Effects of Spatiality on Value-Sensitive Decisions Made by Robot Swarms

Andreagiovanni Reina, Thomas Bose, Vito Trianni and James A. R. Marshall

Abstract Value-sensitive decision-making is an essential task for organisms at all levels of biological complexity and consists of choosing options among a set of alternatives and being rewarded according to the quality value of the chosen option. Provided that the chosen option has an above-threshold quality value, value-sensitive decisions are particularly relevant in case not all of the possible options are available at decision time. This means that the decision-maker may refrain from deciding until a sufficient-quality option becomes available. Value-sensitive collective decisions are interesting for swarm robotics when the options are dispersed in space (e.g., resources in a foraging problem), and may be discovered at different times. However, current design methodologies for collective decision-making often assume a well-mixed system, and clever design workarounds are suggested to deal with a heterogeneous distribution of opinions within the swarm (e.g., due to spatial constraints on the interaction network). Here, we quantify the effects of spatiality in a value-sensitive decision problem involving a swarm of 150 kilobots. We present a macroscopic model of value-sensitive decision-making inspired by house-hunting honeybees, and implement a solution for both a multiagent system and a kilobot swarm. Notably, no workaround is implemented to deal with the spatial distribution of opinions within the swarm. We show how the dynamics presented by the robotic system match or depart from the model predictions in both a qualitative and quantitative way as a result of spatial constraints.

Andreagiovanni Reina, Thomas Bose and James A. R. Marshall
Department of Computer Science, University of Sheffield, S1 4DP, UK, e-mail: a.reina@sheffield.ac.uk,t.bose@sheffield.ac.uk,james.marshall@sheffield.ac.uk

Vito Trianni
ISTC, Italian National Research Council, Rome, Italy, e-mail: vito.trianni@istc.cnr.it

1 Introduction

Engineering large robot swarms with predictable performance is a very challenging problem, which is exacerbated by the spatiality aspects inherent to robotic systems that are widely distributed in space and that feature a highly heterogeneous and dynamic interaction network. For this reason, available design methods for system control resort to space-time models or low-dimensional abstractions [9, 3, 14]. Indeed, even with simple mobility models such as random walks [5], the system dynamics are the more difficult to predict the more the individual actions are influenced by information available only locally. If the state of a robot strongly depends on its spatial location (which in turns determines the interactions with neighbours), it is very likely that the swarm robotic system will present heterogeneities through space that may have a bearing on the macroscopic dynamics. The effects of spatiality are negligible only if the swarm is "well-mixed": in analogy with chemical systems [8], a certain robot state should be uniformly distributed within the swarm, or, in alternative, interactions between any two robots in the swarm should be equally likely. This condition is however not customary in swarm robotics, due to limited motion speed and local communication abilities that prevent sufficient mixing. As a result, the system dynamics may strongly deviate from the predictions of abstract macroscopic models [4, 24].

In collective decision-making problems, spatiality may be determinant for the system dynamics, especially when the decision is the result of the formation of spatial heterogeneities (e.g., in self-organised aggregation [1, 7]). In other cases, it can play against convergence to a coherent outcome due to the formation of spatially isolated clusters that do not sufficiently interact, resulting in a decision deadlock or in long convergence times [24]. The attentive design of the individual robot behaviour can cancel out or even exploit the effects of spatiality [15, 20, 25]. Design methodologies based on well-mixed assumptions propose clever workarounds to deal with spatial constraints, such as limiting the interaction between agents from different populations only when/where the agents populations mix, e.g., at a home location [21].

In this work, we address the design of a collective decision-making behaviour in a swarm robotics scenario characterised by spatial heterogeneities, and analyse the effects of spatiality in the system dynamics. The decision problem falls in the general class of *value sensitive decision-making* [17], that is, it requires that a decision is taken only if there is at least one option that has a sufficiently high quality (i.e., above a given threshold θ_v). In case such a high-quality option is not available, a decision should not be taken in the expectation that a supra-threshold option would become available at a later time (see Sect. 2). Due to spatiality, it may be possible that only low-quality alternatives are discovered first, due to random exploration, so that the decision should be delayed until a high-quality one is eventually found. Quantifying the effects of spatiality in such a decision problem is therefore a fundamental step toward the engineering of swarm robotics systems.

