Chapter VIII

Emergent Specialization in Biologically Inspired Collective Behavior Systems

G. S. Nitschke, Vrije Universiteit, The Netherlands

M. C. Schut, Vrije Universiteit, The Netherlands

A. E. Eiben, Vrije Universiteit, The Netherlands

Abstract

Specialization is observable in many complex adaptive systems and is thought by many to be a fundamental mechanism for achieving optimal efficiency within organizations operating within complex adaptive systems. This chapter presents a survey and critique of collective behavior systems designed using biologically inspired principles. Specifically, we are interested in collective behavior systems where specialization emerges as a result of system dynamics and where emergent specialization is used as a problem solver or means to increase task performance. The chapter presents an argument for developing

design methodologies and principles that facilitate emergent specialization in collective behavior systems. Open problems of current research as well as future research directions are highlighted for the purpose of encouraging the development of such emergent specialization design methodologies.

Introduction

Specialization is observable in many complex adaptive systems[1] and is thought by many to be a fundamental mechanism for achieving optimal efficiency within certain complex adaptive systems. In complex ecological communities, specializations have evolved over time as a means of diversifying the community in order to adapt to the environment (Seligmann, 1999). Over the course of evolutionary time, specialization in biological communities has assumed both morphological (Wenseleers, Ratnieks, & Billen, 2003) and behavioral forms (Bonabeau, Theraulaz, & Deneubourg, 1996). For example, morphologically specialized castes have emerged in certain termite colonies (Noirot & Pasteels, 1987), and honeybees dynamically adapt their foraging behavior for pollen, nectar, and water as a function of individual preference and colony demand (Calderone & Page, 1988). The consequence of such specializations is that labor is efficiently divided between specialized castes[2] and individuals for the benefit of accomplishing group tasks. In such a sense, specialization can be viewed as an adaptive mechanism in a complex adaptive system.

Many artificial complex adaptive systems that exhibit collective behavior have used design principles, which draw their inspiration from examples of specialization in nature. Such examples include complex ecological communities such as social insect colonies (Bonabeau et al., 1996; Bonabeau, Sobkowski, Theraulaz, & Deneubourg, 1997; Calderone et al., 1988; Noirot et al., 1987; Seligmann, 1999; Wenseleers et al., 2003) biological neural networks (Baev, 1997), multi-cellular organisms (Hawthorne, 2001), economies of a nation, companies, corporations, and other business organizations (Abdel-Rahman, 2001; Ng & Yang, 1997; Resnick, 1997). Such biologically inspired design principles are especially prevalent in multi-robot (Potter, Meeden, & Schultz, 2001) swarm intelligence (Bonabeau, Dorigo, & Theraulaz, 1998) and artificial life systems (Nishimura & Takashi, 1997) where it is highly desirable to replicate the success of biological collective behavior systems.

Suppositions of Specialization

Given empirical evidence offered by research in both biological collective behavior systems, and biologically inspired artificial collective behavior systems[3], two key observations can be stated.

- Specialization that assumes either behavioral or morphological forms is often present in biological systems that exhibit collective behavior.
- In biological systems that exhibit collective behavior, specialization is beneficial in that it increases the efficiency of the system, or allows collective behavior tasks to be solved that could not otherwise be solved by individuals within the system.

Given these observations, one can formulate the assumption that specialization is beneficial in biological inspired artificial complex adaptive systems that are designed to solve certain types of collective behavior tasks. Examples of such types of collective behavior tasks are presented in section *Collective Behavior Tasks and Specialization*. In order for this assumption to be proved, this chapter proposes the need to develop *emergent behavior design* methodologies[4]. Such methodologies would dictate design and engineering principles for creating an artificial complex adaptive system capable of solving collective behavior tasks that require or benefit from specialization. Ideally, such methodologies would result in the production of artificial complex adaptive systems that yield emergent yet desired forms of specialization. As in biological systems, this emergent specialization could then be harnessed and used by the system for the benefit of either increasing task performance, or solving certain collective behavior tasks, that could not otherwise be solved.

Chapter Goal and Motivation: Specialization as a Problem Solver

The chapter's scope is a survey and critique of collective behavior systems designed using biologically inspired design principles that use emergent specialization to solve collective behavior tasks. Such design principles include *self-organization*, *learning,* and *evolution* (Brooks, 1990). This chapter presents an argument for utilizing emergent behavioral specialization as a problem solver in biologically inspired artificial complex adaptive systems. Such utilization would be advantageous given the numerous real world applications where specialization is beneficial. Examples of such applications are presented in the section *Collective Behavior Tasks and Specialization*. This chapter's motivation is similar to that cited for the *organic computing* research endeavor (Müller & Sick, 2006). Organic computing has recently achieved some success in investigating the notion of defining and measuring concepts such as emergence and self-organization in large distributed complex adaptive systems. The key idea is to utilize emergent phenomena for the benefit of solving tasks in *organic computing systems*. An organic computing system is a technical system, which adapts dynamically to the current conditions of its environment. It is self-organizing, self-configuring, self-repairing, self-protecting, self-explaining, and context-aware (Müller et al., 2006). Initial research in this area displays great promise, and includes exploiting emergent functionality at the hardware level of visual microprocessors for image recognition tasks (Komann & Fey, 2007), self-organizing, and self-stabilizing role assignment in sensor and actuator networks (Weis, Parzyjegla, Jaeger, & Mühl, 2006), and self-organization of job scheduling and distribution of jobs over nodes in a network (Trumler, Klaus, & Ungerer, 2006).

Chapter Scope: Behavioral Specialization

Another important issue is which type of specialization[5] should be instituted for the benefit of a collective behavior system. We have elected to only survey research literature concerned with *behavioral specialization*. The decision to adopt this focus was based on the discovery that with relatively few exceptions (section: *Types of Specialization*) the majority of research concerning the use of emergent specialization for improving task performance is restricted to simulated systems. This is so, given the obvious engineering challenges and inherent complexity of dynamically creating morphologically specialized robots and computer components, that represent effective solutions to emerging challenges in a physical task environment (Parker & Nathan, 2006; Pfeifer, Iida, & Gomez, 2006; Watson, Ficici, & Pollack, 1999b). Figure 1 presents the scope of the chapter within the dimensions of emergent versus non-emergent phenomena and behavioral versus morphological specialization.

Figure 1. Types of specialization in biologically inspired collective behavior systems. The top left-hand side quadrant defines the scope of this chapter. Specifically, adaptive systems that use heterogeneous or homogenous design approaches with the aim of deriving emergent behavioral specialization for solving collective behavior tasks. See section: Types of Specialization for details.

