Collective decision with 100 Kilobots: speed versus accuracy in binary discrimination problems

Gabriele Valentini, Eliseo Ferrante, Heiko Hamann, Marco Dorigo

► To cite this version:

Gabriele Valentini, Eliseo Ferrante, Heiko Hamann, Marco Dorigo. Collective decision with 100 Kilobots: speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, 2016, 30 (3), pp.553-580. 10.1007/s10458-015-9323-3 . hal-01403730

HAL Id: hal-01403730

<https://hal.science/hal-01403730>

Submitted on 6 Dec 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Université Libre de Bruxelles

*Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*

**Collective Decision with 100 Kilobots
Speed vs Accuracy in Binary
Discrimination Problems**

G. VALENTINI, E. FERRANTE, H. HAMANN, and M.
DORIGO

IRIDIA – Technical Report Series

Technical Report No.
TR/IRIDIA/2015-005

July 2015

IRIDIA – Technical Report Series
ISSN 1781-3794

Published by:

IRIDIA, *Institut de Recherches Interdisciplinaires
et de Développements en Intelligence Artificielle*
UNIVERSITÉ LIBRE DE BRUXELLES
Av F. D. Roosevelt 50, CP 194/6
1050 Bruxelles, Belgium

Technical report number TR/IRIDIA/2015-005

The information provided is the sole responsibility of the authors and does not necessarily reflect the opinion of the members of IRIDIA. The authors take full responsibility for any copyright breaches that may result from publication of this paper in the IRIDIA – Technical Report Series. IRIDIA is not responsible for any use that might be made of data appearing in this publication.

Collective Decision with 100 Kilobots

Speed vs Accuracy in Binary Discrimination Problems

Gabriele Valentini · Eliseo Ferrante ·
Heiko Hamann · Marco Dorigo

Received: date / Accepted: date

Abstract Achieving fast and accurate collective decisions with a large number of simple agents without relying on a central planning unit or on global communication is essential for developing complex collective behaviors. In this paper, we investigate the speed versus accuracy trade-off in collective decision-making in the context of a binary discrimination problem—i.e., how a swarm can collectively determine the best of two options. We describe a novel, fully distributed collective decision-making strategy that only requires agents with minimal capabilities and is faster than previous approaches. We evaluate our strategy experimentally, using a swarm of 100 Kilobots, and we study it theoretically, using both continuum and finite-size models. We find that the main factor affecting the speed versus accuracy trade-off of our strategy is the agents neighborhood size—i.e., the number of agents with whom the current opinion of each agent is shared. The proposed strategy and the associated theoretical framework can be used to design swarms that take collective decisions at a given level of speed and/or accuracy.

Keywords collective decision-making · swarm robotics · majority rule · voter model · self-organization · ordinary differential equations · chemical reaction network · Gillespie algorithm · Kilobot

Gabriele Valentini
IRIDIA, Université Libre de Bruxelles, Brussels, Belgium
Tel.: +32 0 2650 2712
Fax: +32 0 2650 2715
E-mail: gvalenti@ulb.ac.be

Eliseo Ferrante
Laboratory of Socioecology and Social Evolution, KU Leuven, Leuven, Belgium

Heiko Hamann
Department of Computer Science, University of Paderborn, Paderborn, Germany

Marco Dorigo
IRIDIA, Université Libre de Bruxelles, Brussels, Belgium

1 Introduction

The ability to make careful and timely decisions is an essential feature of most artificial systems, being them formed by a single agent or, as in this paper, by a swarm of reactive agents. With the aim of studying problems of swarm robotics [1] within a more principled and general framework, we focus on the design and analysis of collective decision-making—an essential collective behavior that implements agency at the swarm level and is a required component in many applications of swarm robotics [3, 16, 24]. We consider how a swarm of agents with minimal capabilities can solve the problem of finding the best out of two options that are symmetrically distributed in a certain environment [25, 44]. The swarm has to reach such an agreement by local interactions only and by exploiting positive feedback, i.e., by implementing a self-organized process. Contrary to classic multiagent systems formed by rational and more informed agents, the agents in a swarm act highly asynchronously, independently assessing the quality of the two alternatives over and over again, repeatedly changing preference for the best option, and communicating that choice to their local neighbors only. Positive feedback is induced by a mechanism that is generally called positive feedback modulation [11]. It amplifies or reduces the time that individual agents spend on participating to the decision-making process as a function of the option’s quality. Over time, this process biases the opinion of the majority of agents towards the highest quality.

In this paper, we thoroughly study the collective decision-making strategy first proposed in [45]: *Direct Modulation of Majority-based Decisions* (DMMD). In this strategy, modulation of positive feedback is implemented by agents that advertise their opinion for a time that is proportional to the assessed quality. The same modulation of positive feedback is used also in Valentini et al. [44]. The two strategies differ for the decision rule used by the agents. In [44], we used the voter model [4, 18], which is based on agents copying the opinion of a random neighbor at each application of the decision rule. In this paper, we use instead the majority rule [9], which is based on agents changing their opinion to the one shared by the majority of individuals in their neighborhood. Using both a continuous approximation model and a stochastic, finite-size model, we show that the majority rule enables faster although less accurate collective decisions in comparison to the voter model. The speedup allowed us to realize an implementation of the decision-making strategy on a robot swarm composed of 100 Kilobots [31]—simple autonomous robots with limited sensing and actuation capabilities.

In order to understand the dynamics obtained with robot experiments and in order to predict the performance in arbitrary regions of the parameter space, we developed a model of the collective decision-making system that is based on ordinary differential equations (ODEs). We analytically compare the majority rule versus the voter model by means of this model, showing that the majority rule achieves faster decisions at the expense of lower accuracy. This result is confirmed for finite-size systems using a second model, based on a chemical reaction network simulated numerically using the Gillespie algorithm. Using both modeling techniques, we show that the speed-accuracy trade-off [6, 27] of the DMMD strategy is strongly dependent on one key parameter of the system: the neighborhood size of individual agents when applying the majority rule.

The remaining of the paper is organized as follows. In Section 2, we describe our collective decision-making strategy in details. In Section 3, we present the ex-

Fig. 1 Illustration of the probabilistic finite-state machine of the individual agent. Solid and dotted lines represent, respectively, deterministic and stochastic transitions; symbols D_i and E_i with $i \in \{a, b\}$ represent the dissemination states and the exploration states, while symbols MR highlight the application of the majority rule at the end of the dissemination state.

periments we performed with the Kilobots, and the corresponding results. The two mathematical models, ODEs and chemical reaction network, and their predictions, are presented in Section 4 and Section 5, respectively. Section 6, contains a discussion and a comparison with existing approaches in the literature. Conclusions and future research perspective are finally presented in Section 7.

2 Collective Decision-Making Strategy

We design a self-organized decision-making strategy to allow a swarm of agents to discriminate between two options based on their quality. Although the approach is general enough for an arbitrary number of options [45], here we focus on a binary scenario to simplify the description of the DMMD strategy. We refer to the two options as option a and option b . The quality of the two options is denoted with $\rho_i \in (0, 1]$, $i \in \{a, b\}$. Each agent in a swarm has always a preference for an option, either a or b , referred to as the agent's opinion. Furthermore, each agent can be in one of four possible states: *dissemination* states D_a and D_b , *exploration* states E_a and E_b . The resulting probabilistic finite-state machine is shown in Figure 1.

In the dissemination states, each agent communicates its current opinion locally (which is a when in state D_a and b when in state D_b) to other agents in its neighborhood that are also in the dissemination state. This behavior is performed for the entire duration of the dissemination state. Before moving to an exploration state, the agent perceives and collects the opinions of its neighbors. Then, the agent adds its own opinion to this group of opinions and applies the majority rule to determine its next preferred option. Depending on the outcome of the majority rule, the agent switches to one of the two exploration states E_a or E_b (cf. dotted lines in Figure 1). In the case of a tie the agent keeps its current opinion.

In the exploration states, each agent assesses the quality of its currently preferred option, which is a when in state E_a and b when in state E_b . For the entire duration of the exploration state, the agent evaluates the characteristic features that determine the quality associated to its opinion following a given domain-specific routine. The quality-estimation routine depends on the particular target scenario and could involve complex robot behaviors—for example, those necessary to explore a candidate construction site and evaluate its level of safety. Independently of the scenario, the quality estimation routine results in one sample measurement

which is generally subject to noise. The swarm processes noisy measurements by acting as a filter that averages over many individual agent measurements (see also [44]). Once the exploration is completed, the agent switches to the dissemination state that corresponds to its current opinion (cf. solid lines in Figure 1).

A core mechanism of the DMMD strategy, which implements the selection of the best option, is the so-called *modulation of positive feedback* [11, 44]. The agent controller is designed to scale the time spent in the dissemination states proportionally to the quality of the opinions. The time spent disseminating a (respectively, b) is directly proportional to the opinion’s quality $\rho_a g$ ($\rho_b g$) where g is the unbiased dissemination time, a parameter set by the designer. The parameter g represents the average duration of opinion dissemination without considering its modulation and is subject to application-specific considerations. The agents have control of the positive feedback mechanism by adapting the amount of time that they disseminate an opinion. In this way agents influence the frequency with which other agents observe a certain opinion in their dissemination state. As a consequence, observing neighbors that are in favor of the best option is more likely than observing neighbors that are in favor of other, lower quality alternatives. Therefore, the swarm is biased towards achieving a collective decision for the best option. This idea is loosely inspired by the honeybee behavior shown when honeybees search for potential site locations for their new nest [7, 8, 37].

A requirement of this strategy is that the interaction among agents in the dissemination states should be well-mixed or, at least, approximately well-mixed. That is, the probability of any agent to encounter a neighbor of a certain opinion is approximately proportional to the global distribution of opinions in the whole swarm. The well-mixed property is only a weak requirement as it influences the efficiency of the decision-making process but only in extreme cases its efficacy. If the spatial distribution of agents is sufficiently well-mixed, the decision-making strategy is efficient and successful. The more the system deviates from a well-mixed state, the slower the decision-making process is. Only if the spatial distribution of agents is far from well-mixed, then the decision-making process is slowed down considerably by spatial fragmentation of opinions (e.g., formation of clusters of robots with the same opinion) and might even end up in a deadlock, that is, a macroscopic state of indecision far from consensus [5]. In the next section, we explain how this requirement can be fulfilled for the case of autonomous robots.

