Chapter 2
Discrete Consensus Achievement in Artificial Systems

Collective decision making refers to the phenomenon whereby a collective of agents makes a choice in a way that, once made, the choice is no longer attributable to any of the individual agents. This phenomenon is widespread across natural and artificial systems and is studied in a number of different disciplines including psychology, biology, and physics, to name a few. In the case of robot swarms, collective decision-making systems are distinguished between systems for *consensus achievement* and systems for *task allocation*. The first category encompasses systems that aim to establish an agreement among agents on a certain matter. The second category deals with systems that aim to allocate agents, i.e., the available workforce, to a set of tasks with the objective to maximize the performance of the collective. In this chapter, we focus on consensus achievement problems. We define the best-of-*n* problem and a taxonomy of its possible variants. Using this problem-oriented taxonomy, as well as a second taxonomy based on the design methodology, we review the literature of swarm robotics and provide a complete overview of the current state of the art.

2.1 Consensus Achievement

Consensus achievement problems can be further distinguished in two classes depending on the granularity of the choices available to the swarm. When the possible choices of the swarm are finite and countable, we say that the consensus achievement problem is *discrete*. An example of a discrete problem is the selection of the shortest path connecting the entry of a maze with its exit (Szymanski et al. 2006). Alternatively, when the choices of the swarm are infinite and measurable, we say that the consensus achievement problem is *continuous*. An example of a continuous problem is the selection of a common direction of motion by a swarm of agents flocking in a two- or three-dimensional space (Reynolds 1987). Both discrete and continuous consensus achievement problems have already received substantial attention from the scientific community.

G. Valentini, *Achieving Consensus in Robot Swarms*, Studies in Computational Intelligence 706, DOI 10.1007/978-3-319-53609-5_2

Discrete consensus achievement problems have been studied in a number of different contexts. The community of artificial intelligence focused on decision-making approaches for cooperation in teams of agents applying methods from the theory of decentralized partially observable Markov decision processes (Bernstein et al. 2002; Pynadath and Tambe 2002). These approaches, however, rely on sophisticated communication strategies and are suitable only for relatively small teams of agents. Discrete consensus achievement problems have been considered also in the context of the RoboCup soccer competition (Kitano et al. 1997). In this scenario, robots in a team are provided with a predefined set of plays and are required to agree on which play to execute. Different decision-making approaches have been developed to tackle this problem including centralized (Bowling et al. 2004) and decentralized (Kok and Vlassis 2003; Kok et al. 2003) play-selection strategies. Other approaches to consensus achievement over discrete problems have been developed in the context of sensor fusion to perform distributed object classification (Kornienko et al. 2005a, b). Finally, discrete consensus achievement problems are also studied by the community of statistical physics. Example studies include models of collective motion in one-dimensional spaces (Czirók and Vicsek 2000; Czirók et al. 1999; Yates et al. 2009) that describe the marching bands phenomenon of locust swarms (Buhl et al. 2006) as well as models of democratic voting and opinion dynamics (Castellano et al. 2009; Galam 2008).

Continuous consensus achievement problems have been primarily studied in the context of collective motion, that is, flocking (Camazine et al. 2001). Flocking is the phenomenon whereby a collective of agents moves cohesively in a common direction. The selection of a shared direction of motion represents the consensus achievement problem. In swarm robotics, flocking has been studied in the context of both autonomous ground robots (Ferrante et al. 2012; Nembrini et al. 2002; Spears et al. 2004; Turgut et al. 2008) and unmanned aerial vehicles (Hauert et al. 2011; Holland et al. 2005) with a focus on developing algorithms suitable for minimal and unreliable hardware. Apart from flocking, the swarm robotics community focused on spatial aggregation scenarios where robots are required to aggregate in the same region of a continuous space (Garnier et al. 2008; Gauci et al. 2014; Soysal and Şahin 2007; Trianni et al. 2003). The phenomenon of flocking is also studied by the community of statistical physics (Szabó et al. 2006; Vicsek and Zafeiris 2012) with the aim of defining a unifying theory of collective motion that equates several natural systems. A popular example study is provided by the minimalist model of self-driven particles proposed by Vicsek et al. (1995). The community of control theory has intensively studied the problem of consensus achievement (Mesbahi and Egerstedt 2010) with the objective of deriving optimal control strategies and prove their stability. In addition to flocking and tracking (Cao and Ren 2012), the consensus achievement problems studied in control theory include formation control (Ren et al. 2005), agreement on state variables (Hatano and Mesbahi 2005), sensor fusion (Ren and Beard 2008) as well as the selection of motion trajectories (Sartoretti et al. 2014). Finally, continuous consensus achievement problems have been also studied in the context of wireless sensor networks with the aim of developing algorithms for distributed estimation of signals (Schizas et al. 2008a, b).

In the rest of this chapter, we focus on discrete consensus achievement scenarios and we overview a number of research studies that proposed collective decision-making strategies specifically conceived for robot swarms. First, we formally define the best-of-n problem, i.e., a general structure and logic of a decision-making problem that characterizes several application scenarios in swarm robotics. Successively, we review related studies by organizing them in different classes depending on the approach adopted to design the collective decision-making strategy. Finally, we discuss the main differences between the different design approaches.

2.2 The Best-of-*n* Problem

In the swarm robotics literature, a large number of research studies focused on a relatively few application scenarios whose accomplishment requires the swarm to solve a consensus achievement problem (e.g., the shortest-path problem in foraging scenarios, site-selection in aggregation scenarios). These application scenarios have been primarily tackled separately from each other with an application-oriented perspective that resulted in either the development of domain-specific methodologies or the design of black-box controllers (cf. Sect. 2.3). However, we believe that the consensus achievement problems underlying these application scenarios share a similar logic and structure and that they can be abstracted to a unique framework: the *best-of-n problem*.

From an abstract point of view, the best-of-n problem requires a swarm of agents to make a collective decision over which option, out of n available options, offers the best alternative to satisfy the current needs of the swarm. We use the generic term *options* to abstract domain-specific concepts such as foraging patches, aggregation areas, or traveling paths, to name a few, that are related to particular application scenarios. We refer to the different options of the best-of-n problem using natural numbers, $1, \ldots, n$, and we say that the swarm is required to find the option $i \in \{1, \ldots, n\}$ with highest quality. That is, each option $i \in \{1, \ldots, n\}$ is characterized by an option quality ρ_i. Without loss of generality, we consider the quality of each option i to be normalized in the interval (0; 1] and $\rho_i = 1$ to represent the quality of the best option. Again, we use the term *option quality* as an abstraction to represent domain specific features (e.g., the length of a path, the size of an aggregation spot, the quality of food in a foraging patch).

Given a swarm of N agents, we say that the swarm has found a solution to a particular instance of the best-of-n problem as soon as it makes a *collective decision* for any option $i \in \{1, \ldots, n\}$. A collective decision is represented by the establishment of a *large majority* $M \geq (1 - \delta)N$ of agents that favor the same option i, where δ, $0 \leq \delta \ll 0.5$, represents a tolerance threshold set by the designer. In the boundary case with $\delta = 0$, we say that the swarm has reached a *consensus* decision, i.e., all agents favor the same option i. It is worth noting two aspects of a collective decision. On the one hand, a collective decision must satisfy the property of cohesion (Franks et al. 2013); that is, $\delta \ll 0.5$ implies that the opinions within the swarm are not split over different options of the best-of-*n* problem. On the other hand, a collective decision inherits the quality of the associated option i and, therefore, it can be optimal, $\rho_i = 1$, or sub-optimal, $\rho_i < 1$.

In general, the quality of an option is a function of the features of the environment, or of characteristics inherent to the swarm (e.g., the number of agents), or a combination of both factors and possibly multiple attributes (Reid et al. 2015). We distinguish between two factor types that determine the quality of a certain option. On the one hand, the option quality can be determined by an *internal preference* of individual agents for specific attributes characterizing an option. For example, when searching for a new nest site, honeybees instinctively favor candidate sites with a certain volume, exposure, and height from the ground (Camazine et al. 1999) regardless of their distance from the current nest location. This type of factors requires individual agents to directly evaluate the attributes of a certain option and to estimate its quality. On the other hand, the option quality can be determined by an existing bias that does not generate internally to individual agents but from certain features of the environment that indirectly influence the behavior of the swarm. We refer to this type of factors as *environmental bias*. For example, when foraging, ants find the shortest traveling path between a pair of locations as a result of pheromone trails being reinforced more often on the shortest path (Goss et al. 1989). Ants do not measure the length of each path individually. However, the length of a path indirectly influences the amount of pheromone laid over the path by the ants. Environmental bias factors can be interpreted as defining the cost of each option; this environmental cost can affect the selection of the best option by the swarm both positively, when higher quality options have lower costs, or negatively, otherwise (cf. Sect. 3.2.2).

