

Towards a Cognitive Design Pattern for Collective Decision-Making

Andreagiovanni Reina¹, Marco Dorigo¹, and Vito Trianni²

¹ IRIDIA, CoDE, Université Libre de Bruxelles, Brussels, Belgium

² ISTC, Italian National Research Council, Rome, Italy

{areina,mdorigo}@ulb.ac.be, vito.trianni@istc.cnr.it

Abstract. We introduce the concept of *cognitive design pattern* to provide a design methodology for distributed multi-agent systems. A cognitive design pattern is a reusable solution to tackle problems requiring cognitive abilities (e.g., decision-making, attention, categorisation). It provides theoretical models and design guidelines to define the individual control rules in order to obtain a desired behaviour for the multi-agent system as a whole. In this paper, we propose a cognitive design pattern for collective decision-making inspired by the nest-site selection behaviour of honeybee swarms. We illustrate how to apply the pattern to a case study involving spatial factors: the collective selection of the shortest path between two target areas. We analyse the dynamics of the multi-agent system and we show a very good agreement with the predictions of the macroscopic model.

1 Introduction

Several recent studies describe swarm systems as information-processing systems capable of some cognitive ability, which is strongly determined by the interaction patterns among the system components [4,14]. In this paper, we propose to take a similar perspective in the design of large-scale distributed systems. The studies mentioned above suggest that—to a large extent—the cognitive processing of natural decentralised systems takes place in inter-individual interactions, therefore limiting the need to postulate explicit representations within the single units. By viewing artificial swarm systems as distributed cognitive systems, it will be possible to maximise their information processing capability while keeping a low complexity of the individual units. That is, individual units would contribute to the overall system behaviour without having the global picture about the cognitive process they are collectively producing.

Designing such an information-processing system is clearly a complex endeavour, and in general, modelling, predicting and controlling large-scale distributed systems are complex tasks. Therefore, a successful design methodology for such systems should be grounded on solid theoretical premises. We propose a design methodology that leverages the current understanding of cognitive processing in (natural) distributed systems, and that puts this knowledge in use for the

design of artificial ones. Our proposal is based on the concept of *cognitive design patterns*, that is, reusable solutions to tackle problems requiring cognitive processing (e.g., collective decision-making among multiple alternatives). Similarly to common practice in software engineering [5], these design patterns can be used to guide the design and development of distributed cognitive processes, independently of the particular implementation technique. The idea of using design patterns in distributed systems has already been partially explored, but previous studies are not grounded on the theoretical understanding of collective dynamics [6,1,12]. Our proposal aims at providing general solutions grounded on the principled understanding of the basic mechanisms underlying cognitive processing.

In this paper, we propose a cognitive design pattern for collective decision-making. We refer to decision-making as the process of choosing the best option among a (finite, possibly unknown) number of different alternatives. This process requires the estimation of the quality of the available alternatives (possibly with uncertainty) to select the best one. Recent studies have identified optimal decision strategies in decentralised systems [8,13,11]. We propose a cognitive design pattern hinged on these studies, to be applied to artificial distributed systems. Although decision-making has been largely studied in this context [2,7,9,12], general purpose solutions are still missing. Our work aims at filling this gap by providing a design methodology that can be applied to several application domains.

The cognitive design pattern we propose is composed of the following elements: problem, inspiration, solution and case study. The collective decision-making problem and its relevance in artificial distributed systems have been discussed above. The biological inspiration is presented in Section 2, in which we discuss nest-site selection in honeybees and the related theoretical models accounting for the collective decision-making process. From these models, we derive the solution, which is discussed in Section 3 along with the causal relationship between the microscopic and macroscopic levels. To ease the pattern comprehension, we instantiate the design pattern in a case study presented in Section 4. In Section 5, we show the agreement between the collective decision implemented following the design pattern and the macroscopic model predictions. In Section 6, we discuss the proposed methodology and identify directions for further improvements.

