each of which an environmental change occurs. Each demonstration has a total length of 80 minutes.

Appearance of a new best site. We study the situation in which the swarm have reached a consensus on the best site \mathcal{T}_y before a new better site \mathcal{T}_x appears (with $q_x > q_y$). We simulate this event by initialising the swarm with a full commitment to \mathcal{T}_y and running the experiment with both \mathcal{T}_y and \mathcal{T}_x present. In every experiment, there is also a third lower quality site \mathcal{T}_z , with $q_z < q_y$. The results reported in Fig. 3 used quality values $\{q_x, q_y, q_z\} = \{0.8, 0.7, 0.1\}$. Instead, in Fig. 7, the values of q_x and q_y were systematically varied as indicated on the axes (with $q_x = q + \delta$ and $q_y = q$), and q_z was kept to 0.1. The experiments with the Kilobots have studied the case of $\{q_x, q_y, q_z\} = \{0.8, 0.6, 0.4\}$, the results of which are reported in Fig. 5. The appearance of \mathcal{T}_y is studied in the first phase of the demonstration, from minute 0 to 20.

Disappearance of a new best site. We studied the situation in which the best site \mathcal{T}_x , on which the robot reached a consensus, disappears. Therefore, the environment remains with two sites \mathcal{T}_y and \mathcal{T}_z (with $q_y > q_z$). We simulate this event by initialising the swarm with a full commitment to \mathcal{T}_x and running the experiment with only \mathcal{T}_y and \mathcal{T}_z present. In both multiagent and robot experiments, we studied the situation of $\{q_x, q_y, q_z\} = \{0.8, 0.6, 0.4\}$. In the demonstration with the Kilobots, disappearance is studied in the second phase, in the time frame from minute 20 to 40 (Fig. 5). The swarm always adapted to the best remaining site \mathcal{T}_y . The reason is that the robots re-sample the quality of the site they are committed to when they are within sensing range of the site. When the site disappears, their quality estimate will be therefore updated to the value of zero. With zero quality, the robot does not vote for the site (zero frequency of communication) and additionally any other site discovered by the robot will have a better quality. Therefore, when disappearance occurs, adaptation is always achieved.

Quality exchange between the best and the second best site. We studied the situation in which the best site \mathcal{T}_y , on which the robot reached consensus, swaps its quality with the second best site \mathcal{T}_x , resulting in $q_x > q_y > q_z$. Therefore, we simulate this event by initialising the swarm with a full commitment to \mathcal{T}_y with quality estimates that have been initialised using quality q_x , and running the experiment with $\{q_x, q_y, q_z\} = \{0.8, 0.6, 0.4\}$. The results indicated the ability of the swarm to always adapt to such situations. This experimental setup is very similar to the case of appearance, as in both cases robots are initialised with commitment to \mathcal{T}_y and operate in an environment with $n = 3$ sites \mathcal{T}_y with qualities $q_x > q_y > q_z$. However, there is the following difference: in the appearance case, robots are initialised with quality estimates

$\hat{q}_y \sim \mathcal{N}(q_y, \sigma_q)$, instead, in the quality swap case, the quality estimate has an higher mean, $\hat{q}_y \sim \mathcal{N}(q_x, \sigma_q)$, as it simulates a sudden drop in their site \mathcal{T}_y 's quality, simultaneous to the sudden increase of q_x . In the demonstration with the Kilobots, the quality exchange is studied in the third phase, from minute 40 to 80 (Fig. 5).