We provide an implementation for the value-sensitive decision problem following a recently introduced design pattern for decentralised decision-making [21], as detailed in Sect. 3. Although the design pattern provides some guideline to deal with spatiality—at least to a certain extent—, in this work we implement no special workaround apart from parametrising the system in a way to enhance mixing among robots. We study the system behaviour in the special case of binary symmetric decisions, that is, two options are available with the same quality $v \in [v_m, v_M]$, and we analyse when the system is able to break a decision deadlock or remain stuck at indecisions (i.e., in the expectation that a high-quality option would later appear, see Section 4). We provide results for an abstract multiagent system characterised by point-mass particles not physically interacting with each other, and for experiments with a swarm of 150 kilobots. We conclude in Sect. 5 by discussing the relevance of the obtained results for the engineering of large-scale swarm robotics systems.

2 Case study: value-sensitive decision-making by a robot swarm

Problem Description. We consider a case study in which a robot swarm must search and select among options deployed in a bidimensional environment, with an option becoming visible to a robot only in its immediate surroundings. The collective decision is taken when a large fraction x_Q of robots commits to the same option. Each option, O_i is characterised by its quality $v_i \in [v_m, v_M]$. The robots have no a priori information about the decision problem they have to solve, that is, the robots do not know: (i) how many options are available, (ii) where the options are located, and (iii) which is the quality of the options. The swarm is asked to explore the environment in order to identify all the available options and estimate their quality. Finally, the swarm must select the option with the highest quality if the quality is above a given threshold θ_v, otherwise should remain undecided.

The system property of committing or not to a decision as a function of the estimated quality of the options reflects value-sensitivity: The swarm response is *sensitive* to the value of the perceived options' quality. Value-sensitivity is relevant to engineering [21], biology [17] and also neuroscience [18], and is an advantageous property for systems that have to make decisions among an unknown number of options which have to be discovered. This is typical when the discovery of alternatives is an episodic event. In similar conditions, options may become available at various moments in time, for example, some options may need more time to be discovered or may appear later in the environment. In these cases, committing too early to the current best option may preclude the opportunity of selecting a better option that would get discovered later. This situation may be frequent in scenarios where options are deployed in a spatial environment (e.g., nest-sites in house-hunting social insects [23]). Farther options may take longer to get discovered although they might have a better quality than nearer options. This phenomenon is well-known in biology: for instance, ant colonies that change their nest are able to select the best new nest independently of the distance from the old nest [6].

Experimental Setup. The robot used in this study is the kilobot [22], a small and simple robot with limited sensing and actuation capabilities. A kilobot operates at a clock frequency of about 32 Hz (which corresponds to a clock period $\tau_c \simeq 31$ ms). It can move on a flat surface and control its movements through the modulation of the power applied to two vibration motors. The motion speed varies from robot to robot and also depends on the ground friction. In our experiments, a robot can move straight at a speed of 13 mm/s and rotates at 40 °/s, on average. Through IR communication, a robot can exchange 9 bytes messages with neighbours in a limited range d_s which varies in relation to the reflectance of the ground and the brightness of the environment. In our experiments, the communication range was about 100 mm. Finally, each robot is equipped with a RGB LED that we use to let the robot display their current state.

The environment is a circle with a radius of 750 mm and glass ground. Each option is signalled through two static kilobot robots acting as beacons (which hereafter we call simply *beacons* to differentiate them from the robots that compose the swarm). Each beacon broadcasts every second a message with the option ID and its quality. We consider a binary decision problem characterised by two options, and we assign a unique colour to each of them (i.e., red and blue); the colours will help later to visualise each robots commitment state (see Sect. 4). We allocate two beacons for each option to allow the robots to perceive the option messages in a larger area. The two beacons are located at a distance of 150 mm from each other and thus cover an area of $A_O \simeq 0.058\,\text{m}^2$. Therefore, each robot can perceive an option in a very limited portion of space compared to the whole environment which has an area of $A_E = 1.77\,\text{m}^2$. The two options are located at a distance of about 380 mm from the environment center at diametrical opposite position. Thus, the distance between the two options is about 760 mm.

We analyse the robot experiments postprocessing the video of the experiments that we record through four overhead cameras. The cameras have overlapping fields of view and have been calibrated in order to match the coordinates of each tracked robot in a common reference frame to remove duplicate detection. We use a four-camera tracking system to maximise the probability of detecting the LED colour of every robot. In fact, multiple reduced fields of view allow a view on each area that is more orthogonal to the ground plane and reduces the occurrence of robots occluding the view of the LED with their own body.

