

A quantitative micro–macro link for collective decisions: the shortest path discovery/selection example

Andreagiovanni Reina¹ · Roman Miletitch¹ ·
Marco Dorigo¹ · Vito Trianni²

Received: 1 December 2014 / Accepted: 4 May 2015 / Published online: 26 May 2015
© Springer Science+Business Media New York 2015

Abstract In this paper, we study how to obtain a quantitative correspondence between the dynamics of the microscopic implementation of a robot swarm and the dynamics of a macroscopic model of nest-site selection in honeybees. We do so by considering a collective decision-making case study: the shortest path discovery/selection problem. In this case study, obtaining a quantitative correspondence between the microscopic and macroscopic dynamics—the so-called micro–macro link problem—is particularly challenging because the macroscopic model does not take into account the spatial factors inherent to the path discovery/selection problem. We frame this study in the context of a general engineering methodology that prescribes the inclusion of available theoretical knowledge about target macroscopic models into design patterns for the microscopic implementation. The attainment of the micro–macro link presented in this paper represents a necessary step towards the formalisation of a design pattern for collective decision making in distributed systems.

Keywords Collective decision making · Micro–macro link · Shortest path selection · Swarm robotics · Design pattern

Electronic supplementary material The online version of this article (doi:[10.1007/s11721-015-0105-y](https://doi.org/10.1007/s11721-015-0105-y)) contains supplementary material, which is available to authorized users.

✉ Andreagiovanni Reina
areina@ulb.ac.be

Roman Miletitch
rmiletit@ulb.ac.be

Marco Dorigo
mdorigo@ulb.ac.be

Vito Trianni
vito.trianni@istc.cnr.it

¹ IRIDIA, Université Libre de Bruxelles, Brussels, Belgium

² ISTC, Italian National Research Council, Rome, Italy

1 Introduction

The goal of this study is to establish a quantitative correspondence between the dynamics of a microscopic implementation of a swarm robotics system and the dynamics of a macroscopic model of collective decision in honeybees. Our aim is to support a general methodology for the design of behaviours leading to efficient and flexible collective decisions with predictable dynamics, a frequently recurring problem in swarm robotics. Indeed, it is often the case that the solution to a given problem requires some form of consensus among the robots. For instance, coordinated motion, collective transport, aggregation, path selection and resource exploitation are all activities that presuppose a collective decision by the swarm (Trianni and Dorigo 2006; Garnier et al. 2009; Parker and Zhang 2009, 2010; Gutiérrez et al. 2010; Campo et al. 2010, 2011; Ferrante et al. 2012; Wilson et al. 2014). Therefore, sound methodologies for engineering distributed control strategies for collective decision making are of utmost importance.

Normally, given a problem that requires some collective decision, a specific solution is developed (Parker and Zhang 2009; Montes et al. 2010; Scheidler et al. 2015; Valentini et al. 2014; Wilson et al. 2014; Sartoretti et al. 2014). Inspiration from biological and social systems is customary, as it provides a microscopic, individual-level description of the behaviour that can be used as a reference. Then, the proposed solution is analysed in detail through the development of analytical models that reveal the main features of the macroscopic, system-level dynamics, along with various system properties, such as decision accuracy, efficiency, robustness, scalability, and so forth. Whenever the problem changes, the solution is adapted from previous work or invented anew. In both cases, the theoretical analysis of the system dynamics must be repeated, and the system properties re-assessed. In fact, a quantitative correspondence between microscopic and macroscopic descriptions (hereafter referred to as the micro–macro link, see also Hamann and Wörn 2008) is generally difficult to obtain. When starting from the microscopic description, difficulties are due to the several abstractions necessary to develop a treatable macroscopic model. As a result, the micro–macro link holds only qualitatively, and extensive tuning and sensitivity analyses are needed to find the appropriate working regime for the real system.

The problem-specific approach described above can lead to custom-tailored and possibly highly efficient solutions. However, it does not seem suitable as an engineering methodology, given that the macroscopic properties and the micro–macro link must be repeatedly verified for each new implementation. Given the large literature and the extensive theoretical knowledge available on collective decisions and opinion dynamics in social systems (Castellano et al. 2009; Couzin 2009; Vicsek and Zafeiris 2012; Arganda et al. 2012; Baronchelli et al. 2013; Kao et al. 2014), a better approach would be to start from well-understood macroscopic models—which somehow guarantee the attainment of the desired performance/properties—and to (automatically) generate an implementation of the robotic system in such a way as to obtain a one-to-one match with the target model. Unfortunately, obtaining a micro–macro link starting from the macroscopic description of a desired outcome is unworkable in the general case. Some preliminary work in this direction can be found in (Hamann and Wörn 2008; Kazadi 2009; Berman et al. 2009, 2011; Vigelius et al. 2014; Brambilla et al. 2015).

A compromise solution between problem-specific and generic design approaches can be given by resorting to design patterns, which can provide generic solutions to a specific class of problems. Design patterns are good engineering practices to face recurring design problems: they originate from the generalisation of repeatedly employed solutions, and they provide a formalised methodology to simplify system design and promote dissemination

of best practices. First introduced in the architectural domain (Alexander et al. 1977), they have been exploited in several disciplines including software engineering (Gamma et al. 1995). The exploitation of design patterns for the implementation of distributed systems has not received much attention, and only few proposals can be acknowledged (Babaoğlu et al. 2006; Gardelli et al. 2007). A design pattern for distributed systems must provide a set of implementation guidelines in support of the development of the microscopic behaviour, which leads to the desired macroscopic dynamics. Hence, the core of the design pattern comprises both a macroscopic model presenting desired system-level properties and a microscopic description of the individual behaviour sufficient for the implementation. Analytical tools are also provided for validation. Along with tools and guidelines, a design pattern must also provide examples that showcase its usage in simplified situations, as well as in particularly challenging working conditions that require extra care in the implementation.¹

In this paper, we take a step towards the development of a design pattern for collective decision making. In Sect. 2, we introduce a set of high-level implementation guidelines inspired by studies of honeybee nest-site selection, which describe the target macroscopic dynamics (Seeley et al. 2012; Pais et al. 2013). The main contribution of this paper consists in exploiting these guidelines to develop a swarm robotics system featuring an accurate quantitative micro–macro link with the chosen model. The attainment of such a link is a prerequisite to the formal characterisation of the design pattern for the general case of collective decisions in distributed systems, which is out of the scope of the present paper and is deferred to future work. As case study, we have chosen the shortest path discovery/selection problem (Gutiérrez et al. 2010; Montes et al. 2010; Scheidler et al. 2015). This problem requires the identification and collective selection of the shortest path between target areas, to be performed by a (possibly large) swarm of robots. Path discovery/selection is particularly demanding from a decision-making standpoint because spatial factors strongly influence the collective dynamics, as they determine the interaction patterns among the robots. As a consequence, finding a micro–macro link between a non-spatial macroscopic model of decision making and a strongly spatial decision problem is not trivial. A preliminary study and partial results have been presented by Reina et al. (2014) for a very abstract scenario involving agents moving in a one-dimensional space. Here, we present a complete study of a more realistic scenario for agents/robots moving on a flat surface.

In Sect. 3, we provide a solution to the shortest path discovery/selection problem both in an idealised multi-agent simulation that retains the relevant spatial factors but neglects physical interactions (see Sect. 3.2) and in a physics-based simulation of a swarm of e-pucks (Mondada et al. 2009), therefore accounting for limited swarm size and physical interferences (see Sect. 3.3). In Sect. 4, we analyse the implemented behaviour under a variety of different parameterisations and show that in every case there is a precise correspondence between the microscopic and macroscopic descriptions. We extend the analysis towards varying group size and also show the correspondence with Monte Carlo simulations of a macroscopic finite-size model. As discussed in the conclusions (see Sect. 5), the case study presented in this paper contributes to the ongoing formalisation of a design pattern for collective decision making by extending the high-level guidelines provided in Sect. 2 with specific recommendations for dealing with issues deriving from spatiality. Additionally, this paper also provides useful tools to verify the correctness of the implementation.

¹ For a more specific definition of design patterns in the distributed systems domain, the interested reader is referred to the work of Babaoğlu et al. (2006).

2 From macroscopic descriptions to implementation guidelines

As mentioned above, how collectives can achieve consensus is widely studied in many different contexts, and several models have been proposed in the literature (e.g. [Castellano et al. 2009](#); [Vicsek and Zafeiris 2012](#); [Kao et al. 2014](#)). In this paper, we selected a macroscopic model inspired by the nest-site selection behaviour observed in honeybees ([Seeley et al. 2012](#); [Pais et al. 2013](#)). This model has been selected because it possesses properties that are desirable in artificial distributed systems: (i) it attains near-optimal speed-accuracy trade-offs in the selection of the best option ([Marshall et al. 2009](#)), and (ii) it exploits adaptive mechanisms to tune decision speed and to break symmetry deadlocks (e.g. caused by same-quality options). In this section, we first discuss the macroscopic model, then we provide a high-level description of the implementation path prescribed by a design pattern to obtain a quantitative micro–macro link.

2.1 Collective decision through cross-inhibition

An analytical model of the nest-site selection process in honeybee colonies has been developed and confronted with empirical results, confirming the existence of both positive and negative feedback loops that determine the collective decision ([Seeley et al. 2012](#)). The model describes a decision-making process in which only two options are available, referred to as A and B . Each option i is characterised by an objective quality v_i . The collective decision problem consists in identifying and selecting the best option, or any of the equal-best options. The model treats a population of agents that can be either uncommitted (sub-population U with fraction Ψ_U of the total population) or committed to one of the two options (sub-populations A and B , respectively, with fraction Ψ_A and Ψ_B). Population dynamics can be easily described by a system of two coupled ordinary differential equations, plus a mass conservation term:

$$\begin{cases} \dot{\Psi}_A = \gamma_A \Psi_U - \alpha_A \Psi_A + \rho_A \Psi_A \Psi_U - \sigma_B \Psi_A \Psi_B \\ \dot{\Psi}_B = \gamma_B \Psi_U - \alpha_B \Psi_B + \rho_B \Psi_B \Psi_U - \sigma_A \Psi_A \Psi_B \\ \Psi_U + \Psi_A + \Psi_B = 1 \end{cases} \quad (1)$$

The variation of the population fraction Ψ_i , $i \in \{A, B\}$ results from four concurrent processes, which correspond to the four terms of each differential equation in (1):

- (i) Ψ_i increases as uncommitted individuals spontaneously *discover* and become committed to the option i at the rate γ_i ;
- (ii) Ψ_i decreases as individuals committed to option i spontaneously *abandon* it and get uncommitted at the rate α_i ;
- (iii) Ψ_i increases as individuals from population i actively *recruit* uncommitted ones at the rate ρ_i ;
- (iv) Ψ_i decreases as individuals from population i are *inhibited* by individuals of population $j \neq i$ at the rate σ_j .