ST14. Derivation of the mathematical models

At any point in time, let there be n target sites in the environment indexed by \mathcal{T}_i where $i \in \{1, \dots, n\}$. Let S_i be the number of robots committed to \mathcal{T}_i , S_p be the number of polling robots, and S_u be the number of uncommitted robots. As the time progresses, the dynamical variables S_i , S_p and S_u change due to exploration (which is governed by the exploration rule) and social interactions (which are governed by the social interaction patterns). If we define F_i^{ER} and G_i^{SIP} to be the rates of change of a subpopulation S_i due to exploration and social interactions respectively, then

$$\dot{S}_i = F_i^{ER} + G_i^{SIP}, \quad (\text{S1})$$

where $ER = \mathbf{COMP}$ or \mathbf{RES} for **compare** and **resample** rules respectively, and $SIP = \mathbf{DS}$ or \mathbf{CI} for **direct-switching** and **cross-inhibition** patterns respectively. The precise form of these functions also depends on the quality q_i of every target site \mathcal{T}_i , the probability $P_e = k \pi r_e^2$ of a robot encountering a target site in a unit time, and the probability $P_m = k \pi r_s^2$ of a robot being in communication range with another robot in a unit time (with k as a proportionality constant).

Regardless of the exploration rule, an uncommitted robot, upon meeting a target site \mathcal{T}_i , commits to it with probability q_i . Additionally, robot committed to target site \mathcal{T}_i exploring the environment using the **compare** rule may switch its opinion to \mathcal{T}_j upon meeting the site with a probability equal to q_j if $q_j > q_i$. Similarly, if $q_i > q_j$, a robot may switch its opinion from \mathcal{T}_j to \mathcal{T}_i with probability q_i . Therefore,

$$F_i^{\mathbf{COMP}} = P_e q_i S_u + \sum_{j \neq i} (P_e H_{i,j} q_i S_j - P_e H_{j,i} q_j S_i) \quad (\text{S2})$$

where, $H_{i,j} = 1$ if $q_i > q_j$ and 0 otherwise (we assume the **compare** parameter $\epsilon = 0$). By using the **resample** rule, a fixed fraction α of all committed robots may switch their opinion on encountering a different target site with probability equal to its quality. Hence,

$$F_i^{\mathbf{RES}} = P_e q_i S_u + \sum_{j \neq i} (P_e \alpha q_i S_j - P_e \alpha q_j S_i). \quad (\text{S3})$$

All interactions between two robots belonging to different subpopulations result in either no change in the subpopulation dynamics or one robot (recruiter) changing the commitment of the other robot (recruitee). As discussed in the Results section, the mean rate of recruitment non-linearly depends on the number of recruitees in the system due to the voter model (Fig. 4(A)). In particular, for a subpopulation of S_x recruiters and S_y recruitees, we use the Holling Type II function [59] to model the rate of recruitment as,

$$R = \frac{P_m S_y}{1 + P_m S_y} q_x S_x, \quad (\text{S4})$$

where q_x is the quality of the target site \mathcal{T}_x to which the S_x recruiters are committed. The equation demonstrates that the rate of recruitment is proportional to the S_y if $S_y \ll P_m$ and is independent of S_y if $S_y \gg P_m$. We use this functional form to model interactions between the robots hereafter.

An interaction of a committed robot with an uncommitted robot can lead to the uncommitted robot becoming committed. Moreover, if the robots interact using the direct-switching pattern, there are no polling robots, *i.e.*, $S_p = 0$. When a robot from subpopulation S_i meets a robot from subpopulation S_j , either of the robots may switch to the other subpopulation. Hence, the dynamics are given by

$$G_i^{\text{DS}} = \text{discovery} + \sum_{j \neq i} (\text{direct-switch } j \text{ to } i - \text{direct-switch } i \text{ to } j) = \frac{P_m S_u}{1 + P_m S_u} q_i S_i + \sum_{j \neq i} \left(\frac{P_m S_j}{1 + P_m S_j} q_i S_i - \frac{P_m S_i}{1 + P_m S_i} q_j S_j \right) \quad (\text{S5})$$