Given the specification of the environment and of the robots communication capabilities, the swarm size (i.e., the number of robots) determines the average number of neighbours for each robot. Given N robots moving in an environment of size A_E and interacting over a range d_s, the resulting interaction network has average degree $\langle k \rangle = \pi N d_s / A_E$ [24]. We size the swarm to 150 robots which corresponds to having a network with an average interaction degree of $\langle k \rangle \simeq 2.67$. This value allows to have frequent interactions but remains below the percolation threshold $\langle k \rangle \simeq 4,51$

which determines the emergence of a single connected component in the interaction network [24].

3 Top-down implementation through the design pattern

We implemented a swarm robotics system for decentralised decision-making following a design methodology based on the concept of design patterns, as proposed in [21]. The agent behaviour takes inspiration from a mathematical model for honeybee nest-site selection [23]. This model describes the dynamics of a honeybee swarm collectively deciding their future nesting site through a system of ODEs. In [21], the honeybee model is formalised in a design pattern supporting the design of swarm systems (e.g., robot swarms) by linking the macroscopic model parameters to the individual agent behavioural rules.

Preconditions. The decentralised decision strategy that we implemented requires a set of abilities at the individual robot level, as prescribed by the design pattern of [21]. Each robot must be able to:

- explore the environment searching for available options;
- recognise available options once found;
- individually estimate the options' quality;
- exchange with other robots the options' ID and quality.

Each robot modulates its actions as a function of the estimated quality of the opinion to which it is committed to. All these preconditions are met by the kilobot platform, hence the design pattern methodology can be applied.

Individual robot rules. The decision process works as follows. Robots can be committed to an option i (state C_i, and the total number of robots committed to i is N_i) or uncommitted (state C_U, and the total number of uncommitted robots is N_U). An uncommitted robot explores the environment and upon *discovery* of a potential option i it gets committed with probability P_{γ_i}. A robot committed to option i actively *recruits* uncommitted robots (which also become committed to i) with probability P_{ρ_i}. A robot committed to option i sends stop signals to robots committed to option j, with $j \neq i$. The robot that receives the stop signal becomes *inhibited* and reverts to an uncommitted state with probability P_{β_i}. Finally, a robot committed to option i spontaneously *abandons* its commitment and reverts to an uncommitted state with probability P_{α_i}.

Each robot, every second, broadcasts a message with information of its commitment state and, if committed to an option, the estimated option's quality. A robot updates its commitments state every $\tau_u = 400$ clock cycles through one of the four transitions: discovery, abandonment, recruitment and cross-inhibition. The mechanism of state update is described by the probabilistic finite state machine presented in Figure 1. Some of the transitions are available only if in the latest τ_u clock cycles

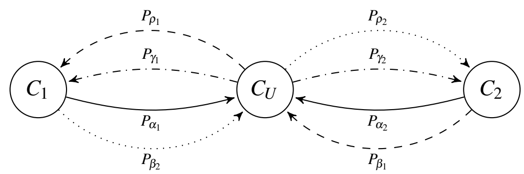

Fig. 1 Robot behaviour described as a probabilistic finite state machine. A robot updates its commitment state every τ_u clock cycles. Solid arrows are spontaneous individual transitions that can happen at any update. Dot-dashed lines are spontaneous individual transitions that can happen only if the robot has found an option in the latest τ_u clock cycles. Dashed and dotted arrows represent interactive transitions that can happen only if the latest message received by the robot in the last τ_u clock cycles is from another robot committed to option O_1 or O_2, respectively.

the robots has encountered an option (necessary for discovery) or received a message from other committed robots (necessary for recruitment and cross-inhibition).

Every robot moves in the environment through an isotropic random walk determined by the alternate sequence of straight motions for $\tau_m = 300$ clock cycles and on-place rotations in a random direction for τ_r clock cycles, which are chosen randomly from a uniform distribution $U(1, 150)$. The random walk is necessary (i) to let uncommitted robot explore the environment to discover potential options, and (ii) to allow robots to mix with each other and thus change their communication neighbourhood [5].

The robot software is available online at `http://diode.group.shef.ac.uk/resources/`.