3 Experiments with the Kilobots

In this section we present a series of robot experiments aimed at validating the robustness of the DMMD strategy to the constraints imposed by the real world. Due to the simplicity of our model, we were able to implement experiments with a relatively large swarm of 100 Kilobots [31]. The Kilobot (shown in Figure 2a) is a low-cost (currently about €120), small-sized (3.3 cm diameter) robot equipped with two independently-controllable vibrating motors for differential drive locomotion, infrared receiver and transmitter for local, close-range communication, RGB LED light emitters, and a light sensor for sensing the intensity of the ambient light. A Kilobot can move with a maximum nominal speed of 1 cm/s and rotational speed of $\pi/4$ rad/s. It can send and receive, within a maximum distance of approximately 20 cm, infrared messages at a rate of up to 30 kb/s.

Fig. 2 The figure shows a) the Kilobot robot highlighting the position of vibrating motors, light sensor, and IR transceiver; and b) the experimental arena partitioned into nest, red site, and blue site with details of both IR and light beacons.

We implemented the DMMD strategy in the following scenario. We built a rectangular arena whose total size is $100 \times 190 \text{ cm}^2$ (see Figure 2b), which is three orders of magnitude larger than the footprint of a single Kilobot. Options *a* and *b* correspond to foraging *sites* of quality ρ_a and ρ_b , respectively. The two sites are $80 \times 45 \text{ cm}^2$ large and are located at the right (site *a*, red) and at the left (site *b*, blue) of the arena. The remaining, central part of the arena is called *nest*. It is $100 \times 100 \text{ cm}^2$ large and it is where the swarm of 100 Kilobots is initially placed. The nest is also the decision-making hub of the swarm, that is, the decision rule is only allowed to be executed within the nest. Robots are initialized in the dissemination state with random position and orientation. At time $t = 0$, the swarm consists of 50 robots with opinion *a* and 50 with opinion *b*; their initial quality estimate is unbiased ($\hat{\rho}_a(0) = \hat{\rho}_b(0) = 1$). Our goal is to have the majority of the swarm foraging from the site associated with the higher quality, in this scenario by definition site *a*. Specifically, the quality of site *a* is twice as high as that of site *b* ($\rho_a = 1$ and $\rho_b = 0.5$). We position a light source on the right of the arena, to provide a landmark that can be used by the robots to navigate and find the three areas. They perform phototaxis when they need to move from site *b* to the nest or from the nest to site *a* and anti-phototaxis in the remaining two cases.

Kilobots can identify the two sites and measure the associated quality by using their infrared sensors. For each site, five additional Kilobots are positioned upside-down under the transparent surface of the arena, at the border between the site and the nest, and act as beacons. These Kilobots continuously communicate locally a message containing the type (*a* or *b*) and the quality (ρ_a or ρ_b) associated to a site. These infrared messages are perceived only within the sites, both due to their local nature (approximately 15 cm) and because we cover the nest area by light occluding paper to prevent robots from sensing this information at the nest.

As defined by the DMMD strategy, robots continuously iterate between periods of exploration and dissemination. Robots explore the site associated with their current opinion by navigating from the nest to that site and measuring its

Fig. 3 The figure shows the finite-state machines that implement the motion control method of the individual robot during the execution of the decision-making strategy: a) FSM used for both dissemination states D_a and D_b ; and b) the two FSMs used for the exploration state E_a (top) and E_b (bottom). Symbols represent low level motion routines, respectively, random walk (RW), phototaxis (PT), and antiphototaxis ($!PT$); colors represent the current robot opinion, respectively, red for opinion a and blue for opinion b .

quality. They then go back to the nest, where they disseminate their current opinion modulating the positive feedback based on the measured quality ρ_i . Finally, they collect the opinions of their neighbors and apply the majority rule potentially changing preference for the best site. As explained in Section 2, the swarm can potentially suffer from opinion fragmentation [5]. For example, the robots might distribute themselves in such a way that all robots with opinion a are positioned close to site a and all robots with opinion b are positioned close to site b . As a consequence, a robot would be more likely to interact with a robot of the same opinion which might cause the decision-making process to enter a deadlock. To maintain the spatial distribution close to a well-mixed distribution, we implemented specialized motion routines that, if performed for a sufficiently long period of time, allow robots to mix well in the nest while disseminating their opinions.

3.1 Robot Control Algorithm

We implemented the DMMD strategy by using the motors, the light sensor, and the infrared transceiver of the Kilobot. Three low-level motion routines—*random walk*, *phototaxis* and *anti-phototaxis*—allow robots to navigate and to explore the environment and to disseminate their opinion. Depending on the current control state and on the current robot opinion, these routines are combined into a probabilistic finite-state machine to implement the behavior in the dissemination states (see Figure 3a) and exploration states (see Figure 3b). In the supplementary material [42], we provide a video highlighting intermediate phases of the robot controller. In the following, we employ the exponential distribution to determine the duration of several sub-routines. We have chosen this distribution due to its large variance that allows us to break the synchrony in the robot motion patterns by introducing noise that improves the mixing of robots (see also Section 6.2 in [45]).

3.1.1 Low-level motion routines

We implemented a correlated random walk in order to improve the mixing of the opinions in the swarm. When performing the random walk, the robot moves forward for an exponentially distributed amount of time; this mostly results in short walks with sporadically longer ones. Then, the robot turns in place for a normally distributed period of time. Phototaxis (respectively, anti-phototaxis) is implemented by letting robots perform oriented motion towards (away from) the light source placed on the right side of the arena. The robots search for the direction with the highest (lowest) light intensity by turning on the spot; once found, they move forward until the ambient light intensity measurement falls outside a tolerance range; when this happens, the robots resumes the on-spot search of the correct direction of motion.

3.1.2 Dissemination states

In both dissemination states D_a and D_b , the robots execute the finite-state machine depicted in Figure 3a. Robots start by performing a random walk in the nest while communicating locally their opinions. The random walk favors the spatial mixing of robots in space and therefore of their opinions. In addition to their current opinion, robots also communicate a randomly generated 16-bit identifier that, with high probability, uniquely identifies the robot in its local neighborhood. This is used to make sure that, at any given time, robots distinguish the opinion of different neighbors. Robots directly modulate positive feedback by spending an exponentially distributed amount of time in the dissemination state. The mean of this exponential distribution is either $\hat{\rho}_a g$ or $\hat{\rho}_b g$, where $\hat{\rho}_i, i \in \{a, b\}$ is the current robot estimate of the option quality. During dissemination, robots might perceive messages from the five robot-beacons positioned at each border between the nest and a site. If such a message is perceived, it means that the robot is mistakenly leaving the nest and it therefore performs either phototaxis or anti-phototaxis in order to return to the nest (see Figure 3a). Oriented motion is performed by the robot for as long as beacon messages are received and proceeds for an additional period of 20 seconds after the last message. This kind of oriented motion allows the robot to keep a distance from the border and to favor a good mixture of robot opinions in space. During the last three seconds before leaving the nest, a robot records the opinions of its neighbors. It then adds its own current opinion to that record, applies the majority rule to determine its next preferred option and, consequently, the next site to explore. We chose a relatively short time for opinion collection in order to reduce the time-correlation of the observed opinions (i.e., robots taking decisions on the basis of outdated information). Nonetheless, this period of time is sufficient for a robot to receive messages from many neighbors as will be clear from the analysis in the next section. Finally, the robot leaves the nest to explore the chosen site.

3.1.3 Exploration states

In states E_a and E_b , robots move to the site associated to their current opinion, performing either phototaxis (towards site a) or anti-phototaxis (towards site b). Once they reach the site, they explore it for an exponentially distributed amount

Fig. 4 The figure shows a series of screen-shots taken from one experiment with a swarm of 100 Kilobots. The screen-shots are taken every 18 minutes of execution.

of time, they record the associated quality (received from the beacons), and then return to the nest. During this time, the robot executes the finite-state machine depicted in Figure 3b (respectively, top for site *a* and bottom for site *b*) in order to stay within the boundaries of the site. We consider this behavior as an abstraction of a quality-estimation routine dependent on the target scenario. For example, the robot might assess during this period how much of a certain resource is available in the site (e.g., construction material), what is the average level in the site of a certain physical feature (e.g., temperature), etc. Additionally, to ensure that robots fully enter the site (i.e., they do not remain in the border region), we implemented the following mechanism. If a robot wants to explore site *a* (respectively, *b*), it performs phototaxis (anti-phototaxis) in two phases. In the first phase, the robot performs phototaxis (anti-phototaxis) until it perceives a message from the beacons, indicating that the robot has crossed the border and entered site *a* (*b*). In the second phase, phototaxis (anti-phototaxis) is continued for as long as messages from the beacons are not received for 5 seconds. Exactly the same mechanism, but reversed, is used by the robots returning to the nest and entering the dissemination state. The second phase also eases the mixing of robot opinions in the nest because robots are programmed to approach the center of the nest.

3.2 Results of Robot Experiments

Our main working hypothesis is that efficiency and accuracy of the decision-making process are affected by the neighborhood size considered when applying the majority rule. The neighborhood size can be directly or indirectly controlled by the experimenter. However, its size could fluctuate over time due to spatial density constraints. In our scenario, we consider two extreme situations. We restrict the maximum neighborhood size to either 4 robots or to 24. The latter case corresponds in practice to no restriction, since the actual number of neighbors perceived by a robot at a given time is rarely greater than 24. We refer to this parameter as the maximum size of the opinion group $\mathcal{G}_{\max}$ and we define it in a way so that it also includes the opinion of the considered robot: $\mathcal{G}_{\max} \in \{5, 25\}$. For each of these two cases, we performed 10 independent runs, each lasting 90 min each (see

Fig. 5 Results of the experiments with the robot and statistical test: a) distributions of the proportion of robots with opinion a over time; b) median and confidence intervals predicted by the GLMM (see explanation in the main text).

supplementary material [42]). Recall that the parameter g determines the duration of the dissemination state without considering positive feedback modulation (i.e., the site quality). The higher the value of g , the longer the robot performs its random walk behavior contributing to the mixing of the opinions, and the longer it takes to the swarm to find a consensus. After a few preliminary test runs, we set $g = 8.4$ min (i.e., about 500 seconds), so as to obtain a proper mixing of the robots opinions while limiting the overall decision time. Some snapshots taken from one of the experiments are depicted in Figure 4, for an explanatory video see [46].