All transition rates— γ_i , α_i , ρ_i , σ_i —are greater than zero. It is worth noting that this model does not require any explicit comparison of the option qualities. The quality value v_i of the two options is instead encoded in the transition rates (i.e. implementing a value-sensitive decision making, see [Pais et al. 2013](#)): different-quality options correspond to biased transition rates, while same-quality options correspond to unbiased transition rates. Overall, the collective decision is based purely on the system dynamics resulting from individual-to-individual interactions.

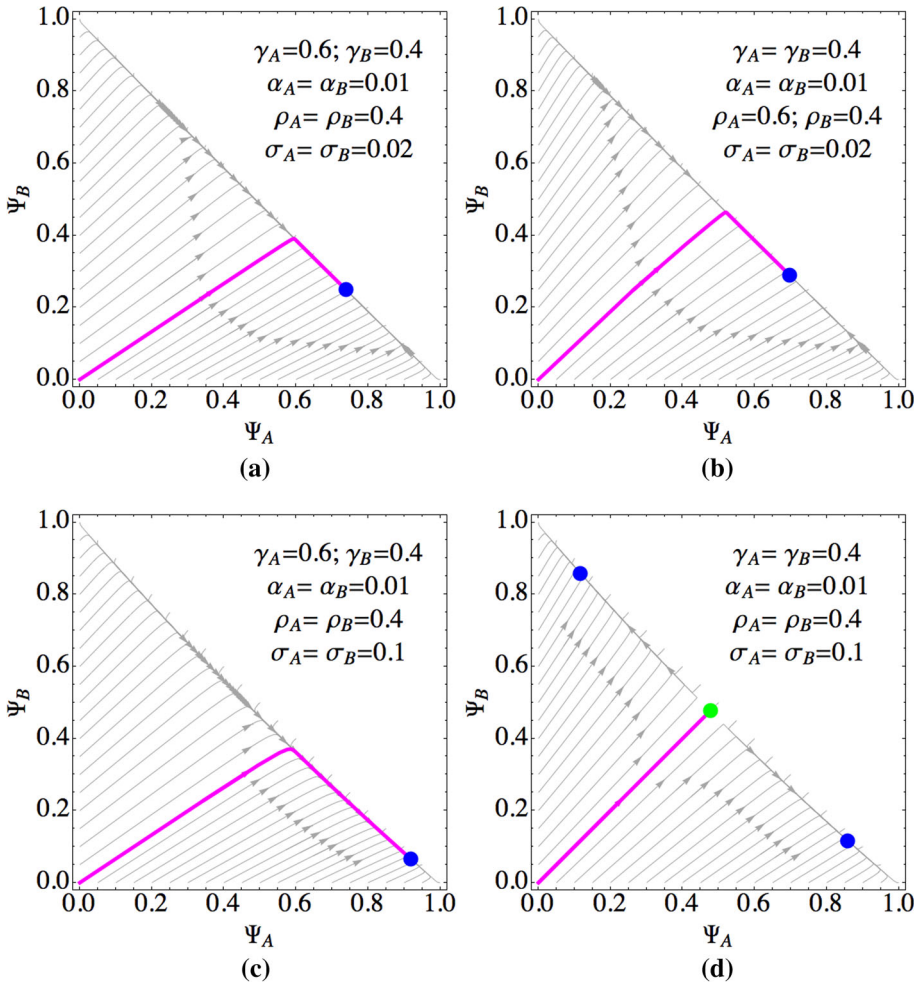


Fig. 1 (Colours online) Effects of different parameterisations on the macroscopic dynamics (trajectories and equilibrium points, shown as *dark (blue) dots* for stable points and *light (green) dots* for unstable saddle points). See text for details

2.2 Implementation guidelines

2.2.1 Working regime

The choice of the parameters of the macroscopic model determines the working regime for the decision-making process. Understanding the macroscopic dynamics leads to a principled choice of the desired parameterisation, which ultimately translates in prescriptions for the implementation.

When the two options have different qualities (e.g. $v_A > v_B$), a biased population distribution is obtained thanks to similarly biased commitment rates. Everything else being equal, a population distribution biased for the better option can be obtained thanks to a higher discovery rate (i.e. $\gamma_A > \gamma_B$, see Fig. 1a) or similarly through recruitment (i.e. $\rho_A > \rho_B$, see

Fig. 1b). Abandonment and cross-inhibition instead reduce the size of a population committed to a given alternative. Abandonment should be small enough to avoid that a large fraction of the population remains uncommitted, possibly biased towards the lower quality (i.e. $\alpha_A < \alpha_B$). Cross-inhibition instead, being proportional to the size of the inhibiting population, contributes to the creation of an unbalanced distribution of individuals between committed populations even for unbiased inhibition rates (i.e. $\sigma_A = \sigma_B$, see Fig. 1c). This is true also for same-quality alternatives (i.e. $v_A = v_B$). In this symmetric case, discovery, abandonment and recruitment are equal and are therefore not sufficient to break the symmetry. However, a sufficient level of cross-inhibition makes the equilibrium point unstable, therefore leading to a symmetry breaking, as shown in Fig. 1d. Through linear stability analysis, it is possible to identify the cross-inhibition level for which the system breaks the deadlock and converges to the choice of one option (see Seeley et al. 2012). The working region is $\{\rho > \alpha, \sigma > \sigma^*\}$, with critical value:

$$\sigma^* = \frac{4\alpha\gamma\rho}{(\rho - \alpha)^2}. \quad (2)$$

As a general guideline for the choice of the transition rates determining the working regime of the macroscopic model, it is advisable to have a parameterisation linked to the option quality, so that commitment is proportional to quality, abandonment inversely proportional, while cross-inhibition can be proportional to the option quality or independent of it, but in any case it must be sufficiently large to ensure convergence (for a detailed analysis, see Pais et al. 2013):

$$\gamma_i, \rho_i \propto v_i, \quad \alpha_i \propto 1/v_i, \quad \sigma_i > \sigma^*, \quad i \in \{A, B\} \quad (3)$$

2.2.2 Individual behaviour

Once a suitable working regime is identified, the four concurrent processes resulting in the macroscopic dynamics of (1) need to be implemented as a multi-agent system. As a general guideline, the agent behaviour should be implemented as the probabilistic finite state machine shown in Fig. 2. Here, the agent can be in three different commitment states that indicate whether the agent is uncommitted (C_U) or committed to either option A or B (C_A or C_B). Two types of transition need to be implemented: spontaneous transitions and interactive transitions. Spontaneous transitions correspond to discovery and abandonment and pertain to the individual agent behaviour. Interactive transitions depend instead on the result of the interaction among agents and are regulated by the probability of encountering agents of population i , which we refer to as P_{ψ_i} : the larger the proportion of agents committed to population i , the larger the probability of encountering one of its members. The transition probabilities between different states determine the outcome of the decision-making process and should be influenced by the option quality v_i to restrict the system dynamics within the working regime discussed above. Therefore, P_{γ_i} and P_{ρ_i} —respectively, the discovery and recruitment probability for option i —should be biased towards the option of higher quality; the abandonment probability P_{α_i} should be small and possibly inversely proportional to the option quality v_i ; finally, cross-inhibition should be governed by a high-enough probability P_{σ_i} . Prior to formalisation of the link between microscopic transition probabilities and macroscopic transition rates, the above high-level guidelines need to be verified for their sufficiency in providing a quantitative micro–macro link. In Sect. 3, we show how these guidelines translate to an actual implementation. Additionally, we discuss the inclusion of spatial factors influencing the interactions between agents, and we indicate possible extensions of the guidelines to deal with dynamic interaction topologies.




Fig. 2 A compact representation as a probabilistic finite state machine of the individual agent behaviour. *Solid arrows* are spontaneous individual transitions, while *dashed arrows* represent interactive transitions that are triggered when an agent encounters an agent of another population (which happens with probability $P\psi_i$ for population i)

2.2.3 Required properties

To obtain a quantitative micro–macro link, two properties are fundamental: state transitions must be memoryless, and the system must be well mixed.

Memorylessness is required by the macroscopic model, which can be derived as a mean-field approximation of a population-level continuous-time Markov process (see Seeley et al. 2012, for details). Memorylessness implies that, at any time, the probability that any agent undergoes a state transitions depends only on the current system state and is independent of the previous state of the system (thus the macroscopic process fulfils the Markov property). This can be achieved by ensuring a memoryless behaviour at the microscopic level, that is, ensuring that state transitions are governed by fixed probabilities per time unit, which results in an exponential distribution of the time intervals between entrance and exit from a state. Note that spatiality may interfere with memorylessness, and particular care must be given to the microscopic implementation, as exemplified in the case study presented in this paper.

The well-mixed property ensures that the probability of interaction between any two agents is constant. This can be achieved by a fully connected interaction topology or by a uniformly random interaction topology. If different populations were segregated, the well-mixed property would not hold anymore, therefore resulting in altered system dynamics (e.g. creation of island of agents with different opinions). To overcome this, interaction among agents should take place only when an unbiased sample of all populations is available.

3 Case study: shortest path discovery/selection

The high-level prescriptions discussed above need to be reified through experimentation in challenging case studies. Among the several factors that may hinder a macroscopic micro–macro link in a swarm robotics context, spatiality and physical interferences are probably the most important ones. Spatial features constrain the ability of interaction among different populations and may easily lead to departures from the ideal well-mixed condition, while the physical embodiment of specific robotic platforms constrains both motion and robot–robot interactions. In some cases, embodiment may also impose a limit on the maximum number of robots that can operate in a given scenario, as beyond this limit physical interferences might impede the robots correct functioning.