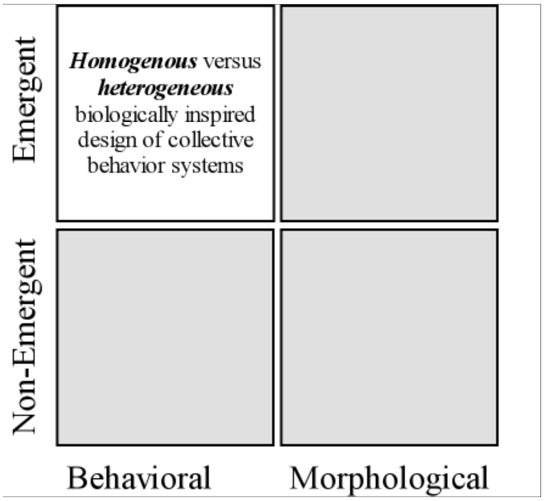

Types of Specialization

Specialization in collective behavior systems has been studied from many different perspectives (Bongard, 2000; Bryant & Miikkulainen, 2003; Blumenthal & Parker, 2004b; Campos, Theraulaz, Bonabeau, & Deneubourg, 2001; Haynes & Sen, 1996b; Nolfi et al., 2003b; Stone & Veloso, 2002; Whiteson, Kohl, Miikkulainen, & Stone, 2003), and is thus often defined in accordance with the goals of researchers conducting the study. Within collective behavior literature, specialization is either studied as an emergent property of the system, or is explicitly pre-programmed into the systems components. With notable exceptions such as Funes, Orme, and Bonabeau (2003), there are few examples of research that successfully specifies, *a priori*, what exactly the behavior of system components should be, in order to produce a specifically desired, yet emergent collective behavior.

Non-Emergent Specialization

Non-emergent specialization is that which is explicitly pre-specified to be apart of the design of system components and global behavior of a system. Such approaches are either static, or utilize learning algorithms so as to ascertain which type of behavioral specialization, selected from a given set, is most appropriate for solving a given task. Such approaches are useful for solving collective behavior tasks that require specialization, where the degree of specialization required can be sufficiently described *a priori* (Arkin & Balch, 1999; Balch, 2002a, 2002b).

Emergent Specialization

Emergent specialization is that which emerges from the interaction of system components in response to a dynamic task that requires varying degrees, or different types of specialization, in order to effectively accomplish. Such approaches have become popular in collective behavior task domains where one does not know, *a priori*, the degree of specialization required to optimally solve the given task (Gautrais, Theraulaz, Deneubourg, & Anderson, 2002; Luke & Spector, 1996; Murciano & Millan, 1997a; Murciano, Millan, & Zamora, 1997b; Potter et al., 2001; Stanley, Bryant, & Miikkulainen, 2005b; Theraulaz, Bonabeau, & Deneubourg, 1998b; Waibel, Floreano, Magnenat, & Keller, 2006). The section *Heterogeneous vs. Homogenous Design of Emergent Specialization* elaborates upon such emergent specialization design approaches.

Morphological vs. Behavioral Specialization

It is possible to further categorize specialization into two distinct classes: *morphological* (Martinoli, Zhang, Prakash, Antonsson, & Olney, 2002; Zhang, Martinoli, & Antonsson, 2003) and *behavioral* (Bonabeau et al., 1997; Li, Martinoli, & Mostafa, 2002).

The term *morphological specialization* is applicable to situated and embodied agents, operating in simulated or physical task environments, with embodiment (sensors and actuators) structured so as to yield an advantage in accomplishing the task (Watson et al., 1998a, 1999b, 2002). Examples of morphological specialization include the evolution of optimal arrangements of sensors and actuators in the design of simulated automobiles (Martinoli et al., 2002; Zhang et al., 2003), evolution of agent morphologies and controllers for various forms of motion in simulated environments (Sims, 2004), evolution of physical electric circuits for control (Thompson, Harvey, & Husbands, 1996), and evolving robot morphology for accomplishing different forms of physical motion (Lipson & Pollack, 2000).

The term *behavioral specialization* is applicable to agents with behaviors that are advantageous for accomplishing specific types of tasks (Balch, 2002a, 2002b; Nolfi & Floreano, 2000; Nolfi & Parisi, 1997). Examples of behavioral specialization include the use of machine learning methods that activate certain behaviors with a particular frequency as a response to dynamically arising tasks (Gautrais et al., 2002).

Collective Behavior Methods for Specialization

There is some agreement among researchers as to the methods for specialization that are appropriate for particular collective behavior tasks. Figure 2 illustrates a categorization of such methods, which are briefly detailed in the following. The categories illustrated in Figure 2 are by no means exhaustive, but rather several examples that have recently received particular research attention.

Figure 2. Collective behavior methods of specialization. See section: Collective Behavior Methods for Specialization for details.

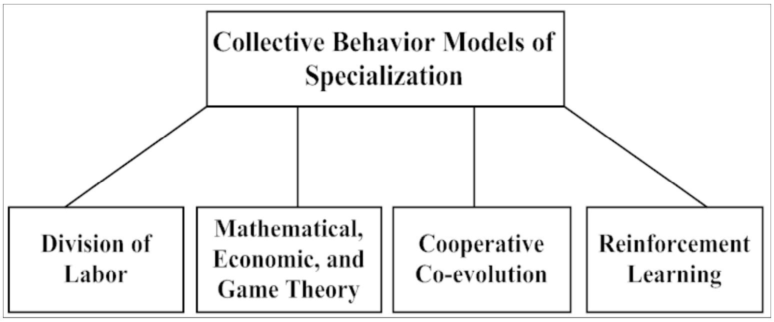

Division of Labor Methods

The use of behavioral threshold and division of labor methods have been investigated within the context of ant-based (Deneubourg, Goss, Pasteels, Fresneau, & Lachaud, 1987) and resource allocation (Bonabeau et al., 1997) methods. Such methods typically utilize feedback signals given to agents of the same caste (Kreiger & Billeter, 2000) in order to encourage the emergence of specialization for a specific task. Many variations of these methods exist (Bonabeau & Theraulaz, 1999; Bonabeau et al., 1996, 1997, 1998; Deneubourg et al., 1987; Robson & Traniello, 1999; Theraulaz, Gervet, & Semenoff, 1991; Theraulaz, Goss, Gervet, & Deneubourg, 1991), including those that use evolutionary algorithms (Tarapore, Floreano, & Keller, 2006; Waibel et al., 2006), and reinforcement learning models (Murciano et al., 1997a, Murciano et al., 1997b) in order to derive threshold values. The goal of such models is typically to optimize global task performance. Such methods are appealing as their evolutionary dynamics and emergent properties can usually be described with a mathematical representation and the results of such models are thus typically amenable to a mathematical analysis (Wu, Di, & Yang, 2003).