The results of the robot experiments are shown in Figure 5. Figure 5a reports the behavior of the proportion of robots with opinion a ($(D_a + E_a)/N$) over time for the two cases: $G_{\max} = 5$ and $G_{\max} = 25$. Qualitatively, we observe that the maximum allowed neighborhood size influences the speed of the decision-making process. To determine whether the observed difference in speed was statistically significant, we fitted a generalized linear mixed model (GLMM) with binomial response, where we considered time as a continuous covariate, $G_{\max}$ as a fixed factor, and $G_{\max}$ nested into the run number as a random factor. In this model, we also included explicitly the interaction of $G_{\max}$ with time as an additional fixed factor, which turned out to be significant (p -value = 0.047). The presence of a significant interaction confirms our qualitative observation: the curves representing the predicted proportion of robots with opinion a as a function of time for the two settings ($G_{\max} = 5$ and $G_{\max} = 25$) do not grow at the same rate. The one with $G_{\max} = 25$ grows faster than the one with $G_{\max} = 5$. These two curves (lines) are shown in Figure 5b, together with the confidence intervals (shaded areas) predicted by the GLMM. As we can see, the system reaches a 90% consensus on a faster with $G_{\max} = 25$ than with $G_{\max} = 5$: with 95% confidence, the system converges to 90% consensus in the interval between $t \approx 55$ and $t \approx 69$ for $G_{\max} = 25$, and in the interval between $t \approx 66$ and $t \approx 80$ for $G_{\max} = 5$.

In both parameter settings, after 90 minutes of execution the swarm always reached a state where the broad majority preferred opinion a , but this almost never coincided with 100% consensus. We identified robot failure as a possible cause of this result: robots occasionally incurred in battery failures, stuck motors, or switched to stand-by due to short circuits caused by collisions with other robots

(0.7 robots per experimental run). Additionally, some robots experienced serious motion difficulties due to poor motor calibration, were unable to reach target areas (i.e., nest, sites), and were thus prevented from changing opinion. Despite these failures, the proposed self-organized decision-making mechanism proved to be very robust by allowing the swarm to always reach a correct collective decision.

4 Ordinary Differential Equations Model

We now deepen our understanding of the direct modulation of majority-based decision strategy. We study the behavior of the systems under the continuous limit approximation ($N \rightarrow \infty$). We also study systematically the impact of the neighborhood size on the speed and accuracy of the decision-making process. For this purpose, we define a system of ordinary differential equations (ODEs) and we analyze it using standard tools of dynamical systems theory. The ODE model describes the dynamics of the expected proportion of agents in the dissemination states (d_a and d_b) and the expected proportion of agents in the exploration states (e_a and e_b). Our mathematical modeling approach relies on two assumptions: *i.* the neighborhood size of agents is constant and *ii.* each agent has always a noiseless quality estimate of its opinion, (even at time $t = 0$). Assumptions *i.* and *ii.* simplify our derivation of the ODE model by allowing us to neglect the effects of random fluctuations of parameters $\mathcal{G}$, ρ_a and ρ_b and to consider instead their mean values.

One essential feature that we need to model is the modulation of positive feedback, that is, the regulation of the time agents spend in the dissemination states. This time is proportional to the quality of the sites ($\rho_a g$ and $\rho_b g$). If these two quantities represent the average time spent by agents to disseminate their opinion, then we can also define the *rates* at which agents move from dissemination to exploration as the inverse of these quantities: $\alpha = (\rho_a g)^{-1}$ and $\beta = (\rho_b g)^{-1}$.

Additionally, to derive our set of differential equations, we need to know the rates at which agents change their opinions. We need to express the probability p_{ab} that an agent with opinion a switches to opinion b as an effect of applying the majority rule for a given group size $\mathcal{G}$ (similarly for probability p_{ba}). In the model, we also need to consider the cases where the application of the majority rule has no effect, that is, no opinion switch is triggered after its application. The probabilities of keeping the same opinion are denoted as p_{aa} and p_{bb} .

First we consider a simplified example to explain how we determined these probabilities (cf. Figure 6). Consider an agent i with opinion a that has two neighbors j, h . Hence, we have $\mathcal{G} = 3$. The probability p_{ab} that this agent switches opinion to b after applying the majority rule is computed by considering all possible combinations of neighbors that form a majority for b . In this simple example with a small group, the only relevant case is when both neighbors j and h have opinion b (denoted by bb). All the other cases, aa , ab , and ba , correspond to a majority of a which leaves agent i unaffected. We define p_a the probability that a neighboring agent has opinion a ; due to symmetry, $(1 - p_a)$ is the probability that a neighboring agent has opinion b . Probability p_a is a function of the proportions d_a and d_b of agents in the dissemination states. Only these agents advertise their opinion and only they can provoke a switch, which gives $p_a = \frac{d_a}{d_a + d_b}$. Given the probability p_a , we can derive p_{ab} as the joint probability $(1 - p_a)^2$ to have two neighbors with opinion b . In the same way, the probability p_{aa} of not provoking

Fig. 6 The figure illustrates the application of the majority rule in a group of $\mathcal{G} = 3$ agents. An agent i with opinion a applies the majority rule over a set of opinions containing its current opinion and the opinions of its two neighbors j and h . In the first three cases, agent i keeps its current preference for option a , in the last case, the agent switches its opinion to option b .

a switch is $p_a^2 + 2p_a(1 - p_a)$, obtained as the sum of the three cases aa , ab , and ba . The derivations of probabilities p_{ij} , $i, j \in \{a, b\}$, is performed by assuming an infinite number of agents ($N \rightarrow \infty$) and a well-mixed distribution of their positions (and therefore their opinions) within the nest. The first assumption is a direct consequence of the continuous nature of the ODE model presented in this section. The second assumption is instead motivated by the requirement of the our strategy to have robots approaching a well-mixed distribution (cf. Section 2), and it is supported by the special motion routines of robots described in Section 3.

The above reasoning to compute probabilities p_{aa} and p_{ab} for a pair of neighbors can be generalized to a generic group size $\mathcal{G}$ using equations

$$p_{aa} = \sum_{i=\lfloor (\mathcal{G}-1)/2 \rfloor}^{\mathcal{G}-1} \binom{\mathcal{G}-1}{i} p_a^i (1 - p_a)^{\mathcal{G}-1-i}, \quad (1)$$

$$p_{ab} = \sum_{i=0}^{\lfloor (\mathcal{G}-1)/2 \rfloor - 1} \binom{\mathcal{G}-1}{i} p_a^i (1 - p_a)^{\mathcal{G}-1-i}. \quad (2)$$

These equations are a discrete integration of a Binomial distribution, where p_a is the success probability, $\mathcal{G} - 1$ the number of trials, and i the number of successes. The rationale is simple. In order to keep opinion a , the number of successes for a needs to be less than half of the neighborhood $(\mathcal{G} - 1)$ minus 1 which is the agent itself. More successes than that provoke a switch. The expressions for probabilities p_{bb} and p_{ba} can be obtained by swapping the power indexes in Equations (1–2).

Table 1 Summary of the parameters used to compare the ODE model to the robot experiments. We use the following legend in the column “Type”: DP (design parameter) is a parameter set by designer; RE (robot experiment) is a parameter estimated from robot experiments; PP (problem parameter) is a parameter that defines the problem to be solved.

<i>Parameter</i>	<i>Value</i>	<i>Type</i>
Quality of Site a	$\rho_a = 1.0$	PP
Quality of Site b	$\rho_b = 0.5$	PP
Maximum group size (robots)	$\mathcal{G}_{\max} \in \{5, 25\}$	DP
Mean group size (ODEs)	$\mathcal{G} \in \{5, 9\}$	RE
Exploration time	$\sigma^{-1} = 6.072$ min	RE
Dissemination time	$g = 8.4$ min	DP

The rates α, β and the probabilities p_{ij} , $i, j \in \{a, b\}$ allow us to finally write the system of four ordinary differential equations that model the proposed strategy

$$\begin{cases} \frac{d}{dt}d_a = \sigma e_a - \alpha d_a, \\ \frac{d}{dt}d_b = \sigma e_b - \beta d_b, \\ \frac{d}{dt}e_a = p_{aa}\alpha d_a + p_{ba}\beta d_b - \sigma e_a, \\ \frac{d}{dt}e_b = p_{ab}\alpha d_a + p_{bb}\beta d_b - \sigma e_b. \end{cases} \quad (3)$$

The dynamics of the proportions of agents in the dissemination states d_a and d_b are modeled by the first two equations. Proportions d_a and d_b increase every time agents return from the corresponding exploration state E_a or E_b , whose rate is determined by σ (see Section 3.2). d_a (respectively, d_b) decreases at a rate α (β), due to agents leaving the dissemination state. The third and fourth equations model the dynamics of the proportion of agents in the exploration states e_a and e_b . The proportion e_a (respectively, e_b) increases every time agents in the dissemination states switch to state E_a (E_b). This happens for all agents in the dissemination state (D_a and D_b) that switch to state E_a (E_b) by applying the majority rule. These rates of switching depend on the probabilities p_{aa} and p_{ba} in Equation (1) and Equation (2) (p_{ba} and p_{bb}). Finally, e_a (e_b) decreases at a rate σ due to agents leaving the exploration state E_a (E_b).