2 Biological Inspiration and Theoretical Models

A remarkable example of collective decision-making is given by honeybee swarms during nest site selection. In spring, honeybee swarms reproduce by colony fission: the queen bee and several thousand workers leave the parent hive and create a cluster in the neighbourhood. Several hundred *scout bees* start searching for new potential nest sites, and return to the swarm to advertise through waggle dances what they have discovered. A number of alternatives may be discovered during the selection process, and a consensus decision is necessary to lead the

whole swarm to the best one. Decision-making is based on peer-to-peer interactions among bees: scouts committed to a potential site recruit other scouts. Additionally, scouts have a certain probability of spontaneously abandoning commitments. As a consequence, a competition between populations committed to different sites takes place, eventually leading to a quorum of individuals committed for the site that is finally chosen. Recently, cross-inhibition between different populations has been discovered. This is implemented through a *stop signal* that scout bees selectively deliver to nest-mates advertising for a different option [13]. A bee receiving several stop signals abandons the recruitment and becomes uncommitted. Thanks to the stop signal, poor-quality sites are quickly abandoned in favour of better ones. Most importantly, the stop signal allows to break decision deadlocks when same-quality alternatives are available, leading to a random decision for one of the two. In this way, the system can optimise the decision making, resulting in a choice that maximises the colony reward.

An analytical model of the nest-site selection process has been developed and confronted with empirical results, confirming the existence of both positive and negative feedback loops that determine the collective decision [13]. The model describes the decision-making process in a binary-choice scenario. The swarm is composed of N individuals (e.g., the scout bees), which can belong to three different groups: uncommitted individuals (population U with size N_U), and individuals committed to one of the alternatives (respectively population A and B , with sizes N_A and N_B). A continuous-time Markov process describes the way in which individuals switch between populations. Four types of transitions are sufficient: *discovery*, *abandonment*, *recruitment* and *cross-inhibition*.

Uncommitted individuals spontaneously *discover* and become committed to the alternative i at the rate γ_i :

$$\begin{aligned} \langle N_U, N_A, N_B \rangle &\xrightarrow{\gamma_A} \langle N_U^-, N_A^+, N_B \rangle \\ \langle N_U, N_A, N_B \rangle &\xrightarrow{\gamma_B} \langle N_U^-, N_A, N_B^+ \rangle \end{aligned} \quad (1)$$

where N_i^+ and N_i^- represent an increment or a decrement in population i . Committed individuals *abandon* the alternative i and thus get uncommitted at the rate α_i :

$$\begin{aligned} \langle N_U, N_A, N_B \rangle &\xrightarrow{\alpha_A} \langle N_U^+, N_A^-, N_B \rangle \\ \langle N_U, N_A, N_B \rangle &\xrightarrow{\alpha_B} \langle N_U^+, N_A, N_B^- \rangle \end{aligned} \quad (2)$$

Individuals from population i actively *recruit* uncommitted ones at the rate $\rho_i N_i/N$ proportional to the recruiting population size:

$$\begin{aligned} \langle N_U, N_A, N_B \rangle &\xrightarrow{\rho_A N_A/N} \langle N_U^-, N_A^+, N_B \rangle \\ \langle N_U, N_A, N_B \rangle &\xrightarrow{\rho_B N_B/N} \langle N_U^-, N_A, N_B^+ \rangle \end{aligned} \quad (3)$$

Finally, individuals from population i actively *inhibit* individuals of population j at the rate $\sigma_i N_i/N$ proportional to the inhibiting population size:

$$\begin{aligned} \langle N_U, N_A, N_B \rangle &\xrightarrow{\sigma_A N_A/N} \langle N_U^+, N_A, N_B^- \rangle \\ \langle N_U, N_A, N_B \rangle &\xrightarrow{\sigma_B N_B/N} \langle N_U^+, N_A^-, N_B \rangle \end{aligned} \quad (4)$$

Here, all transition rates— $\gamma_i, \alpha_i, \rho_i, \sigma_i$ —are greater than zero. It is worth noting that this model does not require any explicit comparison of the alternatives' quality by the single individuals. The quality value of the two alternatives—hereafter labelled v_A and v_B —is instead encoded in the transition rates (e.g., through value-dependent discovery or recruitment rates [11]): different-quality alternatives correspond to biased transition rates, while same-quality alternatives to unbiased ones. Overall, the collective decision is based purely on the system dynamics resulting from individual-to-individual interactions.