In the case in which the option quality is independent of internal agent preferences and of environmental bias, the best-of-*n* problem reduces to a symmetry-breaking problem (de Vries and Biesmeijer 2002; Hamann et al. 2012). In this case, any of the *n* available options of the decision-making problem has the same quality and the goal of the swarm is to collectively choose one of the available options. A symmetry-breaking scenario arises also as a special case of the other classes when two or more options have equal and highest quality.

Finally, depending on the considered application scenario, the option quality is either *dynamic* or *static*. That is, the value of ρ_i may be a function of time. This feature is particularly relevant to guide the choices of designers when defining a collective decision-making strategy. When the option quality is static, the decision-making problem is a non-recurring or rarely recurring problem. In this case designers favor strategies that results in consensus decisions (Montes de Oca et al. 2011; Parker and Zhang 2009; Scheidler et al. 2016). Differently, when the option quality is a function of time (Arvin et al. 2014; Parker and Zhang 2010), designers favor strategies that result in a large majority of agents in the swarm favoring the same option without converging to consensus. In this case, the remaining minority of agents that are not aligned with the current swarm decision keep exploring other options and possibly discover new ones. This approach makes the swarm adaptive to changes in the environment (Schmickl et al. 2009b).

2.3 Overview of Current Design Approaches

The efforts of researchers in the last decade produced a vast literature of studies that spans over a number of application scenarios, design approach, and resulting collective decision-making strategies. In this section, we introduce a taxonomy that will be used in the rest of this chapter to review the most important studies performed in the literature of swarm robotics (see Fig. 2.1).

Particular instances of the best-of-n problem have been tackled using both bottom–up and top–down design approaches (Crespi et al. 2008). In the bottom–up approach, the designer develops the robot controller by hand, following a trial and error procedure where the robot controller is iteratively refined until the swarm behavior fulfills the requirements. In the top–down approach, the controller for individual robots is derived directly from a high-level specification of the desired behavior of the swarm by means of automatic techniques, for example, as a result of an optimization process (Bongard 2013; Nolfi and Floreano 2000) or of a compilation process (Werfel et al. 2014).

In the bottom–up approaches (see Sect. 2.4), the robot controller is usually developed by defining different atomic behaviors that are combined together by the designer to obtain a probabilistic-finite state machine. Each behavior of the robot controller is implemented by a set of control rules that determines (i) how a robot works on a certain task and (ii) how it interacts with its neighbor robots. We distinguish collective decision-making strategies designed by means of a bottom–up process in two categories, cf. Fig. 2.1, depending on how the control rules governing the interaction among robots have been defined. In the first category, that we call opinion-based approaches, robots have an explicit internal representation of their favored option, i.e., an opinion, and the role of the designer is to define the control rules that determine how robots share and change their opinions. Opinion-based approaches are used by the designer to tackle directly a consensus achievement problem rather than specific application scenarios. In the second category, that we call ad hoc approaches, we consider research studies where the control rules governing the interaction between robots have been defined by the designer to address

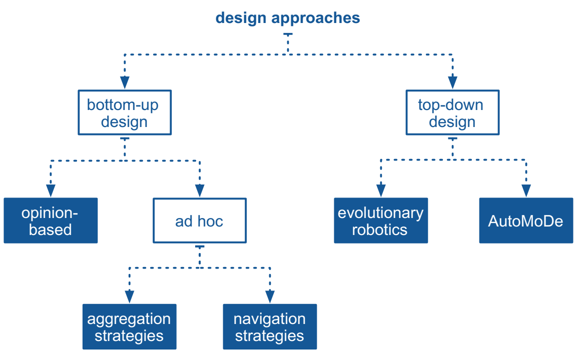

Fig. 2.1 Taxonomy used to review research studies that consider a discrete consensus achievement scenario. Research studies are organized according to their design approach (i.e., bottom–up and top–down) and to how the control rules governing the interaction among robots have been defined

a specific application scenario. As opposed to opinion-based approaches, control strategies belonging to this category are not explicitly designed to solve a consensus achievement problem; nonetheless, their execution by robots of the swarm results in a collective decision. In this category, we consider research studies that focus on the problem of spatial aggregation and on the problem of navigation in unknown environments.

In the top–down approaches (see Sect. 2.5), the robot controller is derived automatically from a high-level description of the desired swarm behavior. We organize research studies adopting a top–down approach in two categories: evolutionary robotics and automatic modular design. Evolutionary robotics (Bongard 2013; Nolfi and Floreano 2000) relies on evolutionary computation methods to obtain a neural network representing the robot controller. As a consequence, this design approach results in black-box controllers. In contrast, automatic modular design (Francesca et al. 2014) relies on optimization processes to combine behaviors chosen from a predefined set and obtain a robot controller that is represented by a probabilistic finite-state machine.

2.4 Bottom–Up Design Approaches

In this section, we consider research studies that developed collective decision-making strategies using a bottom–up design approach.

2.4.1 Opinion-Based Approaches

A large amount of research work has focused on the design of collective decision-making strategies characterized by robots having an explicit representation of their opinions. We refer to these collective decision-making strategies as opinion-based approaches. Using this design approach, robots are required to perform *explicit information transfer* (Ferrante 2013), i.e., to purposely transmit information represented by their opinion to their neighbors. As a consequence, a collective decision-making strategy developed using an opinion-based approach requires robots to have communication capabilities.

One of the first research studies developed with an opinion-based approach is that of Wessnitzer and Melhuish (2003). The authors considered a scenario in which a swarm of robots needs to capture two preys that are moving in the environment (i.e., the option of a best-of-2 problem). To do so, robots are required to collectively choose which prey to hunt first. The authors proposed a collective decision-making strategy based on the majority rule. Initially, each robot favors a prey chosen at random. At each time step, robots apply the majority rule over their neighborhood in order to reconsider and possibly change their opinions. In this case study, the two options of the decision-making problem are characterized by the same quality and thus the decision-making problem requires the swarm to break this symmetry.

Parker and Zhang (2009) developed a collective decision-making strategy by taking inspiration from the house-hunting behavior of social insects (Franks et al. 2002). The authors considered a site-selection scenario where a swarm of robots is required to discriminate between two illuminated sites (i.e., the options of a best-of-2 problem) based on their level of brightness (i.e., option quality defined by an internal preference factor). The proposed control strategy is characterized by three phases. Initially, robots are in the search phase either exploring the environment or waiting in a idle state. Upon discovery of a site and estimating its quality, a robot transits to the deliberation phase. During the deliberation phase, a robot recruits other robots in the search phase by repeatedly sending recruitment messages. The frequency of these messages is proportional to the option quality. Meanwhile, robots estimate the popularity of their favored option and use this information to test if a quorum has been reached. Upon detection of a quorum, robots enter the commitment phase and eventually relocate to the chosen site. The strategy proposed by Parker and Zhang builds on a direct recruitment and a quorum-sensing mechanism inspired by the house-hunting behavior of ants of the *Temnothorax* species. Later, Parker and Zhang (2011) considered a simplified version of this strategy and proposed a rate equation model to study its convergence properties.

Parker and Zhang (2010) proposed a collective decision-making strategy for unary decisions and applied it to the task sequencing problem. In the task sequencing problem, a swarm of robots needs to work sequentially on different tasks. The robots are required to collectively agree on the completion of a certain task prior to begin working on the next task. The task sequencing problem is a best-of-2 problem whose options (i.e., "task complete" or "task incomplete") are characterized by dynamic qualities (i.e., the level of completeness of a task changes over time). The authors proposed a quorum-sensing strategy to address this problem. Robots working on the current task monitor its level of completion (and, therefore, the option quality is due to an internal preference factor); when a robot recognizes the completion of the task, it enters the deliberation phase during which it asks its neighbors if they recognized too the completion of the task. Once a deliberating robot perceives a certain number of neighbors in the deliberation phase (i.e., the quorum), it moves to the committed phase during which it sends commit messages to inform neighbor robots about the completion of the current task. Robots in the deliberation phase that receive a commit message enter the committed phases and responds with an acknowledgement message. Committed robots measure the time passed since the last received acknowledgement and, after a certain time, they begin working on the next task.