4.1 Validation Against Robot Experiments

In order to use the model defined in the system of Equations (3) to study different regions of the parameter space, we first need to check whether the model can qualitatively predict the results obtained with robot experiments. To do, we performed additional robot experiments necessary to estimate the values of the parameters $\mathcal{G}$ and σ of the ODE model. Table 1 lists all parameters used in the ODE model, while Appendix A contains a detailed analysis of the additional robot experiments. We set the group size in the ODE model by rounding the average group size obtained in the robot experiments. This was 8.57 when $\mathcal{G}_{\max} = 25$ and 4.4 when $\mathcal{G}_{\max} = 5$. We therefore set $\mathcal{G} = 5$ and $\mathcal{G} = 9$ in the two cases. The value of g in the robot scenario was set by the designer to $g = 8.4$ min. The

Fig. 7 The figure shows the comparison between robot experiments (box-plots) and predictions of the ODE model in Equations (3) (lines), respectively, for $G_{\max} = 5$ (green) and $G_{\max} = 25$ (purple). In a), we compare robot experiments against the prediction of the ODE model given the estimated parameters; in b), the predictions of the ODE model are instead scaled in time according to $t' = 3t + g$. Parameters: $\sigma = 6.072$, $g = 8.4$, $\rho_a = 1$, $\rho_b = 0.5$, $G_{\max} \in \{5, 25\}$, $\mathcal{G} \in \{5, 9\}$.

mean duration of the exploration state (i.e., the inverse of the rate at which robots transit from the exploration state to the dissemination state) was estimated from data and equals $\sigma^{-1} = 6.072$ min.

The comparison between the system of ODEs and the robot experiments is shown in Figure 7a. As we can see, the trajectories predicted by the model (solid lines) have the same shape but do not match those obtained in the robot experiments (box-plots). Specifically, the ODE model appears to be shifted in time and to evolve at a higher speed. Indeed, the fitting improves if we apply the following time rescaling: $t' = 3t + g$, see Figure 7b. This result suggests that robot experiments are approximately 3 times slower than the dynamics of the ODE model, and shifted by a factor g . The offset g is easily explained by assumption *ii.* of the ODE model: initially, robots do not have a correct estimate of the quality of the two sites but begin the execution with $\rho_a = \rho_b = 1$ (in contrast to assumption *ii.*). Before having a correct quality estimate, robots have to do an initial exploration of the sites, for which they need to wait on average g units of time. Besides, we conjecture that spatial interference among robots might have caused a partial violation of the well-mixed assumption of the model which caused a slow down by a factor of 3. This result could be possibly mitigated by extending (through parameter g) the period of time in which robots perform the random walk designed to minimize this problem (see Section 3.1). Despite this, we obtained correct, qualitative predictions from the ODE model with respect to the asymptotic dynamics of robot experiments.

4.2 Stability of Equilibria

After validating the ODE model with the results from the robot experiments, our next objective is to understand which are all the possible collective decisions that might emerge from the decision-making strategy. To reach this objective, we determine what are all the possible equilibria $\tilde{\gamma} = [\tilde{d}_a, \tilde{d}_b, \tilde{e}_a, \tilde{e}_b]^T$ of the system of

ODEs of Equations (3) and perform a stability analysis. The analysis results in three fixed points

$$\tilde{\gamma}_1 = \left[\frac{g\sigma\rho_a}{1+g\sigma\rho_a}, 0, \frac{1}{1+g\sigma\rho_a}, 0 \right]^T, \quad (4)$$

$$\tilde{\gamma}_2 = \left[0, \frac{g\sigma\rho_b}{1+g\sigma\rho_b}, 0, \frac{1}{1+g\sigma\rho_b} \right]^T, \quad (5)$$

$$\tilde{\gamma}_3 = \frac{1}{\Psi} \left[g\sigma\rho_a\rho_b^2, g\sigma\rho_a^2\rho_b, \rho_b^2, \rho_a^2 \right]^T, \quad (6)$$

where $\Psi = \rho_a^2 + g\sigma\rho_a^2\rho_b + g\sigma\rho_a\rho_b^2 + \rho_b^2$.

Two equilibria, respectively, $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ given in Equation (4) and Equation (5), represent consensus on opinion a and consensus on opinion b . Interestingly, the proportion of agents in exploration and dissemination states predicted by $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ depend only on the exploration and dissemination rates. In turns, this result means that the designer has a tool to fine-tune the desired proportion of agents exploring or disseminating at consensus. This could be of interest during a foraging task [24, 41, 36] to effectively tune the foraging rate, or to aid the calibration of the quorum thresholds [25, 26] when the detection of consensus is necessary to trigger a change in the behavior of the entire swarm (e.g., migration to the selected site). The third equilibrium $\tilde{\gamma}_3$ in Equation (6) corresponds instead to a macroscopic state of indecision where both opinions coexist in the swarm.

A subsequent question that arises is which of these equilibria is asymptotically stable and, more importantly, under which conditions. To answer this question, we linearized the system of ODEs around each equilibrium, calculated the eigenvalues of the corresponding Jacobian matrix, and studied their signs. Note that, due to the conservation of the swarm mass, the system of ODEs in Equations 3 is over-determined. One equation can be omitted, for example the last equation, by rewriting the remaining three using the substitution $e_b = 1 - d_a - d_b - e_a$. Therefore, each equilibrium of the system has only three meaningful eigenvalues. The eigenvalues corresponding to the two consensus equilibria $\tilde{\gamma}_1$ and $\tilde{\gamma}_2$ are

$$\begin{bmatrix} -\frac{1}{g\rho_b} \\ -\sigma \\ \frac{-g\sigma\rho_a\rho_b-\rho_a}{g\rho_a\rho_b} \end{bmatrix}, \begin{bmatrix} -\frac{1}{g\rho_a} \\ -\sigma \\ \frac{-g\sigma\rho_a\rho_b-\rho_b}{g\rho_a\rho_b} \end{bmatrix}, \quad (7)$$

respectively for consensus on option a and for consensus on option b . These eigenvalues depend only on the rates g, σ and on the site qualities ρ_a, ρ_b . Given that these quantities are defined to be always strictly positive, we can conclude that the two consensus equilibria are always asymptotically stable.

The third equilibrium $\tilde{\gamma}_3$ is characterized by eigenvalues with a very complex analytic formulation which prevents us from providing it here for the reader (see supplementary material [42]). Nonetheless, we have performed the stability analysis for this fixed point as well. According to our analysis, for values of $\rho_a, \rho_b \in (0; 1]$, $\rho_a \geq \rho_b$, and for $\sigma, g > 0$, two eigenvalues are always strictly negative while one is always strictly positive. Such a fixed point, which is difficult to visualize due to the high dimensionality of the system, is a saddle point and divides the basin of attraction between trajectories converging to consensus on a and trajectories

converging to consensus on b (see also next section). We can therefore conclude that the macroscopic state of indecision, $\tilde{\gamma}_3$, is not stable.

We finally compare the dynamics of the majority rule against those of the weighted voter model proposed in Valentini et al. [44]. The voter model has simpler asymptotic dynamics as it has only 2 equilibria corresponding to the two consensus decisions. One of the two equilibria is associated with the best opinion a and is asymptotically stable when $\rho_a > \rho_b$. The other equilibrium is unstable. When $\rho_a = \rho_b$, one eigenvalue vanishes for both equilibria that are in this case only Lyapunov stable but not asymptotically stable. Under these conditions, the voter model does not converge to a collective decision but remains indefinitely with the same proportion of opinions a and b with which the swarm was initialized. Therefore, in the limit of $N \rightarrow \infty$, the differences between the voter model and the majority rule are the following. *i*) With the majority rule, convergence to a particular equilibrium depends on the initial conditions (as there are two stable equilibria); whereas the voter model always converges to the best opinion, if one exists ($\rho_a > \rho_b$). *ii*) The majority rule converges, differently from the voter model, to one of the opinions even in the case of symmetric qualities ($\rho_a = \rho_b$). In [44], we show that these properties of the voter model hold only in the deterministic, continuous approximation ($N \rightarrow \infty$), and that they vanish when the influence of finite-size effects is included (see Section 5).

4.3 Speed Versus Accuracy Trade-Off

Our aim in this section is to use the ODE model defined in Equations (3) to analyze how convergence speed and decision accuracy [6, 27] change as a function of a key parameter of our strategy: the group size G . In our terminology, the system has higher accuracy when it can reach consensus on the best opinion (i.e., option a) for a wider range of initial conditions. For all possible initial conditions $d_a(0) \in [0, 1]$, $d_b = 1 - d_a(0)$ we determine the consensus $\tilde{d}_a + \tilde{e}_a \in \{0, 1\}$ that is reached asymptotically from there. We are particularly interested in the border c that separates the two basins of attraction: we converge to $\tilde{d}_a + \tilde{e}_a = 1$ for $d_a(0) \in [c + \epsilon, 1]$ and to $\tilde{d}_a + \tilde{e}_a = 0$ for $d_a(0) \in [0, c - \epsilon]$ where $\epsilon > 0$. Smaller values of c are preferred since they increase the basin of attraction for the best option. Convergence speed is simply measured as the time necessary to reach consensus on any option. To compute this convergence time from the ODE model, we introduce a threshold $\delta = 10^{-3}$ and consider that the system has converged to a collective decision at a certain time t if either $d_a(t) + e_a(t) \geq 1 - \delta$ or $d_b(t) + e_b(t) \leq \delta$. We define convergence time to be the minimum t satisfying this criterion.

The results of this analysis are reported in Figure 8 for decision accuracy and Figure 9 for convergence time. In both figures, the difference between the left and right graphs is the value of the quality parameter ρ_b which determines the difficulty of the decision-making problem. Specifically, a quality of $\rho_b = 0.5$ defines a simpler, more asymmetric discrimination problem where option a is twice as good as option b , whereas $\rho_b = 0.9$ defines a much harder problem where the qualities of the two options are more difficult to distinguish.

In Figure 8a the black, solid line represents the border c between the two intervals of initial conditions (basins of attraction) that lead to different consensus decisions (i.e., asymptotically stable solutions). We observe that this border in-

Fig. 8 The figure shows the results from the speed versus accuracy analysis performed using the ODE model in Equation (3) as a function of the group size G , the initial condition $d_a(0), d_b(0) = 1 - d_a(0)$, and the option quality ρ_b . Figures in the upper row show the border c that divides initial conditions that lead the system to consensus on a (red area) from those with consensus on b (blue area) as a function of the group size G , respectively, a) for $\rho_b = 0.5$ and b) for $\rho_b = 0.9$.

creases roughly logarithmically as a function of the group size $\mathcal{G}$. Higher values of this border indicate a smaller set of initial conditions (red area) that lead the swarm to choose the best option (i.e., site a), and thus lower decision accuracy. The graph shows that the accuracy of the decision-making strategy decreases as a function of the group size. This happens for both easier ($\rho_b = 0.5$, Figure 8a) and more difficult ($\rho_b = 0.9$, Figure 8b) decision-making problems. However, for $\rho_b = 0.9$ this increase is much less noticeable due to the fact that accuracy is already relatively low for small group sizes. Additionally, we can observe that the parity of the group size $\mathcal{G}$ influences the accuracy of the decision-making process. When $\mathcal{G}$ is even the set of initial conditions leading to consensus on option a is smaller than that of the two nearby odd group sizes. This phenomenon, which is more distinct for small group sizes, is characteristic of the majority rule and was reported previously for other systems too [9, 19].