Starting from this stochastic model, it is possible to obtain the continuous-time master equation that describes how the probability of being in each state evolves over time. In the limit of large N , it is possible to extract a mean-field, population-level model of the system dynamics [13]. This takes the form of two coupled ordinary differential equations that describe the dynamics of the fraction $\Psi_i = N_i/N$ of individuals belonging to population $i \in \{U, A, B\}$:

$$\begin{cases} \dot{\Psi}_A = \Psi_U(\gamma_A + \rho_A\Psi_A) - \Psi_A(\alpha_A + \sigma_B\Psi_B) \\ \dot{\Psi}_B = \Psi_U(\gamma_B + \rho_B\Psi_B) - \Psi_B(\alpha_B + \sigma_A\Psi_A) \end{cases} \quad (5)$$

where $\Psi_U = 1 - \Psi_A - \Psi_B$. An extensive analysis of this model showed that the cross-inhibition rates σ_i crucially determine the dynamics of decision-making (see [11]). In case of same-quality alternatives ($v_A = v_B$), the transition rates for different alternatives have the same value (i.e., $\gamma_i = \gamma$, $\rho_i = \rho$, $\alpha_i = \alpha$ and $\sigma_i = \sigma$). In this case, for low rates of cross-inhibition, the system remains deadlocked at indecision with an equal number of individuals committed to either alternative ($\Psi_A = \Psi_B$). Through linear stability analysis [13], it is possible to identify the cross-inhibition level for which the system breaks the deadlock and converges to the choice of one alternative. The working region is $\{\rho > \alpha, \sigma > \sigma^*\}$, with critical value:

$$\sigma^* = \frac{4\alpha\gamma\rho}{(\rho - \alpha)^2}. \quad (6)$$

Additionally, in case of options of different quality, the cross-inhibition rate determines the minimum quality difference between the two alternatives to break the symmetry and make a systematic choice [11]. Therefore, there exist parameterisations of the system that allow to obtain accurate decisions when the quality of the options differs sufficiently, or random decisions when the values are similar.

3 Design Guidelines

The models discussed above provide a link between the individual-level description—given by the continuous time Markov process—and the population-level dynamics—given by the dynamical system in (5). However, the models alone are not sufficient to guide the implementation of a distributed multi-agent system. In fact, the transition rates depend on several factors: they incorporate global knowledge about the populations size, which may not be available to the agents, and they embed spatial and topological factors that partially determine




Fig. 1. Macroscopic dynamics for different parameterisations (trajectories and equilibrium points, shown as light green triangles for stable points and dark blue rhombus for unstable saddle points)

the probability of interaction among agents. For design purposes, it is therefore necessary to identify the causal relationship between microscopic transitions and macroscopic dynamics, and to provide a mechanistic description of the working principles.

The analysis of the macroscopic model in (5) reveals that, when the two alternatives have different value (e.g., $v_A > v_B$), an unbalanced agent distribution between the two populations is obtained thanks to a similarly biased commitment rate. This can be obtained through either discovery or recruitment (e.g., $\gamma_A > \gamma_B$ or $\rho_A > \rho_B$). Different discovery rates directly lead to a distribution of agents between the two alternatives that follows the rate ratio (see for instance Figure 1(a)). Similarly, a difference in the recruitment rates results in a unbalanced distribution even when the discovery rates are equal (see Figure 1(b)). This is due to the positive feedback mechanism that can be noted in (3), which states that the recruitment rate is proportional to the recruiting population size: the larger the population, the stronger the recruitment for the same population. Conversely, cross-inhibition provides a negative feedback loop that reduces the size of a population committed to an alternative, as stated in (4). Crucially, the rate of cross-inhibition is proportional to the size of the inhibiting population, and therefore contributes to the creation of an unbalanced distribution of individuals between committed populations, even for unbiased inhibition rate (i.e., $\sigma_A = \sigma_B$, see Figure 1(c)). 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 (i.e., $\sigma > \sigma^*$) makes the equilibrium point unstable, therefore leading to a symmetry breaking, as shown in Figure 1(d).¹

This mechanistic description clarifies the working regimes and suggests how the transition rates should be chosen to obtain the desired macroscopic behaviour. However, to guide the implementation of a distributed multi-agent

¹ A full characterisation of the parameter space is out of the scope of the present paper, and is object of ongoing studies. The interested reader can find a dynamical systems analysis for the parameter σ in [11].

system it is also necessary to define the main features of the individual agent behaviour and of the agent-to-agent interactions in order to obtain the transition rates in the appropriate range. In doing this, spatial and topological factors need to be suitably taken into account. The definition of these guidelines should provide the minimal requirements to obtain the desired system behaviour, and should incorporate the knowledge gained from the theoretical models.

