

Emergence of Cooperation: State of the Art

Geoff Nitschke

Computational Intelligence Group
Department of Computer Science
Vrije Universiteit Amsterdam
De Boelelaan 1081a, 1081 HV
Amsterdam, The Netherlands
nitschke@cs.vu.nl

Abstract This review presents a review of prevalent results within research pertaining to emergent cooperation in biologically inspired artificial social systems. Results reviewed maintain particular reference to biologically inspired design principles, given that current mathematical and empirical tools have provided only a partial insight into elucidating mechanisms responsible for emergent cooperation, and then only in systems of an abstract nature. This review aims to provide an overview of important and disparate research contributions that investigate utilization of biologically inspired concepts such as emergence, evolution, and self-organization as a means of attaining cooperation in artificial social systems. An introduction and overview of emergent cooperation in artificial life is presented, followed by a survey of emergent cooperation in swarm-based systems, the pursuit-evasion domain, and RoboCup soccer. The final section draws conclusions regarding future directions of emergent cooperation as a problem-solving methodology that is potentially applicable in a wide range of problem domains. Within each of these sections and their respective themes of research, the mechanisms deemed to be responsible for emergent cooperation are elucidated and their key limitations highlighted. The review concludes that current studies in emergent cooperative behavior are limited by a lack of situated and embodied approaches, and by the research infancy of current biologically inspired design approaches. Despite these limiting factors, emergent cooperation maintains considerable future potential in a wide variety of application domains where systems composed of many interacting components must cooperatively perform unanticipated global tasks.

Keywords

Emergence, cooperative behavior,
swarm intelligence, RoboCup soccer,
pursuit-evasion

I Introduction

The global behavior and complexity of biological systems such as ant colonies is considered to be an emergent property of the interactions between the different agents that make up the whole system. Desirable emergent behavior has been observed in many biological systems, though reproducing the conditions leading to the emergence of such behaviors in artificial systems has proved to be difficult, as there is potential for the emergence of undesirable behaviors. It is therefore essential to be able to understand the mechanisms of emergent cooperative behavior in these systems. To date, research that qualitatively measures and evaluates mechanisms that underlie and produce emergent

cooperative behavior in artificial and real-world¹ systems remains largely in the stage of research infancy.

The concept of emergent behavior, which is deemed to be one of the key components of artificial life research, has propagated many ideas about emergent cooperative behavior in biological systems. These ideas have now been adopted by roboticists and computer scientists alike, and have gained prevalence since the rise of decentralized information systems motivated by the proliferation of global systems such as the Internet. Early research in decentralized systems [25, 15, 153] suggested that complexity at a group level might be attainable with very simple individual agents, with no need for central control. For instance, in the mid twentieth century, Grey Walter² and his colleagues studied turtle-like robots equipped with light and touch sensors and very simple behaviors. When placed together, these robots exhibited complex social behavior in response to each other's movements [30]. A derivative of this idea includes the notion of biologically inspired artificial systems, typically designed using an evolutionary computation methodology such that a global organized behavior emerges from interaction of the system's components [1, 18, 22, 39, 84, 98, 99, 104]. It has been argued by many researchers [54, 57, 68, 95, 105, 109, 24, 133–135, 152, 43] that the use of biologically inspired principles such as evolution and emergence in the purposeful design of complex artificial systems is needed in order to replace ineffective preprogrammed and centralized design methodologies. Throughout this review particular reference is made to research that uses biologically inspired design principles such as artificial evolution, self-organization, and emergence as means of attaining cooperative behavior. With a few exceptions, such as the formalization of emergent cooperative behaviors in multi-robot systems developed by Mataric [106], the research in emergent cooperative behavior is restricted to simulated problem domains, given the inherent complexity of applying evolutionary design principles to collective behaviors in groups of real robots [54, 55, 67]. Hence, this review surveys only research pertaining to the study of emergent cooperative behavior using biologically inspired design principles within simulated problem domains. An important future direction that the review emphasizes is the development of algorithms and design methodologies for the synthesis of desired cooperative behavior, where such behavior is applicable to embodied systems. That is, if emergent cooperative behavior in biologically inspired systems were sufficiently understood, purposeful design of cooperative behavior could be applied to benefit a variety of application domains, including telecommunications [44, 150, 41], aerospace and space exploration [28], and multi-robot systems [100, 114, 112, 128].

The remainder of this review is divided into several sections, organized as follows. Section 2 presents an introduction to overview of the topic of emergent cooperation in artificial life as well as providing a historical perspective of traditional game-theoretic and mathematically based approaches to emergent cooperation. This section highlights the constrained nature of cooperation in many of these domains [86, 12, 14, 2, 50] and states the importance of a synthetic approach that uses concepts such as self-organization and emergence in systems that aim to effectively utilize emergent cooperative behavior. Section 3 reviews pertinent research results pertaining to the simulation of emergent cooperation in swarm-based systems, where the goal of such simulations is to reproduce biological phenomena in task domains that include the optimization of network traffic flow, clustering, self-assembly, and cooperative transport. Section 4 details research that has achieved particular success in attaining cooperative behavior in pursuit-evasion and predator-prey systems. Section 5 outlines several research results on emergent cooperative behavior in the simulated league of RoboCup soccer. The final section draws conclusions regarding beneficial future directions of

¹ The emergence of cooperation has also been studied in a variety of fields unrelated to artificial life systems. For example, the book *The evolution of cooperation* by Axelrod [12] included a chapter where Axelrod examined spontaneous instances of cooperation during trench warfare in World War I. It was stated that troops of one side would shell the other side with mortars, but would often do so on a rigid schedule, aiming for a specific point in the other side's trenches, thereby allowing the other side to minimize casualties. The enemy would then reciprocate this action. The generals on both sides were satisfied that shelling was occurring and that the war was progressing satisfactorily, while soldiers in the trenches found a way to cooperatively protect each other.

² Grey Walter was a respected neurophysiologist, who in the late 1940s carried out pioneering research on mobile autonomous robots at the Burden Neurological Institute in Bristol, England, as part of his goal to model brain function.

emergent cooperation as a problem-solving tool in biologically inspired artificial systems. The purpose of the final section is to highlight the key open problems that future research must address in order for the full potential of emergent cooperation to be realized.

As a final introductory note, in the literature various researchers have adopted the use of various nomenclatures that are ambiguous in defining the term *cooperation*. Such terminology often suffers from the frame-of-reference problem [132], as it is typically defined according to the perspectives and interests of the researchers conducting the study. Thus, for the purposes of this review, we are concerned not with a definition of cooperation, but rather with research that uses biologically inspired principles within simulated problem domains as a means of motivating multiple agents to collectively solve a predefined problem of a global nature that could not otherwise be solved by an individual agent.

2 Emergent Cooperation in Artificial Life

Emergent behavior is a key topic in artificial life research, given that artificial life typically adopts a bottom-up³ approach to modeling various forms of collective behavior and emergent social phenomena that are observed in biological systems. Social phenomena that have inspired the design of artificial life systems have included flocking behavior in birds [137], schooling behavior in fish [161], and pursuit and evasion behavior between predators and prey [110], as well as cooperative transportation [100] and the building of nests [20, 38] by ants. The key idea supporting these models is to use biological phenomena as an inspiration for design, and then to analyze and elucidate the mechanisms motivating emergent collective behavior.

Traditionally, before the advent of research fields such as artificial life that exploit concepts such as self-organization and emergence, the concept of cooperative behavior was the emphasis of much research in the field of distributed artificial intelligence, though the definition given to cooperation in such research typically does not relate to the concept of emergent cooperation, as is evident in the field of artificial life. Examples of traditional distributed artificial intelligence approaches to the design of cooperative behavior have included formation control in multi-robot systems [8–10], multi-robot control architectures for cooperative toxic-waste cleanup [126], cooperative box-pushing [127, 129], payload transportation [88], autonomous surveillance of multiple moving targets [128], architectures for establishing and controlling cooperation between multiple agents in industrial systems and hospitals [87, 80], mechanisms for multi-agent negotiation and coalition formation [97, 140], and multi-robot exploration of unknown environments [4], as well as generalized multi-robot control architectures for planning and role assignment in transportation tasks [3, 31].

The primary criticism of such distributed artificial intelligence approaches to cooperation is that they typically apply classical single-agent system artificial intelligence methodologies to the context of distributed systems. In particular, such research usually applies a top-down analysis to the construction of cooperation between agents, for example, via formalizing, or otherwise specifying a priori, the mechanisms that lead to specific types of multi-agent cooperation [86]. Additionally, emergent cooperation is commonly studied within abstract problem domains, such as the iterated prisoner's dilemma (IPD)⁴ [12–14] or other multi-agent scenarios [50] that operate within a game theory⁵ domain. Axelrod [12] explored the conditions under which fundamentally selfish agents were

³ A bottom-up process is usually used in the design of behavior-based systems [27], where this process refers to the incremental development of the system's sophistication from simple to complex.

⁴ The IPD is a two-player game of fundamental importance in game theory. Repetition of the game with the same pair of players changes a payoff matrix, favoring long-term cooperation. This topic was explored in the book *The evolution of cooperation* [12]. In the single-play case, defection is always the preferred strategy. Likewise, if the number of remaining plays in the iterated case is disclosed to the players, cooperation ceases to be appealing. For cooperation to remain appealing, the future must be indeterminate for both players.

⁵ Game theory is a branch of mathematics and economics that analyzes interactions with formalized and structured games. The predicted and actual behavior of individuals in these games as well as optimal strategies is studied. Seemingly different types of interactions can be characterized as having similar incentive structures, thus being examples of a particular game [12].

more likely to cooperate spontaneously. To perform this study, Axelrod used the IPD game, which offers a long-term incentive for cooperation but a short-term incentive for defection. The IPD game has also been extended to account for the emergence of cooperation between more than two players [2], as well as different forms of cooperation, such as non-mutual cooperation, where exhibited altruistic behavior increases the likelihood of cooperative behavior being reciprocated at a later time [6, 125]. Similarly, other problem domains have been concocted for the purpose of studying various forms of emergent social phenomena, not only cooperation. For example, Epstein and Axtell [51] devised what they termed the sugar-scape domain, a grid world with a renewable resource that agents inhabiting the world were dependent upon. The authors demonstrated that a vast range of collective phenomena and social behaviors emerged as a result of spatiotemporal interaction of agents inhabiting the sugar-scape environment where individual agents followed simple local rules.

Aside from classical artificial intelligence applications of cooperation to distributed problem domains, game-theoretic models, or those modeling abstract problem domains such as sugar-scape [51] and other similar models [81–83, 40], research on emergent cooperation using situated and embodied agents [132] has received little attention, due to the inherent complexity of having such agents operate in the real world. In fact, biologically inspired mechanisms for behavioral design have primarily been used in studying emergent cooperative behavior in simulated artificial life systems. Using biologically inspired approaches, researchers attempt to create agents that internally simulate or mimic the social behavior and intelligence found in biological systems. Biologically inspired designs are based on theories drawn from natural sciences, including anthropology, cognitive science, developmental psychology, and ethology. Generally, these theories are used to guide the design of an agent's cognitive, behavioral, motivational, motor, and perceptual systems [56].

Two main arguments and sets of guidelines are commonplace in the literature for drawing inspiration from biological systems. First, numerous researchers contend that nature is the best model for the creation of effective adaptive behavior in artificial agents, embodied or otherwise—meaning that, in order for an artificial agent to be understood by humans, it must be able to interact with its environment in a similar manner to living beings, and it must perceive similar aspects of the environment to those that living beings find salient and relevant [162]. The second argument is that an agent's behavioral design must facilitate examination, evaluation, and refinement of particular theories upon which the design of the agent is based. Using such guidelines, numerous researchers have developed models inspired by biological systems, for the purpose of reproducing the mechanisms and principles, such as self-organization and evolution, that lead to cooperative behavior, or of studying cooperative behavior in biological systems by hypothesizing what kinds of mechanisms lead to what kinds of behaviors, and subsequently testing these hypotheses with an artificial implementation.

Concepts such as self-organization, emergence, and evolution are now thought by many researchers to pose a reasonable alternative to traditionalist artificial intelligence design approaches to multi-agent cooperation. For example, artificial evolution has been used successfully for the derivation of cooperative pursuit strategies in the pursuit-evasion domain [69–72], to attain an ecological equilibrium between groups of predators and prey [117], and to generate effective cooperative behavior in a competitive game scenario played within the simulated league of RoboCup soccer [154, 78, 79, 60, 102]. Also, social insect behavior has proven to be an attractive model for distributed systems of various instantiations [37], in that it enables groups of relatively simple agents to perform relatively difficult tasks. A common problem that confounds such research is that cooperative behaviors that emerge from interacting constituents of the system are difficult to analyze, as it is nontrivial to determine what mechanisms are responsible for what behaviors. The binding theme of this review is that methodologies that are founded upon concepts such as emergence, self-organization, and evolution must be amenable to mathematical and statistical analysis and evaluation in order for resultant emergent behaviors to be effective and meaningful in the context of problem solving in any system, embodied or otherwise.

3 Emergent Cooperation in Swarm-Based Systems

The validity and importance of large artificial swarm systems is clear from drawing parallels to the biological complexity of swarm-based systems such as social insect colonies. In such systems global behavior is considered to be an emergent property of the interactions of the many different components that make up the whole system. Desirable emergent behavior has been observed in many biological systems, though reproducing these conditions in artificial systems has proved to be difficult, as there is potential for the emergence of undesirable behaviors. Certain swarm systems model biological systems that contain hundreds or thousands of agents, such as ant and termite colonies. Social insect colonies present excellent examples of how collectively intelligent systems can be generated by the interaction of a large number of relatively simple agents. Based on the social insect metaphor, *swarm intelligence* has emerged as a novel approach to the design of distributed systems, with emphasis upon flexibility and robustness. For instance, there has been a significant concentration of research on the study of emergent behavior in simulated ant colonies [21, 33, 64–66, 49, 38, 36]. Certain artificial life simulators⁶ and applications, such as *Swarm* [74], *MANTA* [49], *Tierra* [136], and *Avida* [1], have popularized studies of swarm-based systems.