Conversely, if the robots interact using the cross-inhibition pattern, an interaction between committed robots may change the number of polling robots, and an interaction between a committed and a polling robot may lead to the recruitment of the polling robot. In terms of the parameters, we have

$$G_i^{\text{CI}} = \text{discovery} + \text{recruitment of } p \text{ by } i - \sum_{j \neq i} (\text{inhibition of } i \text{ by } j) = \frac{P_m S_u}{1 + P_m S_u} q_i S_i + \gamma \frac{P_m S_p}{1 + P_m S_p} q_i S_i - \sum_{j \neq i} \frac{P_m S_i}{1 + P_m S_i} q_j S_j \quad (\text{S6})$$

and

$$G_p^{\text{CI}} = \sum_i \left(\sum_{j \neq i} \text{inhibition of } i \text{ by } j - \text{recruitment of } p \text{ by } i \right) = \sum_i \left(\sum_{j \neq i} \frac{P_m S_i}{1 + P_m S_i} q_j S_j - \gamma \frac{P_m S_p}{1 + P_m S_p} q_i S_i \right), \quad (\text{S7})$$

where γ is a scaling parameter to take into account that polling robots do not change their commitment directly after receiving a recruitment message. Instead, a polling robot integrates various messages, each making the robot move closer to the target site. When the polling robot reaches the site, it finally changes its commitment state.

Substituting Eqs. (S2)-(S7) in Eq. (S1) gives us the dynamical equations for the four collective behaviours:

- **compare with direct-switching**

$$\begin{aligned}\dot{S}_i &= P_e q_i S_u + \sum_{j \neq i} (P_e H_{i,j} q_i S_j - P_e H_{j,i} q_j S_i) \\ &\quad + \frac{P_m S_u}{1 + P_m S_u} q_i S_i + \sum_{j \neq i} \left(\frac{P_m S_j}{1 + P_m S_j} q_i S_i - \frac{P_m S_i}{1 + P_m S_i} q_j S_j \right) \\ \dot{S}_p &= 0.\end{aligned}\tag{S8}$$

- **compare with cross-inhibition**

$$\begin{aligned}\dot{S}_i &= P_e q_i S_u + \sum_{j \neq i} (P_e H_{i,j} q_i S_j - P_e H_{j,i} q_j S_i) \\ &\quad + \frac{P_m S_u}{1 + P_m S_u} q_i S_i + \gamma \frac{P_m S_p}{1 + P_m S_p} q_i S_i - \sum_{j \neq i} \frac{P_m S_i}{1 + P_m S_i} q_j S_j \\ \dot{S}_p &= \sum_i \left(\sum_{j \neq i} \frac{P_m S_i}{1 + P_m S_i} q_j S_j - \gamma \frac{P_m S_p}{1 + P_m S_p} q_i S_i \right).\end{aligned}\tag{S9}$$

- **resample with direct-switching**

$$\begin{aligned}\dot{S}_i &= P_e q_i S_u + \sum_{j \neq i} (P_e \alpha q_i S_j - P_e \alpha q_j S_i) \\ &\quad + \frac{P_m S_u}{1 + P_m S_u} q_i S_i + \sum_{j \neq i} \left(\frac{P_m S_j}{1 + P_m S_j} q_i S_i - \frac{P_m S_i}{1 + P_m S_i} q_j S_j \right) \\ \dot{S}_p &= 0.\end{aligned}\tag{S10}$$

- **resample with cross-inhibition**

$$\begin{aligned}\dot{S}_i &= P_e q_i S_u + \sum_{j \neq i} (P_e \alpha q_i S_j - P_e \alpha q_j S_i) \\ &\quad + \frac{P_m S_u}{1 + P_m S_u} q_i S_i + \gamma \frac{P_m S_p}{1 + P_m S_p} q_i S_i - \sum_{j \neq i} \frac{P_m S_i}{1 + P_m S_i} q_j S_j \\ \dot{S}_p &= \sum_i \left(\sum_{j \neq i} \frac{P_m S_i}{1 + P_m S_i} q_j S_j - \gamma \frac{P_m S_p}{1 + P_m S_p} q_i S_i \right).\end{aligned}\tag{S11}$$

Note that we make the following assumptions in order to simplify the model. Firstly, the robots are considered to be point particles (zero size) distributed homogeneously in the environment and performing random walk. Secondly, the target sites are assumed to be circular regions with equal diameters which are randomly distributed across the environment. Finally, the transmission of messages between the robots is assumed to be lossless and the estimation of the site qualities by the visiting robots are assumed to be noiseless.