Rather than mimicking biological systems, Montes de Oca et al. (2011) took advantage of the theoretical framework developed in the field of opinion dynamics (Krapivsky and Redner 2003). The authors extended the concept of latent voters introduced by Lambiotte et al. (2009) (i.e., voters stop participating to the decision-making process for a stochastic amount of time after changing opinion) and proposed a collective decision-making strategy referred to as *majority rule with differential*

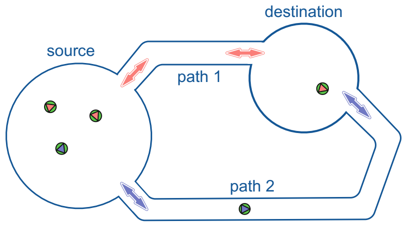

Fig. 2.2 Schematic illustration of the double-bridge problem (Goss et al. 1989). A source and a destination areas are connected by 2 paths: path 2 (*blue arrow*) whose length is approximately twice longer than path 1 (*red arrow*). The figure also shows a swarm of 5 robots (*green circles*), 3 robots with opinion 1 (*red triangle*) and 2 robots with opinion 2 (*blue triangle*)

latency. The proposed strategy is applied to a scenario inspired by the popular *double-bridge* problem (Goss et al. 1989). Robots in a swarm need to transport objects between two locations (i.e., source and destination in Fig. 2.2) connected by two paths of different length (i.e., option quality). Objects are too heavy to be transported by single robots and require a team of 3 robots. During the collective decision-making process, robots repeatedly form teams at the source location. Within a team, robots share with each other their opinion for their favored path and then apply the majority rule (Galam 2008) to determine which path the team should traverse. Then, the team travels back-and-forth along the chosen path before dismantling once back in the source location. Using this strategy, robots do not measure the length of each path; in contrast, the length of a path indirectly biases the frequency of participation of robots to the decision-making process taking place in the source location (i.e., the option quality is defined by an environmental bias). The majority rule with differential latency has been the subject of an extensive theoretical analysis that includes deterministic macroscopic models (Montes de Oca et al. 2011), master equations (Scheidler 2011), statistical model checking (Massink et al. 2013), and Markov chains (see Chap. 4).

The same foraging scenario investigated by Montes de Oca et al. (see Fig. 2.2) has been the subject of other research studies. In Brutschy et al. (2012); Scheidler et al. (2016), the authors extended the control structure underlying the majority rule with differential latency introducing the k-unanimity rule—a novel decision-making mechanism for individual robots. Instead of forming teams and applying the majority rule within each team, robots have a memory of size k where they store the opinions of other robots as they encounter them. A robot using the k-unanimity rule changes its current opinion in favor of a different option only after consecutively encountering k other robots all favoring that other option. The primary benefit of the k-unanimity rule is that it allows the designer to adjust the speed and the accuracy of the collective decision-making strategy by means of the parameter k (Scheidler et al. 2016). The authors studied the dynamics of the k-unanimity rule analytically when applied to decision-making problems with up to $n = 3$ options using a deterministic macroscopic model and a master equation.

Reina et al. (2014, 2015a, b) proposed a collective decision-making strategy inspired by theoretical studies that unify the decision-making behavior of social insects with that of neurons in vertebrate brains (Marshall et al. 2009; Pais et al. 2013; Seeley et al. 2012). The authors considered the problem of finding the shortest path connecting a pair of locations in the environment. In their strategy, robots can be either uncommitted, i.e., without any opinion favoring a particular option, or committed to a certain option, i.e., with an opinion. Uncommitted robots might discover new options in which case they become committed to the discovered option. Committed robots can recruit other robots that have not yet an opinion (i.e., direct recruitment); inhibit the opinion of robots committed to a different option making them become uncommitted (i.e., cross-inhibition); or abandon their current opinion and become uncommitted (i.e., abandonment). In Reina et al. (2014, 2015a), the authors studied the proposed strategy in a foraging scenario with two alternative foraging patches (i.e., the option of a best-of-2 problem); the quality of each option is determined by their distance from a central retrieval area which indirectly influence the behavior of the swarm (i.e., environmental bias). In Reina et al. (2015b), the authors studied the proposed strategy in a different setup: foraging patches are characterized by a quality that the robot can measure (i.e., the internal preference factor) and are positioned at different distances in a way that the best foraging patch is the farthest (i.e., an environmental bias factor that influences negatively the decision-making process). The proposed strategy is supported by both deterministic and stochastic mathematical models (i.e., ordinary differential equations and chemical reaction networks) that link the microscopic parameters of the system to the macroscopic dynamics of the swarm.

2.4.2 Ad Hoc Approaches

In this section, we consider a number of research studies that resulted in the development of control strategies for specific application scenarios, that is, spatial aggregation and navigation in unknown environments. As opposed to opinion-based approaches, the objective of the designers of these control strategies is not to tackle a consensus achievement problem but to address a specific need of the swarm (i.e., aggregation or navigation). Nonetheless, the control strategies reviewed in this section provide a swarm of robots with collective decision-making capabilities.

2.4.2.1 Aggregation-Based Approaches

Aggregation-based approaches are control strategies that make the robots of the swarm aggregate in a common region of the environment forming a cohesive cluster. The opinion of a robot is represented implicitly by its position in space. The primary advantage of an aggregation strategy is represented by the fact that the information regarding a robot opinion can be implicitly transferred to nearby robots without the

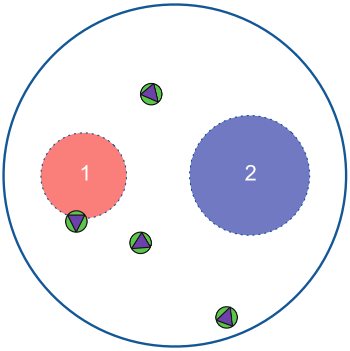

Fig. 2.3 Schematic illustration of the aggregation scenario considered in Campo et al. (2010a); Garnier et al. (2009). The environment is represented by a circular arena enclosed by a wall (*blue line*). The arena is provided with 2 shelters of different size: shelter 1 (*red circle*) has an area approximately twice as smaller than that of shelter 2 (*blue circle*). The figure also shows a swarm of 4 robots (*green circles with purple triangles*) wondering around the environment

need of communication (Ferrante 2013; Sumpter 2010, Chap. 3). Implicit information transfer can be implemented, for example, by means of neighbors observation.

Garnier et al. (2009) considered a behavioral model of self-organized aggregation and studied the emergence of collective decisions. The authors consider an aggregation scenario, cf. Fig. 2.3, where robots are presented with two shelters (i.e., the options of a best-of-2 problem) of different area (i.e., option quality) and are required to select one shelter under which the swarm should aggregate. The proposed control strategy is inspired by the behavior of young larvae of the German cockroach, *Blattella germanica*, (Jeanson et al. 2003, 2005). Robots explore their environment by executing a correlated random walk. When a robot detects the boundary of the arena, it pauses the execution of the random walk and begins the execution of a wall-following behavior. The wall-following behavior is performed for an exponentially distributed period of time after which the robot turns randomly towards the center of the arena. When encountering a shelter, the robot decides whether to stop or not as well as whether to stop for a short or a long period of time as a function of the number of nearby neighbors. Given the number of perceived neighbors, this function returns the probability for a robot to stop and its value has been tuned by the designer to favor the selection of the shelter of bigger area (i.e., shelter 2 in Fig. 2.3). Garnier et al. (2009) studied the proposed strategies in two different setups. In the first setup, the aggregation problem requires to break the symmetry between two shelters of equal size. In the second setup, one shelter is larger than the other. The option quality is determined by an internal preference factor (i.e., the number of perceived neighbors which carries information on the shelter size) and an environmental bias factor (i.e., the shelter size, on the grounds that larger shelters are easier to discover by robots and are of lesser cost). Correll and Martinoli (2011) studied this collective behavior with both Markov chains and difference equations and showed that a collective decision arises only when robots move faster than a minimum speed and are characterized by a sufficiently large communication range.

Campo et al. (2010a) considered the same aggregation scenario of (Garnier et al. 2009) and developed a control strategy taking inspiration from theoretical studies of the aggregation behaviors of cockroaches (Amé et al. 2006). In their strategy the robot controller is composed of 3 phases: exploration, stay under a shelter, and move back to the shelter. Initially, the robots explore the environment by performing a random walk. Once a robot discovers a shelter, it moves randomly within the shelter's area and estimates the density of other robots therein. If during this phase, a robot accidentally exits the shelter, it performs a U-turn aimed at reentering the original shelter. Differently from Garnier et al. (2009), the robots directly decide whether to stay under a shelter or to leave and return to the exploration phase. This decision is stochastic and the probability to leave the shelter is given by a sigmoid function of the estimated density of robots under the shelter. Contrarily to Garnier et al. (2009), the authors tuned the parameters of the sigmoid function with the aim to favor the selection of the smallest shelter that can host the entire swarm (e.g., shelter 1 in Fig. 2.3). As a consequence, the option quality is still determined by a combination of internal preference and environmental bias factors as in Garnier et al. (2009) but, this time, the environmental bias factor (i.e., the size of the shelter) hampers the discovery of the best option (shelter 1 in Fig. 2.3 by the robots of the swarm). A similar aggregation strategy was proposed later by Brambilla et al. (2014) and studied in a binary symmetry-breaking setup. Differently from the sigmoid function used in Campo et al. (2010a), the authors considered a linear function of the number of neighbors to determine the probability with which a robot decides whether to leave a shelter or not.