Drogoul et al. [46–48] presented a simulation model of social organization in an ant colony termed MANTA (*model of an ant-hill activity*). The primary goals of the MANTA application were as follows: first, to model the behaviors of ants at the individual level; second, to test various hypotheses concerning how social structures emerge as a consequence of the behavioral interactions of many individual ants; third, to illustrate that this behavioral model is able to generate functionality such as the division of labor and cooperation when applied in a social context; and fourth, to apply the results to a more general set of systems such as distributed robotic and problem-solving systems where self-organization and emergent collective behavior is a primary focus. The MANTA simulation model was later extended via providing individual ants with a larger set of behaviors, and ant queens were introduced in experiments to reproduce the evolution of a behavioral characteristic known as *sociogenesis*⁷ [159] that is observed in real ants. Preliminary results of Drogoul et al. [46–48] illustrated emergent social structures such as the division of labor within a group of functionally simple artificial ants. The authors also observed the emergence of cooperative behavior beneficial to the colony, similar to social phenomena generally observed among *eusocial*⁸ insects. Results elucidated that the emergence of behavior for the division of labor improved the efficiency of emergent functionality in the population. Such emergent functionality included cooperative foraging and sorting behavior. The authors concluded that the notion of emergent cooperation remains very unclear and difficult to define, and that many of the behaviors viewed as cooperative emerged as a result of the competitive interaction that occurs between individuals in a constrained environment with limited resources.

In the extended version of the MANTA simulation model, Drogoul et al. [49] performed a set of experiments designed to investigate the evolution of a process known as *sociogenesis*, where ant queens needed to interact and cooperate in order for a new ant colony to emerge and survive. Drogoul et al. [49] conducted two types of sociogenesis experiments: those using only a single queen, termed *monogynous sociogeneses*, and those using multiple queens, termed *polygynous sociogeneses*. The hypothesis for the *sociogenesis* experiments was that emergent functionality within a population of agents would be improved by the emergence of a parallel emergent social structure. In this case the

⁶ Such artificial life applications of swarmlike systems typically operate within a modeling or simulation framework such as the Swarm simulator [74]. The initial prototype was developed from a bottom-up design perspective, via the creators' first writing several experiments that the Swarm simulator was to support and then writing the simulator to support these experiments, rather than first developing the simulator and attempting to fit the experiments to that design.

⁷ Sociogenesis, as defined by Wilson [159], is a behavioral process observed in many species of ants, where the newly fertilized queen initiates a new society alone.

⁸ The term *eusocial* describes the most highly developed form of animal societies, such as those of colonial ants, termites, wasps, and bees. Typically there is extensive division of labor and cooperation, with various castes specializing in particular tasks, such as food-gathering, defense, or tending to the young. Reproduction is via an elite group of fertile individuals, assisted by sterile workers [77].

emergent functionality was a cooperative sorting task, and the parallel emergent social structure was the division of labor. In order to facilitate the emergence of cooperation for the sorting task at the colony level, an artificial evolution algorithm was executed, where the selection and reproduction of each new generation was based on the individual genetic design of each artificial ant. Conceptually, each ant was made up of a set of behavioral primitives that described possible actions invoked in response to a set of environmental stimuli. The ants did not use any form of direct communication, but rather used stigmergic communication modeled on pheromone-based communication as used by real ants. During the evolutionary process an indirect form of cooperative behavior emerged, due to particular ants in the population developing specialized behavior for undertaking a specific task. This emergence of specialization removed redundant behaviors in the undertaking of tasks and therefore increased the probability of an ant successfully completing its task. In the polygynous sociogenesis experiments, one of two different results emerged at the end of a successful evolutionary process. The first result was that only a single queen survived, whereas the second result utilized cooperative behavior in that a single queen emerged as the leader of the colony while the other queen ants became analogous to worker ants. An important tradeoff in competitive versus cooperative behavior between the queens was evident from these experiments. This tradeoff proved to be important for the survival of the ant colony as a whole. Cooperative behavior emerged between the queen ants during the initial stage of the colonies' growth; this behavior was manifest in one of the queens taking care of the larvae while others searched for food. The authors concluded that emergent functionality at the colony level was potentially improved via the parallel emergence of a social structure. In the case of these experiments, the emergent functionality was division of labor, and the social structure that concurrently emerged was cooperation.

While the parallels between emergent cooperative behaviors attained under experiments performed using the MANTA simulation model and the emergent behavior observed in real ants makes them intrinsically interesting, what is lacking in this research is a qualitative analysis of the emergent behavior and the mechanisms that lead to cooperative behavior and the concurrent emergence of the division-of-labor functionality. Also, simulations of collective behaviors were limited by the simple and abstract grid-world environment that the artificial ants operated within and only a single case study of emergent collective behaviors has been presented for this artificial ecosystem. Thus, it remains unclear if the approach is applicable to more generalized simulations of emergent social structures that would further test the authors' hypothesis that emergent functionality such as the division of labor facilitates emergent cooperative behavior, which in turn strengthens the performance of the artificial ant colony as a whole.

Aside from simulations that reproduce cooperative behavior in swarm-based systems, certain biological principles that promote cooperative behavior in these systems have also been applied to solving classical artificial intelligence problems. For example, Dorigo et al. [44] and Dorigo and Gambardella [45] applied biological principles from cooperative behavior evident in real ants to solving combinatorial optimization problems such as the traveling salesman problem [32, 33, 44] and the quadratic assignment problem [103, 58].

Dorigo and Gambardella [45] introduced a distributed system algorithm called the *Ant Colony System* that was inspired by the global behavior of biological ant colonies and applied to the traveling salesman problem. The Ant Colony System comprised many agents, specifically artificial ants, each maintaining simple capabilities to mimic the behavior of real ants. The Ant Colony System was inspired by the ability of real ants to find the shortest path from a food source to their nest via the use of pheromone trails promoting cooperative behavior [17, 59]. The task required cooperation in order for the artificial ants to find good solutions to the traveling salesman problem. To facilitate emergent cooperative behavior, the artificial ants used an indirect form of communication mediated by simulated pheromone trails that they deposited on the edges of the traveling salesman problem graph while constructing solutions. Ants use this pheromone information as the medium to communicate information among themselves regarding path length and which path to travel. The execution of, and the emergent cooperative behavior resultant from, the experiments described are as follows. Once all ants have traversed the graph, the best-performing ant deposits its pheromone at

the end of iteration t , thereby defining a preferred route for search in the next iteration of the algorithm. During iteration $t+1$ ants will detect edges belonging to the best traversal of the graph and will elect to traverse these edges with a higher degree of probability. Thus, the cooperative behavior that emerges is a form of autocatalytic behavior where the more the ants follow a trail, the more attractive such trails become as a path for other ants. This process is characterized by a positive feedback loop, where the probability with which an ant chooses a path increases with the number of ants that previously chose the same path. In contrast, whenever an ant visits an edge, it diminishes the amount of pheromone on that edge, thereby making edges less desirable to other ants in the future. This allows for the possibility of an improved future search in the neighborhood of the previous best search.

In their experiments, the authors highlighted the effectiveness of the Ant Colony System applied to the traveling salesman problem by comparing a cooperative search with a noncooperative search. The search executed by a given number of cooperative ants proved superior to the search carried out by the same number of noncooperative ants, each working independently. Specifically, with no cooperation taking place through the pheromone medium, the algorithm slowly derived a suboptimal search solution. When the ants cooperated through the artificial pheromone medium, an optimal solution was quickly converged to, and solutions were not confined to local optima.

The key criticism of this research is that the emergence of cooperative behavior is limited by the constraints of the traveling salesman problem domain. Although similar design principles were applied to solve the quadratic assignment problem [103, 58], the generalization of such design principles to other, less constrained problem domains remains unclear, as experiments were only performed for particular case studies in combinatorial optimization.

Also in the problem domain of artificial ant systems, Perez-Uribe et al. [130] conducted a set of experiments for the purpose of synthesizing cooperative behavior in the context of an artificial evolution process. Simulations were used to study the effects of genetic relatedness and different types of genetic selection in the evolution of cooperation for the accomplishment of a cooperative foraging task. The task was for a group of twenty artificial ants to search for four large and four small food items that, at the beginning of each foraging trial, were randomly scattered in a rectangular grid-world environment. Each foraging trial consisted of two phases; in the first phase each ant activated one of three prespecified behaviors, and in the second phase a group of twenty ants began searching for the food items in their environment. The transportation of the large food items required that two ants cooperate.

The cooperative foraging task was modeled within a mobile robot simulator, with which the authors were able to vary parameters such as the value of food, an ant's genetic specification, and the type of genetic selection and reproduction performed by the artificial evolution process. Changing of these parameters placed selective pressure on the types of cooperative behaviors that were evolved, and the authors argued that the group of mobile robots modeled by the simulator maintained limitations and properties similar to real ants due to their small size. In the evolutionary process, ants were awarded differing fitness scores for either individual or cooperative transportation of food items. Specifically, the total performance of the colony was maximized, in terms of fitness scores, if ants cooperatively transported food items as opposed to acting individually. In the experimental setup, artificial ant colonies were either homogeneous or heterogeneous, where homogeneity was defined by individual ants with genetically similar specifications, whereas heterogeneity was defined by genetically dissimilar individuals. The genetic design of each ant encoded a set of threshold values that were used to determine if a given behavior was activated during each step of a foraging trail. This threshold mechanism was similar to that used by Bonabeau et al. [23] in a model describing division of labor in insect societies.

In their experiments, the authors highlighted that cooperative behaviors were more likely to emerge under a colony level of genetic selection, used for reproduction of homogeneous colonies. In particular, the number of emergent cooperative behaviors was larger in experiments using colony-level genetic selection and homogeneous colonies than with individual-level genetic selection and heterogeneous colonies. The authors stated this to be a result of colony-level selection favoring

individuals that cooperate and not ones that adopt specialized behaviors in foraging for small food items for individual benefit. Experimental results also suggested that genetic relatedness assumed a role in facilitating emergent cooperative behavior, as homogeneous colonies performed better than heterogeneous colonies in the cooperative foraging task. The authors argued that these results maintained biological plausibility, based upon predictions made by certain biologists such as Keller et al. [89, 90], stating that groups should be more efficient when genetic selection acts at the colony level and when there is a high degree of relatedness within groups. However, the same results also yielded no significant difference between homogeneous colonies using colony-level selection and heterogeneous colonies using individual-level selection, indicating that future research should continue to investigate the role of genetic relatedness in facilitating emergent cooperative behavior.

It is clear that modeling emergent cooperative behavior in an artificial ant system with a multi-robot simulator that uses an evolutionary computation methodology is a fruitful approach, since biological social insect systems have a very long generation time and it is inherently difficult to study the evolution of complex social structures such as cooperation. Even though results illustrated that cooperative behavior was more likely to emerge under the colony-level genetic selection within a homogeneous colony, the authors did not clearly state the significance of these results, beyond remarking upon their biological plausibility. It is clear that the inherent complexity of maintaining and analyzing the behavior of large groups of artificial ants within an artificial evolution process justifies the use of a simple task domain, small colonies, and a basic form of genetic-based behavioral encoding. However, the authors did not clearly specify which mechanisms were deemed to be responsible for observed cooperative behavior and the differing degrees of performance between homogeneous and heterogeneous colonies, beyond the conclusion that performance differences were the result of genetic relatedness.

Swarm-bots are a research endeavor concerned with applying biologically inspired design principles in the simulation, and physical construction of, groups of mobile robots that exploit concepts such as emergence, self-assembly, and self-organization in order to accomplish collective goals. Many such collective goals require the use of cooperative behavior. The individual robots are called *s-bots*, and two or more *s-bots* that attach to each other in order to perform a task requiring collective behavior are called a swarm-bot [114]. The key idea of the research is that swarm-bots combine the advantages of swarm intelligence with the flexibility of self-reconfiguration, as they are able to self-assemble and self-organize so as to solve problems that could not otherwise be solved by a single *s-bot*.

As part of the swarm-bot initiative, Nolfi et al. [123] conducted several experiments to address the problem of how a group of *s-bots* could coordinate their movements and actions so as to cooperatively move objects in the environment as far as possible within a given period of time. This research differs from other experiments in the swarm-bot endeavor in that in this case the *s-bots* are given a task that they must cooperate in order to solve. Other swarm-bot research, such as that conducted by Trianni et al. [151], simply maintained the goal of achieving some form of aggregated behavior, which the authors stated would be a prerequisite for various forms of cooperative behavior.

Nolfi et al. [123] conducted experiments designed to facilitate emergent cooperative behavior, where a group of eight *s-bots* were connected to an object, or connected to each other, so as to form a closed structure around an object. The *s-bots* were given the task of moving the object as far as possible in the least amount of time. In the first set of experiments the eight *s-bots* used what the authors termed the *ant formation*, which connected all *s-bots* to the object, but there were no links between the *s-bots* themselves. The resultant collective behavior was dependent upon the weight of the object, such that the *s-bots* cooperatively negotiated to either push or pull the object to their destination. In the second set of experiments, *s-bots* were assembled so as to form a circular structure around the object. The results were similar to those obtained with the ant formation, with the exception that the *s-bot* formation deformed its shape so that some *s-bots* pushed the object, while other *s-bots* pulled it. The mechanism deemed to be primarily responsible for these results were the neural controllers of individual *s-bots*, which evolved the capability to cooperatively coordinate movement when an *s-bot* was connected to another or the object. That is, this cooperation resulted from the inclination for each *s-bot* to follow the direction that the majority of *s-bots* followed at a given time.

It is clear from the interesting nature of these results, and the methods chosen for the evaluation of emergent behavior in the given task domain, that they represent a valuable contribution to the swarm-bot research initiative. Nevertheless, the research lacks formalized methods for the analysis and determination of the mechanisms that led to the successful transportation of objects, meaning that emergent cooperative behavior was only examined from an observational perspective. Also, given that the s-bots are connected to each other or the object at the start of each experiment, the s-bots are forced to cooperate in order to satisfy their individual goals of moving as quickly as possible to a common destination. A form of emergent cooperative behavior not based upon an experimental precondition, but rather based upon a need to solve an unanticipated problem, would have been a more significant contribution to the swarm-bot initiative, especially considering that one of its potential application domains is in real-world search and rescue operations [124].