A well-conceived case study should exemplify the challenges introduced by spatiality and embodiment and propose strategies to address them. We have chosen a decision-making scenario that strongly depends on spatial factors and in which physical interferences may




Fig. 3 Graphical representation of the environment. Target areas A and B are located at distance, respectively, d_A and d_B from the home area (in this fig., $d_A = 2$ m and $d_B = 2.5$ m). All the three areas have radius $R = 0.3$ m. Robots move at constant speed $v = 0.1$ m s⁻¹ and can communicate with neighbours within a range $d_I = 0.6$ m

have strong effects since the robots share the same space: the discovery and selection of the shortest path between two areas in a foraging context (see Sect. 3.1). We first implement an abstract multi-agent simulation (presented in Sect. 3.2) to isolate the challenges introduced by the spatial distribution of target areas while ignoring the target robot embodiment and the resulting physical interactions. Here, agents are modelled as dimensionless particles that do not interfere with each other, so that we can analyse the performance of large swarms, study how the system dynamics vary as a function of the group size and propose specific solutions to tackle spatiality effects. The robotic simulation instead addresses both spatiality and embodiment and requires specific strategies to limit the physical interferences which are discussed in Sect. 3.3. Here, we have chosen a robotic platform—the e-puck robot (Mondada et al. 2009)—that stresses the challenges given by physical interferences due to their relatively small perceptual and interaction range.

Overall, the goal of this study consists in obtaining a good quantitative match between the non-spatial macroscopic model of Sect. 2 and the multi-agent and robotic implementation, as discussed in Sect. 4.

3.1 Problem definition

Foraging is a classic problem in swarm robotics (see Brambilla et al. 2013), and often solutions take inspiration from the food-gathering behaviour observed in social insects. Broadly speaking, agents involved in a foraging task are required to carry out search and retrieval activities. They explore the environment to locate *target areas* that contain the objects to be retrieved (e.g. a food patch). Then, they exploit the chosen path to retrieve objects to a given *home area* (e.g. the nest). Foraging encompasses a wide range of activities, such as exploration, navigation between target areas, retrieval and clustering of objects and collective decision making. All these activities are key components in several real-world applications envisioned as potential swarm robotics scenarios—e.g. search and rescue in disaster zones, collective construction in hazardous environments and nuclear disaster cleaning. Here, we focus on decision making, and we neglect other aspects (e.g. object retrieval) that are out of the scope of the present paper.

In this study, the goal of the swarm is to identify and exploit the shortest path between a home area and any of two target areas, referred to as target area A and target area B (see Fig. 3). The environment is a 2D infinite plane without obstacles, and all areas have circular shape with radius $R = 0.3$ m. The targets are located at a distance from home that varies in the range $d_i \in [1.5, 3.5]$ m, with $i \in \{A, B\}$. The swarm is composed of agents/robots with

local sensing and communication. An agent perceives home and targets only when within the corresponding areas, moves at constant speed of $v = 0.1 \text{ m s}^{-1}$ and communicates only with neighbours up to a maximum distance $d_I = 0.6 \text{ m}$ (indicated by the dotted circle around the robot in Fig. 3). Agents have perfect knowledge of the home location, while target areas need to be located through exploration. The collective decision is taken when the number of agents that have chosen a single target area reaches a quorum Q .

3.2 Implementation of abstract multi-agent simulations

The swarm is composed of dimensionless agents, so that no collisions or physical interferences are possible. Error-free odometry is exploited to track the position of known areas. Agents can be either committed to a specific target, or uncommitted. The committed agents—state C_A and C_B in Fig. 2—keep moving back and forth between home and target thanks to odometry. The uncommitted agents—state C_U —explore the environment to discover potential target areas. Agents can communicate with neighbours and share information about their commitment state and the location of discovered target areas.

3.2.1 Interactive and latent agents

Following the guidelines provided in Sect. 2, we have implemented a microscopic behaviour paying particular attention to the effects of spatiality. Indeed, given the locality of communication and the distance between target areas, agents committed to different targets and uncommitted agents cannot always interact with each other. In particular, different populations of agents are spatially segregated and come into contact only when in the home area. To ensure the well-mixed property, we have therefore decided to limit interactions only when agents are within the home area. When agents leave home, interactions are disallowed. More formally, agents can be in one of two activity states: an interactive state I (e.g. when inside the home area) and a latent state L (e.g. outside the home area). Switching between activity states follows the dynamics prescribed by the implemented behaviour and can be modelled by the simple PFSM shown in Fig. 4 left. We introduce a constant probability P_L for an agent to get latent. Conversely, agents become interactive in different ways depending on their commitment state, as will be detailed below. We refer to P_I as the average probability to become interactive. We can describe the change in the activity state through the PFSM model in the left part of Fig. 4, which predicts that a fraction $P_I/(P_I + P_L)$ of agents can be found on average in the interactive state. We exploit this prediction to tune the value of P_L : to ensure that on average 10% of the agents are interactive within the home area, we set $P_L = 9P_I$.

Given the three possible commitment states prescribed by the design pattern (Fig. 2) and the two activity states, the multi-agent implementation can be described by a PFSM with six states as shown in the right part of Fig. 4. Therefore, the agent committed to target A is interactive when in state C_A^I and latent when in state C_A^L (respectively, C_B^I and C_B^L for the agent committed to target B). Otherwise, the uncommitted agent is interactive in state C_U^I and latent in state C_U^L . The prescribed PFSM of Fig. 2 can be easily retrieved aggregating states by commitment.

3.2.2 Interaction patterns

When agents are in the interactive state, they can exchange short communication messages with neighbours, and on the basis of this communication, they can change their commitment




Fig. 4 Microscopic implementation for the shortest path selection/discovery problem. *Left* PFSM describing the switch between interactive and latent states. *Right* PFSM describing the complete agent behaviour. *White circles* represent interactive states, while *grey circles* represent latent states in which no interaction is possible. *Solid arrows* correspond to spontaneous transitions, while *dashed arrows* represent interactive transitions, and they both correspond to transitions described in the PFSM of Fig. 2. *Dash-dotted arrows* instead identify the transitions between activity states, i.e. the interactive and latent states. See text for further details

state. To obtain the well-mixed property, the distribution of interactive agents in the different commitment states must provide an unbiased representation of the entire population, including latent agents. More precisely, given the fraction Ψ^I (Ψ^L) of agents in the interactive state I (latent state L), we require that:

$$\frac{\Psi_i^I}{\Psi^I} \approx \frac{\Psi_i^L}{\Psi^L} \approx \Psi_i, \quad i \in \{A, B, U\}, \tag{4}$$

where Ψ_i^I and Ψ_i^L represent the fractions of agents that are found in state C_i^I and C_i^L within the entire population. In fact, if changes in the commitment state within the interactive sub-population (fraction Ψ^I) are much faster than changes in the activity state (i.e. agents switching between states I and L), the distribution of commitment states among interactive agents would misrepresent the global population distribution, and therefore, the microscopic and macroscopic dynamics would diverge. We have therefore decided to bind any change in the commitment state resulting from agent–agent interactions to the dynamics of activity change, which ensures the requirements given in (4). In other words, interactions are allowed only upon transitions from the interactive to the latent state: whenever an agent decides to get latent (following the constant probability per time unit P_L), it engages in an interaction with another agent in state I and updates its commitment state accordingly. In this way, changes in commitment state happen at the same rate as changes in the activity state.

Upon interaction, the probability of selecting a partner belonging to each population is proportional to the corresponding fraction of interactive agents. We refer to these probabilities as P_{Ψ_A} , P_{Ψ_B} and P_{Ψ_U} . Thanks to the attentive design of the interactive/latent dynamics discussed above, such probabilities closely represent the global fractions Ψ_A , Ψ_B and Ψ_U . Table 1 indicates all possible transitions from interactive to latent states and also links the transition probabilities to the control parameters P_ρ and P_σ that are introduced in the following. Here, it is worth noting that the sum of all outgoing transitions from any interactive state equals to P_L , which implements the link between changes in activity and commitment state discussed above. The actual change in the commitment state depends on the randomly selected partner, as detailed in the following.

Table 1 Correspondence between the transition probabilities in Fig. 4 and the control parameters of the implemented behaviour, for each interactive state. Considering that $P_{\psi_A} + P_{\psi_B} + P_{\psi_U} = 1$, the sum of all outgoing transitions from any interactive state equals to P_L

From	Transition probability	To
C_U^I	$P_{\psi_A} P_{\rho_A} \rightarrow P_L P_{\psi_A} P_{\rho}$	C_A^L
	$P_{L_U} \rightarrow P_L (P_{\psi_U} + (1 - P_{\rho}) (P_{\psi_A} + P_{\psi_B}))$	C_U^L
	$P_{\psi_B} P_{\rho_B} \rightarrow P_L P_{\psi_B} P_{\rho}$	C_B^L
C_A^I	$P_{\psi_B} P_{\sigma_B} \rightarrow P_L P_{\psi_B} P_{\sigma}$	C_U^L
	$P_{L_A} \rightarrow P_L (P_{\psi_A} + P_{\psi_U} + P_{\psi_B} (1 - P_{\sigma}))$	C_A^I
C_B^I	$P_{\psi_A} P_{\sigma_A} \rightarrow P_L P_{\psi_A} P_{\sigma}$	C_U^L
	$P_{L_B} \rightarrow P_L (P_{\psi_B} + P_{\psi_U} + P_{\psi_A} (1 - P_{\sigma}))$	C_B^I

3.2.3 Motion patterns

Agents always move at constant speed $v = 0.1 \text{ m s}^{-1}$, and their motion direction is determined by a motion vector \mathbf{m} that depends on the agent state. Uncommitted agents explore the environment in search of target areas through a correlated random walk (Bartumeus et al. 2005; Codling et al. 2008). When in the latent state C_U^L , uncommitted agents compute their motion vector \mathbf{m} as the sum of an inertia and a random vector as follows:

$$\mathbf{m} = r_i / \theta_i + 1 / \theta_r \tag{5}$$

where the notation r_i / θ_i indicates a vector in polar coordinates, with length r_i and angle θ_i . The angle θ_i is the agent’s heading direction, θ_r is an angle uniformly drawn in the range $[-\pi, \pi]$, and $r_i = 2$ is the relative strength of the inertia vector determining the random walk. Uncommitted agents switch between latent and interactive states with fixed probabilities per time unit. While searching for target areas (state C_U^L), uncommitted agents stop exploration to return home with constant probability P_{I_U} . The actual transition to the interactive state C_U^I takes place as soon as the agent enters the home area. When in the interactive state C_U^I , agents remain within the home area and move by correlated random walk (5). An interactive uncommitted agent becomes latent and resumes exploration with probability P_{L_U} (see Table 1).

Committed agents move back and forth between home and target exploiting odometry, which is used to update the motion vector \mathbf{m} towards the stored area locations. Similarly to uncommitted agents, committed interactive agents (state C_A^I or C_B^I) remain within the home area and move by correlated random walk (5). They become latent with probability P_{L_A} (P_{L_B} , see also Table 1). When latent (state C_A^L or C_B^L), they travel towards the target area and return home after a full round trip. This is modelled in the PFSM of Fig. 4 by the transition probability P_{I_A} (P_{I_B}).