Macroscopic parameterisation. The macroscopic model of the decision process can be described through a system of stochastic differential equations (SDEs) that describe the changes in the proportion of agents committed to each option. In this paper, we limit the study to binary decisions, hence the model describes the changes of the proportion of agents committed to option O_1 and O_2 as $x_1 = N_1/N$ and $x_2 = N_2/N$, where $N = N_U + N_1 + N_2$ is the total number of robots composing the swarm. The macroscopic model is:

$$\begin{cases} dx_1 = (\gamma_1 x_U - \alpha_1 x_1 + \rho_A x_1 x_U - \beta_2 x_1 x_2)dt + \sigma dW_1(t) \\ dx_2 = (\gamma_2 x_U - \alpha_2 x_2 + \rho_B x_2 x_U - \beta_1 x_1 x_2)dt + \sigma dW_2(t)\,, \\ x_U + x_1 + x_2 = 1 \end{cases} \tag{1}$$

where $\gamma_i, \alpha_i, \rho_i$ and β_i are the transition rates, respectively, of discovery, abandonment, recruitment and cross-inhibition for option i; and σdW_i is a Wiener process which represents additive noise with strength $\sigma \geq 0$. In order to implement a system able of value-sensitive decisions, we use a parameterisation similar to the one proposed in [17]. We make two modifications prescribed by the design pattern [21]: (i) The transition rate of discovery (γ_i) is changed in order to take into account the episodic nature of a discovery; (ii) All transition rates are scaled to meet the maximum speed of the process. The employed parameterisation is:

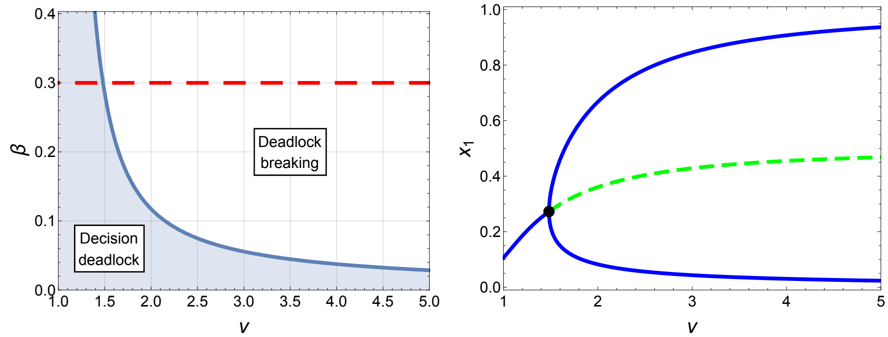

(a) (b)

Fig. 2 Analysis of the macroscopic model of Equation (1) with parameterisation (2) for the binary case $v = v_1 = v_2$. (a) Stability diagram in the parameter space (v, β). In the shaded area, the system has a single attractor with an equal number of committed robots to both options. In the white area, the system has two stable solutions with committed population biased for either of the two options. The horizontal red dashed line shows the selected value of $\beta = 0.3$. (b) Bifurcation diagram for $\beta = 0.3$ as a function of the option quality v. The solid blue lines are stable equilibria, the green dashed line is an unstable saddle point. As desired, the system undergoes a pitchfork bifurcation and breaks the decision deadlock for quality values greater than the treshold $\theta_v \simeq 1.5$.

$$\gamma_i = v_i \cdot P_D, \qquad \alpha_i = v_i^{-1}, \qquad \rho_i = v_i, \qquad \beta_i = \beta, \qquad dt = d\tau/s, \qquad i \in \{1,2\} \tag{2}$$

with P_D the episodic discovery probability, $v_i \in (1,5]$ the option i quality and $s = 0.008$ the temporal scaling. Similarly to [13], we estimate geometrically the probability of encountering an option, $P_D = A_O/A_E \simeq 0.033$, where A_O is the area where an option can be detected by a robot and A_E is the full environment area.