Also as part of the swarm-bot research initiative, Baldassarre et al. [16] presented a set of experiments for investigating emergent cooperation in the form of flocking behaviors. The task was for a group of simulated robots to move in the least amount of time towards a light-source target. An artificial evolution process governing the derivation of robot behaviors over many task trials elucidated emergent forms of situated and specialized behavior that allowed the group to act as a single unit. In many cases the individual robots displayed complementary behaviors in order to form a cooperative group behavior to accomplish their task. Groups consisted of four simulated Khepera robots [113], where all experiments were conducted in simulation using an extended version of the Evorobot simulator [122]. At the beginning of each task trail the four robots were placed in random positions and orientations within a square walled environment, and a light source elsewhere in the environment was switched on. The fitness function of the robot group was based upon how compact the group was with respect to the distances between the robots and upon the average speed of the robots as they moved towards the light source. The fitness of the robot group for a given task trial was determined with respect to each robot's performance in these two aspects of the fitness function. This fitness function produced aggregation of groups, and yielded the emergence of several cooperative strategies. In all executions of the evolutionary process, individuals evolved some form of cooperation to be able to form groups, maintain group coherence, and move uniformly towards the light source. The different group strategies assumed different formations in one of three different classes of strategies, termed *flock*, *amoeba*, and *rose* by the authors. The flock class of group strategies was a particular example where behavioral specializations emerged. This strategy required that different individuals be able to assume and maintain qualitatively different functions in the group. The flock strategy emerged in few executions of the evolutionary process; the simpler set of strategies in the amoeba and rose classes emerged more often, though they were less successful due to their lack of behavioral specialization.

Several forms of cooperative behavior were synthesized via techniques of artificial evolution, though cooperative behaviors using functional specialization performed the best according to the evaluation criteria. The authors argued that functional specialization evolved due to the need to reduce the interference between conflicting subgoals such as the need to turn and move toward the rest of the group and toward the target. The problematic aspect of these experiments was that they aimed to create effective cooperative behaviors purely through the use of artificial evolution. This made analysis of the emergent behaviors difficult, so it is known that behavioral specialization played a key role in the formation of cooperative strategies, but it remains unknown how behavioral specialization emerged in these experiments.

4 Emergent Cooperation in Pursuit-Evasion Systems

The use of biologically inspired design principles for investigating emergent cooperative behavior remains a relatively unexplored area of research in the pursuit-evasion domain [110], as well as for more traditional predator-prey systems [108, 118].

The original version of the pursuit-evasion problem was introduced by Benda et al. [19] and consisted of a grid world containing four pursuers trying to capture a single evader by occupying the four immediate grid spaces around the evader's position. Both the pursuers and the evader were limited to either horizontal or vertical movement at each time step, the movement of the evader was random, and no two agents were allowed to occupy the same grid space at any given simulation time step. The goal of this research was to illustrate emergent cooperative behavior from the interactions of pursuers following simple pursuit strategies. Since this original version, various researchers have used different approaches [94, 101, 141] for the study of collective behavior in the pursuit-evasion domain, where cooperative pursuit strategies are one form of collective behavior studied. Research investigating cooperative pursuit strategies has typically involved studies of cooperative behavior in the context of pursuit strategies that emerge from the interaction of a single pursuer with a single evader operating within a grid-world environment with rules defined according to a game theory model [2, 6, 12, 13, 50, 110], though certain researchers [69, 160, 40] have investigated emergent behavior in the form of cooperative behavior that emerges within a group of pursuers with a need to collectively capture an evader.

For instance, throughout a series of reviews, Haynes and Sen [69–72] compared genetic programming approaches for the evolution of cooperative pursuit strategies. In [73] they proposed a new approach for the development of cooperative strategies derived via genetic programming [95] and tested it within a pursuit-evasion game scenario. The authors argued that the approach differed from existing approaches in that strategies were incrementally constructed via repeatedly evolving and testing them for increasingly difficult pursuit tasks. Additionally the authors argued that their approach relied upon the performance of emergent solutions, rather than domain-specific knowledge.

The experimental setup used a grid world where initially the evader was placed at the center and four pursuers in random positions. A pursuer could see the evader, but not other pursuers, and there was no explicit form of communication between the pursuers. For all experiments, a genetic programming approach called *strongly typed genetic programming* (STGP), devised by Montana [115], was applied to the task of evolving a program that represented a behavior—namely, a pursuit strategy—which was shared by all pursuers in the case of homogeneous teams [69], and was unique to a given pursuer in the case of heterogeneous teams. In order to generate generalized solutions that were not dependent upon initial agent positions, each pursuit strategy in the population of strategies was evaluated by testing it in k randomly generated pursuit-evasion scenarios. The program with the highest percentage of successful pursuit strategies was taken as the fittest. The STGP technique first randomly generated a population of N programs, and then assigned fitness to each after executing and evaluating them in a pursuit-evasion scenario. A subset of the N programs was then selected for propagation of a new population of programs by pairing up the selected programs and swapping random subparts of the programs.

One hypothesis of this research was that evolution of these structures, incrementally evaluated and updated, would produce effective cooperative pursuit strategies for heterogeneous as well as homogeneous teams of pursuers. Homogeneous teams consisted of k pursuers that all shared the same behavioral pursuit strategy (programs), and the evolutionary process would maintain a population of these behavioral strategies (programs). Heterogeneous teams also consisted of k pursuers, but each pursuer utilized a different behavioral strategy. The evolutionary process maintained a population of team-level strategies, where each team-level strategy consisted of some combination of the k behavioral strategies that represented all pursuers in the team.

Haynes and Sen [70, 71] introduced a cooperative coevolutionary process into their experiments, which was designed to facilitate the development of more complex forms of cooperative pursuit strategies in teams of heterogeneous pursuers. The hypothesis was that k different behavioral strategies for controlling the actions of k different pursuers could be combined, through a cooperative coevolution process, to form a cooperative strategy to achieve some predefined global goal. The authors' supposition was that a cooperative-coevolution as opposed to a competitive-coevolution [5], approach would be more effective in the derivation of complex cooperative pursuit strategies.

Haynes and Sen [70, 71] utilized the STGP methodology in order to evolve behavioral strategies that enabled a heterogeneous team of pursuers to cooperatively achieve a common goal. Each pursuer team consisted of k programs, each of which program explicitly represented an individual pursuer, or more precisely, a behavioral aspect of the cooperative team strategy that emerged when the pursuers interacted. Thus entire teams of pursuers were evolved, as opposed to individual pursuers, so that a particular combination of the k programs constituting the team would determine the behavioral strategy of a particular team. Each pursuer always participated in the same team, and fitness was assigned to the team as a whole as a means of addressing the credit assignment problem [63].

Haynes and Sen [72] implemented a series of experiments that evaluated a set of new genetic programming crossover mechanisms for evolving cooperative strategies among a heterogeneous team of pursuers. Results indicated that only one of the new crossover mechanisms, termed *team-uniform* by the authors, evolved a team faster than the traditional crossover mechanisms. The team-uniform crossover mechanism was found to expedite the evolutionary process, as well as facilitate emergent cooperative pursuit strategies with a higher average fitness within heterogeneous teams of pursuers. In several additional experiments, communication between the pursuers was also studied and found to be unnecessary and even detrimental. Specifically, in the experiments that did not use communication, each non-communicating subpopulation converged towards the optimization of a specific function in the team. This resulted in the derivation of cooperative pursuit strategies facilitated by the emergence of pursuers with specialized and complementary pursuit behaviors. Namely, certain pursuers, termed *chasers* by the authors, only chased the evader, while other pursuers, termed *blockers* by the authors, only attempted to block the path of the evader.

The authors compared the evolution of cooperative strategies using homogeneous and heterogeneous teams in experiments testing two types of evaders, those that moved randomly and those that attempted to maintain a maximum distance from the pursuers. Results illustrated that emergent cooperative pursuit strategies outperformed all but one of four preprogrammed heuristic pursuit strategies, which used a greedy search algorithm [69], and that the emergent cooperative strategies of heterogeneous teams outperformed those of homogeneous teams. The authors concluded that their genetic programming approach was an effective means for deriving cooperative behavior, given that it required no explicit communication and minimal domain knowledge.

The key criticism of this series of research is that many questions concerning emergent specialized behavior and how cooperative behavior emerged remain open. For example, it remains unclear which part of the genetic programming tree structure that describes a pursuer's behavior in the case of a heterogeneous approach, or a team's behavior in the case of a homogeneous approach, actually contributes to the cooperative pursuit behavior observed in the experiments. While emergent team-level cooperation in the initial experiments with homogeneous teams and then emergent specialization in the formation of cooperative pursuit strategies with heterogeneous teams were interesting results, the emergence of such behaviors can largely be attributed to the genetic programming implementation and the simple grid-world environment utilized. Also, the application of the genetic programming methodology to other problem domains was not reported upon, so it remains uncertain if the cooperative behaviors would emerge beyond the grid-world implementation.

Similarly to the research of Haynes and Sen [69], Yong and Miiikkulainen [160] investigated the role of behavioral specialization in the evolution of cooperative pursuit strategies in a pursuit-evasion scenario using multiple pursuers and a single evader. This research compared two artificial evolution approaches for the incremental evolution of a neural network architecture, where this architecture controlled the behavior of pursuers. The first approach was a centralized controller, a single neural network that controlled all pursuers, and the second method was a distributed approach where a separate neural network controlled each pursuer in the team. For both of these approaches, an incremental approach to artificial evolution was used, such that evolved neural networks were tested first upon a relatively simple pursuit-evasion task and then upon increasingly complex ones. The incremental evolutionary process proceeded through five stages, where in the simplest stage the evader was stationary, and in each subsequent stage the evader moved progressively faster, until in

the final stage it moved as fast as the pursuers. The authors argued that the advantage of this incremental evolutionary approach was that it prevented the artificial evolution algorithm from converging to a solution in a suboptimal region of the solution space. These approaches for artificial evolution were based on an architecture termed *enforced subpopulations* [116]. This architecture used multiple populations of neurons and at the turn of every generation a single individual; in this case a neuron was selected from each population of neurons in order to construct the neural network for controlling an individual pursuer or a pursuer team. The enforced subpopulations approach to artificial evolution was used to encourage the emergence of specialized behavioral roles in cooperative pursuit behaviors, such as the chaser and blocker behaviors evident in the experiments of Haynes and Sen [72].

The experimental setup was similar to that of Haynes and Sen [69], in that it used a grid world with obstacles, three pursuers, and a single evader, where each agent was able to occupy a single grid space and was able to move in one of four directions at each simulation time step. The goal of any given pursuit-evasion scenario was for two or more pursuers to occupy the grid squares immediately surrounding the evader's position. The fitness of a pursuer team was calculated according to how close they were to the evader at the end of a given pursuit-evasion scenario. The authors' rationale for using this fitness function was that the starting positions of the pursuers should not influence the team's fitness, and hence the time taken for pursuers to capture the evader was not taken into account. Certain experiments also incorporated communication into the behaviors of individual pursuers, where communication was defined as the capability of pursuers to see each other. Thus, neural networks controlling either individual pursuers, or a whole pursuer team, took into account the coordinates of all pursuers in the team in the derivation of cooperative pursuit strategies.

Comparative sets of experiments were performed, using both the centralized and the decentralized approaches to neural network control, and for each of these experiments pursuer teams with and without communication were tested. The results showed that the decentralized approach to evolution without communication yielded pursuer controllers with specific functional roles, such as chasers and blockers [72], where each role contributed to the formation of a cooperative pursuit strategy. Thus, given that each pursuer performed its specific behavioral role, the team was able to effectively capture the evader even though there was no explicit communication to enable this cooperative behavior. Experiments using the decentralized approach to evolution of controllers, in company with communication, produced teams with more flexible behaviors. Several different team-level behaviors emerged, but each lacked the composite forms of behavioral specialization evident in previous experiments, and as a result these team-level behaviors performed worse as pursuit strategies. Specifically, evolution without communication placed strong evolutionary pressure on each pursuer to perform a particular role, whereas evolution with communication utilized variations and combinations of two or more emergent pursuit strategies, so it was not necessary for pursuers to adopt specific roles in order for a pursuit strategy to be successful. Experiments testing the centralized approaches, with and without communication, resulted in the emergence of cooperative pursuit strategies also, though these strategies performed poorly in comparison with the decentralized approaches.

The conclusion was that the distributed approach to the enforced subpopulation methodology for the incremental evolution of neural controllers proved superior in terms of the time taken to evolve good pursuit strategies. Also, the distributed approach, without communication, allowed for the emergence of specialized behavioral roles, such that each subpopulation was optimized for a specific function by the evolutionary process. The authors stated that adaptive niching [61] in the evolutionary process facilitated the emergence of specialized behavioral roles. That is, as one subpopulation started to converge to a particular behavior, other subpopulations that behaved in a complementary manner were rewarded and started to converge to other behavioral roles. In terms of the domain implementation, having all pursuers develop behaviors that converged to complementary behavioral functions contributed to the formation of effective cooperative pursuit strategies, and yielded a higher fitness for the pursuer team as a whole. The authors argued that the distributed enforced subpopulation approach was applicable to any problem domain that can be decomposed

into a sequence of tasks of increasing complexity, as is the case with the theoretically similar SANE reinforcement learning approach [117]. Unfortunately, this approach was not tested beyond the pursuit-evasion grid-world environment using different configurations of pursuer starting positions and obstacles in the environment.

Denzinger and Fuchs [40] investigated the learning of cooperative pursuit behavior, which was achieved via the evolution of a set of appropriate prototypical situation-action pairs. The simulation environment made use of pursuit-evasion scenarios similar to those described by Haynes and Sen [69] and Yong and Miikkulainen [160]. The environment was a grid world containing three pursuers and one evader, and the task in any given pursuit-evasion scenario was for at least two pursuers to position themselves in grid squares adjacent to the evader. The pursuers agent architecture was specified based on the classification of situations with the nearest neighbor rule [34] and a learning mechanism that attempted to generate a set of prototypical situation-action pairs. The pursuers' behavior was derived from that of situation-action pairs in that, when a pursuer was confronted with a new situation, it determined the situation-action pair that was most similar to the given situation in accordance with the nearest neighbor rule.