These equations can be further simplified when describing the adaptation process because fewer subpopulations are present. Initially, all robots are assumed be committed to the previously best target site ($y = 1$). While this is exactly true for the `compare` rule, it is only an approximation for the `resample` rule. However, this approximation is good for low rates of resampling. We report the simplified equations in Eqs. (1)-(4) in the main text. For completeness, we also report here the same simplified Eqs. (1)-(4) in a different format, where we use the probabilities P_e and P_m in place of robot's parameters, as $P_e = k \pi r_e^2$ and $P_m = k \pi r_s^2$. With this formalism the equations are more compact. Also, recall that x and y are the fraction of robots committed to the new best and the previously best target sites, with quality q_x and q_y , respectively. Also, z is the fraction of polling robots. The swarm has a finite size S , and therefore $x + y + z = 1$. The simplified equations are:

- **For compare with direct-switching**

$$\dot{x} = P_e q_x y + \frac{P_m S y}{1 + P_m S y} q_x x - \frac{P_m S x}{1 + P_m S x} q_y y. \quad (\text{S12})$$

- **For compare with cross-inhibition**

$$\begin{aligned} \dot{x} &= P_e q_x y + \gamma \frac{P_m S z}{1 + P_m S z} q_x x - \frac{P_m S x}{1 + P_m S x} q_y y \\ \dot{y} &= -P_e q_x y + \gamma \frac{P_m S z}{1 + P_m S z} q_y y - \frac{P_m S y}{1 + P_m S y} q_x x. \end{aligned} \quad (\text{S13})$$

- **For resample with direct-switching**

$$\dot{x} = P_e \alpha q_x y - P_e \alpha q_y x + \frac{P_m S y}{1 + P_m S y} q_x x - \frac{P_m S x}{1 + P_m S x} q_y y. \quad (\text{S14})$$

- **For resample with cross-inhibition**

$$\begin{aligned} \dot{x} &= P_e \alpha q_x y - P_e \alpha q_y x + \gamma \frac{P_m S z}{1 + P_m S z} q_x x - \frac{P_m S x}{1 + P_m S x} q_y y \\ \dot{y} &= P_e \alpha q_y x - P_e \alpha q_x y + \gamma \frac{P_m S z}{1 + P_m S z} q_y y - \frac{P_m S y}{1 + P_m S y} q_x x. \end{aligned} \quad (\text{S15})$$

ST14.i The possible regimes of the recruitment/inhibition rate

Equation (S4) is the interaction rate that describes the speed at which recruitment and inhibition take place. The functional form of Equation (S4) is such that, if the subpopulation being recruited s_i is very small — or $s_i \ll 1/P_m$ — the term in the summation becomes the $P_m q_j s_i s_j$. This is reasonable because a recruitment occurs when the robots meet (with rate P_m for each pair of robots) and a recruitment message is sent by the recruiting robot (with probability q_j). However, if s_i becomes very large — or $s_i \gg 1/P_m$ — the term in summation becomes $q_j s_j$ which is independent of both P_m and s_i . This can be explained as follows. For very large s_i , some robots of subpopulation s_i are guaranteed to be in range of each recruiting robot. If there are, on average, k robots in range of each recruiting robot, then a message is sent by the recruiting robot to those k robots with a probability q_j . Since, the group is assumed to be homogeneous, each of the robots receiving the recruiting message itself has k robots in its range which can, themselves send recruiting messages. Hence, to a first approximation, a robot receiving a recruitment message from a robot of subpopulation s_j receives k messages. Therefore, the probability of such a robot to get recruited is $1/k$. Hence the number of robots recruited by a single recruiting robot is $q_j \times k \times 1/k = q_j$. Therefore the number of robots recruited by the subpopulation s_j is $q_j s_j$.