Kernbach et al. took inspiration from the thermotactic aggregation behavior of young honeybees, *Apis mellifera* L., (Grodzicki and Caputa 2005), and proposed the BEECLUST algorithm (Kernbach et al. 2009; Schmickl et al. 2009b). The goal of a swarm executing the BEECLUST algorithm is to aggregate around the brightest spot in the environment. For this purpose, a robot moves randomly in the environment; upon encountering another robot, the robot stops moving and measures the local intensity of the ambient light. After waiting for a period of time proportional to the measured light, the robot resumes a random walk. In Schmickl et al. (2009b), the authors studied the BEECLUST algorithm in a setup characterized by two spots (i.e., the options of a best-of-2 problem) of different brightness. The option quality is defined by an internal preference factor (i.e., the brightness measured by each robot) and is also positively influenced by an environmental bias factor (i.e., brighter spots are also characterized by a bigger area which make them easier to discover by the robots of the swarm). Later, Hamann et al. (2012) studied the BEECLUST algorithm in a binary symmetry-breaking setup where both spots are characterized by the same level of brightness. The BEECLUST algorithm has been subject of an extensive theoretical analysis that includes both spatial and non-spatial macroscopic models (Hamann 2013; Hamann et al. 2012; Hereford 2010; Schmickl et al. 2009a). While the resulting decision-making process is robust, it is particularly difficult to model due to the complex dynamics of cluster formation and cluster breakup (Hamann et al. 2012).

More recently, Arvin et al. (2012, 2014) extended the original BEECLUST algorithm by means of a fuzzy controller. In the original BEECLUST algorithm, after the expiration of the waiting period, a robot chooses randomly a new direction of motion. Contrarily, using the extension proposed by Arvin et al., the new direction of motion is determined using a fuzzy controller that maps the magnitude and the bearing of the input signal (in their case, a sound signal) to one out of five predetermined directions of motion (i.e., left, slightly-left, straight, slightly-right, right). The authors studied the extended version of the BEECLUST algorithm in a dynamic, binary decision-making problem defined by two aggregation areas; each area is identified by a sound emitter and the sound magnitudes of the two areas are different and vary over time. As in Schmickl et al. (2009b) where the option quality was determined by the level of brightness, the size of each aggregation area is proportional to the magnitude of the emitted sound. Consequently, the option quality is determined by an internal preference factor and is facilitated by an environmental bias factor. This extension has been shown to improve the aggregation performance of the BEECLUST algorithm (i.e., clusters last for a longer period of time) as well as its robustness to noisy perceptions of the environment.

Mermoud et al. (2010) considered an application scenario where robots of the swarm are required to monitor a certain environment, searching and destroying undesirable artifacts (e.g., pathogens, pollution). Specifically, artifacts correspond to colored spots that are projected on the surface of the arena and can be of two types: “good” or “bad”. The author proposed an aggregation-based strategy that allows robots to collectively perceive the type of a spot and to destroy those spots that have been perceived as bad while safeguarding good spots. Each robot explores the environment by performing a random walk and avoiding obstacles. Once a robot enters a spot, it measure the light intensity to determine the type of the spot. Successively, the robot moves inside the spot area until it detects a border; at this point, the robot decides with a probability that depends on the estimated spot type whether to leave the spot or to remain inside it by performing a U-turn. Within the spot, a robot stops moving and starts to form an aggregate as soon as it perceives one or more other robots evaluating the same spot. When the aggregate reaches a certain size (which is predefined by the experimenter), the spot is collaboratively destroyed and robots resume the exploration of the environment. The achievement of consensus is detected using an external tracking infrastructure which also emulates the destruction of the spot. The scenario proposed by Mermoud et al. corresponds to an infinite series of best-of-2 decision-making problem (i.e., one for each spot) that are tackled in parallel by different subsets of agents of the swarm (i.e., different robot aggregates). The quality of each spot is determined by an internal preference factor that is represented by the measured light. The proposed strategy has been derived following a bottom–up, multi-level modeling methodology that encompasses physics-based simulations, chemical reaction networks, and continuous ODE approximation (Mermoud et al. 2010, 2014).

2.4.2.2 Navigation-Based Approaches

Navigation-based approaches are control strategies that allow a swarm of robots to navigate an environment towards one or more regions of interest. Navigation algorithms have been extensively studied in the swarm robotics literature. However, not all of them provide a swarm with collective decision-making capabilities. For examples, navigation algorithms based on hop-count strategies have been proposed to find the shortest-path connecting a pair of locations (Payton et al. 2001; Szymanski et al. 2006). However, these strategies are incapable of selecting a unique path when there are two or more paths with equal length and thus fail to make a collective decision (Campo et al. 2010b). From a broader perspective, navigation-based approaches include also flocking whereby robots have to agree on a common direction of motion (Ferrante et al. 2012; Nembrini et al. 2002; Spears et al. 2004; Turgut et al. 2008). However, as discussed in the introduction of this chapter, these control strategies are generally studied in experimental setups corresponding to continuous consensus achievement problems (i.e., best-of-∞). In the following, we consider navigation algorithms applied to discrete consensus achievement problems.

Schmickl and Crailsheim took inspiration from the trophallactic behavior of honeybee swarms, *Apis mellifera* L. (Camazine et al. 1998; Korst and Velthuis 1982), and proposed a virtual gradient and navigation strategy that provides a swarm of robots with collective decision-making capabilities. *Trophallaxis* refers to the direct, mouth-to-mouth exchange of food between two honeybees (or other social insects). The authors studied the proposed strategy in a binary aggregation scenario with two spots of different size (Schmickl et al. 2007) and in a foraging scenario reminiscent of the double-bridge problem (Schmickl and Crailsheim 2006, 2008). Robots explore their environment searching for resources (i.e., aggregation spots, foraging patches). Once a robot finds a resource, it loads a certain amount of virtual nectar. As the robot moves in the environment, it spreads and receives virtual nectar to and from other neighboring robots. This behavior allows robots to create a virtual gradient of nectar that can be used by robots to navigate back and forth a pair of locations following the shortest of two paths or to orient towards the largest of two aggregation areas. In both cases, the quality of each option is determined solely by an environmental bias factor (i.e., the length of a path and the size of the aggregation area) which influences the rate of diffusion of virtual nectar. This trophallaxis-inspired strategy has been studied later using models of Brownian motion (Hamann 2010; Hamann and Wörn 2008). The authors defined both a Langevin equation (i.e., a microscopic model) to describe the motion of an individual agent and a Fokker–Planck equation (i.e., a macroscopic model) to model the motion of the entire swarm finding a good qualitative agreement with simulated dynamics of the trophallaxis-inspired strategy.

Garnier et al. (2007) studied a robot control strategy that closely mimics the pheromone-laying behavior characterizing foraging in many ant species (Deneubourg and Goss 1989; Goss et al. 1989). The authors considered a foraging scenario similar to the double-bridge problem where two areas are connected by a pair of paths of equal length (i.e., the options of a best-of-2 symmetry-breaking problem). During robot experiments, pheromone is emulated by means of an external tracking infrastructure interfaced with a light projector that manages both the laying of pheromone and its evaporation. The robots can perceive pheromone trails by means of a pair of light sensors and can recognize the two target areas by means of IR beacons. In the absence of a trail, a robot moves randomly in the environment avoiding obstacles. When perceiving a trail, the robot starts following the trail and depositing pheromone which evaporates with an exponential decay. In their study, the authors show that using this strategy the robots of a swarm are capable to make a consensus decision for one of the two paths. However, the implementation of pheromone-inspired mechanisms on a robotic platform (Fujisawa et al. 2014) still represents a challenge with current technologies which prevents its employment in real-world robotic applications.

Campo et al. (2010b) proposed a solution to the above limitations of pheromone-inspired mechanisms for the case of chain-based navigation systems. In their research work, the robots of the swarm form a pair of chains leading to 2 different locations. Each chain identifies a path and each path has different length (i.e., 2 options with quality defined by an environmental bias factor). The authors proposed a collective decision-making strategy to select the shortest of the 2 paths that is based on virtual pheromones. Robots in a chain can communicate with their 2 immediate neighbors forming a communication network. Virtual ants navigate through the network and lay virtual pheromone eventually leading to the identification and selection of the shortest path.

Gutiérrez et al. (2009) proposed a navigation strategy called *social odometry* that allows a robot of a swarm to keep an estimate of its current location with respect to a certain area of interest. A robot has an estimate of its current location and a measure of confidence about its belief that decreases with the traveled distance. Upon encountering a neighboring robot, they both exchange their location estimates and confidence measures. Successively, each of the two robots updates its current location estimate by averaging its current location with that of its neighbor weighted by the respective measures of confidence. Using social odometry, Gutiérrez et al. (2010) studied a foraging scenario characterized by two foraging patches (i.e., the options of a best-of-2 problem) positioned at different distance (i.e., option quality) from a central retrieval area. The authors find that the weighted mean underlying social odometry favors the selection by the swarm of the closest foraging patch due to the fact that robots traveling to that patch have higher confidence in their location estimates. In this strategy, the option quality is determined by a combination of an internal preference factor with an environmental bias factor. The internal preference is represented by the level of confidence because it is derived by each robot from its own measured distance. The environmental bias is represented by the distance of a patch from the retrieval area because patches that are closer to the retrieval area are easier to discover by robots and are therefore of lesser cost. Due to the presence of noise, social odometry allows a swarm of robots to find consensus on a common foraging patch also in a symmetric setup where the two patches are positioned at the same distance from the retrieval area.