The macroscopic dynamics of the system in (1) can be studied analytically for varying option values [17]. A particularly interesting case is the symmetric condition in which both options have the same quality (i.e., $v_1 = v_2 = v$). In this situation, the two options are equivalent, therefore the swarm must select any of the two but only if their value is higher than θ_v. The study presented in this paper focuses on this case, as it is paradigmatic for evaluating the value-sensitivity property. Figure 2(a) shows the stability diagram as a function of the options' quality v and the cross-inhibition rate β. The diagram shows that there exist two zones determining the macroscopic behaviour: In the grey shaded area, there exist a single attractor corresponding to a decision deadlock, which means that the swarm remains locked at indecision; In the white area, the system presents two stable attractors that correspond to decisions for the one or the other option, and a third unstable saddle point. This diagram can be exploited to select the system parameterisation that provides a desired value-sensitive behaviour. For instance, in order to have a value-sensitive behaviour with $\theta_v = 1.5$, it is necessary to select $\beta = 0.3$ (displayed as a horizontal red dashed line in Figure 2(a)). Figure 2(b) shows the bifurcation diagram for the selected value

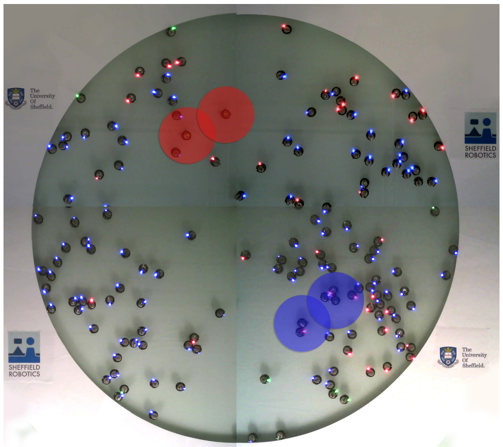

Fig. 3 Screenshot of a 150 kilobots experiment taken from the four-overhead-camera tacking system. The overlaying coloured circles show the two options localised in the environment. The robots light up their LED in a colour that corresponds to their internal commitment state: green for the uncommitted state C_U, red or blue for commitment to the option of the respective colour. This screenshot shows the final state of the swarm, after 30 min, for an experiment with $v_1 = v_2 = 5$. The swarm has a majority of robots committed to the blue option (108), only 35 robots committed to option red and 7 uncommitted robots.

of β, highlighting how the system breaks the decision deadlock when the options' quality v is greater than $\theta_v \simeq 1.5$.

Microscopic parameterisation. The design pattern of [21] explains how to convert the macroscopic parameters in the individual robot probabilities. The main difference between the macroscopic model and the agent implementation concerns the change of the temporal domain. While the macroscopic description is a continous time model, the robots operate, as the most part of electrical devices, in discrete time (i.e., CPU clock cycles). Following the conversion rules of [21], we obtain:

$$P_{\gamma_i} = \gamma_i \cdot T, \quad P_{\alpha_i} = \alpha_i \cdot T, \quad P_{\rho_i} = \rho_i \cdot T, \quad P_{\beta_i} = \beta_i \cdot T, \quad i \in \{1,2\} \tag{3}$$

with T the update timestep length, which is determined by the time between two commitment updates $\tau_u = 400\tau_c$ (where τ_c is the kilobot clock period) and the timescale of the macroscopic model, which is rescaled by the term s. In our experiments, $T = s\tau_u$.

4 Results

In this study, we constrast the dynamics predicted by the macroscopic models with the results obtained from the swarm robotics system which we implemented both in simulation and on real robots. At the macroscopic level, we study the system dynamics through time integration of Equation (1) via a generalised Heun method. Using this numerical scheme the computed solution converges to the Stratonovich solution of the SDEs [10]. Starting from the SDEs model, it is possible to derive the corresponding master equation, which allows the analysis of the finite size effects where the magnitude of the stochastic fluctuations is determined by the finite num-

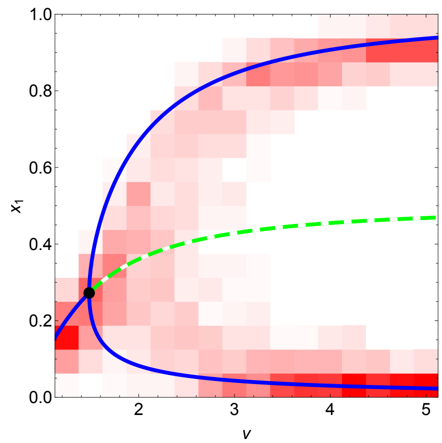

Fig. 4 Comparison between the predicted behaviour of the macroscopic system (as in Fig. 2(b)) and the multiagent simulations (underlaying density histogram). The results of the multiagent simulation correspond to the final distribution of population x_1 in 100 runs of length 5 hours.

ber of robots composing the swarm [21]. We approximate the solution of the master equation through the simulation of the Gillespie algorithm [8].