In the particular case of collective decision making, we define general-purpose design guidelines as follows:

- (i) *discovery*: an uncommitted agent must commit to the alternative i with probability per unit time $P_{\gamma,i}$ (possibly proportional to the value v_i);
- (ii) *abandonment*: an agent committed to the alternative i must become uncommitted with probability $P_{\alpha,i}$ (possibly inversely proportional to v_i);
- (iii) *recruitment*: an uncommitted agent interacting with an agent committed to any alternative i must commit to i with probability $P_{\rho,i}$ (possibly proportional to the value v_i);
- (iv) *cross-inhibition*: an agent committed to alternative i that interacts with an agent committed to alternative $j \neq i$ must become uncommitted with fixed probability P_{σ} ;
- (v) *interaction*: the system must be *well-mixed*, that is, the probability of interaction between any two agents is uniform.

Note that *cross-inhibition* is asymmetric, that is, only one of the two agents (e.g., the one initiating the interaction) can change commitment state, this way respecting the stop signals mechanism. These guidelines are sufficient to generate a collective decision, provided that the agent probabilities of changing the commitment state are in the correct range. In the following section, we follow the design guidelines to implement a multi-agent system.

4 A Simple Spatial Scenario

We introduce a simple, spatial multi-agent scenario to demonstrate the application of the cognitive design pattern for collective decision making. The simplicity of the case study eases the analysis and the comprehension of the implemented mechanisms. At the same time, the studied scenario preserves the relevant ingredients of the collective decision making, therefore can well represent application scenarios where spatiality influences the system dynamics.

We study the collective choice of the shortest path between two alternatives in a 1D space: agents move on a circle and need to collectively select and exploit the shortest path between two *target areas* (see a pictorial representation in Figure 2). Two alternatives are possible: the upper and the lower path, respectively labelled A and B . The angle θ between the target areas defines the decision problem: the best alternative is A for $\theta < \pi$, B for $\theta > \pi$ and any of the two for $\theta = \pi$. To identify and exploit a path, agents need to navigate back and forth between target areas. We assume that agents can move at maximum angular speed $\omega = \pi/18 \text{ s}^{-1}$, and that movements are subject to noise modelled as a Gaussian displacement per arc degree following a $\mathcal{N}(0, \xi)$ distribution,

with $\xi = \pi/4500$. During navigation, agents track the angular distance of the two areas through dead reckoning, and use their estimates to attain previously visited areas without exploring or sensing the environment. However, due to the movement noise, position estimates are subject to cumulative errors. As a consequence, agents may end up with incorrect information and may be unable to attain a target area. Finally, agents have a sensing range of $\beta = \pi/36$ within which they can identify target areas and interact with other agents. All agents start with no knowledge about the target areas, and are therefore uncommitted.

Given the above specifications, we have developed the agent behaviours and interactions following the cognitive design pattern, as stated below.

- (i) *Discovery*: an agent explores the environment through a correlated random walk with persistence rate $\lambda = 0.8$ [3], and gets committed to a path as soon as it stores the position of the two target areas, allowing to navigate back and forth between them on the discovered path. Here, the probability $P_{\gamma,i}$ is determined by spatial factors (such as position and size of the areas) and by the parameters ω and λ . We obtain $\gamma_i \propto v_i$ because shorter paths are easier to discover through random walk.
- (ii) *Abandonment*: an agent abandons its commitment if it fails to attain a target area due to errors in the position estimates. Consequently, it erases the stored locations and resumes exploration. Also in this case, $P_{\alpha,i}$ is not directly under control of the agent's behaviour, but depends on the parameter β and on the movement noise variance ξ . Here, we obtain $\alpha_i \propto 1/v_i$ because lower abandonment rates result from smaller cumulative error on shorter paths: the variance increases proportionally to the path length.
- (iii) *Recruitment*: an uncommitted agent that interacts with an agent committed to alternative i gets recruited with fixed probability P_ρ : it receives the location of the target areas and transforms them in its own reference frame.
- (iv) *Cross-inhibition*: an agent committed to alternative i that interacts with an agent committed to alternative $j \neq i$ becomes uncommitted with fixed probability P_σ : it erases the stored locations and resumes exploration.
- (v) *Interaction*: to provide an equal probability of interaction with agents exploiting different paths, interactions are possible only when agents are within the same target area. Each agent has a maximum of one interaction per time unit. Additionally, agents remain in the target area with




Fig. 2. A graphical representation of the multi-agent scenario. The monodimensional environment is a circle in which the agents move on the circumference line to navigate back and forth between the two target areas.

probability $P_s = 0.9$ per time unit, or until a state change. This helps in creating well-mixed conditions and also increases the interaction rate.