2.5 Top–Down Design Approaches

In this section, we consider research studies that developed collective decision-making strategies using a top–down design approach. All research studies reviewed below make use of automatic optimization approaches to design robot controllers for specific application scenarios.

2.5.1 Evolutionary Robotics

As for most collective behaviors studied in swarm robotics (Brambilla et al. 2013), collective decision-making systems have been also developed by means of automatic design approaches. The typical automatic design approach is represented by evolutionary robotics (Bongard 2013; Nolfi and Floreano 2000) where optimization methods from evolutionary computation (Back et al. 1997) are used to evolve a population of agent controllers following the Darwinian principles of recombination, mutation, and natural selection. Generally, the individual robot controller is represented by an artificial neural network that maps the sensory perceptions of a robot (i.e., input of the neural network) to appropriate actions of its actuators (i.e., output of the neural network). The parameters of the neural network are evolved to tackle a specific application scenario by opportunely defining a fitness function on a per-case base; the fitness function is then used to evaluate the quality of each controller and to drive the evolutionary optimization process.

Evolutionary robotics has been successfully applied to address a number of collective decision-making scenarios. Trianni and Dorigo (2005) evolved a collective behavior that allows a swarm of physically-connected robots to discriminate the type of holes present on the arena surface based on their perceived width and to decide weather to cross the hole (i.e., the hole is sufficiently narrow to be safely crossed) or to avoid it by changing the motion direction (i.e., the hole is too risky to cross). Similarly, Trianni et al. (2007) considered a collective decision-making scenario where a swarm of robots need to collectively evaluate the surrounding environment and determine weather there are physical obstacles that requires cooperation in the form of a self-assembly or, alternatively, if robots can escape obstacles independently of each other.

Francesca et al. (2012, 2014) applied methods from evolutionary robotics to a binary aggregation scenario similar to that studied in Campo et al. (2010a); Garnier et al. (2008, 2009) but with shelters of equal size (i.e., a symmetry-breaking problem). The authors compared the performance of the evolved controller with theoretical predictions of existing mathematical models (Amé et al. 2006); however, their results show a poor agreement between the two models due to the fact that artificial evolution was capable to exploit specific features (e.g., geometric symmetries) present in the simulated environment.

Evolutionary robotics can be successfully applied to the design of collective decision-making systems. However, its use as a design approach suffers of several drawbacks. For example, artificial evolution is a computationally intensive process and the designer is required to perform it for each specific scenario. Artificial evolution may suffer from over-fitting whereby a successfully evolved controller performs well in simulation but poorly on real robots. This phenomenon is also known as the reality gap (Jakobi et al. 1995; Koos et al. 2013). Moreover, artificial evolution does not provide guarantees on the optimality of the resulting robot controller (Bongard 2013). The robot controller, being ultimately a black-box model, is difficult to model and analyze mathematically (Francesca et al. 2012). As a consequence, in general the designer cannot maintain and improve the designed solutions (Matarić and Cliff 1996; Trianni and Nolfi 2011).

2.5.2 Automatic Modular Design

More recently, Francesca et al. (2014) proposed an automatic design method, called AutoMoDe, that provides a white-box alternative to evolutionary robotics. The robot controllers designed using AutoMoDe are behavior-based and have the form of a probabilistic finite-state machine. Robot controllers are obtained by combining a set of predefined modules (e.g., random walk, phototaxis) using an optimization process that, similarly to evolutionary robotics, is driven by a fitness function defined by the designer for each specific scenario.

Using AutoMoDe, Francesca et al. (2014) designed an aggregation strategy for the same scenario as in Campo et al. (2010a); Garnier et al. (2008, 2009). In their experimental setup, the collective decision-making problem corresponds to a binary symmetry-breaking scenario where the swarm needs to select one of two equally good aggregation spots. The resulting robot controller proceeds as follow. A robot starts in the attraction state in which its goal is to get close to other robots. When perceiving an aggregation spot, the robot stops moving. Once stopped, the robot has a fixed probability for unit of time to return to the attraction state and start moving again. Additionally, the robot may transit to the attraction state in the case in which it has been pushed out of the aggregation spot by other robots.

2.6 Discussion

In this chapter, we introduced the reader to several aspects of collective decision making. We followed Brambilla et al. (2013) and organized collective decisions in consensus achievement and task allocation (Gerkey and Matarić 2004) depending on the purpose of the swarm. We showed how decision-making problems requiring the achievement of consensus can be further distinguished in two classes, discrete and continuous, depending on the granularity of the available options. For the case of discrete consensus achievement, we formally defined the structure of the best-of-n problem and showed how this general framework covers a large number of specific application scenarios. Finally, we reviewed the principal research contributions in swarm robotics that focus on discrete consensus achievement problems. We divided our literature review in two parts, bottom–up (see Sect. 2.4) and top–down (see Sect. 2.5) design approaches. For each part, we further distinguished collective decision-making systems and obtained five different categories: opinion-based, aggregation-based, navigation-based, evolutionary robotics, and automatic modular design.

Aggregation-based approaches to collective decision making have the advantage of functioning without the need of communication by exploiting implicit information transfer. However, aggregation as a means of communicating one own opinion provides a viable solution only when the options of the best-of-n problem are clearly separated in space from each other. Similarly, navigation-based approaches can be applied only to scenarios in which the discrete consensus achievement problem requires the swarm to find the shortest-path connecting different locations. In contrast, automatic design approaches as evolutionary robotics and automatic modular design have the potential to be applied to a larger set of consensus achievement scenarios. Evolutionary robotics, however, might suffer from the reality-gap between simulated and real robots. Moreover, it is difficult to derive predictive mathematical models for systems designed using artificial evolution. This latter limitation might also affect automatic modular design depending on the complexity of the resulting probabilistic finite-state machines. Opinion-based approaches offer a more general design methodology that can be applied to different application scenarios. This higher level of generality, however, requires explicit information transfer and can be obtained only at the cost of robot-to-robot communication.

As introduced in Sect. 2.2, the definition of quality of an option in the best-of-n problem depends on the specific application scenario. Nonetheless, we showed that the option quality can be determined by a combination of factors of two types, internal preference of individual agents and environmental bias indirectly affecting the behavior of the swarm. Figure 2.4 illustrates how different combinations of internal

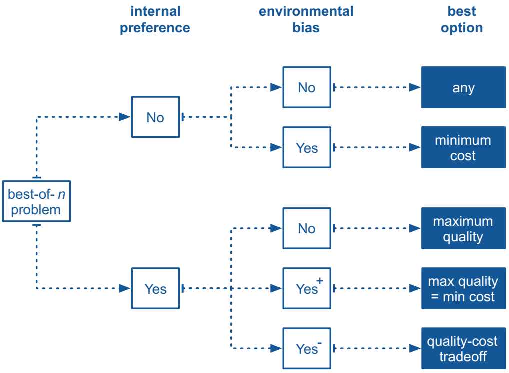

Fig. 2.4 Classification of consensus achievement scenarios corresponding to the best-of-n problem. The schema illustrates how different combinations of internal preference factors and environmental bias factors influences the best option of the decision-making problem

preference and environmental bias factors determine the best option of the best-of-n problem. When the option quality is independent of internal agent preferences and of environmental bias, the best-of-n problem reduces to a symmetry-breaking problem. When the option quality is independent of internal agent preferences and is solely subject to environmental bias, the decision-making problem reduces to finding the option of minimum cost and can be tackled using collective decision-making strategies that do not require agents to directly measure the quality of each option. In the opposite case, i.e., when the option quality depends only on internal preference factors, the best option corresponds to that with highest quality as directly measured by individual agents. Finally, when both factor types coexist, we distinguish between scenarios where the relation between environmental bias and internal preference factors is positive and scenarios in which it is negative. In the first case the best option corresponds to that with highest quality as measured by individual agents which has also minimum cost. In the second case, environmental bias factors influence negatively the option quality and the decision-making problem requires a compromise between option quality and cost.