The pursuer then applied the action associated with the selected pair. The authors argued that this pursuer agent architecture provided a suitable basis for learning cooperative behavior in view of its flexibility. Namely, a pursuer's behavior could be readily changed by modifying, adding, or removing situation-action pairs. The learning of cooperative behavior was defined as searching for an appropriate set of situation-action pairs using a genetic algorithm [76]. The genetic algorithm started with an unfit set of pairs, that is, pairs leading to a set of poor pursuit behaviors for a given set situation. The fitness function defined a comparison procedure for sets of situation-action pairs so that the fitter set could be determined. The authors used several variants of the pursuit-evasion domain [19] where each variant required differing degrees of cooperative behavior among pursuers. The variants on the game include changing the number of pursuers, the boundaries of the grid world, and the communication and observation capabilities of the pursuers. The goal of the experiments was to demonstrate that the approach was versatile in that it allowed a designer of multi-agent systems to specify requirements in terms of representation of situations and possible actions, and then for a satisfactory solution to be evolved automatically.

The agent architecture was able to evolve effective cooperative pursuit strategies for many variants of the pursuit-evasion game, although for problems requiring more complex representations of situations, an agent architecture able to operate in environments other than grid worlds would be required. The use of a grid world allowed for the selection of distinct sets of situation-action values where a finite set of actions and resultant outcomes could be defined. While the emergence of cooperation was simpler to analyze in this grid-world domain, it was limited by its implementation, so the study of mechanisms that facilitated emergent cooperation was limited to trivial situations.

In a variation on the pursuit-evasion domain, Nishimura and Takashi [118] studied the emergence of cooperative behavior in the form of different types of flocking strategies, using a more traditional style predator-prey system [75, 42] that contained large numbers of predators and prey. In this predator-prey system both predators and prey inhabited a simulated two-dimensional grid-world environment and interacted through a succession of pursuit-evasion game scenarios. The game scenarios used a score-based system, and were implemented in the context of an artificial evolution algorithm. The rules of this particular pursuit-evasion game were such that when a predator moved to an adjacent grid square behind a prey, the predator was awarded p points, whereas the prey lost p points, and when a predator moved to a grid square adjacent to a prey, and the two were facing each other, both species lost p points. After receiving a score, individual predators and prey were categorized as either winners or losers. At the end of each generation the species with the higher score was able to reproduce more, and the species with the lower score reproduced less and was consequently diminished. Individual predators and prey in the system were characterized by a set of parameters that controlled their social interaction dynamics and behavior over the course of the evolutionary process. That is, behavioral interaction between predators and prey were formalized as a set of dynamical equations, and adjusting the parameters of those

equations served to yield different individual and collective behaviors over the course of many generations. Since it was difficult for the authors to know which parameter values were relevant to the formation of which types of cooperative behaviors, changes in dynamical equation parameters and state variables were taken into account by the evolutionary process. The offspring of “winner” individuals inherited the behavior of their parents in the form of slightly modified algorithms or parameters. Mutation in the evolutionary process was simulated via the addition of a low level of Gaussian noise to random sets of parameters.

The authors performed several sets of experiments testing various pursuit-evasion game scenarios and different parameter settings. For example, experiments were performed testing predators with adaptive behavior versus prey with fixed behavior, as well as predators with fixed behavior versus prey with adaptive behavior. From the first set of experiments, the authors observed that both predators and prey tended to group in loose probabilistic formations, and that there were certain random swarming dynamics that allowed both predators and prey to maximize their life expectancies. In the second set of experiments, where only the prey were adaptive, cooperative behaviors such as spatially disordered groupings (one-way marching, random swarming, lattice formation, or rotating clusters) emerged. In a third set of experiments that introduced adaptive behavior for predators and prey, similar though more complex forms of these collective behaviors emerged. From their experiments, the authors learned that the predators and prey were able to coexist for the longest time when individuals of both species inherited a parameter responsible for the derivation of a cooperative behavior known as *random swarming*. In particular, the random swarming formation in groups of predators prevented predators from being able to synchronize their headings with groups of prey and thus follow the same prey for extended periods of time. That is, over the course of the evolutionary process, the emergent random swarming formations decreased the chance of predators capturing prey, thereby minimizing the chance of extinction of both predators and prey. Given that the prey then had a lower probability of becoming extinct, they were more readily available to predators as a food source. Thus the predators also benefited from a reduced probability of extinction.

In concluding their research, the authors related their results to natural predator-prey systems, by stating that similar results have been found in theoretical biological studies conducted on schooling fish [7, 26, 35]. These studies also reported cooperative group behaviors that were loose probabilistic formations. Even in such biological studies, though, the relevance of, and mechanisms leading to, emergent cooperative group behaviors have yet to be established. In linking their own experimental results to results of studies in biological predator-prey systems, the authors suggested that instability in predator-prey dynamics, observed in their own experiments, could encourage mutually beneficial coexistence via phenomena such as symbiosis, as evident in certain biological predator-prey systems. That is, in natural systems spatially induced dynamic randomness is important for symbiosis, and in their own experiments the authors demonstrated emergent cooperative behaviors as an unstable dynamic without any explicit organization or structure. Although interesting cooperative behaviors and a stable state of the system were attained by use of a finite set of equations running within an artificial evolution process, the key criticism of this research is that the evolved behaviors were limited by the grid-world environment and were somewhat contrived by adjusting equation parameters prior to the execution of each evolutionary process.

5 Emergent Cooperation in RoboCup Soccer

There is a field of research dedicated to the design and development of multi-robot systems for playing a robotic form of soccer. Various leagues, characterized by the types of robots used as players and the types of game scenarios played, currently exist as research initiatives [85, 92, 93, 145, 158], and each league maintains its own set of technical challenges and engineering accomplishments. Collectively, these robotic soccer systems are known as RoboCup, and they have recently been developed in simulation [102, 107, 142–147] as well as with real robots [11, 92, 139]. For example,

Veloso et al. [156] used a team of small, autonomous two-wheeled robots, and Veloso and Uther [157] used a team of Sony AIBO[®] robot dogs, to play in a *RoboCup soccer tournament*. It is obvious from these latter experiments that robotic systems provide a degree of realism that is never possible in simulation, though, as a complementary research tool, RoboCup simulators allow researchers to readily investigate less tangible research issues such as cooperative behavior, via the implementation of more abstract and complex behaviors. Simulators also have the advantage that it is generally easier to examine the dynamics of group behavior by changing simulation parameters, and it is possible for many experimental trials to be executed in a relatively short time. The RoboCup soccer simulator described by Noda et al. [120] is a virtual soccer field and provides the most commonly used test bed for running experiments to investigate emergent cooperative behavior in various RoboCup game scenarios. Such RoboCup soccer simulators represent a level of abstraction above the low-level perception and action complexities inherent in robotics, and allow researchers to focus on unresolved issues such as emergent team-level behaviors [91].

Much research in the simulated version of RoboCup has focused upon the application of machine learning techniques within constrained experimental scenarios. For example, to date there has not been an implementation of a design methodology that successfully applies emergent cooperative behavior for the consistent benefit of game strategies in a complete RoboCup soccer team. Several researchers have focused on algorithms for describing cooperative behavior between two or three soccer agents in a team, but such behavior is either specified a priori or learned in simplistic game scenarios.

Noda et al. [120] used a RoboCup simulator as a test bed for the learning of cooperative behavior within groups of soccer agents. Learned cooperative behavior took the form of one soccer agent learning when to pass to a teammate and when to shoot the ball at the opponent goal area. The experimental setup included two offensive soccer agents, termed players A and B , and one defensive soccer agent, termed player C . Initially, players A and B were positioned randomly within the penalty area together with player C , and player C was programmed to maintain a position between the ball and the goal area. Player A could not move from its position within the penalty area and had the task of either shooting the ball at the goal area or passing it to player B . Player B was preprogrammed to wait for a pass from player A and then to shoot the ball at the opponent goal area. Player A used a neural network with thirty hidden neurons and a backpropagation method [138] to learn in which situations it was better to cooperate and in which situations it was better not to cooperate, according to the evaluation criteria of the number of goals scored and the time taken to score in a given experiment. For these experiments, cooperation was defined as the situation when player A passed the ball to player B , and noncooperation was defined as the situation when player A shot the ball directly at the opponent goal area. The learning approach was supervised in that over the course of several hundred training scenarios, a coaching agent provided a positive feedback signal when one of the offensive players scored a goal, and a negative feedback signal when a shot aimed at the opponent goal area failed or a time limit expired. Each training scenario consisted of player A randomly selecting when to pass the ball and when to kick it directly at the opponent goal area. Instances when an offensive player successfully scored a goal were used as training data for the neural network. Inputs to the neural network indicated the relative positions of other players, the ball, and the goal, whereas outputs indicated the expected success rates of passing and shooting the ball. Thus, this training data, and the supervised learning technique used, dictated in which instances the offensive players should cooperate and in which instances they should not. The authors illustrated that training the neural network using the backpropagation method allowed the success rates of the shoot and pass actions to increase as player A learned when to pass and when to shoot the ball, depending upon the position of the defensive player relative to the goal area and player B . Learned cooperative behavior of the two offensive players A and B was evaluated in terms of the time taken to score a goal as well as the number of goals scored.

The key criticism of this research is that cooperative behavior was limited to a learned decision-making process for a single soccer agent: the decision to pass or not. The agents, environment, and learning mechanism were kept simple, so that this form of cooperative behavior could be

successfully learned. In their conclusions, the authors justified using a simple neural network learning mechanism in that provided a good starting point for the learning of more complex forms of cooperative behavior that would potentially be applicable to an entire team of soccer agents.

Stone and Veloso [142] introduced a layered learning approach to cooperative behavior, where soccer agents used neural networks to initially learn low-level individual behaviors such as intercepting a ball, and then decision trees [111] to learn higher-level cooperative behaviors such as deciding when and to which soccer agent to pass the ball. The layered learning approach was designed for problem domains where it would be extremely difficult or impossible to find a direct mapping from sensory input to actuator output.

The layered learning approach was implemented within the RoboCup server [120] as a simulated environment, and allowed for a bottom-up definition of soccer agents' capabilities at both the individual and the team level. That is, learned low-level individual behaviors formed the basis for, and were incorporated as part of, higher-level team behaviors. This differed from the other research of Noda et al. [119], in that learned individual behaviors were utilized in a social context involving at least three agents and were thus an important basis for more complex forms of cooperative behavior. The authors implemented a multilayered feedforward neural network as the controller for each soccer agent to first learn the low-level individual skill of intercepting a moving ball. This learned skill was then used by the layered learning approach as the basis for learning the higher-level skill of deciding when and to which soccer agent to pass the ball. Specifically, this higher-level skill involved the ability of a soccer agent to estimate the probability that a pass to another soccer agent would succeed. The authors combined a neural network and a decision tree to demonstrate the feasibility of the layered learning approach, and emphasized that the approach proved empirically successful in a given set of game scenarios. Game scenarios utilized at least three offensive soccer agents that attempted to score a goal, while at least two defensive soccer agents protected the goal area. In these game scenarios, the layered learning approach provided the capability for a group of offensive soccer agents to cooperate via making strategic passes to each other, so that the probability of scoring a successful shot at the opponent's goal area would be maximized.

Stone and Veloso [143] extended the basic behaviors learned in their first set of experiments to higher-level cooperative behavior that was potentially capable of controlling soccer agents throughout an entire RoboCup soccer match. The layered learning approach was used to construct a complete team-level behavior where players decided when to chase the ball and, if intercepted, whether or not to cooperate with other soccer agents by passing the ball. The novel aspect of these experiments was to select the actions of individual soccer agents according to the confidence factors associated with decision tree classifications. The team-level behavior learned was such that a soccer agent moved to intercept the ball when it did not detect any teammates that were likely to reach it more quickly. Each soccer agent used previously learned ball intercept behaviors to first intercept the moving ball, and then a predefined communication protocol to probe teammates so as to ascertain a set of potential receivers for a pass. When a soccer agent decided to pass the ball, the receiver with the highest positive confidence factor was selected. Once a soccer agent decided which teammate to pass to, it communicated its intention to that teammate. The receiving teammate then used its previously learned ball interception skill to intercept the passed ball. If no receiver maintained a positive confidence factor, the soccer agent with the ball was preprogrammed to move with or kick the ball towards the opponent goal area.

In several experiments, this approach for the learning of cooperative behavior was tested using an offensive team playing within a game scenario against an opponent defensive team with a prespecified team-level behavior. The opponent team's behavior was such that a group of soccer agents defended one side of the playing field and did not defend the other side at all. Learned cooperative behaviors exploited the opponents' method of field defense by having particular offensive soccer agents move into an open position on the field prior to receiving a pass. Additionally, other experimental results illustrated that the approach of using a decision tree and confidence factors to make decisions for when and when not to cooperate outperformed both random and preprogrammed approaches in terms of the evaluation criteria of the number of goals

scored and the time for which control of the ball was maintained by the group of three offensive soccer agents.

The research [146] of Stone and Veloso extended their previous work and elaborated upon their approach for having a soccer agent decide whether to cooperate with teammates (pass the ball) or to not cooperate (shoot the ball directly at the opponent goal area). In this research, the authors made the distinction between active and passive soccer agents: an active agent was one that controls the ball, and a passive agent was one that waits for control of the ball. The question addressed in this research was, what actions a passive soccer agent should take in order to improve cooperative behavior within the team. The authors used the layered learning approach to design an action selection mechanism that allowed soccer agents to anticipate if cooperation with a particular teammate would be advantageous. The action selection mechanism allowed passive soccer agents to position themselves with the objective of trying to maximize the chances of a successful pass in the case that the active soccer agent decided to pass. Each passive soccer agent would consider the positions of other passive teammates, opponent soccer agents, and the ball in order to move to a position that would maximize the chances of successful cooperation between itself and the active soccer agent. This action selection mechanism was implemented in several experiments; it proved successful in encouraging cooperative behavior, and outperformed an approach for team control that did not utilize this action selection mechanism. The comparison of performance was in terms of the number of goals scored and the time for which a team maintained control of the ball.