ST15. Stability and bifurcation analysis of the ODEs

We perform stability and bifurcation analysis of the four behaviours described by the models of Eqs. (1)-(4), and we report the bifurcation plots in Fig. S6. While these models differ in the specifics of their dynamics, there are some common features among them. In particular, due to a bifurcation, each of these models can have either one or two stable fixed points depending on the parameters. Let us first consider the parameters where there is only one stable fixed point (say (x_0, y_0, z_0)) in the system. Since this is the only attractor of the system, all trajectories converge to this fixed point. In particular, if x_0 is above the quorum threshold, a system starting at $(x, y, z) = (0, 1, 0)$ would adapt. In Fig. S6, this adaptation region is coloured with a blue background. If x_0 is below the quorum threshold (horizontal dotted line at 80% in Fig. S6), the long-term behaviour of the group is that of indecision (region with a green background). In case the parameters are such that there are two stable fixed points in the system, they are separated by an unstable (saddle) fixed point. Therefore, in addition to the stable fixed point (x_0, y_0, z_0) , there are other two fixed points that we can label as (x_1, y_1, z_1) and (x_2, y_2, z_2) , such

that $x_0 > x_1 > x_2$. Stability analysis and numerical continuation show that (x_1, y_1, z_1) and (x_2, y_2, z_2) are the stable and unstable fixed points, respectively. The bifurcation plots of Fig. S6 show that $x_2 < x_1 < 0.5$ for all tested parameters, and that the initial condition for adaption $(x, y, z) = (0, 1, 0)$ always lies in the basin of attraction of (x_2, y_2, z_2) . This implies that for parameters where two stable fixed points are present (region with an orange background in Fig. S6), our model predicts that the swarm is unable to adapt and does not alter its consensus. In reality, however, the stochastic nature of the dynamics may push the system from one basin to the other. This would lead to some swarms adapting to the changes in the environment even when two stable fixed points are present in the system. However, probability to adapt is expected to decrease for increasing the swarm size (due to smaller fluctuations) and for parameterisations that are distant from the bifurcation point.

The bifurcation analysis for three out of our four models have been conducted using numerical continuation that we perform with the XPPAUT 8.0 software [102]. Instead, for the fourth behaviour, based on the `compare` rule with direct switching pattern, we compute the analytical condition for bifurcation of the swarm-level model using the SymPy Package [103] in Python. We include with the paper a Jupyter notebook to reproduce the analytical results [101] which we also report here. Recall, that the system is governed by the following differential equation,

$$\dot{x} = -\frac{SP_m vx(-x+1)}{SP_mx+1} + \frac{SP_mx(\delta+v)(-x+1)}{SP_m(-x+1)+1} + P_e v(-x+1), \quad (\text{S16})$$

where, $v = q_y$ and $\delta = q_x - q_y$. Note that, $(1-x)$ is a common factor for all terms of the equation. Therefore, the possible fixed points of the system are,

$$x_0 = 1 \quad (\text{S17})$$

$$x_1 = \frac{-SP_e P_m v + SP_m v - \delta + D}{2SP_m(-P_e v + \delta + 2v)} \quad (\text{S18})$$

$$x_2 = \frac{-SP_e P_m v + SP_m v - \delta - D}{2SP_m(-P_e v + \delta + 2v)} \quad (\text{S19})$$

where,

$$D^2 = -(-SP_e P_m v - P_e v)(4S^2 P_e P_m^2 v - 4S^2 P_m^2 \delta - 8S^2 P_m^2 v) + (-S^2 P_e P_m^2 v + S^2 P_m^2 v - SP_m \delta)^2 \quad (\text{S20})$$

Depending on the sign of D^2 , the system might have one or three real fixed points. In other words, $D^2 = 0$ forms the bifurcation condition for the model. In order to obtain this bifurcation condition in terms of the experimental parameters, we substitute $P_e = k\pi r_e^2$, $P_m = k\pi r_s^2$,