Mathematical, Economic, and Game Theory Methods

Linear, non-linear, and dynamic methods based in mathematical, economic, and game theory (Axelrod, 1984; Solow & Szmerekovsky, 2004) have many applications for resource assignment problems in business. For example, the maximum matching algorithm developed by Edmonds (1965) was designed to determine the maximum number of people that can be assigned to tasks in such a way that no person is assigned to more than one task. Thus, it is assumed that each person specializes in performing at most one task. Such methods are advantageous as results can be subject to a formal analysis. However, they are limited by their abstract nature, and assume that the task domain can be mathematically or otherwise formally represented.

Cooperative Co-Evolution Methods

Cooperative co-evolution methods have been implemented both in the context of modified genetic algorithms, for example, *Cooperative Co-evolutionary*

Genetic Algorithms (Potter & DeJong, 2000), and in the context of neuro-evolution methods, for example, *Enforced Sub-Populations* (ESP) (Gomez, 1997). In both cases, the genotype space is decomposed into a set of sub-populations, where each generation, the evolutionary process selects the best performing genotype components from each sub-population so as to construct a complete genotype as a solution. Decomposition of the genotype space into sub-populations, genotype construction from multiple sub-populations, and genotype to phenotype mapping depends upon the approach used. For example, the ESP method encodes separate neurons as genotype components to be distributed between sub-populations, where the composition of neurons encodes a complete neural network. Advantages of such methods include their versatility, and applicability to a broad range of complex, continuous, and noisy task domains. Also, the representation of the genotype space as a set of sub-populations provides a natural representation for many collective behavior tasks, and often effectuates the derivation of specialized phenotypes. A key disadvantage of such approaches is slow derivation of viable solutions in complex task domains due to inherently large search spaces. Also, the genotype representations that produce desired results can typically not be easily interpreted.

Reinforcement Learning Methods

There exists a certain class of reinforcement learning methods that provide periodic feedback signals to agent groups attempting to accomplish a collective behavior task (Sutton & Barto, 1998). A reinforcement signal is either local or global. Local reinforcement signals are calculated by, and given to a single agent, or a caste, upon task accomplishment. Global reinforcement signals are calculated by and given to the entire agent group at the end of a reinforcement learning trial (Li, Martinoli, & Yaser, 2004). The main advantage of reinforcement learning approaches is that agents are able to effectively operate in complex and noisy environments, with incomplete information. However, approaches that utilize only a global reinforcement signal, do not typically effectuate specialization in the group, even if task performance could be increased with specialized agents (Li et al., 2002, 2004). Approaches that utilize local reinforcement signals have been demonstrated as being appropriate for deriving specialized agents (Li et al., 2004), however such approaches suffer from the credit assignment problem (Grefenstette, 1995; Sutton et al., 1998), which potentially leads to sub-optimal collective behavior solutions.

Heterogeneous vs. Homogenous Design of Emergent Specialization

In collective behavior research, approaches to designing emergent specialization usually adopt either *homogeneous* or *heterogeneous* methods for designing system components. Homogeneous approaches utilize a single agent behavior for every agent in a group of agents. Agent behavior may be encoded as one genotype representation, or in some cases simply defined by a given set of parameters, which are copied for each agent in the group (Quinn, Smith, Mayley, & Husbands, 2003). Heterogeneous approaches utilize different behaviors for each agent in a group of agents. The set of different behaviors is sometimes encoded as different populations of genotypes, as in the case of cooperative co-evolutionary genetic algorithms (Parker, 2000). Alternatively, different agent behaviors may simply be represented as different sets of parameters (Campos et al., 2001).

Designing emergent specialization has been studied via specifying homogeneity vs. heterogeneity within both the genotypes and phenotypes of individual agents as well as entire agent groups. Specialization is often closely associated with, and sometimes synonymous with, heterogeneity in collective behavior systems (Balch, 1998; Potter et al., 2001). Heterogeneity can be hardwired or plastic, and may assume either behavioral (Bryant et al., 2003; Noirot et al., 1987; Whiteson et al., 2003), or morphological (O'Riain, Jarvis, Alexander, Buffenstein, & Peeters, 2000; Schultz & Bugajska, 2000; Zhang et al., 2003) forms. Plastic heterogeneity is when a group adapts its degree of heterogeneity as a function of environment and task constraints, where as, hardwired heterogeneity is when the degree of heterogeneity in the group remains static (Li et al., 2002).

Certain researchers have attempted to outline generalized guidelines as when to use either homogeneous or heterogeneous design approaches. For example, Balch (1998) suggested that collective behavior task domains where all individuals are able to perform the task, such as collective gathering, are particularly suited for homogeneous design. Whilst, task domains that explicitly require complementary roles, such as RoboCup soccer, are more suitable for heterogeneous approaches. However, such guidelines, as with studies of specialization, are usually defined according to the goals and perspectives of the researcher. Hence, one can readily find examples of when homogeneity and heterogeneity have been used in a manner incongruent to any given set of design principles or guidelines.

Homogeneous Approaches

In homogeneous approaches, specialization is typically studied at the group level since emergent specialization depends upon the local interactions of cloned behaviors. At the genotype level, the key advantage of a homogeneous approach is that the search space size is kept minimal since an algorithm need only optimize a single behavior. At the phenotype level, homogeneous groups are potentially more adaptive than heterogeneous groups at coping with the loss of group members. Also, homogenous groups typically have greater flexibility in coordinating behaviors so as to produce an effective collective behavior (Stone & Veloso, 1999). The key disadvantage of such approaches is that system homogeneity, either at the genotype or phenotype level, does not facilitate specialization, so it is likely that such collective behavior systems will converge to a non-specialized solution, even if specialization is advantageous in the given task domain.

Heterogeneous Approaches

Heterogeneous approaches typically study emergent specialization at either the local (agent) or global (entire group) level. The key advantage of heterogeneity is that it encourages and facilitates emergent specialization, both at the individual and group level. The key disadvantage of heterogeneous approaches is that the search space is usually (for complex tasks) prohibitively large comparative to homogeneous approaches, since many different agent behaviors need to be optimized or otherwise adapted for task accomplishment.

Collective Behavior Tasks and Specialization

In the design of collective behavior systems, it remains an open research question as to which task domains are most appropriately solved using specialization. However, there is some agreement amongst researchers that if the task can be naturally decomposed into a set of complementary sub-tasks then specialization is often beneficial for increasing collective task performance (Arkin, 1998; Arkin et al., 1999; Balch, 2002a; Balch, 2002b). The following list enumerates several categories of such collective behavior task domains. Each of these task domain categories mandates some degree of collective behavior, where specialization is beneficial for improving task performance. In subsequent sections, specific research examples selected from each of these categories are briefly examined.