Note that we have not specified a direct way to control the transition rates for discovery and abandonment, while recruitment and cross-inhibition are determined by the control probabilities P_ρ and P_σ . We choose fixed probabilities independently of the possible differences in the path lengths. As discussed in Section 3, this should be sufficient to produce a collective choice, provided that the discovery rates are biased toward the best option. Additionally, a sufficient level of cross-inhibition will contribute to make a collective choice and to break decision deadlocks (see Figure 1). To simplify the system analysis, we fix $P_\rho = P_\sigma = P$, which we refer to as the *interaction probability*. We study the system behaviour varying P and θ , while the other parameters are kept constant.

5 Results

To verify the correctness of the design pattern and to study how the collective behaviour changes as a function of the interaction probability P and of the decision problem given by θ , we check the adherence of the multi-agent system with the macroscopic model. To this purpose, we statistically estimate the transition rates γ_i , α_i , ρ_i and σ_i directly from the multi-agent system. Parameter estimation is performed through survival analysis, which permits to estimate how the probability of an event changes over time directly from the experimental data. Using the Nelson-Aalen estimator [10], we computed the hazard curve for the cumulative number of expected events (e.g., the discovery of alternative A):

$$H(t) = \sum_{t_i \leq t} d_i/n_i, \quad (7)$$

where d_i is the number of events recorded at t_i and n_i is the number of events occurring at $t \geq t_i$ (or not occurring at all, e.g., censored cases). In a memoryless process, the instantaneous transition rate is constant over time, therefore events cumulate at a constant rate, which corresponds to the hazard curve being a line. Its slope represents the estimated transition rate (see for instance Figure 3(a)). By employing a survival analysis, we can at the same time estimate all the




Fig. 3. (a) Hazard curve to estimate the transition rate α_A for $\theta = \pi$ and $P = 0.1$. (b) Transition $\sigma_i \Psi_i$ as function of Ψ_i for $i = A, B$, $\theta = \pi$ and $P = 0.1$. Points are unnormalised estimations, and the fitting line slope represents the normalised transition rate.




Fig. 4. Comparison between macroscopic dynamics (trajectories and fixed points, shown as light green triangles for stable points and dark blue rhombus for unstable ones) and the multi-agent simulations (final repartition of agents between the two populations, shown as red empty dots) for the asymmetric case $\theta = 5\pi/4$

transition rates and check the Markov assumption by looking at the shape of the hazard curves. Indeed, even if the agent behaviour is purely memoryless, departure from the Markov assumption is possible due to the spatial factors.

We perform $M = 300$ multi-agent simulations with $N = 400$ agents, and we vary both P and θ . The events we consider are the changes in the commitment status of each agent resulting from discovery, abandonment, recruitment and cross-inhibition, for both alternatives A and B . Discovery and abandonment are transitions spontaneously triggered by an agent, therefore it is possible to directly estimate the parameters γ_i and α_i from the corresponding hazard curve (e.g., Figure 3(a)). Conversely, recruitment and cross-inhibition result from the interaction of agents belonging to different populations, and the transition rates are proportional to the size of the recruiting and inhibiting populations. In this case, it is necessary to first estimate the transition rates for each population fraction (i.e., we estimate $\rho\Psi_i$ and $\sigma\Psi_i$), and then normalize for Ψ_i to obtain ρ and σ (see for instance Figure 3(b)). In this study, we limit the number of different events to consider to $N/10$ by approximating the population fractions within fixed-width windows of 0.025.