The cooperative team-level behaviors described in this series of research reviews were not emergent in the sense that is typically referred to in the artificial life literature, as these cooperative team-level behaviors relied largely upon individual agents learning action selection mechanisms based upon decision tree confidence factors. Cooperative behavior was emergent in the sense that a series of decisions by individual soccer agents regarding whether to pass the ball or not formed a team-level behavior that was more successful in terms of goals scored and the time for which the team maintained control of the ball. In many experiments, the game scenarios tested did not reflect a complete range of scenarios that would be required in an actual RoboCup soccer match, and in certain cases it was unclear if the cooperative behavior exhibited would generalize to a broader class of game situations. Despite this, the layered learning approach provided an excellent methodology for the learning of cooperative behavior in a task environment whose inherent complexity prevented the derivation of a direct mapping from sensors to actuators via the use of more traditional learning methods.

In research on using artificial evolution to derive cooperative behavior within a team of soccer agents, Whiteson et al. [154] compared and evaluated two different neuroevolution approaches to the synthesis of cooperative behavior. These methodologies attempted to derive cooperative behavior within a group of three soccer agents for the *keep-away soccer* task environment [148, 131, 79]. Neuroevolution is an approach that uses genetic algorithms to evolve neural networks, and was designed for the possibility of managing complex control tasks, where learning a direct mapping from sensors to actuators would be extremely difficult. Using these neuroevolution methodologies, soccer agents first learned a small number of subtasks that were then combined, as dictated by an artificial evolution process, so that an overall complex behavior emerged. The authors argued that these neuroevolution approaches were advantageous in that they did not require each soccer agent to learn a direct mapping from sensors to actuators, or to learn a particular means of interaction, in order to derive relatively complex and cooperative behaviors.

In the keep-away soccer task, which was played within a grid-world environment, one team of soccer agents, termed the *keepers*, attempted to maintain possession of the ball, while another team, termed the *takers*, attempted to gain control of the ball. In the experiments performed, the task was to minimize the number of times that the takers gained control of the ball, which occurred whenever a taker was within one grid square of the ball. The objective of the keepers, which mandated cooperative behavior, was to continuously move and pass the ball to other keepers so as to keep the ball away from the takers. The experimental setup was such that three keepers were placed inside a circular area at initially random points that were equidistant from each other. One taker was placed at

the center of the circular area, and the ball was placed in front of a randomly selected keeper. A trial of keep-away soccer proceeded so that keepers received one point for every successful pass, and a trial was completed when a taker intercepted the ball or the ball exited the bounding circle. Over the course of multiple trials, cooperative behavior dictated by a neural network controlling a homogeneous team of three keepers was evolved. The trials included a single taker with heuristic behavior. The aim of the evolved neural controller was to make decisions for each keeper, so as to determine on a case-by-case basis when a keeper was to cooperate with other keepers (passing the ball), and when it was not to cooperate (moving with the ball).

The authors compared two neuroevolution approaches for evolving cooperative behavior among the team of keepers. Both approaches used homogeneous teams, in that each keeper maintained the same neural network controller. In the first approach, genome strings encoded synaptic weights of a population of complete neural network controllers. These genomes were evaluated in each generation; the fittest individuals selectively reproduced and subsequently propagated throughout the evolutionary process. In the second approach, the *enforced subpopulation* method [61] was used to evolve subpopulations of neurons, instead of evolving complete controllers as in the first approach. The enforced subpopulation approach created one subpopulation of neurons for each hidden layer node within a fully connected two-layer feedforward neural network controller that it evolved. Every neuron in a subpopulation was itself a genome that encoded incoming and outgoing weights for a given hidden layer node, where selecting one neuron from each subpopulation formed the hidden layer of a newly derived controller. The authors noted that each subpopulation tended to converge to a behavioral role that maximized the fitness of the networks in which it appeared. So the fitness of a given network was calculated as the average fitness of all neurons that participated in the network. The authors argued that this second approach to neuroevolution was more effective than the first approach tested, as it decomposed the complex problem of finding highly fit controllers into smaller subproblems of finding highly fit neurons.

The authors executed several benchmark tests, and found that the enforced subpopulation approach outperformed other neuroevolution algorithms [29, 62] as well as several reinforcement-learning methods [149] in this keep-away soccer task environment. Results also showed that both approaches evolved a successful neural network controller, though the enforced subpopulation approach performed significantly better in terms of the evaluation measures defined for the keep-away soccer task and in facilitating emergent cooperative behavior. In the first approach, cooperative behavior was not encouraged in that an intercepting keeper was evaluated only by how quickly it could reach the ball, though in keep-away soccer a good interceptor will approach the ball from an appropriate angle in order to make its next pass easier. That is, using the first approach, an evolved keeper approached the ball directly, whereas in the enforced subpopulation approach, ball interception was learned together with passing. Thus, the enforced subpopulation approach was better able to facilitate emergent cooperation in that different keepers evolved complementary behaviors, which aided in the formation of a coherent and effective form of cooperative behavior. Although the enforced subpopulation approach proved superior in these experiments, an obvious criticism of this approach is that for more difficult tasks—for example, those not executed in a grid-world environment—the solution space would be too large for an artificial evolution algorithm to search effectively and construct an appropriate controller.

Hsu and Gustafson [78, 79] also investigated a methodology for facilitating emergent cooperative behavior, using the keep-away soccer task. The methodology combined layered learning [142] and genetic programming [95] approaches. The authors argued that by using a layered learning approach to genetic programming, as opposed to a pure genetic programming approach [95], team-level behaviors such as cooperation could readily be derived. In complex problem domains that necessitate many low-level operations, such as RoboCup soccer, it would be intractable to derive complex and desirable forms of cooperative behavior purely via the use of genetic programming or a genetic algorithm.

Each experimental setup used three keepers, a ball, and a taker, located within a rectangular grid-world environment. The taker was able to move twice as fast as the keepers, and the ball could move,

when passed, at twice the speed of the taker. In all experiments, the taker used a preprogrammed ball interception skill. For all experiments, the collective goal of the keepers was to minimize the number of times the taker gained control of the ball in a given trial. To accomplish this, the design methodology was decomposed into two behavioral layers. The first layer controlled passing the ball accurately to other keepers with no taker present, and the second layer controlled moving with the ball and passing between keepers, with a taker present. Cooperative behavior within the team of keepers was derived using a genetic programming technique, where a single program represented the behavioral strategy of each keeper, and over the course of an evolutionary process, cooperative team-level behavior was acquired using the layered learning approach. Forty percent of the artificial evolution process was executed within the first layer of the layered learning part of the methodology, so the keepers first acquired the skill of accurate passing without the taker present. The fitness function for a keeper's accurate passing skill awarded fitness proportional to the number of passes made to within three grid-squares of another keeper. The remaining sixty percent of the artificial evolution process was executed within the second layer of the layered learning aspect of the design methodology, where the goal was to evolve a team-level behavior that utilized cooperation in order to minimize the number of times the taker gained control of the ball. The fitness function for this layer awarded fitness inversely proportional to the number of times the ball was intercepted by the taker.

In several experiments, the authors compared a standard genetic programming approach [95] with their methodology that combined layered learning and genetic programming. These experiments highlighted that the layered learning approach was able to more quickly evolve cooperative behavioral strategies within the team of keepers, and with a higher fitness than the standard genetic programming approach. The authors argued that the layered learning approach allowed for a workable decomposition of a complex problem into many readily solvable subproblems, and that for each of these subproblems, corresponding fitness functions were readily identifiable. Using the layered learning approach, team-level behavior was formed via the evolution of complementary keeper strategies, such that when the three keepers interacted with each other, an overall cooperative behavior emerged. This cooperative team-level behavior was derived in a bottom-up manner, where keepers first learned the skills necessary to cooperate as a pair of players, and then as a team of three players. The authors attributed the success of their design methodology, which combined layered learning and genetic programming approaches, to the ability of layered learning to provide for the incremental learning of individual keeper behaviors, and the ability of the genetic programming approach to effectively combine these behaviors in order to form a team-level cooperative behavior. Specifically, the methodology maintained the functional capability to first learn low-level behavioral skills such as accurate ball passing between keepers, and then learn to minimize the number of times the taker intercepted the ball during keeper passing maneuvers. The genetic programming process would then refine and compose these learned forms of cooperative behavior that operated between two keepers, into a team-level cooperative behavior.

The key criticism of this research is that only homogeneous teams were evolved, and team level cooperative behavior was derived from the use of only two layers in the layered learning approach. Specifically, only two low-level behaviors were used in the derivation of team-level behaviors. Also, the use of a grid world placed severe limitations on the form of cooperative team-level behavior that could be evolved.

In a similar theme of research, Luke et al. [102] implemented genetic programming techniques within a RoboCup simulator, in an attempt to evolve cooperative behavior within an entire team of eleven soccer agents. The performance of different genetic programming techniques were compared based on the derivation of cooperative behavior, where such behavior was evaluated according to the criteria of the number of goals scored by the team, the number of successful passes, and the period of time for which the team maintained control of the ball. The authors' argument for using genetic programming to evolve cooperative behavior at the team level was that genetic programming uses evolutionary computation to derive symbolic functions and algorithms that operate effectively in unpredictable and dynamic problem domains, whereas learning techniques such as neural networks

and decision trees are designed not to develop algorithmic behaviors but to learn nonlinear functions over a discrete set of variables. Thus, these learning techniques only operate effectively in abstract and constrained problem domains, and are not necessarily well suited to the derivation of complex forms of multi-agent cooperation [142–146].

Luke et al. [102] used the *strongly typed genetic programming* (STGP) technique [115] to simultaneously test entire teams of soccer agents against each other in competitive coevolution scenarios. Using this approach, each genome in a population of genomes specified an entire team of soccer agents, where a single genome comprised many STGP trees. Each STGP tree specified the particular behavior of each soccer agent in the team, and subtrees within each STGP tree specified various aspects of an agent's behavior. Certain low-level behavioral functions such as ball interception and kicking behaviors were preprogrammed, as the STGP technique was unsuccessful at evolving these low-level behaviors. The STGP technique was used, though, to evolve symbolic functions to determine the probability of a successful pass to a teammate, or a shot at the opponent goal area. The STGP technique was also used to compare the evolution of cooperative behavior within homogeneous versus heterogeneous teams of soccer agents. In a homogeneous team each agent was specified using the same STGP tree, and thus maintained the same behavior throughout the course of a game scenario. That is, at the end of each generation in the evolutionary process, the fittest genome was selected to represent the behavior of each agent in the team. In contrast, in a heterogeneous team, each agent developed and followed a unique behavior, derived from a combination of STGP subtrees taken from a set of the fittest genomes in the current generation of the evolutionary process. In certain experiments, the authors also introduced a special crossover operator, termed *root crossover*, which swapped whole STGP trees instead of subtrees. This genetic operator effectively allowed the “trading of players” between different genomes in the population during selective reproduction in the evolutionary process. The root crossover operator was designed with the intention of hastening the dissemination of cooperative team-level strategies throughout the genome population.

Initial competitive coevolution game scenarios executed using the evaluation criteria of the number of goals scored by the team, the number of successful passes, and the period of time for which the team maintained control of the ball converged to very poor solutions with no form of cooperative behavior evident. In subsequent experiments, team fitness was evaluated based only upon the number of goals achieved by a given team. As a result, each population in the competitive coevolution process converged to good solutions in several game scenarios, as well as a number of suboptimal solutions in other game scenarios. One suboptimal solution that emerged was termed “kiddy-soccer” by the authors, as it entailed all soccer agents belonging to a team chasing the ball and attempting to kick it into their opponents' goal area. In early competitive coevolution game scenarios, this strategy gained dominance in the evolutionary process, as opponent teams had not yet evolved cooperative team-level behaviors that allowed for an effective defense of the goal area. After many generations of the coevolutionary process, cooperative behavior emerged within each of the competing teams that effectively combined offensive and defensive team-level strategies. The authors also noted that in the formation of such cooperative team-level behaviors, different soccer agents assumed complementary behavioral roles. For example, one cooperative strategy that emerged entailed some soccer agents maintaining a position close to their own goal area when not close to ball, while the remaining soccer agents of the team maintained other positions in an attempt to gain control of the ball. This particular team-level behavior served to reduce the chance of success of long distance goal shots by opponents, which was previously the behavioral strategy primarily responsible for high team fitness. Eventually, teams on both sides evolved equally effective cooperative team-level behaviors that comprised offensive and defensive strategies. In the final set of evolutionary runs, cooperative team-level behaviors were evolved such that different groups of soccer agents within each team cooperated with each other in a complementary manner via simultaneously defending the goal area and dispersing throughout the field to hold certain positions so as to increase the chance of receiving a pass from fellow soccer agents. Such cooperative behaviors prevented the soccer agents from interfering with each other as had occurred in early

evolutionary runs, where many agents were closely gathered about the ball in an attempt to gain control of it.

The key problem with these experiments was that they relied purely upon a competitive coevolution process and the functionality of genetic programming in order to produce cooperative behavior within a team of soccer agents. This meant that in order to evolve team-level cooperative behavior within a feasible amount of time, several constraints had to be placed on the artificial coevolution process, such as limited population sizes, and teams composed of functionally simple agents. Additionally, cooperative behaviors that emerged under the coevolutionary process could only be analyzed from a purely observational perspective. That is, fitness comparisons between the competing teams only illustrated progress and counterprogress of emergent cooperative behaviors, and did not correspond to true evolutionary progress [52, 53], because the fitness landscapes of both teams were continuously changing due to the Red Queen effect [155]. Also, even though the cooperative behaviors that were reported upon had the advantage of having emerged purely in response to opponent team behavior and not as a result of design constraints specified a priori, it was difficult to envisage how such a competitive coevolutionary process could be utilized for the synthesis of desired forms of cooperative behavior, for unanticipated tasks, in problem domains that mandate real-time behavioral adaptation.