- **Collective Gathering** (Bonabeau et al., 1998; Perez-Uribe, Floreano, & Keller, 2003).
- **Collective Construction** (Murciano et al., 1997a, 1997b; Theraulaz & Bonabeau, 1995).
- **Collective Resource Distribution and Allocation** (Bonabeau et al., 1996, 1997; Campos et al., 2001; Theraulaz et al., 1998a, 1998b).
- **Multi-Agent Computer Games** (Bryant et al., 2003), (Stanley & Miikkulainen, 2002; Stanley et al., 2005b).
- **RoboCup Soccer** (Luke, Farris, Jackson, & Hendler, 1998; Luke et al., 1996; Stone et al., 1999.
- **Predator-Prey and Collective Herding Behaviors** (Blumenthal et al., 2004b), (Blumenthal & Parker, 2004a, 2004c; Luke et al., 1996; Potter et al., 2001).
- **Moving in Formation and Cooperative Transportation Tasks** (Kube & Bonabeau, 1999; Nolfi et al., 2003b; Quinn et al., 2003).

Collective Gathering

Collective gathering is a task domain characterized by the social insect metaphor. That is, collective gathering tasks seek to emulate the success and efficiency of social insects in gathering resources. Collective gathering tasks have been studied in the context of both physical multi-robot systems (Kreiger et al., 2000; Mataric, 1997) and simulated multi-robot systems (Ijspeert, Martinoli, Billard, & Gambardella, 2001), as well as more abstract artificial life simulations (Bongard, 2000; Deneubourg, Theraulaz, & Beckers, 1991; Perez-Uribe et al., 2003). The collective gathering task domain requires that a group of agents search for, collect, and transport resources in the environment from their initial locations to some particular part of the environment. Such gathering tasks typically require that the group of agents allocate their labor efforts to particular sub-tasks so as to derive a collective behavior that maximizes the quantity of resources gathered[6]. Collective gathering tasks are typically viewed as optimization problems and have been traditionally studied with mathematical or otherwise analytical methods (Bonabeau et al., 1996; Gautrais et al., 2002; Theraulaz et al., 1998a).

Learning Behavioral Specialization for Stick Pulling

The research of Li et al. (2004) addressed the important issue of attempting to specify the concepts of heterogeneity and specialization in a formal definition, so as emergent heterogeneity and specialization[7] would be measurable within the larger context of collective behavior and distributed systems research. In a case study that compared centralized and distributed learning methods, the authors qualitatively measured the diversity and specialization of a simulated multi-robot system given a stick-pulling task that mandated specialized and cooperative behavior. One research goal was to investigate the impact of diversity, in the form of heterogeneity in behaviors, upon emergent specialization and in turn the impact of specialization on task performance.

In all experiments, the authors presented a learning method that effectively operated within a multi-robot simulator, where specialization emerged as a function of task constraints and environmental conditions regardless of whether local or global reinforcement signals were used. The authors' explanation for this result was that if behavioral diversity (heterogeneity) is beneficial to task performance, then the learning method facilitates emergent specialization as a means of taking advantage of this behavioral diversity.

The key criticism of this research is the dependency between emergent specialization and the learning method used, and consequently the methods applicability to more generalized optimization tasks. Results supported a hypothesis that if behavioral diversity in a group was beneficial to task performance, then specialization was likely to emerge and increase accordingly with behavioral diversity and task performance. However, these results largely depended upon the type of learning method, the model of the task environment, robot controller parameters that defined membership to a caste, and the task related parameters that the learning method sought to optimize. Thus, the degree to which emergent specialization depended upon the underlying adaptation process remains an open question. Also, the system designer needed to select task environment parameters for the learning method. This cast doubt upon the possibility of applying the learning method to more complex and dynamic task environments, where pertinent task environment parameters that the learning method would require in order to encourage diversity, specialization, and increased task performance, could not be identified *a priori*.

Furthermore, the number of castes composing a group was determined by the system designer and not by the adaptive process. Experiments that analyzed emergent caste formation would be necessary in order to effectively ascertain the relationship between heterogeneity, specialization, and collective behavior task performance. An adaptive process where a particular number of castes emerge in response to simulation environment and task constraints would make such a process applicable to complex task environments where task challenges dynamically arise.

Collective Construction

Collective Construction is a task domain characterized by the social insect metaphor. That is, collective construction tasks seek to emulate the success and efficiency of social insects in gathering resources. Collective construction tasks have mainly been studied in the context of artificial life simulations (Bonabeau, Theraulaz, Arpin, & Sardet, 1994; Murciano et al., 1997a, 1997b; Theraulaz & Bonabeau, 1995. Collective construction is typically viewed as an extension of the collective gathering task, in that it requires the agents to construct a particular structure, with gathered resources, at a home area of the environment. Specialization is typically required for building complex structures from many different types of component resources.

Reinforcement Learning for Specialization in Collective Construction

Murciano et al. (1997a) and Murciano et al. (1997b) applied *reinforcement learning* (RL) methods to a group of homogeneous agents operating in a discrete simulation environment. A collective gathering task mandated that individual agents derive specialized behavior in order to then derive an optimal collective behavior.

The authors used a RL method that independently modified action selection parameters within the controller of each agent. The RL method used either *global* or *local* RL signals so as to effectuate the learning of specialized behaviors. Behavioral specialization took the form of an agent learning to consistently select one action from a set of possible actions. The global RL signal measured group performance, and the local RL signal measured individual performance. The global RL signal was given at the end of a RL trial, where the signal was equal for all agents in the group. The local RL signal was given to individual agents, where the signal was calculated in terms of the agents own successes or failures. Murciano et al. (1997b) conducted experiments that tested the impact of local versus global RL signals upon the learning of specialized behaviors in a homogenous group of agents with no communication. The goal of these experiments was for agents to specialize via learning to gather specific object types so as to construct complex objects. Thus, when agents interacted an effective collective gathering and construction behavior emerged. Group task performance was measured as the number of complex objects assembled in a given RL trial. In the same experimental setup (Murciano et al., 1997a) conducted experiments that utilized only global RL signals for the purpose of facilitating emergent specialization within a homogeneous group of communicating agents. The task of individual agents and the group was to maximize the number of objects gathered over the course of a RL trial. The goal of experiments was for agents to specialize to different behaviors so as communication would facilitate the collective gathering of an optimal number of objects.

One criticism of the research of Murciano et al., (1997b), and Murciano et al. (1997a) derives from the use of RL signals in effectuating specialized behavior. Experimental results indicated that a global RL signal successfully motivated emergent specialization, given the assumption that all agents contribute equally to the task, and the signal was translated so as it could be meaningfully interpreted by each agent in a homogenous group. This casts doubt upon the applicability of global RL signals to heterogeneous groups. Likewise, the applicability of local RL signals was not tested in complex task domains that provided more realistic simulations of multi-robot systems. The possibility of applying the RL method to facilitate specialization in continuous simulation and physical task domains seems unlikely given the sparse reinforcement limitations of global RL signals and the noisy nature of local RL signals (Sutton et al., 1998) that inhibit learning. One aim of the research was to demonstrate that specialization emerges as a function of task constraints on the environment and agent group, irrespective of the type of reinforcement signal used. Achieving scalability in the learning of behavioral specialization is especially prevalent for tasks that require an increasing degree of heterogeneity, and complexity in collective behavior, as a response to dynamically emerging task challenges. However, the scalability of the RL method as a mechanism for encouraging behavioral specialization remains unclear since only two group sizes (10 and 30 agents), and a discrete environment of one size (54 x 54 grid cells) was tested. Also, the impact of more dynamic versions of the simulation environment upon the RL algorithm, were not tested. That is, only one redistribution of objects, during given RL trials, was tested.