The swarm robotics system has been implemented and analysed both through multiagent simulations and through a 150 kilobot swarm. The multiagent system is implemented in MASON [12] and simulates point-mass particles that move in a bidimensional plane. The scenario, the agent behaviour and its parameterisation are coherent with the kilobot implementation described in Sections 2 and 3. However, in this system, noise and collisions among agents are not taken into account. The kilobot swarm implementation, instead, allows us to study the dynamics of a real physical system that includes all the aspects inherent to robotic experimentation. Figure 3 shows a screenshot of one experiment, while two videos of the robot experiments are available online[1]

A first assessment of the effects of spatiality can be obtained looking at the asymptotic dynamics of the system of Eq. (1) using the parameterisation (2), in comparison with the dynamics of the multiagent simulations. We performed a large set of simulations by extensively varying the quality value v, and Figure 4 contrasts the macroscopic bifurcation diagram as a function of the options' quality v with the final distribution of the simulated swarm at convergence. We can appreciate a good qualitative agreement of the multiagent dynamics with the stable attractors of Eq. (1). The existence of a bifurcation is well predicted by the macroscopic model, although the multiagent simulations appear to break the symmetry for slightly higher values of v. Figure 5 shows the comparison of the temporal dynamics of the macroscopic models and of the swarm robotics systems (both real and simulated). The plots show aggregated results of several experiments for two options' quality val-

[1] `https://www.youtube.com/watch?v=Gdy5o18y5lg` and `https://www.youtube.com/watch?v=EJtcpuj1Q5o`

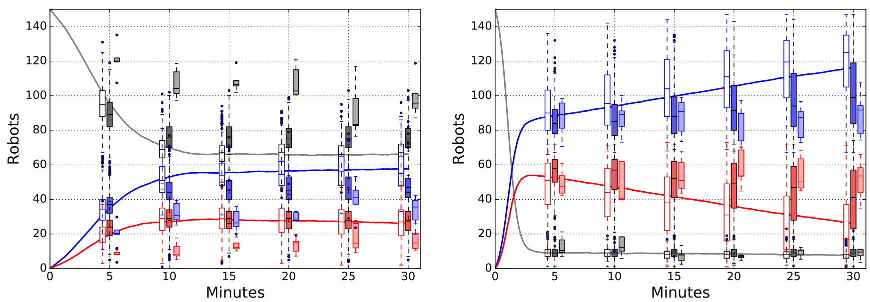

(a) $v = 1.5$ (b) $v = 5$

Fig. 5 Comparison of the temporal dynamics of the macroscopic models and the swarm robotics system. The solid lines show the dynamics of Equation (1) with noise strength $\sigma = 0.0032$ (average over 500 runs). The empty boxplots are the dynamics of the master equation describing the process with $N = 150$ agents (the solution is approximated with 500 runs of the Gillespie algorithm). The darker boxplots are the results of the spatial multiagent simulations (500 runs). Finally, the lighter boxplots are the results of 150-robots experiments (5 runs). Colours represent the three subpopulations: gray the uncommitted robots, blue/red the committed robots. Since both options have the same quality, the results show each time as blue the selected option and as red the *discarded* option.

ues: $v = 1.5$ and $v = 5$. The results show that both the multiagent simulations and the kilobot swarm have slightly different dynamics compared to the macroscopic description (SDE and Gillespie simulations). For $v = 1.5$, the system is close to the critical bifurcation point, and the dynamics are not quantitatively matched by the spatial system (see Fig. 5(a)). For $v = 5$ instead, the asymptotic behaviour is very similar, but the dynamics of the swarm robotics system are slower (see Fig. 5(b)). A slower convergence is the consequence of spatiality that leads to a drift of the system from a well-mixed condition. We observe indeed a slightly heterogeneous distribution of the commitment state among robots, with the emergence of clusters of robots with uniform commitment and a slow mixing of the two populations. Such a slowing-down effect is typical of consensus problems on regular lattices [2], and appears also with mobile agents in case the mobility pattern is not sufficient to produce an adeguate mixing [24]. Additionally, collisions among kilobots represent an additional factor that slows down the diffusion of robots in the environment, which justifies the slower dynamics detected during the kilobots experiments with respect to multiagent simulations, as shown in Fig. 5.