The taxonomy illustrated in Fig. 2.4 provides us with different means to interpret the swarm robotics literature. We conclude this chapter by reorganizing the research studies reviewed in Sects. 2.4 and 2.5 according to this new taxonomy as shown in Table 2.1. We distinguish research studies in five different categories determined by the specific coupling between internal preference factors and environmental bias

Table 2.1 Classification of the swarm robotics literature according to the combination of factors that determines the quality of the options of the best-of-n problem

Internal preference	Environmental bias	Research lines/studies
No	No	**i**. Wessnitzer and Melhuish (2003) **ii**. Garnier et al. (2007) **iii**. Garnier et al. (2009), Brambilla et al. (2014) **iv**. Hamann et al. (2012), Hamann (2013) **v**. Francesca et al. (2012, 2014)
	Yes	**i**. Hamann (2010); Hamann and Wörn (2008); Schmickl and Crailsheim (2006, 2008); Schmickl et al. (2007) **ii**. Campo et al. (2010b) **iii**. Massink et al. (2013); Montes de Oca et al. (2011); Scheidler (2011) **iv**. Brutschy et al. (2012); Scheidler et al. (2016) **v**. Reina et al. (2014, 2015a)
Yes	No	**i**. Parker and Zhang (2009, 2010, 2011) **ii**. Mermoud et al. (2010, 2014)
	Yes, positive bias	**i**. Garnier et al. (2009) **ii**. Schmickl et al. (2009a, b), Arvin et al. (2012, 2014) **iii**. Gutiérrez et al. (2010)
	Yes, negative bias	**i**. Campo et al. (2010a) **ii**. Reina et al. (2015b)

factors that shapes the best-of-n problem. For each category, we further group the literature in separate lines of research, where each line of research is centered around a particular collective decision-making strategy. In our endeavor, we could not assign some research studies based on evolutionary robotics to a distinct category owing to the difficulty of understanding the precise functioning of the underlying neural networks. The first two categories in Table 2.1, namely, symmetry-breaking problems and problems where the option quality is defined only by environmental bias factors, are covered by a large number of research studies. Moreover, these studies are distributed in five separate lines of research for each category. Differently, the swarm robotics literature has considered a significantly smaller number of studies that focused on the best-of-n problem in the case in which the option quality is determined by an internal preference factor. The majority of these studies considered application scenarios where the option quality is either independent of environmental bias factors or it is positively influenced by them. Note that collective decision-making strategies developed for the former of these two categories directly apply to the latter due to the positive influence of environmental bias factors. The last category, i.e., research studies considering application scenarios where the internal agent preferences are negatively influenced by environmental biases, is the less developed area of study in the literature of discrete consensus achievement with only two research contributions. From the perspective of the designer, this category represents application scenarios with the highest level of complexity that require design solutions able to compensate the negative influence of environmental bias.

References

J.-M. Amé, J. Halloy, C. Rivault, C. Detrain, J.-L. Deneubourg, Collegial decision making based on social amplification leads to optimal group formation. Proc. Natl. Acad. Sci. **103**(15), 5835–5840 (2006)

F. Arvin, A.E. Turgut, S. Yue, Fuzzy-based aggregation with a mobile robot swarm, ed. By M. Dorigo, M. Birattari, C. Blum, A.L. Christensen, A.P. Engelbrecht, R. Groß, T. Stützle. *Swarm Intelligence*, vol. 7461. LNCS (Springer, 2012), pp. 346–347

F. Arvin, A.E. Turgut, F. Bazyari, K.B. Arikan, N. Bellotto, S. Yue, Cue-based aggregation with a mobile robot swarm: a novel fuzzy-based method. Adapt. Behav. **22**(3), 189–206 (2014)

T. Back, D.B. Fogel, Z. Michalewicz (eds.), *Handbook of Evolutionary Computation* (IOP Publishing Ltd., Bristol, 1997)

D.S. Bernstein, R. Givan, N. Immerman, S. Zilberstein, The complexity of decentralized control of markov decision processes. Math. Oper. Res. **27**(4), 819–840 (2002)

J.C. Bongard, Evolutionary robotics. Commun. ACM **56**(8), 74–83 (2013)

M.H. Bowling, B. Browning, M.M. Veloso, Plays as effective multiagent plans enabling opponent-adaptive play selection, in *Proceedings of the 14th International Conference on Automated Planning & Scheduling* (AAAI Press, 2004), pp. 376–383

M. Brambilla, E. Ferrante, M. Birattari, M. Dorigo, Swarm robotics: a review from the swarm engineering perspective. Swarm Intell. **7**(1), 1–41 (2013)

M. Brambilla, A. Brutschy, M. Dorigo, M. Birattari, Property-driven design for robot swarms: a design method based on prescriptive modeling and model checking. ACM Trans. Auton. Adapt. Syst. **9**(4), 17:1–17:28 (2014)

A. Brutschy, A. Scheidler, E. Ferrante, M. Dorigo, M. Birattari, “Can ants inspire robots?” Self-organized decision making in robotic swarms, in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (IEEE Press, 2012), pp. 4272–4273

J. Buhl, D.J.T. Sumpter, I.D. Couzin, J.J. Hale, E. Despland, E.R. Miller, S.J. Simpson, From disorder to order in marching locusts. Science **312**(5778), 1402–1406 (2006)

S. Camazine, K. Crailsheim, N. Hrassnigg, G.E. Robinson, B. Leonhard, H. Kropiunigg, Protein trophallaxis and the regulation of pollen foraging by honey bees (Apis mellifera l.). Apidologie **29**(1–2), 113–126 (1998)

S. Camazine, P.K. Visscher, J. Finley, R.S. Vetter, House-hunting by honey bee swarms: collective decisions and individual behaviors. Insectes Sociaux **46**(4), 348–360 (1999)

S. Camazine, J.-L. Deneubourg, N.R. Franks, J. Sneyd, G. Theraulaz, E. Bonabeau, *Self-Organization in Biological Systems* (Princeton University Press, Princeton, 2001)

A. Campo, S. Garnier, O. Dédriche, M. Zekkri, M. Dorigo, Self-organized discrimination of resources. PLoS ONE **6**(5), e19888 (2010a)

A. Campo, Á. Gutiérrez, S. Nouyan, C. Pinciroli, V. Longchamp, S. Garnier, M. Dorigo, Artificial pheromone for path selection by a foraging swarm of robots. Biol. Cybern. **103**(5), 339–352 (2010b)

Y. Cao, W. Ren, Distributed coordinated tracking with reduced interaction via a variable structure approach. IEEE Trans. Autom. Control **57**(1), 33–48 (2012)

C. Castellano, S. Fortunato, V. Loreto, Statistical physics of social dynamics. Rev. Mod. Phys. **81**, 591–646 (2009)

N. Correll, A. Martinoli, Modeling and designing self-organized aggregation in a swarm of miniature robots. Int. J. Robot. Res. **30**(5), 615–626 (2011)

V. Crespi, A. Galstyan, K. Lerman, Top-down vs bottom-up methodologies in multi-agent system design. Auton. Robots **24**(3), 303–313 (2008)

A. Czirók, T. Vicsek, Collective behavior of interacting self-propelled particles. Phys. A: Stat. Mech. Appl. **281**(1–4), 17–29 (2000)

A. Czirók, A.-L. Barabási, T. Vicsek, Collective motion of self-propelled particles: kinetic phase transition in one dimension. Phys. Rev. Lett. **82**, 209–212 (1999)

H. de Vries, J.C. Biesmeijer, Self-organization in collective honeybee foraging: emergence of symmetry breaking, cross inhibition and equal harvest-rate distribution. Behav. Ecol. Sociobiol. **51**(6), 557–569 (2002)

J.-L. Deneubourg, S. Goss, Collective patterns and decision-making. Ethol. Ecol. Evol. **1**(4), 295–311 (1989)

E. Ferrante, *Information transfer in a flocking robot swarm*. Ph.D. thesis, Université libre de Bruxelles, Brussels, BE (2013)

E. Ferrante, A.E. Turgut, C. Huepe, A. Stranieri, C. Pinciroli, M. Dorigo, Self-organized flocking with a mobile robot swarm: a novel motion control method. Adapt. Behav. **20**(6), 460–477 (2012)

G. Francesca, M. Brambilla, V. Trianni, M. Dorigo, M. Birattari, Analysing an evolved robotic behaviour using a biological model of collegial decision making, ed. By T. Ziemke, C. Balkenius, J. Hallam. *From Animals to Animats 12*, vol. 7426, LNCS (Springer, 2012), pp. 381–390

G. Francesca, M. Brambilla, A. Brutschy, V. Trianni, M. Birattari, AutoMoDe: a novel approach to the automatic design of control software for robot swarms. Swarm Intell. **8**(2), 89–112 (2014)

N.R. Franks, S.C. Pratt, E.B. Mallon, N.F. Britton, D.J.T. Sumpter, Information flow, opinion polling and collective intelligence in house-hunting social insects. Philos. Trans. R. Soc. B: Biol. Sci. **357**(1427), 1567–1583 (2002)

N.R. Franks, T.O. Richardson, N. Stroeymeyt, R.W. Kirby, W.M.D. Amos, P.M. Hogan, J.A.R. Marshall, T. Schlegel, Speed-cohesion trade-offs in collective decision making in ants and the concept of precision in animal behaviour. Anim. Behav. **85**(6), 1233–1244 (2013)