Finally, the RL method assumed that the given task environment could be abstracted to the form of a multi-objective function which could be optimized. In this case the function was represented as a set of agent affinities that determined an agent's propensity to adopt particular behavioral roles. This severely limited the applicability of the RL method to more general and complex task environments.

Collective Resource Distribution and Allocation

In a series of research endeavors inspired by social insects (Bonabeau et al., 1996, 1997; Campos et al., 2001; Theraulaz et al., 1998a, 1998b), studied emergent specialization using response threshold methods in simulations of homogenous agent groups that were implemented within the context of mathematical frameworks.

Division of Labor for Dynamic Task Allocation

Theraulaz et al. (1998a) extended a previous formalization for the regulation of division of labor (Bonabeau et al., 1996) in simulated social insect colonies so as to include a *reinforcement learning* process. A formal variable response threshold method was implemented for purpose of facilitating emergent specialization in the form of division of labor. The authors highlighted similarities between their results and observations made within biological social systems where specialist workers were dynamically allocated based upon sub-task demand within a collective behavior task (O'Donnell, 1998).

Division of Labor for Dynamic Flow Shop Scheduling

Campos et al. (2001) introduced a division of labor method and applied it as a method for assigning resources within a dynamic flow shop scheduling task. The task entailed assigning trucks to paint booths in a factory, where trucks moved along an assembly line at a given pace. The color of a truck was predetermined by customer order. Three minutes was needed to paint a truck, but an additional three minutes was required if the color of a paint booth was to be changed for the truck. There was also a cost associated with paint changeover for a booth. A division of labor method was applied to minimize the number of such changeovers. Such paint fit-and-finish tasks are traditional bottleneck problems that can significantly reduce production throughput and thus require optimal solutions (Morley & Ekberg, 1998).

Division of Labor as a Function of Group Size

Gautrais et al., (2002) implemented a variable response threshold method to demonstrate that increasing agent group size and demand for tasks generated specialized agents. As with previous research (Bonabeau et al., 1996; Theraulaz et al., 1998a, 1998b), the response threshold method provided each agent in a group with an internal threshold for activating a particular behavior. Each agent's response threshold was influenced by the level of demand for a particular task, and agents allocated themselves so as to satisfy demand for these tasks. The authors' main conclusion was that their response threshold method demonstrated emergent specialization to be function of group size in the given resource allocation task, where group sizes exceeding a critical threshold value contained specialized agents, and group sizes below the critical threshold value contained only unspecialized agents. These findings were corroborated by similar findings in empirical theoretical biology studies (Robson et al., 1999).

Division of Labor Methods for Collective Resource Distribution and Allocation: Comments

Such response threshold methods represent a very simple, yet powerful, self-regulating feedback system that assigns the appropriate numbers of agents to different tasks. It is obvious that the study of such biologically inspired formalizations of specialization are worthy of future research attention given their applicability to a broad range of optimization tasks including dynamic scheduling and resource allocation. The methods of Bonabeau et al. (1997), Campos et al. (2001), Gautrais et al. (2002), and Theraulaz et al. (1998a) were prevalent in that they eloquently demonstrated how behavioral specialization emerged as a result of self-regulating task assignment and accomplishment, for which there exists a large amount of corroborating biological literature and empirical evidence (Chen, 1937a, 1937b; Deneubourg et al., 1987; O'Donnell, 1998; Robson et al., 1999; Theraulaz, Gervet, & Semenoff, 1991).

The main appeal of this set of research examples was their successful modeling of specialized behavior in the form a set of equations. These equations were successfully applied as a method for regulating the specialization of agents to specific tasks, in order to optimally accomplish a collective behavior task. However, in many cases the adaptive nature of response threshold regulation was never tested for more than one group or environment size, and more than two tasks. Also, the removal of specialized agents to test the adaptation process was limited to two agents. This was an important aspect of the adaptive nature of response thresholds, since if task allocation becomes too dynamic, or oscillatory, it is conceivable that the advantages of specialization could be lost as an agent spends all of its time switching between tasks, and consequently never dedicates enough time to accomplish a given task.

In each case, a simple set of experiments illustrated the importance and necessity of utilizing models of biological social behavior as a step towards understanding such social behavior, and then applying the underlying techniques, namely response thresholds, as a means of designing problem solving methods for optimization tasks. The main advantage of division of labor methods is their eloquence and simplicity of formal specification. Also, such methods yield results that are amenable to a mathematical or formal analysis. However, such methods are also limited to task domains that can be completely represented via the mechanics of a mathematical method. This makes the contributions of such methods limited to optimization tasks that can be formally represented, or to supporting empirical results evident in related biological literature.

Multi-Agent Computer Games

The application of biologically inspired methods to multi-agent computer games (Fogel, Hays, & Johnson, 2004; Laird & vanLent, 2000) has recently achieved particular success and gained popularity. For example, there has been particular research interest in the creation of adaptive interactive multi-agent first-person shooter games (Cole, Louis, & Miles, 2004; Hong & Cho, 2004; Stanley et al., 2005b), as well as strategy games (Bryant et al., 2003; Revello & McCartney, 2002; Yannakakis, Levine, & Hallam, 2004) using artificial evolution and learning as design methods for agent behavior. However, the study of specialized game playing behaviors, in teams of agents, has received relatively little research attention. Specialization is beneficial since it is often necessary for teams of agents to formulate collective behavior solutions in order to effectively challenge a human player, where an increasingly difficult level of agent performance is expected as game time progresses.

Legion-I: Neuro-Evolution for Adaptive Teams

Bryant et al. (2003) utilized the *Enforced Sub-Populations* (ESP) neuro-evolution method (Gomez, 1997) for the derivation of collective behavior in a multi-agent strategy game called *Legion-I*. The research hypothesis was that a team of homogeneous agents, where agents were capable of adopting different behavioral roles would be advantageous in terms of task performance, comparative to heterogeneous groups, composed of agents with static complementary behaviors. These experiments highlighted the effectiveness of the ESP method for deriving a dynamic form of emergent behavioral specialization motivated by division of labor. Results supported the hypothesis that for the Legion game, a homogenous team, where individuals could dynamically switch between specialized behaviors was effective. However, the analysis of emergent specialization was only at a behavioral level, so one could not readily ascertain the relationship between behavioral specialization and the evolved genotypes responsible for such behaviors. This would make an exploration of the mechanisms responsible for emergent specialization resulting from division of labor problematic. The task environment used a discrete simulation environment popular in multi-agent strategy games, but this was not sufficiently complex or dynamic in order to adequately test and support suppositions stating the advantages of behavioral specialization in homogenous teams. Also, the task performance of homogenous groups was not compared with heterogeneous groups. Valuable insight into the capabilities of homogenous versus heterogeneous agent groups for facilitating emergent specialization, could be gained by a comparison between groups represented by one neural controller, versus each agent within a group being represented by a different neural controller.