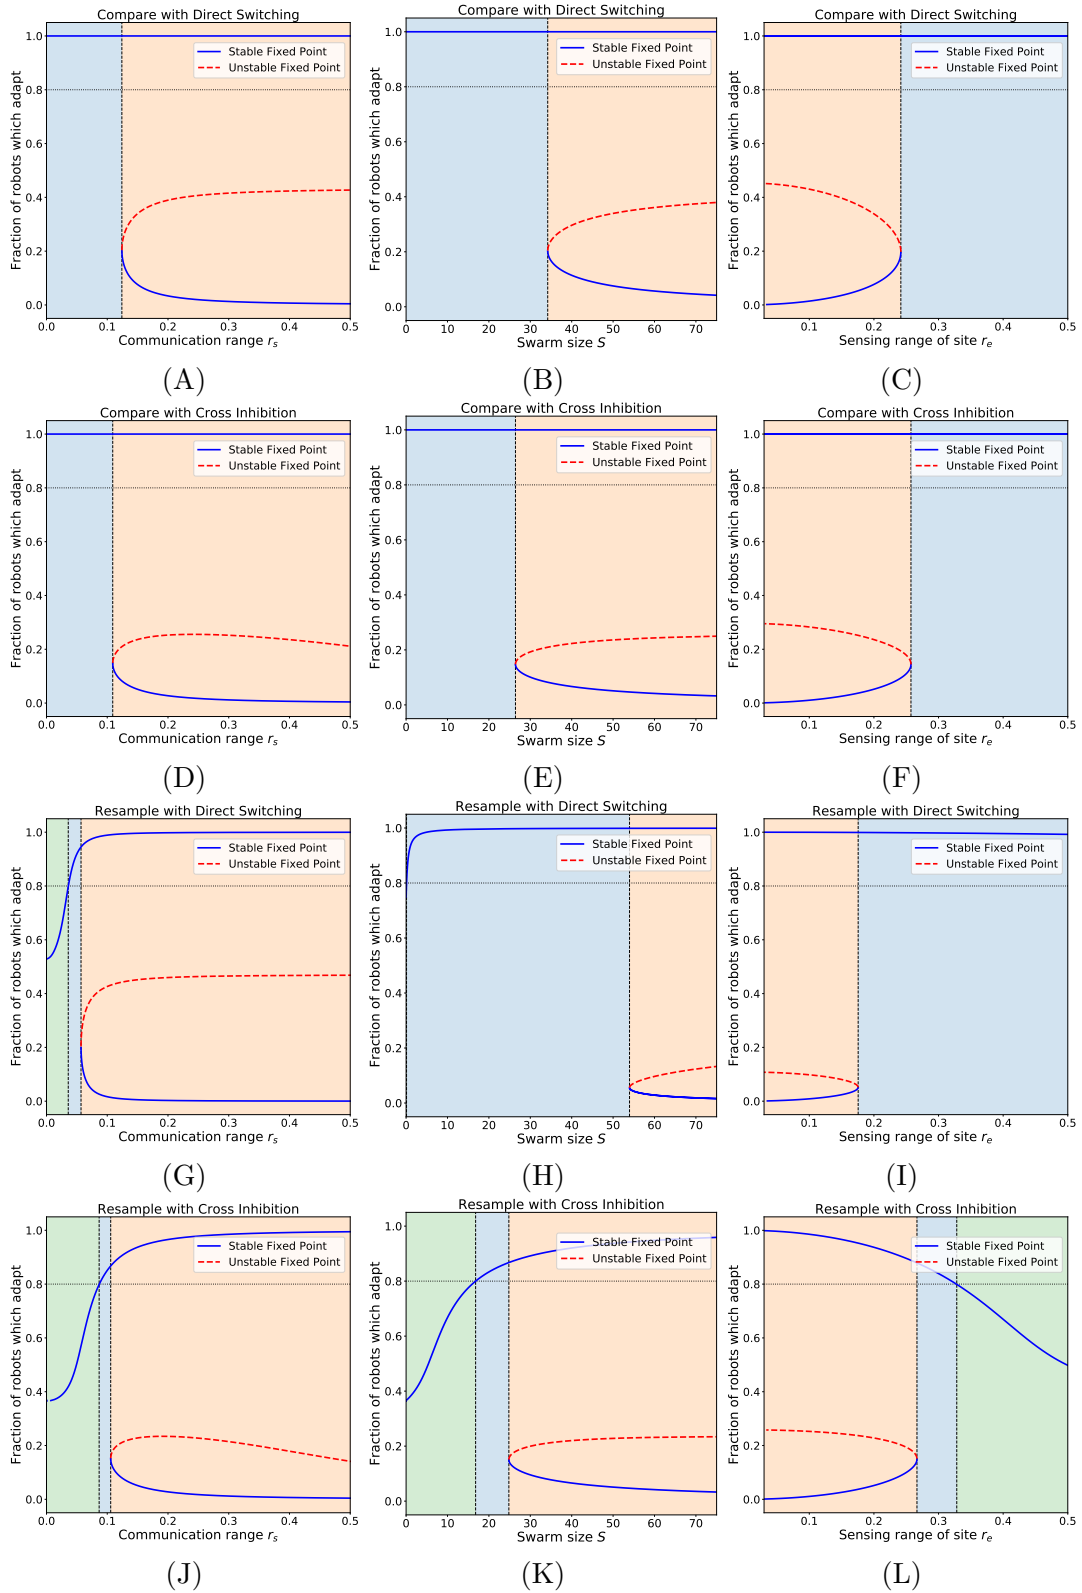


Figure S6: Bifurcation plots for the four behaviours resulting from the combination of the two exploration rules—**compare** and **resample**—and the two social interaction patterns—direct switching and cross-inhibition. We plot the fixed points of the system with respect to the communication range r_s (first column), the swarm size S (second column), and the sensing range of target sites r_e (third column). There are three system's regimes marked with different background colours: adaptation (blue), indecision (green), and stagnation (orange); see full discussion in Text ST15. Fixed parameters: $S = 50$, $r_e = 0.2$, $r_s = 0.1$, $q_x = 0.9$, $q_y = 0.8$, $k = 1$, and $\epsilon = 0$. For panels (D)-(F), $\gamma = 0.95$; for panels (G)-(I), $\alpha = 0.1$; and for panels (J)-(L), $\gamma = 0.9$ and $\alpha = 0.9$.

$q_x = v + \delta$ and $q_y = v$ to obtain,

$$\begin{aligned}
& - \left(-Sk^2\pi^2q_yr_e^2r_s^2 - k\pi q_yr_e^2 \right) \left(4S^2k^3\pi^3q_yr_e^2r_s^4 - 8S^2k^2\pi^2q_yr_s^4 - 4S^2k^2\pi^2r_s^4(q_x - q_y) \right) \\
& + \left(-S^2k^3\pi^3q_yr_e^2r_s^4 + S^2k^2\pi^2q_yr_s^4 - Sk\pi r_s^2(q_x - q_y) \right)^2 = 0.
\end{aligned} \tag{S21}$$

ST16. More localised information diffuses less

The model predicts that the ability of the swarm to adapt decreases by reducing the robot's sensing range r_e (Fig. S6, right column). This results hints at the fact that adaptation can be more difficult when the information is highly localised and accessible to individuals in a limited area. Such a model prediction, to be confirmed with experiments, needs further investigation, out of the scope of this study. Intuitively, we can understand that by increasing the robot's sensing range r_e , the swarm reduces the few-vs-many ratio, and therefore it can exploit a larger minority to change the opinion of the rest of the swarm. Increasing robot's sensing range r_e to reduce the few-vs-many ratio is analogous to including a number of stubborn robots that do not change opinion. This strategy has been empirically investigated in [22], and through our study we can explain it via mathematical modelling (see Fig. 4). Indeed, increasing the number of stubborn robots, increases the rate by which individual are exposed to new target sites. In our system, robots are instead exposed to new target sites through individual exploration. Therefore, either increasing the number of stubborn or the sensing range r_e , the final effect is the same: the few-vs-many ratio reduces and adaptation is facilitated.