R. Fujisawa, S. Dobata, K. Sugawara, F. Matsuno, Designing pheromone communication in swarm robotics: group foraging behavior mediated by chemical substance. Swarm Intell. **8**(3), 227–246 (2014)

S. Galam, Sociophysics: a review of Galam models. Int. J. Mod. Phys. C **19**(03), 409–440 (2008)

S. Garnier, F. Tache, M. Combe, A. Grimal, G. Theraulaz, Alice in pheromone land: an experimental setup for the study of ant-like robots. Proc. IEEE Swarm Intell. Symp. SIS **2007**, 37–44 (2007)

S. Garnier, C. Jost, J. Gautrais, M. Asadpour, G. Caprari, R. Jeanson, A. Grimal, G. Theraulaz, The embodiment of cockroach aggregation behavior in a group of micro-robots. Artif. Life **14**(4), 387–408 (2008)

S. Garnier, J. Gautrais, M. Asadpour, C. Jost, G. Theraulaz, Self-organized aggregation triggers collective decision making in a group of cockroach-like robots. Adapt. Behav. **17**(2), 109–133 (2009)

M. Gauci, J. Chen, W. Li, T.J. Dodd, R. Groß, Self-organized aggregation without computation. Int. J. Robot. Res. **33**(8), 1145–1161 (2014)

B.P. Gerkey, M.J. Matarić, A formal analysis and taxonomy of task allocation in multi-robot systems. Int. J. Robot. Res. **23**(9), 939–954 (2004)

S. Goss, S. Aron, J.-L. Deneubourg, J.M. Pasteels, Self-organized shortcuts in the argentine ant. Naturwissenschaften **76**(12), 579–581 (1989)

P. Grodzicki, M. Caputa, Social versus individual behaviour: a comparative approach to thermal behaviour of the honeybee (Apis mellifera L.) and the american cockroach (Periplaneta americana L.). J. Insect Physiol. **51**(3), 315–322 (2005)

Á. Gutiérrez, A. Campo, F. Santos, F. Monasterio-Huelin Maciá, M. Dorigo, Social odometry: imitation based odometry in collective robotics. Int. J. Adv. Robot. Syst. **6**(2), 129–136 (2009)

Á. Gutiérrez, A. Campo, F. Monasterio-Huelin, L. Magdalena, M. Dorigo, Collective decision-making based on social odometry. Neural Comput. Appl. **19**(6), 807–823 (2010)

H. Hamann, *Space-time continuous models of swarm robotic systems: Supporting global-to-local programming*, vol. 9, Cognitive Systems Monographs (Springer, Berlin, 2010)

H. Hamann, Towards swarm calculus: urn models of collective decisions and universal properties of swarm performance. Swarm Intell. **7**(2–3), 145–172 (2013)

H. Hamann, H. Wörn, A framework of space-time continuous models for algorithm design in swarm robotics. Swarm Intell. **2**(2–4), 209–239 (2008)

H. Hamann, T. Schmickl, H. Wörn, K. Crailsheim, Analysis of emergent symmetry breaking in collective decision making. Neural Comput. Appl. **21**(2), 207–218 (2012)

Y. Hatano, M. Mesbahi, Agreement over random networks. IEEE Trans. Autom. Control **50**(11), 1867–1872 (2005)

S. Hauert, S. Leven, M. Varga, F. Ruini, A. Cangelosi, J.C. Zufferey, D. Floreano, Reynolds flocking in reality with fixed-wing robots: communication range vs. maximum turning rate, in *Proceedings of the 2011 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2011*, pp. 5015–5020 (2011)

J.M. Hereford, Analysis of a new swarm search algorithm based on trophallaxis, in *2010 IEEE Congress on Evolutionary Computation, CEC* (IEEE Press, 2010), pp. 1–8

O. Holland, J. Woods, R.D. Nardi, A. Clark, Beyond swarm intelligence: the UltraSwarm, in *Proceedings of the 2005 IEEE Swarm Intelligence Symposium, SIS 2005*, pp. 217–224 (2005)

N. Jakobi, P. Husbands, I. Harvey, Noise and the reality gap: the use of simulation in evolutionary robotics, ed. By F. Morán, A. Moreno, J.J. Merelo, P. Chacón. *Proceedings of the Third European Conference on Advances in Artificial Life*, vol. 929. LNCS (Springer, 1995), pp. 704–720

R. Jeanson, S. Blanco, R. Fournier, J.-L. Deneubourg, V. Fourcassié, G. Theraulaz, A model of animal movements in a bounded space. J. Theor. Biol. **225**(4), 443–451 (2003)

R. Jeanson, C. Rivault, J.-L. Deneubourg, S. Blanco, R. Fournier, C. Jost, G. Theraulaz, Self-organized aggregation in cockroaches. Anim. Behav. **69**(1), 169–180 (2005)

S. Kernbach, R. Thenius, O. Kernbach, T. Schmickl, Re-embodiment of honeybee aggregation behavior in an artificial micro-robotic system. Adapt. Behav. **17**(3), 237–259 (2009)

H. Kitano, M. Asada, Y. Kuniyoshi, I. Noda, E. Osawa, H. Matsubara, RoboCup: a challenge problem for AI. AI Mag. **18**(1), 73–85 (1997)

J.R. Kok, N. Vlassis, Distributed decision making of robotic agents, in *Proceedings of the 8th Annual Conference of the Advanced School for Computing and Imaging*, pp. 318–325 (2003)

J.R. Kok, M.T. Spaan, N. Vlassis, Multi-robot decision making using coordination graphs, in *Proceedings of the 11th International Conference on Advanced Robotics, ICAR*, pp. 1124–1129 (2003)

S. Koos, J.-B. Mouret, S. Doncieux, The transferability approach: crossing the reality gap in evolutionary robotics. IEEE Trans. Evol. Comput. **17**(1), 122–145 (2013)

S. Kornienko, O. Kornienko, C. Constantinescu, M. Pradier, P. Levi, Cognitive micro-agents: individual and collective perception in microrobotic swarm, in *Proceedings of the IJCAI-05 Workshop on Agents in real-time and dynamic environments, Edinburgh, UK*, pp. 33–42 (2005a)

S. Kornienko, O. Kornienko, P. Levi, Minimalistic approach towards communication and perception in microrobotic swarms, in *2005 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pp. 2228–2234 (2005b)

P.J.A.M. Korst, H.H.W. Velthuis, The nature of trophallaxis in honeybees. Insectes Sociaux **29**(2), 209–221 (1982)

P.L. Krapivsky, S. Redner, Dynamics of majority rule in two-state interacting spin systems. Phys. Rev. Lett. **90**, 238701 (2003)

R. Lambiotte, J. Saramäki, V.D. Blondel, Dynamics of latent voters. Phys. Rev. E **79**, 046107 (2009)

J.A.R. Marshall, R. Bogacz, A. Dornhaus, R. P̃lanqué, T. Kovacs, N.R. Franks, On optimal decision-making in brains and social insect colonies. J. R. Soc. Interface **6**(40), 1065–1074 (2009)

M. Massink, M. Brambilla, D. Latella, M. Dorigo, M. Birattari, On the use of Bio-PEPA for modelling and analysing collective behaviours in swarm robotics. Swarm Intell. **7**(2–3), 201–228 (2013)

M. Matarić, D. Cliff, Challenges in evolving controllers for physical robots. Robot. Auton. Syst. **19**(1), 67–83 (1996)

G. Mermoud, L. Matthey, W.C. Evans, A. Martinoli, Aggregation-mediated collective perception and action in a group of miniature robots, in *Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 2 - Volume 2, AAMAS '10* (IFAAMAS, 2010), pp. 599–606

G. Mermoud, U. Upadhyay, W.C. Evans, A. Martinoli, Top-down vs. bottom-up model-based methodologies for distributed control: a comparative experimental study, ed. By O. Khatib, V. Kumar, G. Sukhatme.*Experimental Robotics*, vol. 79. STAR (Springer, 2014), pp. 615–629

M. Mesbahi, M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*, Princeton Series in Applied Mathematics (Princeton University Press, Princeton, 2010)

M.A. Montes de Oca, E. Ferrante, A. Scheidler, C. Pinciroli, M. Birattari, M. Dorigo, Majority-rule opinion dynamics with differential latency: a mechanism for self-organized collective decision-making. Swarm Intell. **5**, 305–327 (2011)

J. Nembrini, A. Winfield, C. Melhuish, Minimalist coherent swarming of wireless networked autonomous mobile robots, in *Proceedings of the Seventh International Conference on Simulation of Adaptive Behavior on From Animals to Animats, ICSAB* (MIT Press, 2002), pp. 373–382

S. Nolfi, D. Floreano, *Evolutionary Robotics. The Biology, Intelligence, and Technology of Self-organizing Machines* (MIT Press, Cambridge, 2000)

D. Pais, P.M. Hogan, T. Schlegel, N.R. Franks, N.E. Leonard, J.A.R. Marshall, A mechanism for value-sensitive decision-making. PLoS ONE **8**(9), 1–9 (2013)