As already mentioned, to provide the well-mixed property, the probability to switch between active and interactive states must be comparable across the agents’ commitments states: in this way, the agents in state I are an unbiased sample of the populations A , B and U . To this end, we tune the probability P_{I_U} to obtain a match between the average time of return to the home area of uncommitted agents with the average round-trip time of the committed agents. Given the agent speed v and the distance range of target areas, the average round-trip time is $\tau_m = 2\bar{d}/v = 50 \text{ s}$, where $\bar{d} = 2.5 \text{ m}$ is the average target area distance. Therefore, we fix the probability to become interactive per time unit to $P_{I_U} = 1/\tau_m = 0.02 \text{ s}$.

Note that in the present case study, different distances of the target correspond to different round-trip times: the closer the target area, the shorter the time needed to return home, the higher the frequency with which agents become interactive. This bias towards closer target areas can be exploited for decision purposes, as the rate of becoming interactive positively biases the system dynamics towards the selection of the shortest path. Starting from these considerations, the implementation of the four concurrent processes prescribed by the design pattern—discovery, abandonment, recruitment and cross-inhibition—can be easily obtained.

3.2.4 Discovery

Uncommitted agents discover target areas through random exploration while in the latent state C_U^L . Whenever an agent stumbles upon a target area, it stores the area location and gets committed to it. In the PFSM model of Fig. 4, this event corresponds to the transition between C_U^L and C_A^L (C_B^L), which happens with a probability P_{γ_A} (P_{γ_B}). Discovery events result from correlated random walks that start from the home area: the closer the target area, the higher the discovery probability. Therefore, at a macroscopic level, the discovery rate is biased towards shorter paths as prescribed by the design pattern.

3.2.5 Abandonment

Committed agents may spontaneously abandon their commitment and revert to an uncommitted state with constant probability P_α . Only agents in the latent state C_A^L (C_B^L) are allowed to abandon commitment and become uncommitted in order to resume exploration in state C_U^L , as shown in Fig. 4. Given that longer paths imply longer travel times, the macroscopic abandonment rate is larger for longer paths, which is in agreement with the prescriptions of the design pattern. After abandonment, agents return home and from there retrieve exploration. In this way, the dynamics of abandonment do not interfere with discovery, ensuring that every discovery event results from a random exploration that originates from the home location. This allows to preserve the memoryless property of the agent behaviour as required by the design pattern.

3.2.6 Recruitment

An uncommitted agent (state C_U^I), upon interaction with a committed agent (state C_A^I or C_B^I), gets recruited and thus committed to the other agent's target area with constant probability P_ρ . Given the interaction pattern discussed above, uncommitted agents get recruited only when becoming latent, therefore with a constant probability P_L . Given that the interactions with committed agents are bound to the probability P_{ψ_A} (P_{ψ_B}), the overall recruitment probability is given by $P_L P_{\psi_A} P_\rho$ ($P_L P_{\psi_B} P_\rho$), as shown in Table 1. Given that the interactive populations are slightly biased by shorter paths, at the macroscopic level recruitment for closer targets is slightly higher, as prescribed by the design pattern. The commitment message sent by a committed agent contains the information relative to the corresponding target area location, which can be used by the recruited agent for navigation. The transferred information is the angle and distance of the target location relative to the recruiter. The receiver combines this information with the recruiter's relative location and orientation, and through triangulation computes the target location.

Table 2 Control parameters and chosen value or value range

Parameter	Description	Value
v	Agent speed	0.1 m s^{-1}
d_I	Interaction radius	0.6 m
r_i	Relative strength of the inertia vector	2
P_{IU}	Probability of abandoning exploration	$v/2\bar{d} = 0.02 \text{ s}^{-1}$
P_L	Probability of getting latent	$3P_{IU}$
P_α	Probability of abandoning commitment	0.005 s^{-1}
P_ρ	Probability of recruitment	[0, 1]
P_σ	Probability of cross-inhibition	[0, 1]

In this study, all parameters are fixed but the recruitment and cross-inhibition probabilities

3.2.7 Cross-inhibition

An interactive agent committed to target A (B) in state C_A^I (C_B^I), upon interaction with an agent in state C_B^I (C_A^I), gets cross-inhibited and reverts to an uncommitted state with constant probability P_σ . Recall that interactions take place with probability P_L and that the probability of interacting with an agent from population A (B) is P_{ψ_A} (P_{ψ_B}). It follows that the overall cross-inhibition probability is $P_L P_{\psi_B} P_\sigma$ ($P_L P_{\psi_A} P_\sigma$, see also Table 1). Upon interaction, an agent recognises that the partner is committed to a different target area by measuring the distance between the target area location internally stored and the area location communicated by the partner. If the distance is greater than the target area radius R , cross-inhibition takes place.

3.2.8 Control parameters

The implemented behaviour has several control parameters that can be varied to obtain slightly different macroscopic dynamics. Some parameters have been arbitrarily chosen, while others are tuned to preserve the properties prescribed by the design pattern. In this study, we fix all parameters but P_ρ and P_σ , which determine the interaction pattern between agents. Table 2 summarises the parameters we introduced and the values that have been chosen.

3.3 Implementation of physics-based swarm robotics simulations

To study how the agent embodiment influences the system dynamics, we implemented the decision process on a swarm of simulated robots. The robotic platform of reference is the e-puck robot (Mondada et al. 2009), equipped with a range-and-bearing board (Gutiérrez et al. 2009) to allow short-range localised communication and an embedded computer running Linux.² Experimentation is conducted exploiting the ARGoS simulator (Pinciroli et al. 2012). The robot behaviour is implemented in accordance with the multi-agent behaviour described in Sect. 3.2. However, the robot characteristics and the physical embodiment introduce several constraints that require ad hoc modifications of the behaviour. We group the introduced modifications for (i) sensors and actuators, (ii) random walk, (iii) obstacle avoidance, and (iv) path exploitation.

² http://www.gctronic.com/doc/index.php/Overo_Extension.

3.3.1 Sensors and actuators

In the multi-agent implementation, agents have an abstract perception of the environment as well as abstract interactions with neighbours. In a robotic system, the same functionalities are implemented employing the available sensors and actuators. More specifically, the robotic implementation needs specific solutions for area localisation and robot–robot communication.

E-pucks detect and discriminate areas according to the floor grey level, as measured through the infrared ground sensor. We paint the floor with three grey levels: one grey level for the target areas, a second grey level for the home area and a third one for the empty area. In this way, robots can perceive an area when physically over it through their ground sensor and differentiate between types of area by looking at the different sensor activations resulting from different grey scales.

Communication is implemented through a combination of range-and-bearing infrared communication (IR) and WiFi, in a way similar to the range-and-bearing implementation of Roberts et al. (2009). Even if robots communicate via WiFi, interactions are kept local by delivering messages only to robots within IR range (we limit this range to $R = 30$ cm). Robots constantly broadcast via the IR channel their IP address, so that only upon reception of an IR message, a robot may deliver a WiFi message to the corresponding IP address. On the receiver side, WiFi messages are filtered out when they do not have an IR counterpart. In this way, WiFi communication gets enhanced by the localisation aspect: besides receiving the message content, a robot may localise the sender position using the IR signal strength and angle. In addition, by keeping interactions local, the risk of a communication channel overload is limited even when operating with large groups. However, communication with high robot densities or cluttered environments is limited to the subset of the closest neighbours because IR communication works only in line-of-sight. This may have a bearing on the macroscopic dynamics, as discussed below. Range-and-bearing communication is also used to establish a common frame of reference among two communicating agents, given that robots cannot recognise the relative heading of neighbours. We adopt here the solution described by Gutiérrez et al. (2010), based on sharing the relative bearing among the interacting robots. In this study we consider noiseless sensors and actuators. In a real robot implementation, the odometry error of the physical robot could be compensated for exploiting robot–robot communication as done in (Gutiérrez et al. 2010).

3.3.2 Random walk

Differently from the multi-agent system, e-pucks employ a differential-drive motion system that does not allow immediate changes of direction. Therefore, during random walk, if the movement vector \mathbf{m} varies at each control step as described in Eq. (5), the robot might not be able to change its position due to the time required to rotate on place for heading towards \mathbf{m} . To work out this issue while keeping the implementation as close as possible to the multi-agent system, a robot computes \mathbf{m} as in (5) and keeps unaltered the desired motion vector \mathbf{m} for $w = 5$ control steps before the next update. This gives the robot a sufficient time to rotate and move along the desired direction, while keeping a frequent update of the motion vector \mathbf{m} .

3.3.3 Obstacle avoidance

Due to their physical embodiment, robots must avoid collisions with each other and consequently alter their ideal motion trajectories. To avoid collisions, a robot computes its motion

vector \mathbf{m} —either from random walk (5) or from directed movements towards target areas—and sums to this an obstacle avoidance vector \mathbf{o} . The latter is computed by taking into account neighbouring robots' locations that are perceived in a range of 20 cm over the IR communication channel and calculating a sum vector in the opposite direction. Finally, the robot motion vector \mathbf{m}_o is computed as:

$$\mathbf{m}_o = \mathbf{m} + 2\mathbf{o}, \quad (6)$$

where a larger weight is given to the obstacle avoidance component to ensure collision-free motion.

3.3.4 Path exploitation

Obstacle avoidance alone is not sufficient to allow a smooth navigation back and forth between home and target areas. Indeed, groups of robots moving towards opposite locations interfere with each other. To reduce such interferences, we designed the robot trajectories in a round trip to create a double-line motion, letting robots keep the right with respect to the robots travelling in the opposite direction. To achieve this organised motion, committed robots going to a location rotate clockwise their motion vector \mathbf{m} by an angle $\theta \in [0^\circ, 30^\circ]$, with θ linearly decreasing as a function of the distance to the target location.

3.3.5 Expected effects of embodiment

The physics-based implementation of the robotic simulations allows us to investigate the effects of physical embodiment on the collective dynamics. The solutions described above deal with part of the constraints and interferences caused by the robot embodiment, but do not completely solve them. The most important consideration is that the robotic system cannot work with high robot densities, as motion between and within target areas would be strongly altered or completely disrupted. Therefore, we limit our study to groups of 50 robots. With this group size, we expect only a mild divergence from the ideal motion patterns. Additional constraints are given by the line-of-sight communication, which may prevent well-mixed interactions between the interactive robots. However, given that the average number of interacting robots is limited to 10% of the group size according to the multi-agent implementation, minor departures from a well-mixed condition are expected. A further analysis of the effects of embodiment in the robotics implementation is presented in Sect. 4.2.