Figures 9a and 9b show through heatmaps how the time necessary to reach a decision varies as a function of the group size $\mathcal{G}$ and of initial conditions $d_a(0), d_b(0) = 1 - d_a(0)$. The black lines provide the border c between consensus on a and consensus on b . As we can see, the consensus time increases with the proximity to the border c . Figure 9c and Figure 9d detail instead the shape of consensus time for selected values of the group size $\mathcal{G}$. The color of the lines represents the asymptotic result of the decision-making process, respectively, red for consensus on a and blue for consensus on b . As we can see, the consensus time is higher when the initial proportion $d_a(0)$ of agents favoring option a is closer to the border c between the basins of attraction that divides initial conditions leading to consensus on a from those leading to consensus on b (i.e., where lines turn from blue to red in Figures 9c and 9d). Additionally, we observe that increasing the group size $\mathcal{G}$ speeds up the decision-making process for a wide range of initial conditions $d_a(0)$. This speedup is approximately halved every time we double the number of neighbors in the group (c.f. the speedup given by $\mathcal{G} = 9$ with respect to $\mathcal{G} = 5$ with that given by $\mathcal{G} = 17$ with respect to $\mathcal{G} = 9$).

Fig. 9 The figure shows the results from the speed versus accuracy analysis performed using the ODE model in Equation (3) as a function of the group size G , the initial condition $d_a(0), d_b(0) = 1 - d_a(0)$, and the option quality ρ_b . The heatmaps in the upper row show the consensus time (min) for group size $G \in \{3, 25\}$ and initial condition $d_a(0) \in [0, 1]$, respectively, a) for $\rho_b = 0.5$ and b) for $\rho_b = 0.9$. Black solid lines represent the border points c for each value of G . In the lower row, figures show the consensus time over initial conditions $d_a(0)$ for group size $G \in \{5, 9, 17\}$, respectively, c) for $\rho_b = 0.5$ and d) for $\rho_b = 0.9$. Red and blue lines represent initial conditions with consensus on option a and option b .

The results given in Figure 8 and Figure 9 reveal the crucial trade-off between convergence speed and decision accuracy of the DMMD strategy. We can increase convergence speed by increasing the group size G at the cost of lower accuracy. Similarly, we can have higher accuracy at the cost of lower convergence speed. This behavior is particularly evident for simple decision-making problems (e.g., $\rho_b = 0.5$). For more difficult discrimination problems (e.g., $\rho_b = 0.9$), the group size G has a lower influence on the decision accuracy while the swarm can still benefit in terms of convergence speed.

5 Chemical Reaction Network

In Section 4 we studied the asymptotic properties of the DMMD strategy using the continuous limit approximation ($N \rightarrow \infty$). Real-world swarm systems, however, are composed of a large but finite number of agents. In many of these systems, finite size crucially influences the system's dynamics so that predictions based on continuous approximations might be of limited use [40, 44]. A number of different

Algorithm 1: GILLESPIE ALGORITHM($N, g, \sigma, \rho_a, \rho_b, D_a(0)$)

```

1 Initialize agents in the swarm:  $D_a = D_a(0), D_b = N - D_a(0), E_a = 0, E_b = 0$ 
2 Initialize time:  $t = 0$ 
3 repeat
4   Compute the total reaction rate:  $\kappa = \alpha D_a + \beta D_b + \sigma E_a + \sigma E_b$ 
5   Generate an exponentially distributed time  $t'$  with rate  $\kappa$  and set  $t = t + t'$ 
6   Randomly choose the next reaction to occur with probabilities
       $(\alpha D_a / \kappa, \beta D_b / \kappa, \sigma E_a / \kappa, \sigma E_b / \kappa)$ 
7   If Reaction 8 (Reaction 9) occurs, the outcome of its right-hand side  $\mathfrak{E}_a | \mathfrak{E}_b$  is
      determined by probabilities  $q_{aa}$  and  $q_{ab}$  ( $q_{ba}$  and  $q_{bb}$ ).
8   Update  $D_a, D_b, E_a, E_b$  according to the outcome of steps 6–7
9 until  $D_a + E_a \in \{0, N\}$ ;

```

modeling techniques exist to deal with finite-size effects, such as Markov chains [38, 13, 41, 43] and master equations [22, 21, 44, 47, 36, 17]. Here, we use the formalism of (chemical) master equations which are derived from a chemical reaction network [15].

Chemical master equations are stochastic differential equations modeling the dynamics of coupled chemical reactions among a set of molecules. Using this formalism to model a multiagent system, agents in different states are represented by molecules of different types, while state transitions of individual agents are represented by chemical reactions with certain rates. Chemical master equations are often hard if not impossible to solve analytically. For this reason, we base our study on numerical simulations using the Gillespie algorithm [12]. The Gillespie algorithm—also known as Stochastic Simulation Algorithm—generates statistically correct trajectories of a master equation which can be used to approximate its exact solution.

Given a swarm of N agents, we use symbols D_a and D_b to denote the number of agents in the dissemination states and symbols E_a and E_b to denote the number of agents in the exploration states. Additionally, we refer to an individual agent in one of these states using symbols $\mathfrak{D}_a$ and $\mathfrak{D}_b$ for opinion dissemination and symbols $\mathfrak{E}_a$ and $\mathfrak{E}_b$ for exploration, respectively. The proposed decision-making strategy is modeled by the chemical reactions

The above set of reactions is sufficient to define a master equation as described by van Kampen [15]. According to Reactions (8–9), each agent in a dissemination state (either $\mathfrak{D}_a$ or $\mathfrak{D}_b$) switches to an exploration state (either $\mathfrak{E}_a$ or $\mathfrak{E}_b$) at a constant rate. Specifically, at rate $\alpha = (\rho_a g)^{-1}$ if the agent is in state $\mathfrak{D}_a$ or at rate $\beta = (\rho_b g)^{-1}$ otherwise. Reactions (10–11) model instead the transition of agents from an exploration state (either $\mathfrak{E}_a$ or $\mathfrak{E}_b$) to a dissemination state (either $\mathfrak{D}_a$ or $\mathfrak{D}_b$) which happens at a constant rate σ .

In the Gillespie algorithm [12], the evolution in time of the numbers of agents D_a, D_b, E_a , and E_b is obtained by iteratively performing two steps: *i*) determine

the time of the next reaction and *ii*) determine which is the reaction that occurs and consequently update the macroscopic state D_a, D_b, E_a, E_b of the system. Since the execution time of chemical reactions is modeled by exponentially distributed times [15], we have that the time before the next occurrence of any reaction is also exponentially distributed. Specifically, this time is computed as the minimum of a set of exponentially distributed variables which is still exponentially distributed with a rate κ equal to the sum of the individual reactions rates (see lines 4-5 in Algorithm 1). The specific reaction that occurs is randomly determined with probabilities equal to the ratio between each reaction rate and the overall rate κ (see line 6). If Reaction (10) occurs, respectively Reaction (11), the outcome is uniquely determined. We have that the number E_a of agents exploring option a (E_b) decreases by one unit and the number D_a of agents disseminating opinion a (D_b) increases by one unit. However, if Reaction (8) occurs, respectively Reaction (9), the outcome is determined by an additional probabilistic experiment. The number D_a of agents (D_b) decreases by one unit, and the type of agents increasing by one unit is E_a with probability q_{aa} (q_{ba}) or E_b with probability q_{ab} (q_{bb}).

This additional step is required because Reactions (8-9) are in fact “meta-reactions” that expand in a larger reaction set having one entry for each possible configuration of an agent neighborhood during the application of the majority rule. Probabilities q_{aa}, q_{ab}, q_{ba} , and q_{bb} are the discrete equivalent of the switching probabilities in Equations (1-2) used for the continuous ODE model. In contrast to the Binomial distribution used in Section 4, in the discrete case we use an hypergeometric distribution, which yields probabilities

$$q_{aa} = \sum_{i=\lfloor (\mathcal{G}-1)/2 \rfloor}^{\mathcal{G}-1} \frac{\binom{D_a-1}{i} \binom{D_b}{\mathcal{G}-i-1}}{\binom{D_a+D_b}{\mathcal{G}-1}}, \quad (12)$$

$$q_{ab} = \sum_{i=0}^{\lfloor (\mathcal{G}-1)/2 \rfloor - 1} \frac{\binom{D_a-1}{i} \binom{D_b}{\mathcal{G}-i-1}}{\binom{D_a+D_b}{\mathcal{G}-1}}. \quad (13)$$

Probabilities q_{aa} and q_{ab} are a discrete integration of an hypergeometric distribution, whereby D_a and D_b are the number of success states and failure states in the population, $D_a + D_b$ is the population size, $\mathcal{G} - 1$ is the number of trials, and i the actual number of successes. The expressions for probabilities q_{bb} and q_{ba} can be obtained by swapping the number of successes i with the number of failures $\mathcal{G} - i - 1$ in Equations (12-13).

After simulating trajectories of the master equation using the Gillespie algorithm, we compute the exit probability E_N , that is, the probability that a swarm of N agents reaches consensus on opinion a , and the average consensus time T_N , that is, the time necessary to reach consensus on any option. In all of our studies, we use the nominal parameters that characterized the robot experiments: $N = 100$, $\sigma = 6.072$, $g = 8.4$, and $\rho_a = 1$. In the reminder of this section, we validate the chemical reaction network model against the results of the robot experiments, we perform a thorough analysis of the speed versus accuracy trade-off (as we did for the ODE model in Section 4.3), and we compare the proposed DMMD strategy against the voter model previously described in [44].