6 Conclusions and Future Directions

In order to draw conclusions for this review, it is important to note that in the research field of artificial life alone, there exists a vast range of disparate research endeavors that could be considered as exploring some aspect of emergent cooperation in artificial social systems. Thus, for the purposes of this review, we were not concerned with finding a definition for the term “emergent cooperative behavior,” or providing an exhaustive compilation of research results on artificial social systems, but rather with identifying a set of pertinent research examples that used biologically inspired design principles as a means of motivating multiple agents to collectively solve a predefined problem of a global nature that could not otherwise be solved by an individual agent. The pertinent research examples were identified and selected based upon results where emergent cooperative behavior had been achieved using biologically inspired design methodologies that made use of concepts such as self-organization, learning, and evolution. The research results reviewed were from three disparate problem domains that facilitated and benefited from emergent cooperation. These problem domains were swarm-based systems, pursuit and evasion, and RoboCup soccer.

The binding theme of the review argued, from pertinent results in these three problem domains, that the majority of emergent cooperative behavior research utilizes neither situated or embodied approaches, but only abstract task domains that limit the derivation of cooperative behavior to simple or trivial forms. However, given that the mechanisms leading to emergent cooperation in biological systems such as social insect colonies largely remain a mystery, the use of abstract task domains constructed in order to achieve simple forms of cooperation is justified. It is evident from the literature that the use of various forms of simulated artificial social systems is deemed by many researchers to be an effective approach for investigating emergent cooperation, in that such simulations provide a means for studying the conditions under which cooperation emerges, and the effects of parametric changes can be seen in a relatively short space of time. Unfortunately, the application of biologically inspired design principles to current artificial social system simulations lacks proven methodologies that allow for effective analysis and evaluation of the mechanisms that motivate the emergence of desired forms of cooperation. Also, the transfer of these simulated biological mechanisms to algorithms that are applicable to real-world artificial systems, such as decentralized control systems or multi-robot systems, is not yet plausible.

Additionally, the use of concepts such as evolution, self-organization, and learning was highlighted in many cases as being an effective means for the derivation of cooperative behavior, though the use of many biologically inspired design methods such as artificial evolution is still in a

stage of research infancy, so emergent cooperation synthesized using these design principles is currently limited to simple forms involving only a few participants.

Given this general evaluation of the literature, in the results outlined and research reviewed, several key open problems were identified. In all cases, these open problems were not imposed by the nature of the problem domains themselves, but rather by the infancy of the biologically inspired design mechanisms and the distinct lack of analytical methods and techniques. In each set of results surveyed, researchers were using different approaches and development platforms for the synthesis of cooperative behavior, as well as different methods for the interpretation, evaluation, and analysis of emergent cooperative behavior under relatively similar task environments.

It is obvious that if emergent cooperative behavior is to be used to any great benefit in real-world artificial systems potentially consisting of thousands of embodied agents, as for example in biological swarm-based systems, then it is important that future research address particular open problems evident from current research results. Specifically, if the notion of emergent cooperation is to gain any maturity and credence as a viable means of problem solving in any kind of artificial system that is currently evident in various real-world technological forms, then the results yielded must be quantifiable and comparable with those of traditional methods of achieving collective goals. Ideally, proven design methodologies for achieving desired emergent cooperation must be scalable and transferable to a counterpart situated and embodied application domain, so such methodologies would need to be defined by algorithms and methods of analysis that are equally applicable in the physical world. Notable exceptions that partially address this issue include the research of Dorigo et al. [44] and Stone and Veloso [142]. In the latter research, high-level multi-agent behaviors, originally developed in simulation, were subsequently transferred to robot-control code, and only certain run-time functions were redefined [156, 146]. In this case, though, cooperative behavior was primarily defined a priori in that each robot was equipped with the knowledge required to play in any position in several different types of group formations. The utilization of biologically inspired design principles or task domains contrived in order to promote with certainty the emergence of particular types of cooperative behaviors was found to be a common aspect of many research endeavors, especially in RoboCup where there was a particular emphasis on validation of behaviors in a physical system. This is regrettable, since many of the researchers that utilize biologically inspired design principles as a means of deriving cooperative behavior also subscribe to the notion that the world is its own best model. Hence, it is deemed that the most promising research avenues for significant future progress are those that attempt to define structured and interdisciplinary approaches to developing theories, design methodologies, and evaluation methods for emergent cooperation, both in simulation and in physical systems.

Future research avenues would also do well to exploit the advantages of biological design principles in the derivation of cooperative forms of behavior that potentially benefit an artificial social system as a whole. For example, advantages such as redundancy, scalability, and minimalist component design have been utilized to great lengths in swarm-based systems research to achieve various forms of collective behavior. Many of these systems, though, (with notable exceptions such as the Ant System [44]), are used only to demonstrate concepts such as self-organization, emergence, and their apparent contribution to “swarm intelligence,” not to formulating and addressing effective forms of cooperative behavior that can be evaluated or otherwise conform to a standardized benchmark. Additionally, when emergent cooperation in biologically inspired artificial systems is achieved, it is rarely tested concurrently in a real-world problem domain, and the results are not compared with those of more traditional approaches that do not utilize a biologically inspired design approach, or even no cooperation at all, to achieve group- or global-level goals. The comparison of results using emergent cooperation with those attained using more classical distributed artificial intelligence design approaches is an aspect that is missing from many current research endeavors, and should form a greater part of future research if the notion of emergent cooperation as a means of problem solving is to gain credibility.

Another open problem is that there currently exists no standardized benchmark or method for evaluating or otherwise classifying emergent cooperative behavior. Admittedly, in the RoboCup

soccer problem domain, the annual RoboCup tournaments for both simulated [121] and physical [93, 96] robot teams provide an effective method for evaluating the performance of a given team-level behavior or control algorithm. Also, having robot teams play each other provides a method for competitive performance evaluation that allows for the improvement of behavior without the need for an absolute performance measure. Such teams, though, typically make use of preprogrammed behaviors and various optimization algorithms as means of improving performance, and do not use the concept of emergent cooperation as a means of attaining team-level solutions. Furthermore, in teams that used biologically inspired design principles to achieve emergent cooperation, there was a clear lack of analytical tools, making it difficult to determine if team success could be attributed to cooperative behavior or to other elements in game play, such as a particular combination of individual behaviors, specialization of behavioral skills, or (in the case of the simulated league) the nature of the environment itself.

Results from emergent cooperative behavior studies in the pursuit-evasion domain have also suffered from the problem of trivialized models and design approaches as well as a lack of evaluation benchmarks and methods for performance comparison with counterpart situated and embodied systems. Emergent cooperation in the pursuit-evasion domain has obvious real-world applications, such as the formulation of military, reconnaissance, or search and rescue strategies in environments for which there are relatively few possibilities of specifying cooperative behaviors a priori. However, for emergent cooperation to be effectively applied as a means of problem solving in such task domains, and for any effective multi-robot system to be implemented, it would first be necessary for a cooperative behavior design methodology to be based upon theories and proofs of convergence that would clearly derive a form of cooperative behavior that could be evaluated according to a threshold for quality.

Given the early stage of research in artificial social systems that use biologically inspired design principles as a means of achieving emergent cooperation in order to solve system-level problems, it is understandable that standardized methods for deriving, testing, proving the convergence of, and evaluating emergent cooperative behavior do not yet exist. Nevertheless, from a review and evaluation of pertinent research results in the RoboCup soccer, pursuit-evasion, and swarm intelligence problem domains, it is obvious that much success has already been achieved using relatively simple synthetic approaches for the design of emergent cooperation and using a disparate array of preliminary methods for behavioral analysis. Thus, it is plausible that if particular key problems highlighted throughout this review are focused upon as subjects of future research, then the concept of emergent cooperation will no longer be restricted to simple task domains and abstract methodologies that are necessarily restricted to simulation.

References

1. Adami, C., & Brown, T. (1994). *Evolutionary learning in the 2D artificial life system "Avida"* (Tech. Rep. 106-38). Pasadena, CA: Kellogg Radiation Lab, California Institute of Technology.
2. Akiyama, E., & Kunihiko, K. (1995). Evolution of cooperation, differentiation, complexity, and diversity in an iterated three-person game. *Artificial Life*, 2(1), 293–304.
3. Alami, R., Fleury, S., Herrb, M., Igrand, F., & Robert, F. (1998). Multi-robot cooperation in the MARTHA project. *IEEE Robotics and Automation Magazine*, 5(1), 36–47.
4. Amat, J., Mantaras, R., & Sierra, C. (1995). Cooperative autonomous low-cost robots for exploring unknown environments. In O. Khatib & J. Salisbury (Eds.), *Proceedings of the Fourth International Symposium on Experimental Robotics* (pp. 40–49). Stanford, CA: Springer-Verlag.
5. Angeline, P., & Pollack, J. (1993). Competitive environments evolve better solutions for complex tasks. In S. Forrest (Ed.), *Proceedings of the Fifth International Conference on Genetic Algorithms* (pp. 264–270). San Francisco: Morgan Kaufmann.
6. Angeline, P. (1994). An alternate interpretation of the iterated prisoner's dilemma and the evolution of non-mutual cooperation. In R. Brooks & P. Maes (Eds.), *Proceedings of the Fourth Workshop on Artificial Life* (pp. 353–358). Cambridge, MA: MIT Press.
7. Aoki, I. (1982). A simulation study on the schooling mechanism in fish. *Bulletin of the Japan Society of Scientific Fisheries*, 48, 1081–1088.

8. Arkin, R. (1998). *Behavior based robotics*. Cambridge, MA: MIT Press.
9. Arkin, R., & Balch, T. (1999). Cooperative multi-agent robotic systems. In D. Kortenkamp, R. Bonasso, & R. Murphy (Eds.), *Artificial intelligence and mobile robots*. Cambridge, MA: MIT Press.
10. Arkin, R., & Balch, T. (1999). Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, 14(6), 926–939.
11. Asada, M., Uchibe, E., Noda, S., Tawaratsumida, S., & Hosoda, K. (1994). A vision based reinforcement learning for coordination of soccer playing behaviors. In *Proceedings of AAAI Workshop on AI, ALife and Entertainment* (pp. 16–21). Seattle, WA: AAAI Press.
12. Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.
13. Axelrod, R. (1987). Evolution strategies in the iterated prisoner's dilemma. In L. Davis (Ed.), *Genetic algorithms and simulated annealing*. San Francisco, CA: Morgan Kaufmann.
14. Axelrod, R., & Hamilton, W. (1981). The evolution of cooperation. *Science*, 211(4489), 1390–1396.
15. Avouris, N., & Gasser, L. (1992). *Distributed artificial intelligence: Theory and praxis*. Dordrecht, The Netherlands: Kluwer Academic.
16. Baldassarre, G., Nolfi, S., & Parisi, D. (2002). Evolving mobile robots able to display collective behaviors. In C. Hemelrijk & E. Bonabeau (Eds.), *Proceedings of the International Workshop on Self-Organization and Evolution of Social Behavior* (pp. 11–22). Zurich: University of Zurich Press.
17. Beckers, R., Deneubourg, J., & Goss, S. (1992). Trails and U-turns in the selection of the shortest path by the ant *Lasius niger*. *Journal of Theoretical Biology*, 159(1), 397–415.
18. Beckers, R., Holland, O., & Deneubourg, J. (1994). From local actions to global tasks: Stigmergy and collective robotics. In R. Brooks & P. Maes (Eds.), *Artificial life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 181–189). Cambridge, MA: MIT Press.
19. Benda, M., Jagannathan, V., & Dodhiawalla, R. (1986). *On optimal cooperation of knowledge sources* (Tech. Rep. BCS-G2010-28). Bellevue, WA: Boeing AI Center, Boeing Computer Services.
20. Bonabeau, E., Theraulaz, G., Arpin, E., & Sardet, E. (1994). The building behavior of lattice swarms. In R. Brooks & P. Maes (Eds.), *Artificial life IV: Proceedings of the Fourth International Workshop on the Synthesis and Simulation of Living Systems* (pp. 307–312). Cambridge, MA: MIT Press.
21. Bonabeau, E., Theraulaz, G., Deneubourg, J., & Camazine, S. (1997). Self-organization in social insects. *Trends in Ecology and Evolution*, 12(1), 188–193.
22. Bonabeau, E., Dorigo, M., & Theraulaz, G. (1998). *Swarm intelligence: From natural to artificial systems*. Oxford, UK: Oxford University Press.
23. Bonabeau, E., Theraulaz, G., Deneubourg, J., & Camazine, S. (1998). Group and mass recruitment in ant colonies: The influence of contact rates. *Journal of Theoretical Biology*, 195(1), 157–166.
24. Bongard, J., & Pfeifer, R. (2003). Evolving complete agents using artificial ontogeny. In F. Hara & R. Pfeifer (Eds.), *Morpho-functional machines: The new species (designing embodied intelligence)*. Berlin: Springer-Verlag.
25. Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. Cambridge, MA: MIT Press.
26. Breder, C. (1954). Equations ascriptive of fish school and other animal aggregations. *Ecology*, 35(3), 129–211.
27. Brooks, R. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, 6, 3–15.
28. Brooks, R., & Flynn, A. (1989). Fast, cheap and out of control: A robot invasion of the solar system. *Journal of the British Interplanetary Society*, 1, 478–485.
29. Bryant, B., & Miikkulainen, R. (2003). Neuroevolution for adaptive teams. In *Proceedings of the 2003 Congress on Evolutionary Computation*. Piscataway, NJ: IEEE Press.
30. Cao, Y., Fukunaga, A., & Kahng, A. (1997). Cooperative mobile robotics: Antecedents and directions. *Autonomous Robots*, 4, 7–27.
31. Chaimowicz, L., Sugar, T., Kumar, V., & Campos, M. (2001). An architecture for tightly coupled multi-robot cooperation. In *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 2292–2297). Seoul: IEEE Press.