4 Results

We evaluate the correctness of the microscopic implementation in both the multi-agent and the swarm robotics simulation. To this end, we compare the dynamics of the macroscopic model with the results of the microscopic simulations. In this study, the parameterisation of the macroscopic model is estimated directly from the simulations and we investigate the decision-making dynamics for varying probabilities P_ρ and P_σ and for a set of different decision problems for varying option quality v_i (i.e. varying the target areas distance d_i , $i \in \{A, B\}$).

4.1 Abstract multi-agent simulations

The effects of spatiality can be appreciated by looking at the decision-making dynamics of multi-agent simulations by varying the position of target areas. We show results for a set

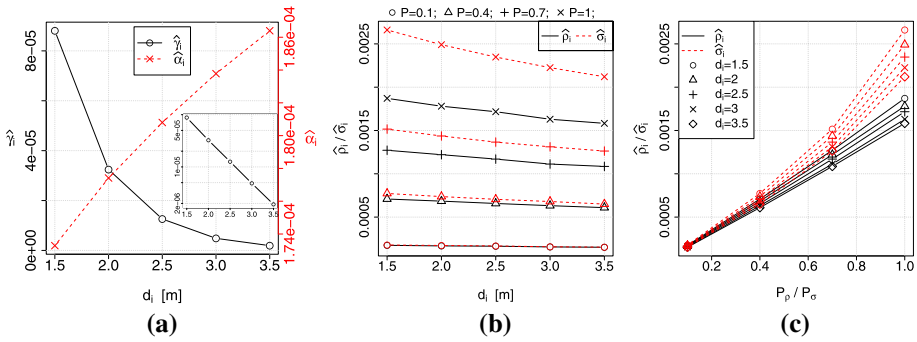


Fig. 5 Macroscopic transition rates estimated from the multi-agent simulation. **a** Estimated rates for discovery ($\hat{\gamma}_i$) and abandonment ($\hat{\alpha}_i$) plotted against the target area distance d_i . Note the exponential decay of the discovery rate ($\hat{\gamma}_i = \gamma_0 e^{-\lambda d_i}$, $\lambda = 1.97$), as indicated by the inset where the y-axis is plotted in logarithmic scale. **b, c** Estimated rates for recruitment ($\hat{\rho}_i$) and cross-inhibition ($\hat{\sigma}_i$) plotted against target area distance d_i and control probability P_ρ/P_σ . See text for details

of decision problems in which the distance of target B is fixed ($d_B = 2.5$ m), while the distance of target A systematically varies ($d_A \in \{1.5, 2, 2.5, 3, 3.5\}$ m). When not stated otherwise, simulations are run for a total population of $N = 500$ agents. We first discuss the relation between estimated transition rates and the target area distances. Then, we evaluate the micro–macro link comparing the final population distributions in multi-agent simulations with the macroscopic model attractors, and we verify the presence of the same type of phase transitions at both levels. Finally, we extend the analysis to varying group size N to identify finite-size effects on both microscopic and macroscopic dynamics.

4.1.1 Estimation of the macroscopic transition rates

To verify the existence of the micro–macro link, it is necessary to relate the microscopic implementation to the macroscopic model. This is possible by estimating the transition rates of the macroscopic model directly from the multi-agent simulation. We do so through survival analysis, which provides powerful nonparametric methods (Nelson 1969) to estimate how the probability of events changes over time directly from the experimental data (for a detailed explanation of the methods, see Appendix). We designed a set of experiments to estimate all the macroscopic transition rates of the ODE system in (1)— $\gamma_i, \alpha_i, \rho_i, \sigma_i, i \in \{A, B\}$ —for each possible distance of the target areas. For the interactive transitions, we also varied the control parameters P_ρ and P_σ in the set $\{0.1, 0.4, 0.7, 1\}$.

Figure 5 shows the result of the parameter estimation. In agreement with the implementation choices, the estimated discovery rate $\hat{\gamma}_i$ decreases with the distance d_i of the target area i , while the estimated abandonment rate $\hat{\alpha}_i$ follows the opposite trend (see Fig. 5a). Note that discovery follows an exponential decay with increasing distance at an estimated decay rate $\lambda = 1.97$, as shown in the inset of Fig. 5a. This is a result of the motion pattern implemented for uncommitted agents, which is governed by exponentially distributed exploratory trips. Such an exponential decay implies that farther targets are more difficult to discriminate than closer ones through discovery, and agent–agent interactions are useful to break possible deadlocks and lead to consistent decisions.

The interaction rates of recruitment $\hat{\rho}$ and cross-inhibition $\hat{\sigma}$ have been estimated for different values of the control probabilities P_ρ and P_σ , respectively. Also in this case,

the rates are inversely proportional to the distance of the target area to which the interacting population is committed (see Fig. 5b), in agreement with the implementation choices. Indeed, such a bias in the transition rates derives from the shorter round trips performed by agents committed to closer targets. As a result, biased rates favour convergence on shorter paths, confirming the correctness of the multi-agent implementation. Note that the effect of distance on the transition rates is linear and is important for larger values of the control probability but nearly negligible for smaller values, as shown in Fig. 5b. The estimated transition rates vary approximately linearly also with respect to the control probability of the individual agents, P_ρ and P_σ , as shown in Fig. 5c. A slight departure from linearity is visible for cross-inhibition, which is mainly due to errors in the estimation procedure given by the fact that returning home of committed agents is not a purely memoryless process.

4.1.2 Final distribution

Given the estimated macroscopic parameters for varying distance d_A and varying control parameters P_ρ and P_σ as resulting from the previous analysis, we can evaluate the micro–macro link and compare the dynamics displayed by the two description levels. For each decision problem obtained varying target distances and control probabilities, we perform 500 runs and we compare the population distribution after 5×10^3 s (with an integration timestep of 0.1 s) with the vector field and phase portrait of the ODE system (1) obtained using the estimated macroscopic parameterisation. Figure 6 shows the correspondence between microscopic and macroscopic dynamics for selected decision problems. The complete set of results for every decision problem we investigated is available in Figures S1–S5 in the online supplementary material. As can be observed in Fig. 6, we obtained a very good agreement between the macroscopic dynamics and the multi-agent simulations, which confirms the existence of a quantitative micro–macro link as a result of the correct implementation of the multi-agent behaviour. The agreement is noticeable not only when the microscopic dynamics have stabilised (i.e. population distributions around the macroscopic stable point, see Fig. 6a, b, d), but also in unstable transitory states which precisely follow the macroscopic vector field, as shown in Fig. 6c. Here, the macroscopic dynamics predict a quick convergence to a one-dimensional manifold, followed by a slower diffusion towards the one or the other attractor (Pais et al. 2013). These dynamics are well reproduced by the multi-agent simulations, which show several points scattered around the one-dimensional manifold.

Analysing the system dynamics, it is possible to notice that for low values of P_ρ a decision is not taken and a large majority of the agents remains uncommitted (see Fig. 6a). This behaviour results from a small positive feedback from recruitment which cannot balance the spontaneous abandonment rate (see also Fig. 5). In fact, by increasing the value of P_ρ , the system reliably converges towards the best option (see Fig. 6b). As shown in the inset, 100% of the runs reach the quorum $Q = 0.75$ for the best option. Cross-inhibition has a different role in the system dynamics: it speeds up the decision process and avoids deadlocks at indecision (Pais et al. 2013). In particular, for equally distant target areas, the dynamics with low values of P_σ are slower and the system is often found to be still at indecision after 5×10^3 s. As shown by the inset of Fig. 6c, in this time frame only about half of the runs reached the decision quorum for one or the other option. As soon as we increase P_σ , the system always converges to a large majority of agents committed to either one of the two alternatives (see Fig. 6d).

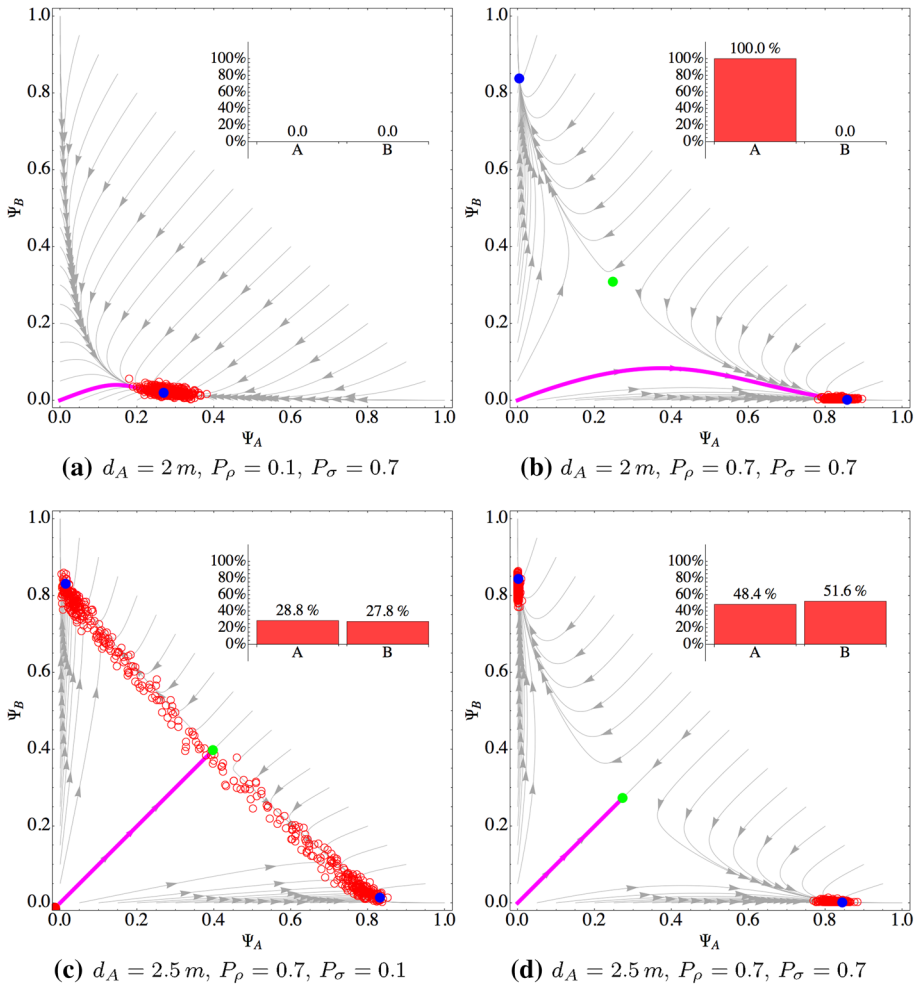


Fig. 6 (Colours online) Comparison of the microscopic and macroscopic dynamics for different decision problems and control probabilities. The macroscopic dynamics are displayed through the phase portrait (*grey arrows* show the vector field, *filled circles* are equilibrium points—*dark blue*: stable; *light green*: unstable). The *bold magenta arrow* represents the trajectory starting from a fully uncommitted population ($\Psi_A = \Psi_B = 0$). The microscopic dynamics are displayed as a scatterplot representing the final distribution of 500 independent runs (*red empty circles*). The *insets* show the decision pattern for a quorum $Q = 0.75$, indicating the percentage of runs that resulted in an above-quorum fraction of agents committed to either alternative