Fig. 10 The figure shows the comparison between robot experiments (box-plots) and the predictions of the chemical reaction network in Equations (8–11) approximated by the Gillespie algorithm (shaded areas). The shaded areas correspond to a confidence region computed using the 25th and the 75th percentiles of 1000 independent executions of Algorithm 1 with time rescaled according to $t' = 3t + g$. a) reports the robot scenario with $\mathcal{G}_{\max} = 5$ compared with the Gillespie algorithm with $\mathcal{G} = 5$ and b) the robot scenario with $\mathcal{G}_{\max} = 25$ compared with the Gillespie algorithm with $\mathcal{G} = 9$. Parameters: $N = 100$, $\sigma = 6.072$, $g = 8.4$, $\rho_a = 1$, $\rho_b = 0.5$.

5.1 Validation Against Robot Experiments

We validate the chemical reaction network defined to model our decision-making strategy against the results of the robot experiments. The results of this validation are shown in Figure 10 as a function of time, where shaded areas provide a confidence region between the 25th and the 75th percentile predicted by the model and box-plots give the outcome of robot experiments. The results of the Gillespie algorithm depicted in the figure are obtained using the same time rescaling $t' = 3t + g$ used in Section 4.1 (refer to that section for the rationale). Remarkably, the predictions of the simulated chemical reaction network fit the robot experiments very well, for both group sizes $\mathcal{G} = 5$ (Figure 10a) and $\mathcal{G} = 9$ (Figure 10b). In contrast to the ODE model, the chemical reaction network can also very well predict the variance of the system. Both in data from robot experiments and in the prediction of the model, the variance is higher for intermediate values of time, and lower at the beginning and at the end of the execution of the system. Given that the Gillespie algorithm, that we recall approximates a chemical master equation, gives qualitatively correct predictions of the dynamics of our system, we can proceed to study the speed-accuracy trade-off in different regions of the parameter space.

5.2 Speed Versus Accuracy Trade-Off

The results of the speed versus accuracy analysis performed by approximating the chemical master equation are reported in two separate figures: Figure 11 reports the accuracy of the system by showing the exit probability E_N as a function of the group size $\mathcal{G}$ and of the initial condition $D_a(0)$; Figure 12 reports the convergence speed of the decision-making strategy by showing the time T_N necessary to reach consensus as a function of the same two parameters. Throughout this analysis we keep the same color notation used in the figures of Section 4.3 for the ODE model with the purpose to simplify the comparison of the results.

Fig. 11 The figure shows the results of the accuracy part of the speed versus accuracy analysis performed using the chemical reaction network defined by Reactions (8–11). The x-axis refers to the group size $G \in \{3, \dots, 25\}$ dimension, the y-axis to the initial condition $D_a(0) \in \{0, \dots, N\}$, $D_b(0) = N - D_a(0)$ dimension, while the column of the plot refers to two different problem difficulties, encoded in the option quality $\rho_b \in \{0.5, 0.9\}$. The heatmaps show: the exit probability E_N to reach consensus on option a , respectively a) for $\rho_b = 0.5$ and b) for $\rho_b = 0.9$. Panels c) and d) are zoomed-in versions of Panels a) and b) focusing on the region that separates the two basins of attraction. In all panels, red encodes $E_N = 1$ while blue encodes $E_N = 0$. Finite-size predictions have been approximated using 2.5×10^4 independent executions of the Gillespie algorithm for each point. Parameters: $\sigma = 6.072$, $g = 8.4$, $N = 100$.

For what concern accuracy, the outcome of the analysis with the Gillespie algorithm is in accordance with that obtained with the ODE model: the system is more accurate for lower values of the group size G , particularly for easier decision-making problems ($\rho_b = 0.5$, Figures 11a–11c). The main difference between the continuous and finite-size analysis is that in the latter case we do not have anymore a clear border dividing the two basins of attractions for different consensus decision. In contrast, we obtain a border that gather all points having equal probability to converge to either option ($E_N = 0.5$). Under this line, the probability to converge to option a smoothly decreases to 0, above this line, it increases to 1 (Figures 11c–11d). This behavior is a direct consequence of finite-size effects modeled by the chemical reaction network and ignored in the ODE model. Where the ODE model predicts a macroscopic state of indecision, we have instead that the system converges anyway to a consensus in the finite-size model. Additionally, the results in Figure 11 also show the same pattern in the decision accuracy that we

Fig. 12 The figure shows the results of the speed part of the speed versus accuracy analysis performed using the chemical reaction network defined by Reactions (8–11) for a swarm of $N = 100$ agents as a function of the group size $G \in \{3, \dots, 25\}$, the initial condition $D_a(0) \in \{0, \dots, N\}$, $D_b(0) = N - D_a(0)$, and option quality $\rho_b \in \{0.5, 0.9\}$. The heatmaps in the upper row show the average consensus time T_N (min) for group size $G \in \{3, 25\}$ and initial condition $d_a(0) \in [0, 1]$, respectively, a) for $\rho_b = 0.5$ and b) for $\rho_b = 0.9$. Black solid lines represent the border points c for each value of $\mathcal{G}$. In the lower row, figures show the consensus time over initial conditions $d_a(0)$ for group size $G \in \{5, 9, 17\}$, respectively, c) for $\rho_b = 0.5$ and d) for $\rho_b = 0.9$. Red and blue lines represent initial conditions with consensus on option a and option b . Finite-size predictions have been approximated using 2.5×10^4 independent executions of the Gillespie algorithm for each point. Parameters: $\sigma = 6.072$, $g = 8.4$, $N = 100$.

observed with the ODE model where even groups $\mathcal{G}$ were less accurate than odd groups (cf. Section 4.3).

The results of the analysis of the convergence speed are shown in Figure 12. In agreement with the prediction obtained with the ODE model, we have that the system is faster for higher values of the group size $\mathcal{G}$. This is particularly true for initial conditions that are closer to the state of indecision (i.e., the border line for $E_N = 0.5$) as shown in Figure 12a and Figure 12b. The primary difference between the current finite-size analysis and that of the continuous approximation in Section 4 is evident when looking at the shape of the consensus time as a function of the initial condition. Figure 12c and Figure 12d show these results: the curve of the consensus time T_N has a much smoother shape around the point of indecision as opposed to the exponentially increasing curves shown in Figure 9c and Figure 9d. Additionally, we also observe that the value of T_N predicted by the chemical reaction network is lower than that predicted by the ODE model for

Fig. 13 The figure shows the results of the speed versus accuracy comparison between the majority rule and the voter model. Data shown have been generated using Algorithm 1 for the majority rule and the Gillespie algorithm in [44] for the voter model. Results are given for a swarm of $N = 100$ agents as a function of the initial condition $D_a(0) \in \{0, \dots, N\}$, $D_b(0) = N - D_a(0)$ and option quality $\rho_b \in \{0.5, 0.7, 0.9, 0.99\}$. In a), we show the difference $E_N^{VM} - E_N^{MR}$ between the exit probability of the voter model, E_N^{VM} , and that of the majority rule, E_N^{MR} . In b), we show the time ratio T_N^{VM}/T_N^{MR} between the time necessary to reach consensus on any option using the voter model, T_N^{VM} , and using the majority rule, T_N^{MR} . Red horizontal dotted lines represent equivalent points, respectively, $E_N^{VM} - E_N^{MR} = 0$ for a), and $T_N^{VM}/T_N^{MR} = 1$ for b); symbols (square, diamond) visualize the configuration corresponding to the robot scenario. Parameters: $N = 100$, $\sigma = 6.072$, $g = 8.4$, $\mathcal{G} = 5$.

almost all initial conditions. As already mentioned above, the discrepancies in the behavior of the system are a direct consequence of the finiteness of the swarm size.

Overall, the speed versus accuracy analysis presented here provides the same message as the continuous analysis in Section 4.3. By increasing the group size $\mathcal{G}$, the swarm benefits in terms of convergence speed at the cost of a lower decision accuracy. This loss in accuracy is stronger in easier decision-making problems (e.g., $\rho_b = 0.5$) while it is mitigated in more difficult discrimination problems (e.g., $\rho_b = 0.9$). Conversely, the benefits concerning convergence speed obtained by increasing the group size are relatively unaffected by the difficulty of the problem. With respect to the continuous approximation provided by the ODE model, the analysis of the chemical master equations allowed us to better quantify the performance of the system by catching the stochastic effects resulting from a finite swarm.

5.3 Comparison with the Voter Model

We conclude the analysis of finite-size effects by performing again a speed versus accuracy study but this time we compare the performance of the DMMD strategy against the one of the weighted voter model [44]. Let us recall that the primary difference between the two decision-making strategies is given by the decision rule utilized by individual agents—in the voter model agents decide by copying the opinion of a random neighbor. The difference in decision accuracy between the two strategies is reported in Figure 13a for increasingly difficult decision-making problems ($\rho_b \in \{0.5, 0.7, 0.9, 0.99\}$). The majority rule is less accurate than the voter model when lines are greater than 0, is equally accurate for values equal to 0,

and more accurate for values greater than 0. As it can be noticed, the voter model is in general more accurate than the majority rule for initial conditions $D_a(0) < 50$. Conversely, the accuracy of the majority rule reaches that of the voter model for initial conditions $D_a(0) > 50$ and it even outperforms the voter model for the hardest decision problem ($\rho_b = 0.99$). This behavior is a direct consequence of the stability analysis provided at the end of Section 4.2: since the voter model has only one asymptotically stable state, when considering finite-size effects, we have that its dynamics converge with high probability to the best option for a larger set of initial conditions (cf. Valentini et al. [44]). In contrast, the dynamics of the majority rule strongly depend on the initial conditions. Figure 13b reports instead the comparison between the two strategies in terms of convergence time depicted as the ratio between the consensus time of the voter model, T_N^{VM} , over that of the majority rule, T_N^{MR} . We can observe that the majority rule considerably speeds up the decision-making process for all considered parameters. The difference in speed ranges from almost two-fold for the easy problem with $\rho_b = 0.5$ up to twenty-fold for the difficult problem with $\rho_b = 0.99$.