32. Coloni, A., Dorigo, M., & Maniezzo, V. (1991). Distributed optimization by ant colonies. In F. Varela & P. Bourguin (Eds.), *Proceedings of the European Conference on Artificial Life* (pp. 134–142). Amsterdam: Elsevier Academic.
33. Coloni, A., Dorigo, M., & Maniezzo, V. (1992). An investigation of some properties of an algorithm. In R. Manner & B. Manderick (Eds.), *Proceedings of the Parallel Problem Solving from Nature Conference* (pp. 509–520). Amsterdam: Elsevier Academic.
34. Cover, T., & Hart, P. (1967). Nearest neighbor pattern classification. *IEEE Transactions on Information Theory*, *13*, 21–27.
35. Cullen, J., Shaw, E., & Baldwin, H. (1965). Methods for measuring the three dimensional structure of fish schools. *Animal Behavior*, *13*(1), 534–543.
36. Deneubourg, J., & Goss, S. (1989). Collective patterns and decision-making. *Ecology, Ethology, and Evolution*, *1*, 295–311.
37. Deneubourg, J., Goss, S., Pasteels, J., Fresneau, D., & Lachaud, J. (1987). Self-organization mechanisms in ant societies (II): Learning in foraging and division of labor. *Behavior in Social Insects*, *54*(1), 177–196.
38. Deneubourg, J., Goss, S., Franks, A., Detrain, C., & Chretien, L. (1990). The dynamics of collective sorting: Robot-like ants and ant-like robots. In J. Meyer & S. Wilson (Eds.), *From animals to animats: Proceedings of the First International Conference on Simulation of Adaptive Behavior* (pp. 356–363). Cambridge, MA: MIT Press.
39. Deneubourg, J., Theraulaz, G., & Beckers, R. (1991). Swarm-made architectures. In F. Varela & P. Bourguin (Eds.), *Proceedings of the European Conference on Artificial Life* (pp. 123–133). Amsterdam: Elsevier Academic.
40. Denzinger, J., & Fuchs, M. (1996). Experiments in learning prototypical situations for variants of the pursuit game. In *Proceedings of the Second International Conference on Multi-Agent Systems* (pp. 48–55). Kyoto: AAAI Press.
41. Di Caro, G., & Dorigo, M. (1998). AntNet: Distributed stigmergic control for communications networks. *Journal of Artificial Intelligence Research*, *9*, 317–365.
42. Dieckmann, U., Marrow, P., & Law, R. (1995). Evolutionary cycling in predator-prey interactions: Population dynamics and the Red Queen. *Journal of Theoretical Biology*, *176*, 91–102.
43. Dorigo, M., Colombetti, M., & Patel, M. (1995). Evolutionary learning for intelligent automation: A case study. *Intelligent Automation and Soft Computing Journal*, *1*, 29–42.
44. Dorigo, M., Maniezzo, V., & Coloni, A. (1996). The ant system: Optimization by a colony of cooperating agents. *IEEE Transactions on Systems, Man, and Cybernetics—Part B*, *26*(1), 1–13.
45. Dorigo, M., & Gambardella, L. (1997). Ant Colony System: A cooperative learning approach to the traveling salesman problem. *IEEE Transactions on Evolutionary Computation*, *1*, 53–66.
46. Drogoul, A., Ferber, J., Corbara, B., & Fresneau, D. (1992). A behavioral simulation model for the study of emergent social structures. In F. Varela & P. Bourguin (Eds.), *Toward a practice of autonomous systems* (pp. 161–170). Cambridge, MA: MIT Press.
47. Drogoul, A., Corbara, B., & Fresneau, D. (1992). Applying etho-modeling to social organization in ants. In J. Billen (Ed.), *Biology and evolution of social insects* (pp. 375–383). Leuven, Belgium: Leuven University Press.
48. Drogoul, A., & Ferber, J. (1992). Using reactive multi-agent systems in simulation and problem solving. In N. Avouris & L. Gasser (Eds.), *Distributed artificial intelligence: Theory and praxis* (pp. 53–80). Brussels, Belgium: Kluwer Academic.
49. Drogoul, A., Corbara, B., & Lalande, S. (1995). MANTA: New experimental results on the emergence of (artificial) ant societies. In N. Gilbert & R. Conte (Eds.), *Artificial societies: The computer simulation of social life* (pp. 190–211). London: University College of London Press.
50. Dugatkin, L. (1990). N-person games and the evolution of cooperation: A model based on predator inspection in fish. *Journal of Theoretical Biology*, *142*(1), 123–135.
51. Epstein, J., & Axtell, R. (1996). *Growing artificial societies: Social science from the bottom up*. Cambridge, MA: MIT Press.
52. Floreano, D., & Nolfi, S. (1997). God save the Red Queen! Competition in co-evolutionary robotics. In J. Koza, D. Kalyanmoy, M. Dorigo, D. Fogel, M. Fogel, M. Garzon, H. Iba & R. Riolo (Eds.), *Genetic programming 1997: Proceedings of the Second Annual Conference*. San Francisco: Morgan Kaufmann.

53. Floreano, D., & Nolfi, S. (1997). Adaptive behavior in competing co-evolving species. In P. Husbands & I. Harvey (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*. Cambridge, MA: MIT Press.
54. Floreano, D., & Nolfi, S. (2000). *Evolutionary Robotics*. Cambridge, MA: MIT Press.
55. Floreano, D., & Urzelai, J. (2000). Evolutionary robots with on-line self-organization and behavioral fitness. *Neural Networks*, 13(1), 431–443.
56. Fong, T., Nourbakhsh, I., & Dautenhahn, K. (2003). A survey of socially interactive robots. *Robotics and Autonomous Systems*, 42, 143–166.
57. Frutiger, D., Bongard, J., & Iida, F. (2002). Iterative product engineering: Evolutionary robot design. In P. Bidaud & F. Amar (Eds.), *Proceedings of the Fifth International Conference on Climbing and Walking Robots* (pp. 619–629). Paris: Professional Engineering Publishing.
58. Gambardella, L., Taillard, E., & Dorigo, M. (1999). Ant colonies for the quadratic assignment problem. *Journal of the Operational Research Society*, 50, 167–176.
59. Goss, S., Aron, S., Deneubourg, J., & Pasteels, J. (1989). Self-organized shortcuts in the Argentine ant. *Naturwissenschaften*, 76(1), 579–581.
60. Gomez, F., & Miikkulainen, R. (1997). Incremental evolution of complex general behavior. *Adaptive Behavior*, 5(1), 317–342.
61. Gomez, F., & Miikkulainen, R. (2001). *Robust non-linear control through neuroevolution (Tech. Rep. AI-02-292)*. Austin, TX: Computer Science Department, University of Texas.
62. Gomez, F. (2003). *Robust non-linear control through neuroevolution. PhD thesis*. Austin, TX: Computer Science Department, University of Texas.
63. Grefenstette, J. (1988). Credit assignment in rule discovery systems. *Machine Learning*, 3(3), 225–246.
64. Gutowitz, H. (1993). Complexity-seeking ants. In J. Deneubourg & S. Goss (Eds.), *Proceedings of the Third European Conference on Artificial Life* (pp. 10–16). Cambridge, MA: MIT Press.
65. Hackwood, S., & Beni, G. (1991). Self-organizing sensors by deterministic annealing. In T. Yamazaki (Ed.), *Proceedings of the International Conference on Intelligent Robots and Systems* (pp. 1177–1183). Osaka, Japan: IEEE Press.
66. Hackwood, S., & Beni, G. (1992). Self-organization of sensors for swarm intelligence. In Y. Ishida (Ed.), *Proceedings of the International Conference on Robotics and Automation* (pp. 819–829). Nice, France: IEEE Press.
67. Harvey, I. (1997). Artificial evolution and real robots. In S. Masanori (Ed.), *Proceedings of International Symposium on Artificial Life and Robotics* (pp. 18–20). Berlin: Springer-Verlag.
68. Harvey, I., Husbands, P., Cliff, D., Thompson, A., & Jakobi, N. (1997). Evolutionary robotics: The Sussex approach. *Robotics and Autonomous Systems*, 20(4), 205–224.
69. Haynes, T., & Sen, S. (1995). Evolving cooperation strategies. In V. Lesser (Ed.), *Proceedings of the First International Conference on Multi-Agent Systems* (pp. 450–459). Cambridge, MA: MIT Press.
70. Haynes, T., & Sen, S. (1996). Co-adaptation in a team. *International Journal of Computational Intelligence and Organizations*, 1(4), 1–20.
71. Haynes, T., & Sen, S. (1996). Evolving behavioral strategies in predators and prey. In G. Weib & S. Sen (Eds.), *Adaptation and learning in multi-agent systems* (pp. 113–126). Berlin: Springer-Verlag.
72. Haynes, T., & Sen, S. (1997). Crossover operators for evolving a team. In J. Koza, K. Deb, M. Dorigo, D. Fogel, M. Garzon, H. Iba, & R. Riolo (Eds.), *Proceedings of Genetic Programming 1997: The Second Annual Conference* (pp. 162–167). San Francisco: Morgan Kaufmann.
73. Haynes, T., Wainwright, R., Sen, S., & Schoenefeld, D. (1995). Strongly typed genetic programming in evolving cooperation strategies. In L. Eshelman (Ed.), *Proceedings of the Sixth International Conference on Genetic Algorithms* (pp. 271–278). San Francisco: Morgan Kaufmann.
74. Hiebeler, D. (1994). The swarm simulation system and individual-based modeling. In *Proceedings of Decision Support 2001: Advanced technology for natural resource management* (pp. 20–26). Base Services Canada.
75. Hirsch, M., & Smale, S. (1974). *Differential equations, dynamical systems, and linear algebra*. Orlando, FL: Academic Press.
76. Holland, J. (1992). *Adaptation in natural and artificial systems*. Cambridge, MA: MIT Press.

77. House Books Limited. (2000). *Dictionary of Biology*. Oxford, UK: Oxford University Press.
78. Hsu, W., & Gustafson, S. (2001). Layered learning in genetic programming for a cooperative robot soccer problem. In J. Miller, M. Tomassini, P. Lanzi, C. Ryan, A. Tetamanzi, & W. Langdon (Eds.), *Proceedings of the Fourth European Conference on Genetic Programming* (pp. 291–301). Como, Italy: Springer-Verlag.
79. Hsu, W., & Gustafson, S. (2002). Genetic programming and multi-agent layered learning by reinforcements. In W. Langdon, E. Cantupaz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M. Potter, A. Schultz, J. Miller, E. Burke, & N. Jonoska (Eds.), *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 764–771). San Francisco: Morgan Kaufmann.
80. Huang, J., Jennings, N., & Fox, J. (1995). An agent-based approach to health care management. *Applied Artificial Intelligence*, 9(4), 401–420.
81. Iba, H. (1996). Emergent cooperation for multiple agents using genetic programming. In H. Voigt, W. Ebeling, I. Rechenberg, & H. Schwefel (Eds.), *Parallel problem solving from nature* (pp. 23–41). Berlin: Springer-Verlag.
82. Iba, H. (1997). Evolving communicating agents based on genetic programming. In Z. Michalewicz, T. Baeck, & X. Yao (Eds.), *Proceedings of the IEEE International Conference on Evolutionary Computation* (pp. 13–16). Piscataway, NJ: IEEE Press.
83. Iba, H. (1998). Evolutionary learning of communicating agents. *Journal of Information Science*, 108(1), 181–205.
84. Ijspeert, A., Martinoli, A., Billard, A., & Gambardella, L. (1999). *Cooperation through the exploration of local interactions in autonomous collective robotics: The stick pulling experiment (Tech. Rep.)*. Lausanne, Switzerland: Swiss Federal Institute of Technology.
85. Jamzad, M., Foroughnasirai, A., Kazemi, M., Ghorbani, R., Chiniforushan, E., & Chitsaz, H. (2000). Middle sized soccer robots: ARVAND. In *RoboCup-99: Robot Soccer World Cup III* (pp. 61–73). Berlin: Springer-Verlag.
86. Jennings, N. (1995). Controlling cooperative problem solving in industrial multi-agent systems using joint intentions. *Artificial Intelligence*, 75(2), 195–240.
87. Jennings, N., Corera, J., & Laresgoiti, I. (1995). Developing industrial multi-agent systems. In V. Lesser (Ed.), *Proceedings of the First International Conference on Multi-Agent Systems* (pp. 423–430). Cambridge, MA: MIT Press.
88. Johnson, P., & Bay, J. (1994). Distributed control of simulated autonomous mobile robot collective in payload transportation. *Autonomous Robots*, 2(1), 43–64.
89. Keller, L., & Ross, K. (1998). Selfish genes: A green beard in the red fire ant. *Nature*, 394, 573–575.
90. Keller, L., & Chapuisat, M. (1999). Cooperation among selfish individuals in insect colonies. *Bioscience*, 49, 899–909.
91. Kitano, H., Kuniyoshi, Y., Noda, I., Asada, M., Matsubara, H., & Osawa, E. (1997). RoboCup: A challenge problem for AI. *AI Magazine*, 18(1), 73–85.
92. Kitano, H., Fujita, M., Zrehen, S., & Kageyama, K. (1998). Sony legged robot for RoboCup challenge. In *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 2605–2612). Leuven, Belgium: IEEE Press.
93. Kitano, H., & Asada, M. (2000). The RoboCup humanoid challenge as the millennium challenge for advanced robotics. *Advanced Robotics*, 13(1), 723–736.
94. Korf, R. (1992). A simple solution to pursuit games. In *Working papers of the Eleventh International Workshop on DAI* (pp. 183–194). Geneva: AAAI Press.
95. Koza, J. (1992). Evolution of subsumption using genetic programming. In F. Varela & P. Bourguine (Eds.), *Proceedings of the First European Conference on Artificial Life* (pp. 110–119). Cambridge, MA: MIT Press.
96. Kim, J., Ed. (1996). *Proceedings of the Micro-Robot World Cup Soccer Tournament*. Taejon, Korea: Elsevier Science.
97. Kraus, S. (1997). Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94(1), 80–95.