4.1.3 Bifurcation

To further validate the existence of a precise quantitative micro–macro link, we tested whether the macro and the micro systems undergo the same phase transition at the same predicted value σ^* (see Sect. 2.2). For an unbiased decision problem (i.e. $d_A = d_B$), transition rates for any of the two options are equal, and in such a completely symmetric situation, the system risks getting stuck at indecision. As discussed in Sect. 2.2.1, there exists a value of sigma, σ^* , for which the system breaks the symmetry and converges to either one of the two options. Using the estimated transition rates for $d_A = d_B = 1.5m$ and $P_\rho = 1$, we computed




Fig. 7 (Colours online) Phase transition predicted by the macroscopic model (*lines*) and resulting from microscopic multi-agent simulations (*colour shades*) for two unbiased decision problems: **a** $d_A = d_B = 1.5$ m and **b** $d_A = d_B = 2.5$ m. We show the proportion of agents committed to option A, Ψ_A , varying the control probability P_σ . The macroscopic bifurcation diagram illustrates stable states as *dark (blue) solid lines* and unstable state as *light (green) dashed lines*. The distribution of Ψ_A obtained from multi-agent simulations is illustrated through a density histogram in which more frequent values are represented as *darker (red) boxes*

the estimated critical value $\hat{\sigma}^* = 3.96261 \times 10^{-5}$. Note that the estimated macroscopic transition rate $\hat{\sigma}$ is linearly related to the individual-level cross-inhibition probability P_σ , as shown in Fig. 5c, especially for small values of P_σ :

$$\hat{\sigma} = K(d_A, d_B)P_\sigma, \tag{7}$$

where $K(d_A, d_B)$ is a constant value that embeds spatial factors related to the distance of both target areas, as well as the constant probability P_L of engaging in an interaction when agents are interactive (see Table 1). Since we have control only on the parameter P_σ , we discount $K(d_A, d_B)$ from $\hat{\sigma}^*$. The pitchfork bifurcation is predicted by the macroscopic model at the value $P_{\sigma^*} = \hat{\sigma}^*/K(d_A, d_B) = 1.4089 \times 10^{-2}$ (see Fig. 7a). We have therefore performed a bifurcation analysis with the multi-agent simulation, varying P_σ in the range $[0, 0.1[$ (with a 0.005 step increment). For each condition, we performed 500 runs lasting 2×10^4 s, a sufficient time to ensure convergence. Figure 7a illustrates the final distribution of Ψ_A as a density histogram, which results in a very good agreement with the macroscopic bifurcation diagram. Figure 7b shows similar results for a different parameterisation: $d_A = d_B = 2.5$ m and $P_\rho = 1$.

4.1.4 Finite-size effects

A limitation of the macroscopic model (1) is the infinite-size approximation, which turns into the impossibility to precisely analyse the system behaviour for varying group size N . Similarly to Valentini et al. (2014), we numerically study the finite-size effects through Monte Carlo simulations of a macroscopic finite-size model exploiting the Gillespie algorithm (Gillespie 1976), a widely used method to study the behaviour of continuous-time, well-mixed, memoryless processes. Figure 8 shows that also for small values of N the finite-size dynamics are in agreement with the infinite size ODE system (1). However, for small

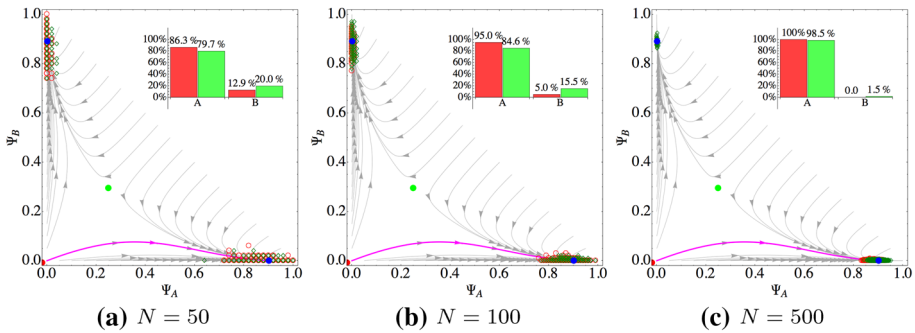



Fig. 8 (Colours online) Comparison of the macroscopic and microscopic dynamics including macroscopic simulations of finite-size effects. The experiment parameterisation is $d_A = 2\text{ m}$, $d_B = 2.5\text{ m}$, $P_\rho = P_\sigma = 1$. The population distribution for 2000 independent runs is displayed as *red empty circles* for multi-agent simulations, and as *green diamonds* for the Gillespie simulations. The *inset* shows the percentage of runs in which a decision was made for one or the other option (quorum $Q = 0.75$)

Fig. 9 (Colours online) Exit probability for multi-agent (*black solid lines*) and Gillespie simulations (*red dashed line*)



size N , the macroscopic system is subject to larger random fluctuations which result in more frequent wrong decisions especially in decision problems with a small difference between options (see the inset histogram in Fig. 8).

To better quantify the accuracy of the micro–macro link including finite-size effects, we introduce the *exit probability*, a measure that indicates the percentage of runs that terminate with a proportion of agents committed for the best (or equally best) option greater than a quorum $Q = 0.75$. Figure 9 shows the estimated exit probability for varying group size for both multi-agent and Gillespie simulations. Also in this case, we recognise a good agreement between macroscopic dynamics and multi-agent simulation results. Multi-agent simulations have a slightly higher exit probability than predicted by the macroscopic model. This is to be accounted to the larger time delay in reporting discoveries of farthest targets, which slightly bias the decision problem towards the closest option. Such effects cannot be grasped by the macroscopic dynamics starting from a fully uncommitted population, but could be accounted for by an appropriately measured bias in the starting conditions (e.g. starting with a small population fraction committed to the closest target). In future studies, by introducing such a bias in the macroscopic finite-size models, better predictions could be achieved.

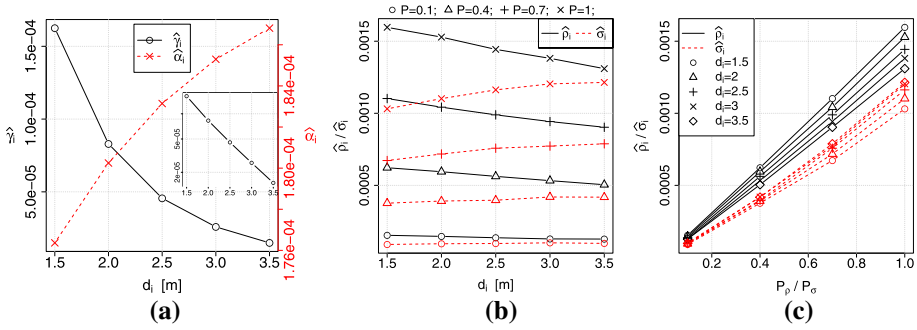


Fig. 10 Macroscopic transition rates estimated from the robotics simulations. See also the caption of Fig. 5 for details

4.2 Physics-based swarm robotics simulations

The swarm robotics implementation aims to investigate the effect of embodiment on the macroscopic dynamics. In Sect. 3.3, we have discussed several solutions introduced to reduce interferences and deal with physical interactions among robots and between robots and environment. We have therefore run several simulations using the ARGoS framework (Pinciroli et al. 2012) to study the microscopic dynamics. Also in this case, we study several decision problems fixing the distance of target B to $d_B = 2.5$ m and systematically varying the distance of target A , $d_A \in \{1.5, 2, 2.5, 3, 3.5\}$ m. Simulations are performed with groups of $N = 50$ robots.

4.2.1 Estimation of the macroscopic transition rates

For each experimental condition, we estimated the macroscopic transition rates through survival analysis, following the same methodology used for multi-agent simulations and detailed in Appendix (see Fig. 10). The estimated discovery ($\hat{\gamma}$) and abandonment ($\hat{\alpha}$) rates vary with target distance in a similar way to the multi-agent simulations, as shown in Fig. 10a. Abandonment increases with distance while discovery exponentially decays at a decay rate $\lambda = 1.27$ (see also the figure’s inset). Abandonment rates are also in a good quantitative agreement with the multi-agent simulations, while discovery rates present larger values. This is a result of the difference in the correlated random walk by agents and robots, the former having a smaller correlation in the motion direction due to the ability of instantaneous turning.

The estimated recruitment ($\hat{\rho}_i$) and cross-inhibition ($\hat{\sigma}_i$) rates depend on the distance d_i , as well as on the control probabilities P_ρ and P_σ , respectively, as shown by Fig. 10b, c. While recruitment correctly decays with distance, similarly to the multi-agent implementation, we notice that the cross-inhibition rate slightly increases with distance. This is mainly the result of physical interferences among robots committed to the same option, which are less frequent for longer paths because the density of robots over the path is reduced, therefore leading to higher transition rates. Indeed, notwithstanding the ad hoc motion pattern for path exploitation discussed in Sect. 3.3, physical interferences cannot be completely neutralised, the stronger their effect the higher the robot density, resulting in a penalty for motion on shorter paths.