The above analysis shows that, for application scenarios where the time available to reach a collective decision is critical, such as the Kilobot scenario described in Section 3, the majority rule proposed in this paper is a much more practical design solution than the voter model. In our specific scenario, by using the majority rule instead of the voter model, we obtained a speed up of 1.89 while keeping the accuracy at the same level as in the voter model. This key difference is extremely relevant when considering the limitations in energy autonomy of the available robotic platform. Although the voter model is more accurate than the majority rule for most initial conditions $D_a(0) < 50$, this difference is considerably reduced for values of $D_a(0) \approx 50$ and vanishes for $D_a(0) > 50$. That is, when the designer has means to initialize the swarm with an approximately uniform distribution of opinion, the accuracy of the majority rule is close to that of the voter model.

6 Discussion and related work

The collective decision-making strategy and the scenario studied in this paper are inspired by the collective behavior of social insects, such as ants and honeybees [20, 7, 37, 39]. Specifically, the scenario was inspired by the site selection problem often faced by honeybee swarms [7, 37], and was tackled by a swarm of 100 Kilobots [31]. The same and similar robots have successfully been used in swarms sized up to thousands to complete tasks such as aggregation [16], collective transport [32] and self-assembly [33]. However, the site-selection scenario discussed here and in our previous paper [45], is the first experiment in which a large swarm of robots has tackled a collective discrimination problem. The decision-making strategy that we have designed included a modulation of positive feedback mechanism that is loosely inspired by the waggle dance communication language observed in honeybees [8].

Collective decision-making systems have already received substantial attention from the engineering community. For example, the control theory community has intensively studied the problem of consensus achievement [14, 29, 30, 23, 34]. However, this research line mostly focuses on continuous decision-making problems, that is, problems with infinite number of alternatives which do not require discrimination based on quality assessment. Differently, we focus on engineering lit-

erature that has tackled discrete discrimination problems in the remaining of this section. We organize this literature into two categories. In the first category, we discuss decision-making strategies that are related to our work but do not directly apply to the scenario considered in this paper. These strategies do not include a mechanism for the direct modulation of positive feedback. Instead, features of the environment (e.g., the length of a path in a shortest-path problem) are used to indirectly bias the collective decisions towards consensus on the best option. Such strategies are not easily transferred to scenarios that are relevant here. In the second category, the studied strategies include a direct modulation mechanism of positive feedback as done in this paper. This modulation is used to engineer the bias towards the best option directly in the controller.

In the first category, we find works on collective decision-making inspired by the aggregation behavior of cockroaches [10, 3]. In these works, a specific feature of the environment—the size of the aggregation site—represents the quality of each option. The authors implement the decision-making process by two individual decision rules (the first proposed in [10] and the second in [3]) that allow agents to adjust the probability of staying within a site as a function of the site’s area. Free parameters of the strategy are determined empirically or through a genetic algorithm, and are thus environment-specific. In [24, 35, 41], the authors use a majority rule-based strategy that does not include direct modulation of positive feedback. The goal of the swarm is to find the shortest path connecting the starting location to a site. When agents meet in the starting location, they form teams of 3 agents and they apply the majority rule. However, differently from our strategy, the dissemination time does not depend on the quality (path length) of the option because agents are assumed to be incapable of measuring the path length. Instead, modulation is indirectly provided by the asymmetry in the environment: the shorter the distance of a particular path, the more frequently agents return to the nest and disseminate the corresponding option. The work in [36, 2] uses a similar, environmentally-provided mechanism for the modulation of positive feedback. Differently from [24], the path selection is specifically represented by a double-bridge experiment and the majority rule is substituted by the so called k -unanimity rule. When using the k -unanimity rule, an agent switches to a particular option only after observing it k times in a row in other neighboring agents. Similarly to the group size $\mathcal{G}$ in our strategy, the parameter k can be used by the designer to regulate the speed and the accuracy of the collective decision.

In the second category, we describe scientific contributions that make use of a mechanism that allows agents to actively and directly modulate positive feedback. Reina et al. [28] developed a decision-making strategy inspired by mathematical models that generalize the collective decision-making behavior of social insects (ant colonies, honeybee swarms) and that of neurons in vertebrate brains [20]. Their modeling and design method is sufficiently generic to model either environmentally-induced or internally-designed positive feedback modulation. However, the authors focused so far only on the study of environmentally-induced modulation mechanism while no experiment have been performed using direct modulation as done in this paper. The main differences with respect to our approach are that their decision-making strategy allows agents to be in an uncommitted state (not favoring any of the two options) and that the individual decision rule implements recruitment and cross-inhibition. In Parker and Zhang [25], the authors consider the best-of- n decision-making problem in an aggregation task inspired by

the house-hunting behavior observed in ant colonies. In contrast to the majority rule used here, agents directly recruit other agents and perform explicit search and commitment phases. The decision-making strategy also includes a mechanism for consensus awareness by the swarm that is implemented following a quorum-sensing procedure. Finally, in our earlier work on the voter model [44] we proposed a decision-making strategy that utilizes the same mechanism for the direct modulation of positive feedback as here. The weighted voter model can be considered as a simplified version of the algorithm proposed by Parker and Zhang [25] where the initial exploration of the environment and the quorum-sensing mechanism are neglected. This feature allowed us to simplify the comparison between the voter model/recruitment rule and the majority rule presented above.

7 Conclusions and Future Work

In this paper, we have proposed a collective decision-making strategy—direct modulation of majority-based decisions (DMMD)—that allows a swarm of agents to choose the best of two options. Following the DMMD strategy, individual agents in the swarm couple the use of the majority rule with a mechanism for direct modulation of positive feedback to implement collective agreement for the best option. Specifically, agents iteratively alternate periods of opinion dissemination, where they advertise their preference for particular alternatives of the problem, with periods of exploration, where they gather information from the environment concerning the quality of their current opinion. The information gathered from the environment is utilized by individual agents to modulate their efforts of opinion promotion, that is, amplifying or reducing the time spent in the dissemination state during which they advertise a particular option. At the end of the dissemination period, agents reconsider their current opinion by adopting the opinion favored by the majority of their neighbors. This coupling of positive feedback modulation and the majority rule introduces a bias in the agents opinions that steers the swarm towards a collective decision for the best option.

We have shown that the DMMD strategy can be successfully implemented to let a swarm of 100 Kilobots tackle a binary foraging/site-selection problem. We have validated our decision-making strategy by performing more than 20 independent repetitions, equivalent to ≈ 35 hours of robot experiments. The results of the robot experiments prove that: *i.* the DMMD strategy has sufficiently low requirements that allow its implementation on robots with very limited perception and actuation capabilities; *ii.* it is fast enough to implement a feasible collective decision-making process within the robots' limited energy autonomy; and *iii.* it is robust to failures of real robotic hardware.

Along with robot experiments, we have defined a mathematical framework to analyze the performance of the decision-making strategy over broader regions of the parameter space. We have investigated the limiting dynamics ($N \rightarrow \infty$) of the DMMD strategy using an ordinary differential equation model and finite-size effects ($N < \infty$) using a chemical reaction network approximated with the Gillespie algorithm. Both mathematical models have been validated against data from robot experiments showing good qualitative agreement based on the mere requirement of linear time rescaling ($t' = 3t + g$). Using this mathematical framework, we proved that consensus decisions are the only asymptotically stable solutions of

the system. We investigated the trade-off between the convergence speed and the decision accuracy that arises when varying the average neighborhood size of agents applying the majority rule. The primary result of this analysis is that quicker collective decisions can be obtained with larger neighborhood sizes (i.e., higher robot swarm densities) at the cost of a lower probability to reach the optimal decision. Additionally, we observed that the parity of the group of opinions utilized in the majority rule influences this trade-off as well, with odd group sizes having greater chances to choose the best option than even groups. Finally, we compared the performance of the DMMD strategy with those of the weighted voter model [44] which was previously proposed to address the same scenario. With respect to the voter model, the use of the majority rule speeds up the decision-making process considerably (e.g., $1.89\times$ speed-up in the considered robot scenario) although it is characterized by a lower accuracy in all but the harder decision-making problems.

In our study we have focused on the minimal case of binary decision-making problems. This simplification allowed us to perform an extensive analysis of the proposed strategy based on data from robot experiments and based on the mathematical analysis of continuous and finite-size models. Nonetheless, the DMMD strategy already applies to the general case of more than two options as discussed in [45]. Currently, we are extending our entire mathematical framework in order to account for this generalized scenario and we plan to keep validating our results through large-scale robot experiments. An additional simplification in our study is the symmetry between options in relation to the environment. As reviewed in Section 6, asymmetries in the environment might introduce a bias in the decision-making process. In cases when this bias affects the dissemination of the best option negatively, an individual agent’s direct modulation of opinion promotion might not suffice to implement convergence on the best option. As a future line of research, we plan to analyze the robustness of the modulation mechanism to asymmetries in the environment. In addition, we plan to extend the proposed strategy with a more powerful modulation mechanism that is capable to balance a potentially negative influence of the environment on the decision-making process.

Acknowledgements This work has been partially supported by the European Research Council through the ERC Advanced Grant “E-SWARM: Engineering Swarm Intelligence Systems” (contract 246939) and by the EU-H2020-FET project ‘flora robotica’, no. 640959. Marco Dorigo acknowledges support from the Belgian F.R.S.–FNRS. Eliseo Ferrante acknowledges support from the Fund for Scientific Research (FWO), Flanders, Belgium.

A Appendix

We performed additional robot experiments to estimate the values of the average group size $\mathcal{G}$ in the two settings and that of the time σ^{-1} necessary to exploration a site. Each Kilobot records internally its series of exploration times and that of the number of neighbors at decision time. After an entire experiment, the acquired data is downloaded from the robots using a wired connection. We had to limit the number of experiments for data acquisition because it is very time-consuming. Specifically, we performed four runs: two runs to measure the actual average group size in the two settings ($\mathcal{G}_{\max} = 5$ and $\mathcal{G}_{\max} = 25$); and two runs to measure the average time required to complete the exploration of a site, again in the two settings.