NERO: Neuro-Evolution of Augmenting Topologies

Stanley et al. (2005b), Stanley, Bryant, Karpov, and Miikkulainen, (2006), and Stanley, Bryant, and Miikkulainen (2005a) introduced a neuro-evolution method for the online evolution of neural controllers that operated in the context of an interactive multi-agent computer game called *Neuro-Evolving Robotic Operatives* (NERO). NERO is a first-person perspective shooter game, where a human player competes with teams of agents, and agents compete against each other. The rtNEAT neuro-evolution method was used for evolving increasing complex agent neural controllers using a process known as *complexification*. This was an extension of the *Neuro-Evolution of Augmenting Topologies* (NEAT) method (Stanley et al., 2002) that operated using online evolution. The authors demonstrated the effectiveness of the rtNEAT method for dynamically adapting agent controllers within a team playing against other agent teams or a human player in real time. Agent game playing behavior became increasingly sophisticated over successive generations as a result of changing neural network topological structure as well as evolving network connection weights. As an extension of the NEAT method, rtNEAT used online evolution to yield impressive results in terms of facilitating effectively competitive collective behaviors in the game playing time of NERO. The NEAT and rtNEAT methods successfully implemented a speciated representation of the genotype space, and a distance measure for genotype similarities, that provided a clear method for relating observed behaviors with a given set of genotypes.

However, the specialized controllers evolved were primarily determined by a training phase of NERO. Agent teams evolved specializations that were suitable for a given environment. Given that simulation environments were the same for both training and a subsequent battle phase, it remains unclear how suitable evolved teams would be for generalized collective behavior games. The true potential and beneficial nature of the rtNEAT method for evolving specialized behaviors in an online evolutionary process, for purpose of increasing team task performance, was not tested in other simulated multi-robot task domains. In realistic collective behavior tasks where the environment is dynamic and its structure and layout are not known *a priori*, training phases would only be partially effective since controllers trained in a simulation of the environment would simply be representing a best guess behavior. Currently, it remains unclear if rtNEAT could be successfully applied to collective behaviors tasks where there is a significant disparity between a training simulation and a subsequent *actual simulation* (called the battle phase in NERO). Such an issue is especially prevalent if online evolution of controllers is to eventually be applied for accomplishing multi-robot tasks, with time and energy constraints, in dynamic and complex physical task environments.

RoboCup Soccer

A distinct relation to multi-agent game research is RoboCup (Kitano & Asada, 2000). RoboCup is a research field dedicated to the design and development of multi-robot systems for the purpose of playing a robotic form of soccer. It is widely recognized as a specific test bed for machine learning algorithms, and engineering challenges (Noda & Stone, 2001). The very nature of the RoboCup game demands the existence of several types of behavioral specialization, in the form of different player roles. Such behaviors must be complementary and able to interact in such a way so as to produce a desired global behavior. That is, a team strategy that wins the game in a competitive scenario. Several researchers have focused on machine learning, evolutionary computation, and neuro-evolution methods that derive task accomplishing collective behaviors within groups of two or three soccer agents. Although, specialized behaviors of individual soccer agents was either specified *a priori* or was derived in simplistic game scenarios (Hsu & Gustafson, 2001, 2002; Luke et al., 1998; Matsubara, Noda, & Hiraki, 1996; Stone et al., 1998, 1998b, 2002; Whiteson et al., 2003. Each of these research examples has been critiqued elsewhere (Nitschke, 2005).

Pursuit-Evasion

Pursuit-evasion is a collective behavior task that is commonly used within artificial life research to test both non-adaptive (typically game theoretic) and adaptive (typically learning and evolution) methods for agent controller design. The task requires that multiple pursuer agents derive a collective behavior for the capture of one or more evading agents (Haynes & Sen, 1996a). The investigation of emergent specialization remains a relatively unexplored area of research in the pursuit-evasion domain (Luke et al., 1996), the collective herding variation (Potter et al., 2001), as well as more traditional predator-prey systems (Nishimura et al., 1997).

Evolving Pursuit-Evasion Behavior with Hexapod Robots

Blumenthal et al. (2004a, 2004b, 2004c) expanded previous work via combining a *punctuated anytime learning* (Blumenthal & Parker, 2006; Parker, 2000) method with an evolutionary algorithm within a co-evolution scenario. Although not the main research focus, this work addressed the issue of using morphological differences in agents in order to effectuate the derivation of behavioral specialization, and consequently a collective prey-capture behavior. The co-evolution scenario operated within a simulated multi-robot system of five hexapod robots where the goal was to derive an effective prey-capture behavior within four predator robots, and a predator-evasion behavior within one prey robot. This study effectively illustrated the derivation of prey-capture behavior based upon specialized behaviors that utilized differences in simulated hexapod robot morphology. Such as, the least maneuverable robot adopting a passive defensive position, whilst the fastest and most maneuverable robots adopted proactive pursuit behaviors.

However, the morphological differences between the robots were simple, leading one to speculate that a higher degree of complexity in specialized behavior may have emerged if differences in sensors and controller structure were included along with a greater disparity in actuator capabilities. Also, the prey was always initially placed at the center of the simulation environment, which made it easier for predators to form an effective prey capture behavior, and influenced the types of prey-capture behaviors that could emerge. Though not explicitly stated as a being a goal of this research, a valuable contribution to this research, would have been a methodological study that described a mapping or set of principles linking types of sensor and actuator capabilities to resulting forms of emergent behavioral specialization. Such a study could potentially form the basis of multi-robot system design methodologies that use evolution and learning mechanisms that capitalize on morphology in order to produce desired collective behaviors for solving a given task.

Evolving Herding Behavior in a Multi-Robot System

The research of Potter et al., (2001) investigated the evolution of homogeneous vs. heterogeneous controllers within a simulated multi-robot system that was given a collective herding task. A group of Nomad 200s were simulated within the *TeamBots* simulator (Balch, 1998). The research hypothesis was that as task difficulty increased, heterogeneity and specialization become essential for successful task accomplishment. Heterogeneity was defined as the number of different behaviors one robot could select from, as well as the number of behaviors in the group. This hypothesis was tested with experiments that introduced a predator into the environment. The goal was to encourage the emergence of specialized defensive behaviors in addition to herding behaviors. Experiments effectively illustrated that emergent behavioral specialization, for the benefit of collective behavior task performance, could be facilitated in a heterogeneous team of agents. Furthermore, results supported a hypothesis that constructing a collective behavior task such that multiple behaviors are required, increases the need for heterogeneity, and in turn specialization. However, the inducement of emergent specialization via increasing the number of behaviors required, and not simply task complexity, was only investigated within a single case study.