We first consider the asymmetric case of $\theta = 5\pi/4$. In this case, the decision problem should lead to the systematic choice of the alternative B . In Figure 4, we show the dynamics of the macroscopic model of equation (5) with the parameters estimated from the multi-agent system. We note that for low values of the interaction probability P there exists a single stable fixed point for $\Psi_B > 0.8$, and all trajectories converge to it. The agreement between the multi-agent system and the macroscopic dynamics is remarkable: the final distribution of agents from the simulations perfectly matches the model predictions (see Figure 4(a) and 4(b)). For higher values of P , the macroscopic dynamics show that the system undergoes a bifurcation, and a second stable fixed point appears that corresponds to the choice of the inferior option (see Figure 4(c) and 4(d)). The basin of attraction of the inferior fixed point is however smaller, and the trajectories starting from the origin (i.e., when all agents are uncommitted) always lead to




Fig. 5. (a) Convergence time for symmetric and asymmetric cases as a function of the interaction probability P , with decision threshold at $\Psi_i = 0.7$ and $t_{max} = 2000s$. (b) Transition rates as a function of interaction probability P for the symmetric case, with $\theta = \pi$.

the selection of the best option. This is confirmed by the multi-agent simulations, as all the repetitions resulted in a systematic choice of the alternative B . The bifurcation observed in the macroscopic dynamics appears when cross-inhibition is sufficiently strong compared to the other transition rates. On the one hand, this may lead to errors in the decision making if the system happens to be in the basin of attraction of the inferior choice. On the other hand, as also noted in [11], larger cross-inhibition rates lead to increased decision speed (see Figure 5(a)). Similar dynamics can be observed for different values of $\theta < \pi$, for which we observed that smaller differences between the alternatives sometimes lead to the wrong choice, as predicted by the macroscopic model (see supplementary material <http://iridia.ulb.ac.be/supp/IridiaSupp2014-005/sm.pdf>).

In Figure 6 we show the case for $\theta = \pi$, which corresponds to equal alternatives and—potentially—to a decision deadlock. This is actually the case for very low values of the interaction probability (e.g., $P = 0.01$ shown in Figure 6(a)). In this case, the model predicts a single stable fixed point for $\Psi_A = \Psi_B$, in agreement with the multi-agent simulations that equally remain deadlocked at indecision (see the red dot-cloud around the fixed point in Figure 6(a)). However, a phase transition is observed for increasing interaction probability, corresponding to a higher cross-inhibition rate and therefore to the ability of breaking the symmetry: two stable solutions appear indicating a collective choice for either A or B . The accordance between the multi-agent simulations and the macroscopic dynamics is very good also in this case, as shown in Figure 6(b)-(d). Note also that the macroscopic dynamics are highly symmetric, in accordance with the underlying multi-agent system. Similarly to the asymmetric case, we observe that higher values of P lead to a more definite choice of one or the other option, and that the convergence speed is also increased (see Figure 5(a)).

Figure 5(b) shows how the estimated transition rates vary with respect to the interaction probability P for the symmetric case. While discovery and abandonment remain roughly constant, both ρ and σ increase quasi-linearly with P , indicating that a higher probability of interaction among agents directly translates in increased recruitment and cross-inhibition rates. We note that the estimated cross-inhibition rate is initially below the critical value ($\sigma < \sigma^*$) for small interaction probabilities ($P < 0.07$). These are actually the values at which the multi-agent simulations remain deadlocked at indecision. For larger P , cross-inhibition is sufficiently high and the collective decision is efficiently performed.




Fig. 6. Macroscopic dynamics vs. multi-agent simulations for the symmetric case $\theta = \pi$

6 Conclusions

Overall, the results we obtained confirm that the multi-agent system that we implemented following the guidelines given by the cognitive design pattern corresponds very well to the reference macroscopic model. This result is non-trivial due to the differences between the two levels: while the model is continuous and deterministic, the multi-agent system, due to the finite swarm size, is discrete and stochastic. We have exploited the understanding of the basic mechanisms underlying the collective decision-making process in order to perform important implementation choices, such as the use of a fixed interaction probability independent of the decision problem to be faced. This minimalistic choice would not have been safe without knowledge about the system dynamics and about the correspondence between microscopic rules and macroscopic behaviour. Similarly, the implementation choices to grant a uniform interaction probability among agents—and therefore the *well-mixed* property—are also a result of the design pattern guidelines, which allowed to pinpoint the important aspects to be considered (e.g., the need to limit peer-to-peer interactions in a location containing a good sample of the population distribution). The parameter estimation we performed and the subsequent analysis for varying P and θ suggest that our implementation results in a well-behaved system not violating any assumption, despite the spatial factors that hinder the adherence to a Markov process.