Figure 14a reports the probability mass function of the actual neighborhood size $P(\mathcal{G})$ estimated from a single experimental run for each setting. When $\mathcal{G}_{\max} = 25$ (purple histograms, 652 samples), the average group size estimated using this mass function was 8.57, while it was 4.4 for $\mathcal{G}_{\max} = 5$ (green histograms, 682 samples). We have therefore a difference of almost a

Fig. 14 The figure shows the results from the statistical analysis performed on the additional robot experiments: a) probability mass function of the group size G for the two parameter settings $G_{\max} = 5$ (green) and $G_{\max} = 25$ (purple); b) probability density function of the time σ^{-1} necessary for a robot to explore a site (dotted line gives the average value of σ^{-1}).

factor of two between the two averages. Concerning the exploration time, we have graphically shown in [45] that the probability density functions resulting from $G_{\max} = 5$ (504 samples) and $G_{\max} = 25$ (602 samples) were very similar (as one would expect). To further investigate this point, we performed a two-sample Kolmogorov-Smirnov test. The null hypothesis that the two samples come from a different distribution could not be rejected (p-value = 0.5364), which supports our original conclusion that data sets are consistent with each other. We therefore merged the two data sets to improve our estimate of the exploration time. Figure 14b shows the probability density function of the time σ^{-1} (where σ is a rate, see Section 4) a robot spends to complete the exploration of a site. The average exploration time is 6.072 min (dotted line).

References

1. Brambilla, M., Ferrante, E., Birattari, M., Dorigo, M.: Swarm robotics: a review from the swarm engineering perspective. *Swarm Intelligence* **7**(1), 1–41 (2013)
2. Brutschy, A., Scheidler, A., Ferrante, E., Dorigo, M., Birattari, M.: “Can ants inspire robots?” Self-organized decision making in robotic swarms. In: *IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS)*, pp. 4272–4273. IEEE Press (2012)
3. Campo, A., Garnier, S., Dédriche, O., Zekkri, M., Dorigo, M.: Self-organized discrimination of resources. *PLoS ONE* **6**(5), e19,888 (2010)
4. Clifford, P., Sudbury, A.: A model for spatial conflict. *Biometrika* **60**(3), 581–588 (1973)
5. Deffuant, G., Neau, D., Amblard, F., Weisbuch, G.: Mixing beliefs among interacting agents. *Advances in Complex Systems* **3**(01n04), 87–98 (2000)
6. Franks, N.R., Dornhaus, A., Fitzsimmons, J.P., Stevens, M.: Speed versus accuracy in collective decision making. *Proc. R. Soc. Lond. B* **270**, 2457–2463 (2003)
7. Franks, N.R., Pratt, S.C., Mallon, E.B., Britton, N.F., Sumpter, D.J.T.: Information flow, opinion polling and collective intelligence in house-hunting social insects. *Philosophical Transactions of the Royal Society of London. Series B: Biological Sciences* **357**(1427), 1567–1583 (2002)
8. von Frisch, K.: The dance language and orientation of bees. Harvard University Press (1967)
9. Galam, S.: Majority rule, hierarchical structures, and democratic totalitarianism: A statistical approach. *Journal of Mathematical Psychology* **30**(4), 426–434 (1986)
10. Garnier, S., Gautrais, J., Asadpour, M., Jost, C., Theraulaz, G.: Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior* **17**(2), 109–133 (2009)
11. Garnier, S., Gautrais, J., Theraulaz, G.: The biological principles of swarm intelligence. *Swarm Intelligence* **1**(1), 3–31 (2007)

12. Gillespie, D.T.: Exact stochastic simulation of coupled chemical reactions. *The Journal of Physical Chemistry* **81**(25), 2340–2361 (1977)
13. Hamann, H.: Towards swarm calculus: Urn models of collective decisions and universal properties of swarm performance. *Swarm Intelligence* **7**(2-3), 145–172 (2013)
14. Hatano, Y., Mesbahi, M.: Agreement over random networks. *IEEE Transactions on Automatic Control* **50**(11), 1867–1872 (2005)
15. van Kampen, N.G.: *Stochastic processes in physics and chemistry*. Elsevier, Amsterdam, NL (1992)
16. Kernbach, S., Thenius, R., Kernbach, O., Schmickl, T.: Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic system. *Adaptive Behavior* **17**(3), 237–259 (2009)
17. Lerman, K., Martinoli, A., Galstyan, A.: A review of probabilistic macroscopic models for swarm robotic systems. In: E. Şahin, W. Spears (eds.) *Swarm Robotics, Lecture Notes in Computer Science*, vol. 3342, pp. 143–152. Springer (2005)
18. Liggett, T.M.: Stochastic interacting systems: contact, voter and exclusion processes, *Grundlehren der mathematischen Wissenschaften*, vol. 324. Springer (1999)
19. List, C.: Democracy in animal groups: a political science perspective. *Trends in Ecology & Evolution* **19**(4), 168–169 (2004)
20. Marshall, J.A., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., Franks, N.R.: On optimal decision-making in brains and social insect colonies. *Journal of The Royal Society Interface* **6**(40), 1065–1074 (2009)
21. Martinoli, A., Easton, K., Agassounon, W.: Modeling swarm robotic systems: a case study in collaborative distributed manipulation. *The International Journal of Robotics Research* **23**(4-5), 415–436 (2004)
22. Martinoli, A., Ijspeert, A., Mondada, F.: Understanding collective aggregation mechanisms: From probabilistic modelling to experiments with real robots. *Robotics and Autonomous Systems* **29**(1), 51–63 (1999)
23. Mesbahi, M., Egerstedt, M.: *Graph theoretic methods in multiagent networks*. Princeton University Press (2010)
24. Montes de Oca, M., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., Dorigo, M.: Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. *Swarm Intelligence* **5**, 305–327 (2011)
25. Parker, C.A.C., Zhang, H.: Cooperative decision-making in decentralized multiple-robot systems: The best-of-n problem. *IEEE/ASME Transactions on Mechatronics* **14**(2), 240–251 (2009)
26. Parker, C.A.C., Zhang, H.: Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence* **4**, 199–220 (2010)
27. Passino, K.M., Seeley, T.D.: Modeling and analysis of nest-site selection by honeybee swarms: the speed and accuracy trade-off. *Behavioral Ecology and Sociobiology* **59**(3), 427–442 (2006)
28. Reina, A., Miletitch, R., Dorigo, M., Trianni, V.: A quantitative micro–macro link for collective decisions: the shortest path discovery/selection example. *Swarm Intelligence* **9**(2-3), 75–102 (2015)
29. Ren, W., Beard, R., Atkins, E.: A survey of consensus problems in multi-agent coordination. In: *Proceedings of the 2005 American Control Conference*, pp. 1859–1864 vol. 3. IEEE Press (2005)
30. Ren, W., Beard, R.W.: *Distributed consensus in multi-vehicle cooperative control: theory and applications*. Communications and control engineering. Springer (2008)
31. Rubenstein, M., Ahler, C., Hoff, N., Cabrera, A., Nagpal, R.: Kilobot: A low cost robot with scalable operations designed for collective behaviors. *Robotics and Autonomous Systems* **62**(7), 966–975 (2014)
32. Rubenstein, M., Cabrera, A., Werfel, J., Habibi, G., McLurkin, J., Nagpal, R.: Collective transport of complex objects by simple robots: Theory and experiments. In: T. Ito, C. Jonker, M. Gini, O. Shehory (eds.) *Proceedings of the 12th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '13*, pp. 47–54. IFAAMAS (2013)
33. Rubenstein, M., Cornejo, A., Nagpal, R.: Programmable self-assembly in a thousand-robot swarm. *Science* **345**(6198), 795–799 (2014)
34. Sartoretti, G., Hongler, M.O., de Oliveira, M., Mondada, F.: Decentralized self-selection of swarm trajectories: from dynamical systems theory to robotic implementation. *Swarm Intelligence* **8**(4), 329–351 (2014)

35. Scheidler, A.: Dynamics of majority rule with differential latencies. *Phys. Rev. E* **83**, 031,116 (2011)
36. Scheidler, A., Brutschy, A., Ferrante, E., Dorigo, M.: The k-unanimity rule for self-organized decision-making in swarms of robots. *IEEE Transactions on Cybernetics* **In press** (2015)
37. Seeley, T.D.: *Honeybee Democracy*. Princeton University Press (2010)
38. Soysal, O., Şahin, E.: A macroscopic model for self-organized aggregation in swarm robotic systems. In: E. Şahin, W.M. Spears, A.F. Winfield (eds.) *Swarm Robotics, LNCS*, vol. 4433, pp. 27–42. Springer (2007)
39. Sumpter, D.J.: *Collective animal behavior*. Princeton University Press (2010)
40. Toral, R., Tessone, C.J.: Finite size effects in the dynamics of opinion formation. *Communications in Computational Physics* **2**(2), 177–195 (2007)
41. Valentini, G., Birattari, M., Dorigo, M.: Majority rule with differential latency: An absorbing Markov chain to model consensus. In: *Proceedings of the European Conference on Complex Systems 2012, Springer Proceedings in Complexity*, pp. 651–658. Springer (2013)
42. Valentini, G., Ferrante, E., Hamann, H., Dorigo, M.: Collective decision with 100 kilobots: Speed vs accuracy in binary discrimination problems (2015). URL <http://iridia.ulb.ac.be/supp/IridiaSupp2015-005/>. Supplementary material
43. Valentini, G., Hamann, H.: Time-variant feedback processes in collective decision-making systems: influence and effect of dynamic neighborhood sizes. *Swarm Intelligence* **9**(2-3), 153–176 (2015)
44. Valentini, G., Hamann, H., Dorigo, M.: Self-organized collective decision making: The weighted voter model. In: A. Lomuscio, P. Scerri, A. Bazzan, M. Huhns (eds.) *Proceedings of the 13th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '14*, pp. 45–52. IFAAMAS (2014)
45. Valentini, G., Hamann, H., Dorigo, M.: Efficient decision-making in a self-organizing robot swarm: On the speed versus accuracy trade-off. In: R. Bordini, E. Elkind, G. Weiss, P. Yolum (eds.) *Proceedings of the 14th International Conference on Autonomous Agents and Multiagent Systems, AAMAS '15*, pp. 1305–1314. IFAAMAS (2015)
46. Valentini, G., Hamann, H., Dorigo, M.: Self-organized collective decisions in a robot swarm. In: *Proceedings of the 29th AAAI Conference on Artificial Intelligence, AI Video Competition*. AAAI Press (2015). http://youtu.be/5lz_HnOLBW4
47. Vigeliuss, M., Meyer, B., Pascoe, G.: Multiscale modelling and analysis of collective decision making in swarm robotics. *PLoS ONE* **9**(11), e111,542 (2014)