The key criticism lies in the comparison of homogenous and heterogeneous groups for deriving collective herding behaviors. Particularly, it is unclear why the authors opted to use only two genotype populations to represent a group of three shepherds in the heterogeneous design approach. The impact of homogeneity and heterogeneity on emergent specialization was not validated with larger groups of shepherds. Also, only one increment in the complexity of the task environment was tested. That is, the addition of the predator to the collective herding task. Complete validation of the authors' hypothesis that specialization emerges not as a consequence of task complexity, but rather as a result of the number of behaviors required to solve the task, would require several comparative case studies. Such studies would need to test tasks of varying degrees of difficulty versus tasks that require numerous complementary and potentially specialized behaviors. Such a comprehensive study would yield a valuable contribution to ones understanding of the relation between heterogeneous and homogenous design approaches, task performance, task complexity, and emergent specialization.

Moving in Formation and Cooperative Transportation Tasks

Certain collective behavior research endeavors, mainly in the fields of artificial life and multi-robot systems, have aimed to model and reproduce various forms of social phenomena that are observable in biological systems (Reynolds, 1987; Zaera, Cliff, & Bruten, 1996). Coordinated movement and cooperative transport is sometimes studied within the context of a gathering task, and has been studied separately in both physical and simulated environments. Cooperative transport is inspired by biological prey retrieval models, which present many examples of the value of specialization, such as the pushing vs. pulling behaviors exhibited in stigmatic coordination that allows several ants to transport a large prey (Kube et al., 1999). Such inspiration was used by the research of Dorigo et al. (2004) and Nolfi, Baldassarre, and Parisi (2003a), which described the evolution of coordinated motion, and self-assembly in a simulated multi-robot system for the purpose of cooperatively transporting objects. Similarly, the research of (Nolfi et al., 2003a) described the evolution of particular group formations in a simulated multi-robot system, which allowed efficient forms of coordinated group movement across an environment towards a light or sound source. The research of Baldassarre, Nolfi, and Parisi (2003), Dorigo et al. (2004), and Nolfi et al. (2003a) has been reviewed in related work (Nitschke, 2005), and is thus not described here.

Future Directions

Consequent of the literature reviewed, we deem the most viable future research direction to be the development of structured and principled emergent behavior design methodologies. From a broad range of methods that utilize emergent specialization for solving collective behavior tasks, a lack of a unifying set of design principles (methodologies) that link the workings of each of these methods, was highlighted. Such design methodologies would provide definitions and measures of specialization, and allow researchers to construct collective behavior systems that facilitate desired forms of emergent specialization that solve given tasks. If emergent specialization is to be utilized as a problem solver in systems that are designed using biologically inspired principles such as evolution and learning, then the concept of specialization must be defined, so as it can be identified and used in a problem solving process. In order to validate design methodologies that identify, measure, and harness emergent specialization as a problem solving tool in artificial complex adaptive systems, several considerations must be made.

1. Given the disparate and disjoint nature of biologically inspired and collective behavior research, validation of emergent specialization design methodologies would be experimental, and not necessarily constructed from a set of mathematical or otherwise theoretical suppositions that are proved.
2. Such methodologies would need to encapsulate the various types of specialization that benefit particular types of collective behavior tasks. These types would be identified through extensive experimentation.
3. Such methodologies would need to use specialization that can be identified and categorized, either dynamically by the design method, or *a priori* by a human designer. Importantly, dynamic identification of the type and degree of specialization required for a given task by a method would greatly increase the applicability of the method. That is, such a method would by applicable to complex task environments where specific task challenges dynamically arise in the environment and the exact nature of tasks cannot be described ahead of time.

Hence, if emergent specialized behavior is to be used as a means of deriving solutions to complex and dynamic task challenges in both simulated and physical collective behavior systems[8] then future research is obliged to look towards addressing the considerations delineated herein.

Conclusion

In drawing conclusions for this chapter, it is important to note that the chapter's goal was not to present an exhaustive list of research relating to emergent specialization, but rather to identify and present a set of pertinent research examples that use biologically inspired design approaches for the purpose of facilitating emergent behavioral specialization. Such research examples were selected based upon results that indicated emergent behavioral specialization as being beneficial for solving collective behavior tasks.

The binding theme of the chapter argued, that the majority of collective behavior research is currently analyzed and evaluated from empirical data gathered and emergent behavioral specialization observed, without analytical methods for identifying the means and causes of emergent specialization. An obvious reason for this is that the use of biologically inspired concepts such as evolution, self-organization, and learning as design methods is still in a phase of research infancy. Consequently, emergent specialization derived using such biologically inspired design concepts is currently constrained to simple forms. Given this general evaluation of prevalent literature, we identified several unresolved issues that inhibit the development of biologically inspired design methodologies that synthesize emergent specialization in solving collective behavior tasks.

1. It was evident that many researchers deem the simulation of collective behavior systems to be an effective approach for investigating emergent behavioral specialization, given that simulations provide a convenient means for studying the conditions under which specialization emerges. For example, the effects of parametric changes can be observed in a relatively short space of time. However, with notable exceptions, such as *SwarmBots* (Dorigo et al., 2004), the identification and transference of mechanisms motivating emergent specialization observed in simulation to counter-part algorithms operating in physical collective behavior systems such as multi-robot systems, is not yet plausible. In the case of *SwarmBots* (Dorigo et al., 2004), a simple task environment made the

transference to a physical environment possible, and emergent specialization was not necessarily a problem solver for dynamic challenges in the environment, but rather a solution to a given task that was emergent but not necessarily desired.