98. Kube, C., & Zhang, H. (1993). Collective robotics: From social insects to robots. *Adaptive Behavior*, 2(1), 189–218.
99. Kube, C., & Zhang, H. (1997). Task modeling in collective robotics. *Autonomous Robots*, 4(1), 53–72.
100. Kube, C., & Bonabeau, E. (1999). Cooperative transport by ants and robots. *Robotics and Autonomous Systems*, 1, 20–29.
101. Levy, R., & Rosenschein, J. (1992). Game theoretic approach to distributed artificial intelligence and the pursuit problem. In *Decentralized AI III* (pp. 129–146). Kaiserslautern, Germany: Springer-Verlag.
102. Luke, S., Hohn, C., Farris, J., Jackson, G., & Hendler, J. (1998). Co-evolving soccer softbot team coordination with genetic programming. In H. Kitano (Ed.), *RoboCup-97: Robot Soccer World Cup I* (pp. 398–411). Berlin: Springer-Verlag.
103. Maniezzo, V., & Colorni, A. (1994). The ant system applied to the quadratic assignment problem. *IEEE Transactions on Knowledge and Data Engineering*, 11(5), 769–778.
104. Mataric, M. (1992). Designing emergent behaviors: From local interactions to collective intelligence. In J. Meyer, H. Roitblat, & S. Wilson (Eds.), *From animals to animats 2: Second International Conference on the Simulation of Adaptive Behavior* (pp. 432–441). Cambridge, MA: MIT Press.
105. Mataric, M., & Cliff, D. (1996). Challenges in evolving controllers for physical robots. *Robotics and Autonomous Systems*, 19(1), 67–83.
106. Mataric, M. (1997). Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, 9(1), 62–78.
107. Matsubara, H., Noda, I., & Hiraki, K. (1996). Learning of cooperative actions in multi-agent systems: A case study of pass and play in soccer. In S. Sandip (Ed.), *Adaptation, co-evolution, and learning in multi-agent systems: Papers from the 1996 AAAI Spring Symposium* (pp. 63–67). Boston: AAAI Press.
108. McCauley, E., Wilson, W., & deRoos, A. (1993). Dynamics of age-structured and spatially structured predator-prey interactions: Individual-based models and population-level formulations. *American Naturalist*, 142(3), 412–442.
109. Meeden, L., & Kumar, D. (1998). Trends in evolutionary robotics. In L. Jain & T. Fukuda (Eds.), *Soft computing for intelligent robotic systems* (pp. 215–233). New York: Physica-Verlag.
110. Miller, G., & Cliff, D. (1994). *Co-evolution of pursuit and evasion I: Biological and game-theoretic foundations (Tech. Rep. CSR311)*. Brighton, UK: School of Cognitive and Computing Sciences University of Sussex.
111. Mitchell, T. (1997). *Machine learning*. Berkshire, UK: McGraw-Hill.
112. Mitsumoto, N., Fukuda, T., Shimojima, K., & Ogawa, A. (1995). Micro autonomous robotic system and biologically inspired swarm strategy as a multi-agent robotic system. In *Proceedings of the IEEE International Conference on Robotics and Automation* (pp. 2187–2192). Piscataway, NJ: IEEE Press.
113. Mondada, F., Franzi, E., & Jenne, P. (1993). Mobile robot miniaturization: A tool for investigation in control algorithms. In T. Yoshikawa & F. Miyazaki (Eds.), *Proceedings of the Third International Symposium on Experimental Robotics* (pp. 501–513). Berlin: Springer-Verlag.
114. Mondada, F., Floreano, D., Guignard, A., Deneubourg, J., Gambardella, L., Nolfi, S., & Dorigo, M. (2002). *Search for rescue: An application for the SWARM-BOT self-assembling robot concept* (Tech. Rep. I2S-LSA-STI). Lausanne, Switzerland: Swiss Federal Institute of Technology.
115. Montana, D. (1995). Strongly typed genetic programming. *Evolutionary Computation*, 3(2), 199–230.
116. Moriarty, D., & Miikkulainen, R. (1997). Forming neural networks through adaptive co-evolution. *Evolutionary Computation*, 5(4), 373–399.
117. Moriarty, D. (1997). *Symbiotic evolution of neural networks in sequential decision tasks*. Ph.D. dissertation. Austin, TX: Department of Computer Sciences The University of Texas.
118. Nishimura, S., & Takashi, I. (1997). Emergence of collective strategies in a predator-prey game model. *Artificial Life*, 3(1), 243–260.
119. Noda, I., Matsubara, H., & Hiraki, K. (1996). Learning cooperative behavior in multi-agent environment: A case study of choice of play-plans in soccer. In N. Foo & R. Goebel (Eds.), *Proceedings of the Fourth Pacific Rim International Conference on Artificial Intelligence* (pp. 570–579). Cairns, Australia: Springer-Verlag.

120. Noda, I., Matsubara, H., Hiraki, K., & Frank, I. (1998). Soccer Server: A tool for research on multi-agent systems. *Applied Artificial Intelligence*, 12(1), 233–250.
121. Noda, I., & Stone, P. (2001). The RoboCup Soccer Server and CMUnited clients: Implemented infrastructure for MAS research. *Autonomous Agents and Multi-Agent Systems*, 7(1), 101–120.
122. Nolfi, S. (2000). *Evorobot 1.1 user manual* (Tech. rep.). Rome: Institute of Cognitive Sciences and Technologies National Research Council.
123. Nolfi, S., Baldassarre, G., & Parisi, D. (2003). Evolution of collective behaviour in a team of physically linked robots. In R. Gunther, A. Guillot, & J. Meyer (Eds.), *Applications of evolutionary computing* (pp. 581–592). Heidelberg, Germany: Springer-Verlag.
124. Nolfi, S., Deneubourg, J., Floreano, D., Gambardella, L., Mondada, F., & Dorigo, M. (2003). Swarm-Bots: Swarm of mobile robots able to self-assemble and self-organize. *Ecrium New*, 53, 25–26.
125. Oliphant, M. (1994). Evolving cooperation in the non-iterated prisoner's dilemma: The importance of spatial organization. In R. Brooks & P. Maes (Eds.) *Proceedings of the Fourth Artificial Life Workshop* (pp. 349–352). Cambridge, MA: MIT Press.
126. Parker, L. (1994). An experiment in mobile robotic cooperation. In *Proceedings of the ASCE Specialty Conference on Robotics for Challenging Environments* (pp. 131–139). Cambridge, MA: MIT Press.
127. Parker, L. (1994) L-ALLIANCE: Task-oriented multi-robot learning in behavior-based systems. *Advanced Robotics*, 11(4), 305–322.
128. Parker, L. (1999). Cooperative robotics for multi-target observation. *Intelligent Automation and Soft Computing*, 5(1), 5–19.
129. Parker, L. (1999). Adaptive heterogeneous multi-robot teams. *Neuro-computing*, 28(1), 75–92.
130. Perez-Uribe, A., Floreano, D., & Keller, L. (2003). Effects of group composition and level of selection in the evolution of cooperation in artificial ants. In W. Banzhaf, T. Christaller, P. Dittrich, J. Kim, & J. Ziegler (Eds.), *Advances of artificial life: Proceedings of the Seventh European Conference on Artificial Life* (pp. 128–137). Berlin: Springer-Verlag.
131. Di Pietro, A., While, L., & Barone, L. (2002). Learning in RoboCup keep-away using evolutionary algorithms. In W. Langdon, E. Cantupaz, K. Mathias, R. Roy, D. Davis, R. Poli, K. Balakrishnan, V. Honavar, G. Rudolph, J. Wegener, L. Bull, M. Potter, A. Schultz, J. Miller, E. Burke, & N. Jonoska (Eds.), *GECCO 2002: Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 1065–1072). San Francisco: Morgan Kaufmann.
132. Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. Cambridge, MA: MIT Press.
133. Quinn, M., Smith, L., Mayley, G., & Husbands, P. (2002). Evolving team behavior for real robots. In *International Workshop on Biologically-Inspired Robotics: The legacy of W. Grey Walter*. Bristol, UK: Hewlett Packard Labs.
134. Quinn, M., Smith, L., Mayley, G., & Husbands, P. (2002). *Evolving formation movement for a homogeneous multi-robot system: Teamwork and role-allocation with real robots* (Tech. rep. 515). Brighton, UK: School of Cognitive and Computing Sciences University of Sussex.
135. Quinn, M., Smith, L., Mayley, G., & Husbands, P. (2003). Evolving controllers for a homogeneous system of physical robots: Structured cooperation with minimal sensors. *Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 361(1), 2321–2344.
136. Ray, T. (2001). *Overview of Tierra at ATR* (Tech. rep. ATR-HIP-15). Kyoto, Japan: Advanced Telecommunications Research Institute, Human Information Processing Research Laboratories.
137. Reynolds, C. (1987). Flocks, herds and schools: A distributed behavioral model. *Computer Graphics*, 21(4), 25–36.
138. Rumelhart, D., Hinton, G., & Williams, R. (1986). Learning internal representations by error propagation. In D. Rumelhart & J. McClelland (Eds.), *Parallel distributed processing, Volume 1: Foundations* (pp. 318–362). Cambridge, MA: MIT Press.
139. Sahota, M., Mackworth, A., Barman, R., & Kingdon, S. (1995). Real-time control of soccer playing robots using off-board vision: The Dynamite test bed. In *IEEE International Conference on Systems, Man, and Cybernetics* (pp. 3690–3663). Vancouver, BC: IEEE Press.
140. Sandholm, T., & Lesser, V. (1997). Coalitions among computationally bounded agents. *Artificial Intelligence*, 94(1), 99–137.

141. Singh, M. (1990). Towards a theory of situated know-how. In L. Aiello (Ed.), *Proceedings of the Ninth European Conference on Artificial Intelligence* (pp. 604–609). Stockholm, Sweden: Springer-Verlag.
142. Stone, P., & Veloso, M. (1998). Towards collaborative and adversarial learning: A case study in robotic soccer. *Evolution and Learning in Multi-Agent Systems*, 48(1), 83–104.
143. Stone, P., & Veloso, M. (1998). Using decision tree confidence factors for multi-agent control. In H. Kitano (Ed.), *Proceedings of the Second International Conference on Autonomous Agents* (pp. 110–116). ACM/SIGART.
144. Stone, P., & Veloso, M. (1998). A layered approach to learning client behaviors in the RoboCup Soccer Server. *Applied Artificial Intelligence*, 12, 165–188.
145. Stone, P., & Veloso, M. (1998). The CMUNITED-97 simulator team. In H. Kitano (Ed.), *RoboCup-97: Robot soccer world cup* (pp. 389–397). Berlin: Springer-Verlag.
146. Stone, P., & Veloso, M. (1999). Task decomposition, dynamic role assignment, and low bandwidth communication for real time strategic teamwork. *Artificial Intelligence*, 110(2), 241–273.
147. Stone, M., Veloso, M., & Bowling, M. (1999). Anticipation as a key for collaboration in a team of agents: A case study in robotic soccer. In G. McKee & P. Schenker (Eds.), *Proceedings of SPIE Sensor Fusion and Decentralized Control in Robotic Systems II* (pp. 134–141). Bellingham, WA: SPIE Press.
148. Stone, P., & Sutton, R. (2002). Scaling reinforcement learning toward RoboCup Soccer. In *Proceedings of the Eighteenth International Conference on Machine Learning* (pp. 537–544). San Francisco: Morgan Kaufmann.
149. Sutton, R., & Barto A. (1998). An introduction to reinforcement learning. Cambridge, MA: MIT Press.
150. Theraulaz, G., & Bonabeau, E. (1995). Modeling the collective building of complex architectures in social insects with lattice swarms. *Journal of Theoretical Biology*, 177(1), 381–400.
151. Trianni, V., Gross, R., Labella, T., Sahin, E., & Dorigo M. (2003). Evolving aggregation behaviors in a swarm of robots. In W. Banzhaf, T. Christaller, P. Dittrich, J. Kim, & J. Ziegler (Eds.), *Advances in artificial life: Proceedings of the Seventh European Conference on Artificial Life* (pp. 865–874). Dortmund, Germany: Springer-Verlag.
152. Watson, R., Ficici, S., & Pollack, J. (1999). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. In P. Angeline, Z. Michalewicz, M. Schoenauer, X. Yao, & A. Zalzalá (Eds.), *Proceedings of the IEEE Congress on Evolutionary Computation* (pp. 335–342). Berlin: Springer-Verlag.
153. Werner, G., & Dyer, M. (1993). Evolution of herding behavior in artificial animals. In J. Meyer, H. Roitblat, & S. Wilson (Eds.), *Proceedings of Second International Conference on Simulation of Adaptive Behavior* (pp. 393–399). Cambridge, MA: MIT Press.
154. Whiteson, S., Kohl, N., Miikkulainen, R., & Stone, P. (2003). Evolving keep-away soccer players through task decomposition. In *Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 356–368). Chicago: Springer-Verlag.
155. Van Valen, L. (1973). A new evolutionary law. *Evolution Theory*, 1, 1–30.
156. Veloso, M., Bowling, M., Achim, S., Han, K., & Stone, P. (1999). The CMUnited-98 champion small robot team. In M. Asada & H. Kitano (Eds.), *RoboCup-98: Robot soccer world cup II* (pp. 77–92). Berlin, Germany: Springer Verlag.
157. Veloso, M., & Uther, W. (1995). The CM Trio-98 Sony legged robot team. In M. Asada & H. Kitano (Eds.), *RoboCup-98: Robot soccer world cup II* (pp. 491–497). Berlin: Springer-Verlag.
158. Veloso, M., Bowling, M., & Stone, P. (2000). The CMUnited-98 champion small-robot team. *Advanced Robotics*, 13(1), 753–767.
159. Wilson, E. (1985). The sociogenesis of insect colonies. *Science*, 228, 1489–1495.
160. Yong, C., & Miikkulainen, R. (2001). *Cooperative co-evolution of multi-agent systems* (Tech. rep. AI01-287). Austin, TX: Department of Computer Sciences University of Texas.
161. Zaera, N., Cliff, D., & Bruten, J. (1996). *(Not) evolving collective behaviors in synthetic fish* (Tech. rep.). Bristol, UK: Hewlett-Packard Laboratories.
162. Zlatev, J. (2001). The epigenesis of meaning in human beings, and possibly in robots. *Minds and Machines*, 11(2), 155–195.