The main feature of the design pattern we have developed consists in the possibility to perform decisions with minimal complexity at the individual level. Indeed, the only requirement is that agents can interact and recognise that their peers have a different opinion. Quality comparison is not necessary, which allows to implement the system in a large number of possible applications.

The present paper represents the very first step toward the definition and formalisation of cognitive design patterns for swarm systems. Several aspects must be investigated further in order to provide a proper engineering methodology. For what concerns collective decision-making mechanisms, future work will be dedicated to the characterisation of the full parameter space, in order to identify the parameter ranges that result in desired macroscopic dynamics. Additionally, we will characterise the relationship between individual-level parameters (e.g., the interaction probabilities P_ρ and P_σ) and the corresponding transition

rates (e.g., ρ and σ). Another issue to be considered is the effect of spatial and topological constraints on the collective decision process. In some preliminary studies, we have observed that spatiality influences the macroscopic dynamics (e.g., violating the well-mixed condition), and therefore needs to be characterised properly in order to provide guidelines re-usable in multiple domains.

Finally, following software engineering common practices [5], to let a solution become a design pattern we need to apply such solution to at least three different problems. In particular, we aim, as future work, to implement this cognitive design pattern in the fields of cognitive radio networks and language games.

References

1. Babaoğlu, O., Canright, G., Deutsch, A., Di Caro, G., Ducatelle, F., Gambardella, L.M., Ganguly, N., Jelasity, M., Montemanni, R., Montresor, A., Urnes, T.: Design patterns from biology for distributed computing. *Transactions on Adaptive and Autonomous Systems* 1(1), 26–66 (2006)
2. Campo, A., Garnier, S., Dédriche, O., Zekkri, M., Dorigo, M.: Self-Organized Discrimination of Resources. *PLoS One* 6(5), e19888 (2011)
3. Codling, E.A., Plank, M.J., Benhamou, S.: Random walk models in biology. *Journal of the Royal Society, Interface* 5(25), 813–834 (2008)
4. Couzin, I.: Collective cognition in animal groups. *Trends in Cognitive Sciences* 13(1), 36–43 (2009)
5. Gamma, E., Helm, R., Johnson, R., Vlissides, J.: *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional (1995)
6. Gardelli, L., Viroli, M., Omicini, A.: Design patterns for self-organising systems. In: Burkhard, H.-D., Lindemann, G., Verbrugge, R., Varga, L.Z. (eds.) *CEEMAS 2007. LNCS (LNAI)*, vol. 4696, pp. 123–132. Springer, Heidelberg (2007)
7. Hamann, H.: Towards swarm calculus: Urn models of collective decisions and universal properties of swarm performance. *Swarm Intelligence* 7(2-3), 145–172 (2013)
8. Marshall, J.A.R., Bogacz, R., Dornhaus, A., Planqué, R., Kovacs, T., Franks, N.R.: On optimal decision-making in brains and social insect colonies. *Journal of the Royal Society, Interface* 6(40), 1065–1074 (2009)
9. Montes, M., Ferrante, E., Scheidler, A., Pinciroli, C., Birattari, M., Dorigo, M.: Majority-rule opinion dynamics with differential latency: A mechanism for self-organized collective decision-making. *Swarm Intelligence* 5(3-4), 305–327 (2010)
10. Nelson, W.: Hazard plotting for incomplete failure data. *Journal of Quality Technology* 1, 27–52 (1969)
11. Pais, D., Hogan, P.M., Schlegel, T., Franks, N.R., Leonard, N.E., Marshall, J.A.R.: A mechanism for value-sensitive decision-making. *PLoS One* 8(9), e73216 (2013)
12. Parker, C.A.C., Zhang, H.: Cooperative decision-making in decentralized multiple-robot systems: the best-of-N problem. *IEEE Transactions on Mechatronics* 14(2), 240–251 (2009)
13. Seeley, T.D., Visscher, P.K., Schlegel, T., Hogan, P.M., Franks, N.R., Marshall, J.A.R.: Stop signals provide cross inhibition in collective decision-making by honeybee swarms. *Science* 335(6064), 108–111 (2012)
14. Trianni, V., Tuci, E., Passino, K.M., Marshall, J.A.R.: *Swarm Cognition: An interdisciplinary approach to the study of self-organising biological collectives*. *Swarm Intelligence* 5(1), 3–18 (2010)