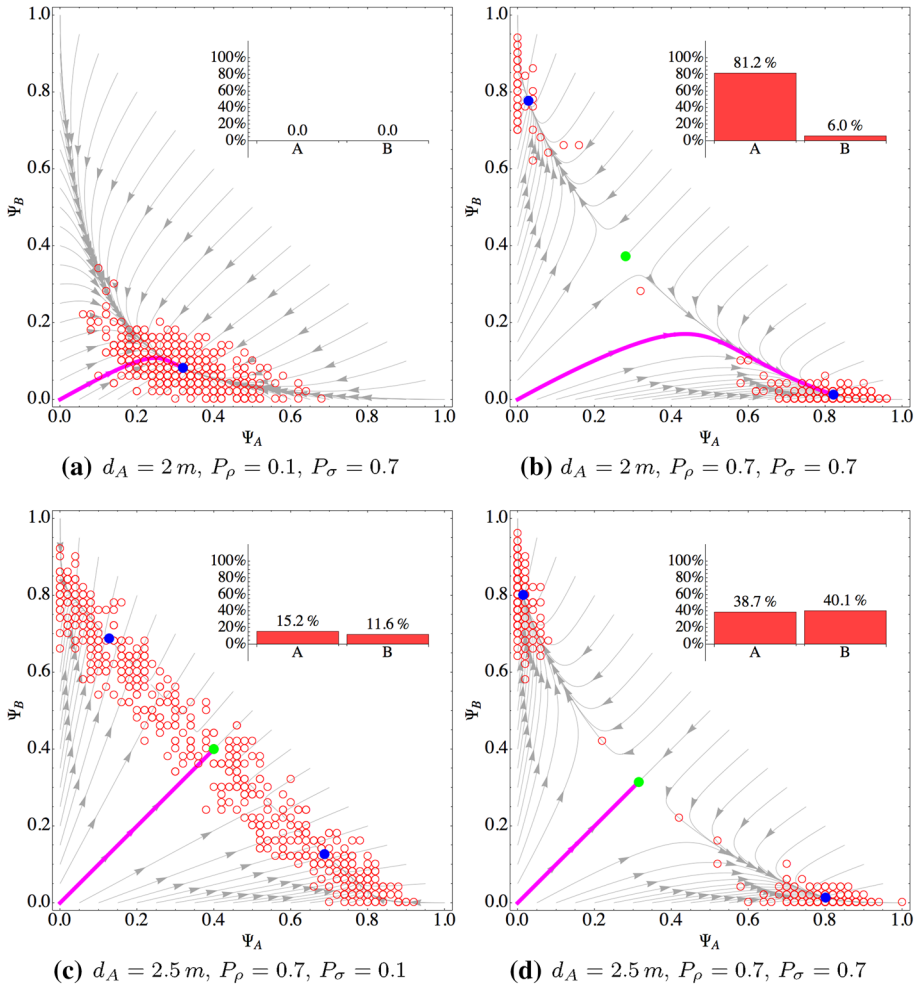


Fig. 11 (Colours online) Comparison of the microscopic and macroscopic dynamics for different decision problems and control probabilities. See the caption of Fig. 6 for details

4.2.2 Final distribution

Exploiting the estimated macroscopic transition rates, we can analyse the micro–macro link also for the robotics simulations. We have therefore matched the macroscopic dynamics with the final population distribution from 500 independent runs of the robotics simulations. Fig. 11 shows four selected conditions (corresponding to the same decision problems and control parameters used in Fig. 6). The complete set of results for every decision problem we investigated is available in Figures S6–S10 in the online supplementary material. Also in this case, the agreement between macroscopic dynamics and microscopic simulations is remarkable. With respect to multi-agent simulations, we note a larger scatter of data points, which is due to larger random fluctuations related to the physical embodiment of robots, as well as to finite-size effects given the relatively low group size. Apart from this, the microscopic implementation accurately matches the macroscopic dynamics. We also

recognise that the decision process is not compromised by a cross-inhibition rate increasing with distance. Indeed, the additional negative feedback is counterbalanced by the positive feedback given by recruitment, which allows to achieve rather accurate decisions.

5 Conclusions

The problem of correctly engineering large-scale distributed systems so as to obtain desired macroscopic dynamics is a complex one, and no general solution exists to date. In this paper, we have demonstrated how to obtain a quantitative relationship between a desired macroscopic model of collective decision making borrowed from studies of nest-site selection in honeybees (Seeley et al. 2012) and a shortest path discovery/selection problem within both multi-agent and robotics simulations. This work is particularly relevant in the perspective of formalising a design pattern for collective decisions based on the mechanisms observed in honeybees. The high-level guidelines provided in Sect. 2 have been verified in a particularly challenging example, and the relationship between microscopic parameters and macroscopic transition rates was obtained directly from the experimental data, showing that in many cases an approximately linear relation holds, with few exceptions that can be ascribed to the effects of spatiality and embodiment. Additionally, specific solutions have been proposed to deal with both spatiality and embodiment which deserve further investigations before proceeding to the formalisation of a design pattern for the general case.

For instance, the distinction between interactive and latent states can be generalised to any scenario in which—due to spatiality or other constraints—agent–agent interactions are not always possible, so that a necessity to differentiate between activity states arises. The lesson learned from the present case study is the importance of implementing spontaneous transitions between the activity states with comparable dynamics across different populations of committed and uncommitted agents. This should ensure that the interactive agents are an overall unbiased representation of the entire population. What needs to be further investigated is the amount of bias that can be introduced before disrupting the micro–macro link.

Another important implementation choice consists in linking the possibilities of interaction to the dynamics of activity change. In this respect, the lesson learned is the importance of having similar dynamics between interaction rates and activity changes, in order to maintain unbiased proportions within the interactive population. Also in this case, it would be important to ascertain whether alternative implementation patterns exist in order to speed up the decision process beyond the constraints imposed by activity dynamics.

Finally, it is worth noticing that spatial factors are well managed also thanks to well-designed motion patterns and transitions between states of the PFSM individual-level model. The lesson learned is that a proper design of the individual behaviour should lead to memoryless processes whenever possible. In this way, the microscopic implementation can be easily linked to the macroscopic description. When spatiality provides some form of bias, it is important to design the individual behaviour to neutralise such a bias (e.g. returning to the home location after an abandonment is mandatory to avoid a bias in the discovery of the abandoned alternative). Neutralising spatiality effects may have costs (e.g. it slows down discovery of potential alternatives), but these are costs that need to be paid to obtain a quantitative micro–macro link.

The problem studied in this paper and the solution proposed have some specificities that do not allow to generalise towards any collective decision-making problem. However, this case study can help in understanding the effects of spatiality on the decision dynamics. In this study, the quality of the alternatives is not directly available to the agents and decisions are bound

to biases related to spatial factors (i.e. different latencies related to different options, see also [Montes et al. 2010](#); [Scheidler et al. 2015](#)). As a consequence, not all possible macroscopic dynamics can be obtained, because the macroscopic transition rates can only be modulated by the chosen control parameters but not completely controlled. Different dynamics could be achieved in case each agent can individually estimate (with noise) the quality of the available option to contribute to the collective choice.

Future work will be tailored to formalise the suggestions given in this paper into a design pattern for collective decision making. More formal guidelines and methods to deal with spatiality issues—or, more generally, with the existence of latent states—will be provided, together with a formal relationship between microscopic control parameters and macroscopic transition rates. As a consequence, it would become possible to select the relevant parameterisation of the system at the macroscopic description level and immediately derive the corresponding parameterisation of the individual behaviour that leads to the desired collective outcome. This would lead to a complete micro–macro link for collective decision making allowing both top-down design and bottom-up verification and analysis.

Acknowledgments This work was partially supported by the European Research Council through the ERC Advanced Grant “E-SWARM: Engineering Swarm Intelligence Systems” (contract 246939). Vito Trianni acknowledges support from the EU-FP7 Project “DICE: Distributed Cognition Engineering” funded by the European Commissions FP7 People Programme under the Marie Curie Career Integration Grant scheme (Project ID: 631297). Marco Dorigo acknowledges support from the Belgian F.R.S.-FNRS. We thank Gabriele Valentini for sharing the code of the Gillespie algorithm used for the study of macroscopic finite-size effects and Carlo Pinciroli for the support in the implementation of the swarm robotics simulations.

Appendix

Survival analysis

We estimated the transition rates of the macroscopic model from the multi-agent/swarm robotics simulations through survival analysis. In this appendix, we shortly introduce this methodology and we present the Nelson–Aalen estimator that we employed for our analysis. Finally, we describe how to estimate the transition rates directly from experimental data. We believe that showing the application of this statistical tool for the analysis of a swarm robotics experiment may be of interest to the community.

Survival analysis

Survival analysis is a branch of statistics that offers tools to estimate the change over time of the probability of an event from experimental data. Survival analysis has been initially introduced in medicine to estimate the probability of survival (or death) of an organism under some treatment. Subsequently, these tools generalised to the estimate of any transition probability between populations. Nowadays, survival analysis is employed in several fields, such as economics—e.g. to estimate the probability of a stock market crash—or mechanical engineering—e.g. to estimate the probability of engine failures. In this work, we apply survival analysis to estimate the transition rates of the macroscopic model of collective decision making.

Other works have used this type of analysis to estimate the parameters of multi-agent systems behaviour. [Jeanson et al. \(2003\)](#) use survival analysis to estimate the probability




Fig. 12 **a** Cumulative discovery probability over time for target area at distance $d = 1.5$ m. The slope of the resulting line corresponds to the discovery rate estimate. **b** Cross-inhibition rate σ_A estimates over population fraction Ψ_A for target area distances $d_A = d_B = 2.5$ m and probability $P_\sigma = 0.1$

with which cockroaches change their behaviour. Garnier et al. (2008) and Reina et al. (2014) employed survival analysis to compute transition rates of artificial agents behaviour.

Hazard curve

We consider the three populations $\{A, B, U\}$ in accordance with the agent’s commitment state described in Sect. 2. To compute the rate at which agents switch (transit between) their commitment state, we log the number of timesteps t interlaying between two commitment switches (transitions) and the relative type of event causing the switch (e.g. discovery or recruitment). At the end of an experiment, we log the timesteps t from the last commitment switch as censored event, which indicates that after t timesteps no transition happened. We use the Nelson–Aalen estimator (Nelson 1969) to compute the hazard curve $H(t)$ from the collected experimental data. The hazard curve $H(t)$ shows the cumulative probability of events occurring until time t and is computed as follows:

$$H(t) = \sum_{t_i \leq t} d_i / n_i, \tag{8}$$

where d_i is the number of events recorded at t_i , and n_i is the number of events occurring (or censored) at time $t \geq t_i$. In a memoryless system, the probability of an event does not change over time, therefore the curve of the cumulative probability as function of time corresponds to a line with a slope equal to the constant event rate (i.e. the transition rate). Assuming our system as memoryless, we compute the transition rate by linear fitting the hazard curve with a line passing through the origin. Additionally, the quality of the fitting can be used to verify the correctness of a memoryless implementation.