5 Discussions

The study we have presented highlights the complexity of dealing with spatiality in the engineering of a swarm robotics system. Although the specific decision problem we tackle seems ideal for obtaining a quantitative match between macroscopic models and experimental system—due to the agents/robots living (and mixing) in the same space—a non-negligible deviation from the model predictions is observed in both the multiagent simulations and the experiments with kilobots, especially in the transitory dynamics, which are slower for the spatial system. The design pattern for decentralised decision-making [21] is key for achieving a good qualitative match, as demonstrated with the extensive multiagent simulations we performed. Similarly to the results presented in this study, other work [26] obtained a qualitative match between spatial systems' dynamics and non-spatial mathematical models. However, a quantitative micro-macro link requires some additional workaround to better approximate a well-mixed system. For instance, the introduction of a distinction between latent and interactive agents as suggested in [21] could be key to allow a better mixing of the populations. In future work, we will implement such workaround and evaluate the extent to which the well-mixed condition is attained. More generally, effects of spatiality should be included in the macroscopic models, possibly resorting to heterogeneous mean-field approximations [16, 8]. In this way, the design pattern methodology could be enriched with tools to deal with spatially heterogeneous systems, and also for systems interacting on networks with heterogeneous topology (e.g., scale-free networks) [16]. Another possible approach to obtain a quantitative match of the system dynamics influenced by spatiality consists in including such factors into the macroscopic models [19, 9, 3].

An additional problem we recognise is given by collisions among robots, which further reduce mobility and limit mixing within the system, as already noted in previous studies [11, 24]. The kilobot platform does not provide means for efficient collision avoidance, and the high density that characterises experimentation with large groups plays against the population mixing. Indeed, besides being well-mixed, large-scale systems need also be "diluted" to ensure a good micro-macro link [8]. Efforts to provide guidelines to deal with less dilute systems will greatly benefit the engineering practice for swarm robotics systems. Indeed, we plan to investigate through further studies the effects of density on robot mobility in order to provide a model with diffusion coefficient as a function of the robot density.

The experimentation we performed in this study is limited to the symmetric binary decision case, which we deem sufficient to identify the relevant dynamics and compare with the model predictions. However, the implementation we provide is agnostic on the number of possible alternatives, and on the relative difference in quality, as already demonstrated in [21]. This means that the implemented system would work out-of-the-box also with an increasing number of alternatives. In such a best-of-n scenario, however, the macroscopic dynamics may be influenced by the number of available options, and the specific parameterisation we selected (i.e., the transitions rate strengths, as suggested by [17]) may need to be adjusted to produce value-sensitive decisions for any given number of options. Analytical studies in this direction are ongoing, and tests with a swarm robotics system will allow the validation of the design method beyond the binary case also for large-scale physical systems.

Acknowledgements This work was partially supported by the European Research Council through the ERC Consolidator Grant "DiODe: Distributed Algorithms for Optimal Decision-Making" (contract 647704). Vito Trianni acknowledges support from the project DICE (FP7 Marie Curie Career Integration Grant, ID: 631297). Finally, the authors thank Michael Port for the valuable help in building the infrastructure necessary to conduct the robot experiments.