C.A.C. Parker, H. Zhang, Cooperative decision-making in decentralized multiple-robot systems: the best-of-n problem. IEEE/ASME Trans. Mechatron. **14**(2), 240–251 (2009)

C.A.C. Parker, H. Zhang, Collective unary decision-making by decentralized multiple-robot systems applied to the task-sequencing problem. Swarm Intell. **4**, 199–220 (2010)

C.A.C. Parker, H. Zhang, Biologically inspired collective comparisons by robotic swarms. Int. J. Robot. Res. **30**(5), 524–535 (2011)

D. Payton, M. Daily, R. Estowski, M. Howard, C. Lee, Pheromone robotics. Auton. Robots **11**(3), 319–324 (2001)

D.V. Pynadath, M. Tambe, The communicative multiagent team decision problem: analyzing teamwork theories and models. J. Artif. Intell. Res. **16**(1), 389–423 (2002)

C.R. Reid, S. Garnier, M. Beekman, T. Latty, Information integration and multiattribute decision making in non-neuronal organisms. Anim. Behav. **100**, 44–50 (2015)

A. Reina, M. Dorigo, V. Trianni, Towards a cognitive design pattern for collective decision-making, ed. By M. Dorigo, M. Birattari, S. Garnier, H. Hamann, M.A. Montes de Oca, C. Solnon, T. Stützle. *Swarm Intelligence*, vol. 8667. LNCS (Springer, 2014), pp. 194–205

A. Reina, R. Miletitch, M. Dorigo, V. Trianni, A quantitative micro-macro link for collective decisions: the shortest path discovery/selection example. Swarm Intell. **9**(2–3), 75–102 (2015a)

A. Reina, G. Valentini, C. Fernández-Oto, M. Dorigo, V. Trianni, A design pattern for decentralised decision making. PLoS ONE **10**(10), e0140950 (2015b)

W. Ren, R.W. Beard, *Distributed consensus in multi-vehicle cooperative control: theory and applications*, Communications and control engineering (Springer, London, 2008)

W. Ren, R. W. Beard, E.M. Atkins, A survey of consensus problems in multi-agent coordination, in *Proceedings of the 2005 American Control Conference*, vol. 3 (IEEE Press, 2005), pp. 1859–1864

C.W. Reynolds, Flocks, herds and schools: a distributed behavioral model. ACM SIGGRAPH Comput. Graph. **21**(4), 25–34 (1987)

G. Sartoretti, M.-O. Hongler, M. de Oliveira, F. Mondada, Decentralized self-selection of swarm trajectories: from dynamical systems theory to robotic implementation. Swarm Intell. **8**(4), 329–351 (2014)

A. Scheidler, Dynamics of majority rule with differential latencies. Phys. Rev. E **83**, 031116 (2011)

A. Scheidler, A. Brutschy, E. Ferrante, M. Dorigo, The k-unanimity rule for self-organized decision-making in swarms of robots. IEEE Trans. Cybern. **46**(5), 1175–1188 (2016)

I. Schizas, G. Giannakis, S. Roumeliotis, A. Ribeiro, Consensus in ad hoc WSNs with noisy links - Part II: distributed estimation and smoothing of random signals. IEEE Trans. Signal Proc. **56**(4), 1650–1666 (2008a)

I. Schizas, A. Ribeiro, G. Giannakis, Consensus in ad hoc WSNs with noisy links - Part I: distributed estimation of deterministic signals. IEEE Trans. Signal Proc. **56**(1), 350–364 (2008b)

T. Schmickl, K. Crailsheim, Trophallaxis among swarm-robots: a biologically inspired strategy for swarm robotics. First IEEE/RAS-EMBS Int. Conf. Biomed. Robot. Biomech. BioRob **2006**, 377–382 (2006)

T. Schmickl, K. Crailsheim, Trophallaxis within a robotic swarm: bio-inspired communication among robots in a swarm. Auton. Robots **25**(1–2), 171–188 (2008)

T. Schmickl, C. Möslinger, K. Crailsheim, Collective perception in a robot swarm, ed. By E. Şahin, W.M. Spears, A.F. Winfield. *Swarm Robotics*, vol. 4433. LNCS (Springer, 2007), pp. 144–157

T. Schmickl, H. Hamann, H. Wörn, K. Crailsheim, Two different approaches to a macroscopic model of a bio-inspired robotic swarm. Robot. Auton. Syst. **57**(9), 913–921 (2009a)

T. Schmickl, R. Thenius, C. Moeslinger, G. Radspieler, S. Kernbach, M. Szymanski, K. Crailsheim, Get in touch: cooperative decision making based on robot-to-robot collisions. Auton. Agents Multi-Agent Syst. **18**(1), 133–155 (2009b)

T.D. Seeley, P.K. Visscher, T. Schlegel, P.M. Hogan, N.R. Franks, J.A.R. Marshall, Stop signals provide cross inhibition in collective decision-making by honeybee swarms. Science **335**(6064), 108–111 (2012)

O. Soysal, E. Şahin, A macroscopic model for self-organized aggregation in swarm robotic systems, ed. By E. Şahin, W.M. Spears, A.F.T. Winfield. *Swarm Robotics*, vol. 4433. LNCS (Springer, 2007), pp. 27–42

W.M. Spears, D.F. Spears, J.C. Hamann, R. Heil, Distributed, physics-based control of swarms of vehicles. Auton. Robots **17**(2), 137–162 (2004)

D.J.T. Sumpter, *Collective Animal Behavior* (Princeton University Press, Princeton, 2010)

B. Szabó, G.J. Szöllösi, B. Gönci, Z. Jurányi, D. Selmeczi, T. Vicsek, Phase transition in the collective migration of tissue cells: experiment and model. Phys. Rev. E **74**, 061908 (2006)

M. Szymanski, T. Breitling, J. Seyfried, H. Wörn, Distributed shortest-path finding by a micro-robot swarm, ed. By M. Dorigo, L.M. Gambardella, M. Birattari, A. Martinoli, R. Poli, T. Stützle. *Ant Colony Optimization and Swarm Intelligence*, vol. 4150. LNCS (Springer, 2006), pp. 404–411

V. Trianni, M. Dorigo, Emergent collective decisions in a swarm of robots, in *Proceedings of the IEEE Swarm Intelligence Symposium, SIS 2005* (IEEE Press, 2005), pp. 241–248

V. Trianni, S. Nolfi, Engineering the evolution of self-organizing behaviors in swarm robotics: a case study. Artif. Life **17**(3), 183–202 (2011)

V. Trianni, R. Groß, T.H. Labella, E. Şahin, M. Dorigo, Evolving aggregation behaviors in a swarm of robots, ed. By W. Banzhaf, J. Ziegler, T. Christaller, P. Dittrich, J.T. Kim. *Advances in Artificial Life*, vol. 2801. LNCS (Springer, 2003), pp. 865–874

V. Trianni, C. Ampatzis, A.L. Christensen, E. Tuci, M. Dorigo, S. Nolfi, From solitary to collective behaviours: decision making and cooperation, ed. By F. Almeida Costa, L.M. Rocha, E. Costa, I. Harvey, A. Coutinho. *Advances in Artificial Life*, vol. 4648. LNCS (Springer, 2007), pp. 575–584

A.E. Turgut, H. Çelikkanat, F. Gökçe, E. Şahin, Self-organized flocking in mobile robot swarms. Swarm Intell. **2**(2), 97–120 (2008)

T. Vicsek, A. Zafeiris, Collective motion. Phys. Rep. **517**(3–4), 71–140 (2012)

T. Vicsek, A. Czirók, E. Ben-Jacob, I. Cohen, O. Shochet, Novel type of phase transition in a system of self-driven particles. Phys. Rev. Lett. **75**, 1226–1229 (1995)

J. Werfel, K. Petersen, R. Nagpal, Designing collective behavior in a termite-inspired robot construction team. Science **343**(6172), 754–758 (2014)

J. Wessnitzer, C. Melhuish, Collective decision-making and behaviour transitions in distributed ad hoc wireless networks of mobile robots: target-hunting, ed. By W. Banzhaf, J. Ziegler, T. Christaller, P. Dittrich, J.T. Kim. *Advances in Artificial Life*, vol. 2801. LNCS (Springer, 2003), pp. 893–902

C.A. Yates, R. Erban, C. Escudero, I.D. Couzin, J. Buhl, I.G. Kevrekidis, P.K. Maini, D.J.T. Sumpter, Inherent noise can facilitate coherence in collective swarm motion. Proc. Nat. Acad. Sci. **106**(14), 5464–5469 (2009)

http://www.springer.com/978-3-319-53608-8

Achieving Consensus in Robot Swarms
Design and Analysis of Strategies for the best-of-n
Problem
Valentini, G.
2017, XIV, 146 p. 46 illus., 37 illus. in color., Hardcover
ISBN: 978-3-319-53608-8