Rate estimation

The rate of spontaneous transitions, in this study discovery and abandonment, can be directly estimated by computing the slope of the hazard curve, as detailed above. For instance, Fig. 12a

shows the hazard curve of the discovery rate computed from multi-agent experiments with target area at distance $d = 1.5$ m. Differently, rate estimates of transitions consequent to an interaction—in this work, recruitment and cross-inhibition—include the probability of interaction with an agent of the other population, which we call hereafter the *interacting population*. This probability changes during the process as the interacting population size changes and must be taken into account to estimate a constant transition rate independently from the interacting population size. Therefore, we first compute one aggregate transition rate for every interacting population size, and then, we normalise the rates for the interacting population fraction. For instance, from Eq. (1), the cross-inhibition rate for the population B is $(-\sigma_A \Psi_A)$, which includes the size of the interacting population A that delivers the inhibition signal. Through survival analysis, we compute the aggregate rate $(\sigma_A \Psi_A)$ for varying values of Ψ_A in the range $]0, 1[$. Then, by linear fitting, we discount from the aggregate rates the interacting population fraction to obtain σ_A . Figure 12b shows a set of 50 aggregate rate estimates plotted as function of Ψ_A , the slope of the fitted line corresponds to the estimate of σ_A .

To estimate the transition rates with constant population sizes, we run ad hoc experiments where the population sizes are fixed. In these experiments, agents follow the normal behaviour and, in case of commitment state transitions, they only log the event but do not change commitment state. In this way, we can quickly gather a large amount of data for every population size and parameterisation.

References

- Alexander, C., Ishikawa, S., & Silverstein, M. (1977). *A pattern language: Towns, buildings, construction*. New York: Oxford University Press.
- Arganda, S., Pérez-Escudero, A., & de Polavieja, G. (2012). A common rule for decision making in animal collectives across species. *Proceedings of the National Academy of Sciences*, 109(50), 20508–20513.
- Babaoglu, O., Canright, G., Deutsch, A., Di Caro, G. A., Ducatelle, F., Gambardella, L. M., et al. (2006). Design patterns from biology for distributed computing. *Transactions on Adaptive and Autonomous Systems*, 1(1), 26–66.
- Baronchelli, A., Ferrer-i Cancho, R., Pastor-Satorras, R., Chater, N., & Christiansen, M. H. (2013). Networks in cognitive science. *Trends in Cognitive Sciences*, 17(7), 348–360.
- Bartumeus, F., Da Luz, M. G. E., Viswanathan, G. M., & Catalan, J. (2005). Animal search strategies: A quantitative random-walk analysis. *Ecology*, 86(11), 3078–3087.
- Berman, S., Halasz, A., Hsieh, M. A., & Kumar, V. (2009). Optimized stochastic policies for task allocation in swarms of robots. *IEEE Transactions on Robotics*, 25(4), 927–937.
- Berman, S., Kumar, V., & Nagpal, R. (2011). 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 2011)*, (pp. 378–385). IEEE.
- Brambilla, M., Ferrante, E., Birattari, M., & Dorigo, M. (2013). Swarm robotics: A review from the swarm engineering perspective. *Swarm Intelligence*, 7(1), 1–41.
- Brambilla, M., Brutschy, A., Dorigo, M., & Birattari, M. (2015). Property-driven design for robot swarms. *ACM Transactions on Autonomous and Adaptive Systems*, 9(4), 17:1–17:28.
- Campo, A., Gutiérrez, A., Nouyan, S., Pinciroli, C., Longchamp, V., Garnier, S., et al. (2010). Artificial pheromone for path selection by a foraging swarm of robots. *Biological Cybernetics*, 103(5), 339–352.
- Campo, A., Garnier, S., Dédriche, O., Zekkri, M., & Dorigo, M. (2011). Self-organized discrimination of resources. *PLoS One*, 6(5), e19888.
- Castellano, C., Fortunato, S., & Loreto, V. (2009). Statistical physics of social dynamics. *Reviews of Modern Physics*, 81(2), 591–646.
- Codling, E. A., Plank, M. J., & Benhamou, S. (2008). Random walk models in biology. *Journal of the Royal Society, Interface*, 5(25), 813–834.
- Couzin, I. D. (2009). Collective cognition in animal groups. *Trends in Cognitive Sciences*, 13(1), 36–43.

- Ferrante, E., Turgut, A. E., Huepe, C., Stranieri, A., Pinciroli, C., & Dorigo, M. (2012). Self-organized flocking with a mobile robot swarm: A novel motion control method. *Adaptive Behavior*, 20(6), 460–477.
- Gamma, E., Helm, R., Johnson, R., & Vlissides, J. (1995). *Design patterns: Elements of reusable object-oriented software*. Boston, MA: Addison-Wesley Professional.
- Gardelli, L., Viroli, M., & Omicini, A. (2007). Design patterns for self-organising systems. In H. D. Burkhard, G. Lindemann, R. Verbrugge, & L. Z. Varge (Eds.), *Multi-agent systems and applications V, volume 4696 of LNCS* (pp. 123–132). Springer: Berlin.
- Garnier, S., Jost, C., Gautrais, J., Asadpour, M., Caprari, G., Jeanson, R., et al. (2008). The embodiment of cockroach aggregation behavior in a group of micro-robots. *Artificial Life*, 14(4), 387–408.
- Garnier, S., Gautrais, J., Asadpour, M., Jost, C., & Theraulaz, G. (2009). Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. *Adaptive Behavior*, 17(2), 109–133.
- Gillespie, D. T. (1976). A general method for numerically simulating stochastic time evolution of coupled chemical reactions. *Journal of Computational Physics*, 22, 403–434.
- Gutiérrez, A., Campo, A., Dorigo, M., Donate, J., Monasterio-Huelin, F., & Magdalena, L. (2009). Open e-puck range and bearing miniaturized board for local communication in swarm robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*, (pp. 3111–3116). IEEE Press.
- Gutiérrez, A., Campo, A., Monasterio-Huelin, F., Magdalena, L., & Dorigo, M. (2010). Collective decision-making based on social odometry. *Neural Computing & Applications*, 19(6), 807–823.
- Hamann, H., & Wörn, H. (2008). A framework of space-time continuous models for algorithm design in swarm robotics. *Swarm Intelligence*, 2(2), 209–239.
- Jeanson, R., Blanco, S., Fournier, R., Deneubourg, J.-L., Fourcassié, V., & Theraulaz, G. (2003). A model of animal movements in a bounded space. *Journal of Theoretical Biology*, 225(4), 443–451.
- Kao, A. B., Miller, N., Torney, C., Hartnett, A., & Couzin, I. D. (2014). Collective learning and optimal consensus decisions in social animal groups. *PLoS Computational Biology*, 10(8), e1003762.
- Kazadi, S. (2009). Model independence in swarm robotics. *International Journal of Intelligent Computing and Cybernetics*, 2(4), 672–694.
- Marshall, J. A. R., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., & Franks, N. R. (2009). On optimal decision-making in brains and social insect colonies. *Journal of the Royal Society, Interface*, 6(40), 1065–1074.
- Mondada, F., Bonani, M., Raemy, X., Pugh, J., Cianci, C., Klapotcz, A., Magnenat, S., Zufferey, J.-C., Floreano, D., & Martinoli, A. (2009). The e-puck, a robot designed for education in engineering. In *Proceedings of the 9th conference on autonomous robot systems and competitions*, volume 1(1), (pp. 59–65). IPCB, Castelo Branco, Portugal.
- Montes, M., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., & Dorigo, M. (2010). Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence*, 5(3–4), 305–327.
- Nelson, W. (1969). Hazard plotting for incomplete failure data. *Journal of Quality Technology*, 1, 27–52.
- Pais, D., Hogan, P. M., Schlegel, T., Franks, N. R., Leonard, N. E., & Marshall, J. A. R. (2013). A mechanism for value-sensitive decision-making. *PLoS One*, 8(9), e73216.
- Parker, C. A. C., & Zhang, H. (2009). Cooperative decision-making in decentralized multiple-robot systems: The best-of-N problem. *IEEE Transactions on Mechatronics*, 14(2), 240–251.
- Parker, C. A. C., & Zhang, H. (2010). Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. *Swarm Intelligence*, 4(3), 199–220.
- Pinciroli, C., Trianni, V., O’Grady, R., Pini, G., Brutschy, A., Brambilla, M., et al. (2012). ARGoS: A modular, parallel, multi-engine simulator for multi-robot systems. *Swarm Intelligence*, 6(4), 271–295.
- Reina, A., Dorigo, M., & Trianni, V. (2014). Towards a cognitive design pattern for collective decision-making. In M. Dorigo, et al. (Eds.), *Proceedings of 9th International Conference on Swarm Intelligence (ANTS) Volume 8667 of LNCS* (pp. 194–205). Springer Verlag, Berlin, Germany.
- Roberts, J. F., Stirling, T. S., Zufferey, J.-C., & Floreano, D. (2009). 2.5d infrared range and bearing system for collective robotics. In *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. IROS 2009*, (pp. 3659–3664). IEEE Press.
- Sartoretto, G., Hongler, M.-O., de Oliveira, M. E., & Mondada, F. (2014). Decentralized self-selection of swarm trajectories: From dynamical systems theory to robotic implementation. *Swarm Intelligence*, 8(4), 329–351.
- Scheidler, A., Brutschy, A., Ferrante, E., & Dorigo, M. (2015). The k-unanimity rule for self-organized decision making in swarms of robots. *IEEE Transactions on Cybernetics* (in press).
- Seeley, T. D., Visscher, P. K., Schlegel, T., Hogan, P. M., Franks, N. R., & Marshall, J. A. R. (2012). Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science*, 335(6064), 108–111.

- Trianni, V., & Dorigo, M. (2006). Self-organisation and communication in groups of simulated and physical robots. *Biological Cybernetics*, *95*(3), 213–231.
- Valentini, G., Hamann, H., & Dorigo, M. (2014). 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). International Foundation for Autonomous Agents and Multiagent Systems.
- Vicsek, T., & Zafeiris, A. (2012). Collective motion. *Physics Reports*, *517*(3–4), 71–140.
- Vigelius, M., Meyer, B., & Pascoe, G. (2014). Multiscale modelling and analysis of collective decision making in swarm robotics. *PloS One*, *9*(11), e111542.
- Wilson, S., Pavlic, T. P., Kumar, G. P., Buffin, A., Pratt, S. C., & Berman, S. (2014). Design of ant-inspired stochastic control policies for collective transport by robotic swarms. *Swarm Intelligence*, *8*(4), 303–327.