References

1. J.-M. Amé, J. Halloy, C. Rivault, C. Detrain, and J.-L. Deneubourg. Collegial decision making based on social amplification leads to optimal group formation. *Proceedings of the National Academy of Sciences*, 103(15):5835–5840, 2006.
2. A. Baronchelli, L. Dall'Asta, A. Barrat, and V. Loreto. Topology-induced coarsening in language games. *Physical review E, Statistical, nonlinear, and soft matter physics*, 73(1):015102, 2006.
3. S. Berman, V. Kumar, and R. Nagpal. Design of control policies for spatially inhomogeneous robot swarms with application to commercial pollination. In *Proceedings of the 2011 IEEE International Conference on Robotics and Automation (ICRA)*, pages 378–385. IEEE Press, 2011.
4. N. Correll and A. Martinoli. Collective Inspection of Regular Structures using a Swarm of Miniature Robots. In *The 9th International Symposium on Experimental Robotics (ISER) (Springer Tracts in Advanced Robotics)*, volume 21, pages 375–385. Springer, Berlin, Germany, 2006.
5. C. Dimidov, G. Oriolo, and V. Trianni. Random walks in swarm robotics: an experiment with kilobots. In M. Dorigo et al., editor, *Proceedings of the 10th International Conference on Swarm Intelligence (ANTS 2016)*, volume 9882 of *LNCS*, pages 185–196. Springer Verlag, Berlin, Germany, 2016.
6. N. R. Franks, K. A. Hardcastle, S. Collins, F. D. Smith, K. M. Sullivan, E. J. Robinson, and A. B. Sendova-Franks. Can ant colonies choose a far-and-away better nest over an in-the-way poor one? *Animal Behaviour*, 76(2):323–334, 2008.
7. S. Garnier, C. Jost, J. Gautrais, M. Asadpour, G. Caprari, R. Jeanson, A. Grimal, and G. Theraulaz. The embodiment of cockroach aggregation behavior in a group of micro-robots. *Artificial life*, 14(4):387–408, 2008.
8. D. T. Gillespie, A. Hellander, and L. R. Petzold. Perspective: Stochastic algorithms for chemical kinetics. *The Journal of Chemical Physics*, 138(17):170901–15, 2013.
9. H. Hamann and H. Wörn. A framework of spacetime continuous models for algorithm design in swarm robotics. *Swarm Intelligence*, 2(2-4):209–239, 2008.
10. P. E. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Stochastic Modelling and Applied Probability*. Springer, Berlin, Germany, 1992.
11. K. Lerman and A. Galstyan. Mathematical Model of Foraging in a Group of Robots: Effect of Interference. *Autonomous Robots*, 13(2):127–141, 2002.
12. S. Luke, C. Cioffi-Revilla, L. Panait, K. Sullivan, and G. Balan. Mason: A multiagent simulation environment. *Simulation: Transactions of the society for Modeling and Simulation International*, 81(7):517–527, 2005.
13. A. Martinoli, K. Easton, and W. Agassounon. Modeling Swarm Robotic Systems: A Case Study in Collaborative Distributed Manipulation. *International Journal of Robotics Research*, 23(4):415–436, 2004. Special Issue on Experimental Robotics, B. Siciliano, editor.
14. N. Michael and V. Kumar. Control of ensembles of aerial robots. *Proceedings of the IEEE*, 99(9):1587–1602, 2011.
15. M. Montes, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, and M. Dorigo. Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3-4):305–327, 2010.

16. P. Moretti, S. Liu, A. Baronchelli, and R. Pastor-Satorras. Heterogenous mean-field analysis of a generalized voter-like model on networks. *The European Physical Journal B*, 85(3):1–6, 2012.
17. D. Pais, P. M. Hogan, T. Schlegel, N. R. Franks, N. E. Leonard, and J. A. R. Marshall. A mechanism for value-sensitive decision-making. *PLoS ONE*, 8(9):e73216, 2013.
18. A. Pirrone, T. Stafford, and J. A. R. Marshall. When natural selection should optimise speed-accuracy trade-offs. *Frontiers in Neuroscience*, 8(73), 2014.
19. A. Prorok, N. Corell, and A. Martinoli. Multi-level spatial modeling for stochastic distributed robotic systems. *The International Journal of Robotics Research*, 30(5):574–589, 2011.
20. A. Reina, R. Miletitch, M. Dorigo, and V. Trianni. A quantitative micro-macro link for collective decisions: The shortest path discovery/selection example. *Swarm Intelligence*, 9(2–3):75–102, 2015.
21. A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, and V. Trianni. A design pattern for decentralised decision making. *PLoS ONE*, 10(10):e0140950, 2015.
22. M. Rubenstein, C. Ahler, N. Hoff, A. Cabrera, and R. Nagpal. Kilobot: A low cost robot with scalable operations designed for collective behaviors. *Robotics and Autonomous Systems*, 62(7):966–975, 2014.
23. T. D. Seeley, P. K. Visscher, T. Schlegel, P. M. Hogan, N. R. Franks, and J. A. R. Marshall. Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064):108–11, 2012.
24. V. Trianni, D. De Simone, A. Reina, and A. Baronchelli. Emergence of consensus in a multi-robot network: from abstract models to empirical validation. *Robotics and Automation Letters, IEEE*, PP(99):1–1, 2016.
25. G. Valentini, E. Ferrante, H. Hamann, and M. Dorigo. Collective decision with 100 kilobots: Speed versus accuracy in binary discrimination problems. *Autonomous Agents and Multi-Agent Systems*, 30(3):553–580, 2016.
26. G. Valentini and H. Hamann. Time-variant feedback processes in collective decision-making systems: influence and effect of dynamic neighborhood sizes. *Swarm Intelligence*, 9(2–3):153–176, 2015.