2. In the pertinent research examples reviewed, the complete potential of biologically inspired design, and the advantages of emergent specialization were not always effectively exploited. For example, many collective behavior systems, with notable exceptions such as division of labor methods applied to optimization tasks (Bonabeau et al., 1997), were simply attempting to synthesize emergent specialization, or to demonstrate the veracity of concepts such as self-organization, learning, and evolution for deriving novel agent behaviors. Such concepts were rarely applied to methods that derived emergent specialization as a means of increasing task performance or accomplishing unforeseen challenges in collective behavior tasks.
3. There is currently no standardized benchmark or research test-bed for testing, interpreting, evaluating, and classifying emergent specialized behavior. RoboCup was included as an honorable mention in the chapter, given that it provides an effective platform for testing and evaluating various forms of collective and individual behavior, emergent or otherwise, implemented either within an agent simulator or a physical multi-robot system. That is, collective behavior is simply evaluated within a competitive game scenario, so collective behavior performance is determined according to the evaluation criteria of the game. Another exception is collective gathering and dynamic scheduling in distributed systems, which can be represented as optimization tasks. In this case, standardized benchmarks exist in the form of performance results yielded by classical adaptive approaches. This makes the results of biologically inspired and classical methods to such tasks comparable. However, with exceptions such as Bonabeau et al. (1997) and Theraulaz and Bonabeau (1995) many optimization tasks do not benefit from the use of emergent behavioral specialization. Thus, the testing, interpretation, and evaluation of emergent specialized behavior within the context of collective behavior systems, is currently conducted according to the performance benchmarks of the researcher's own experimental simulation platform. This means that the experimental results can only be compared within the context of their own simulation environment. The development of emergent specialization design methodologies that could be equally ap-

plied to physical collective behavior systems would remove this critical constraint.

Given these open research issues, one may conclude that if the notion of emergent specialization as a problem solver for collective behavior tasks is to gain any maturity and credibility, then collective behavior systems must be built upon proven *emergent specialization design methodologies*. Ideally, such methodologies must be proven for convergence to desired forms of collective behavior (achieved as a consequence of emergent specialization), scalable and transferable to a counterpart situated and embodied collective behavior task environments.

References

Abdel-Rahman, H. (2001). When do cities specialize in production. *Reg Sci Urban Econ*, *26*(1), 1-22.

Arkin, R. (1998). *Behavior based robotics*. Cambridge, MA: MIT Press.

Arkin, R., & Balch, T. (1999). Behavior-based formation control for multi-robot teams. *IEEE Transactions on Robotics and Automation*, *14*(6), 926-939.

Axelrod, R. (1984). *The evolution of cooperation*. New York: Basic Books.

Baev, K. (1997). *Biological neural networks*. Berlin, Germany: Birkuser.

Balch, T. (2002b). Taxonomies of multi-robot task and reward. In T. Balch & E. Parker (Eds.), *Robot teams: From diversity to polymorphism* (pp. 23-35). Natick: A K Peters.

Balch, T. (2002a). Measuring robot group diversity. In T. Balch & E. Parker (Eds.), *Robot teams: From diversity to polymorphism* (pp. 93-135). Natick: A K Peters.

Balch, T. (1998). *Behavioral diversity in learning robot teams. PhD thesis.* Atlanta: College of Computing, Georgia Institute of Technology.

Baldassarre, G., Nolfi, S., & Parisi, D. (2003). Evolving mobile robots able to display collective behavior. *Artificial Life*, *9*(3), 255-267.

Beni, G. (2004). From swarm intelligence to swarm robotics. *Proceedings of the 1st International Workshop on Swarm Robotics* (p. 1-9). Santa Monica, CA: Springer.

Blumenthal, J., & Parker, G. (2006). Benchmarking punctuated anytime learning for evolving a multi-agent teams binary controllers. *World Automation Congress 2006.* Budapest, Hungary: IEEE Press.

Blumenthal, J., & Parker, G. (2004a). Co-evolving team capture strategies for dissimilar robots. *AAAI Artificial Multi-Agent Learning Symposium* (pp. 15-23). Arlington, VA: AAAI Press.

Blumenthal, J., & Parker, G. (2004b). Competing sample sizes for the co-evolution of heterogeneous agents. *Proceedings of the 2004 IEEE/RSJ International Conference on Intelligent Robots and Systems* (pp. 1438-1443). Sendai, Japan: IEEE Press.

Blumenthal, J., & Parker, G. (2004c). Punctuated anytime learning for evolving multi-agent capture strategies. *Proceedings of the Congress on Evolutionary Computation* (pp. 1820-1827). Portland, OR: IEEE Press.

Bonabeau, E., & Theraulaz, G. (1999). Role and variability of response thresholds in the regulation of division of labor in insect societies. In J. Deneubourg & J. Pasteels (Eds.), *Information processing in social insects* (p. 141-163). Basel, Switzerland: Springer Verlag.

Bonabeau, E., Dorigo, M., & Theraulaz, G. (1998). *Swarm intelligence: From natural to artificial systems.* Oxford, England: Oxford University Press.

Bonabeau, E., Sobkowski, A., Theraulaz, G., & Deneubourg, J. (1997). Adaptive task allocation inspired by a model of division of labor in social insects. In D. Lundh, B. Olsson, & A. Narayanan (Eds.), *Bio-computing and emergent computation* (pp. 36-45). Singapore: World Scientific.

Bonabeau, E., Theraulaz, G., & Deneubourg, J. (1996). Quantitative study of the fixed threshold model for the regulation of division of labor in insect societies. *Proceedings of the Royal Society of London B*, *263*(1), 1565-1569.

Bonabeau, E., Theraulaz, G., Arpin, E., & Sardet, E. (1994). The building behavior of lattice swarms. *Artificial Life IV: Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems* (pp. 307-312). Cambridge, MA: MIT Press.

Bongard, J. (2000). The legion system: A novel approach to evolving heterogeneity for collective problem solving. *Proceedings of Eurogp-2000* (pp. 16-28). Edinburgh, UK: Springer-Verlag.

Brooks, R. (1990). Elephants don't play chess. *Robotics and Autonomous Systems*, *6*(1), 3-15.

Bryant, B., & Miikkulainen, R. (2003). Neuro-evolution for adaptive teams. *Proceedings of the 2003 Congress on Evolutionary Computation* (pp. 2194-2201). Canberra, Australia: IEEE Press.

Calderone, N., & Page, R. (1988). Genotypic variability in age polyethism and task specialization in the honeybee. *Apis mellifera. Behav. Ecol. Sociobiol*, *22*(1), 17-25.

Campos, C., Theraulaz, G., Bonabeau, E., & Deneubourg, J. (2001). Dynamic scheduling and division of labor in social insects. *Adaptive Behavior*, *8*(2), 83-94.

Chen, C. (1937a). The leaders and followers among the ants in nest-building. *Physiol. Zool.*, *10*(1), 437-455.

Chen, C. (1937b). Social modification of the activity of ants in nest-building. *Physiol. Zool.*, *10*(1), 420-436.

Cole, N., Louis, S., & Miles, C. (2004). Using a genetic algorithm to tune first-person shooter bots. *Proceedings of the 2004 Congress on Evolutionary Computation* (Vol. 1, pp. 139-145). Piscataway: IEEE Press.

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. In J. Pasteels & J. Deneubourg (Eds.), *From individual to collective behavior in social insects* (pp. 177-196). Basel, Switzerland: Birkhauser.

Deneubourg, J., Theraulaz, G., & Beckers, R. (1991). Swarm-made architectures. *Proceedings of the European Conference on Artificial Life* (pp. 123-133). Amsterdam, Holland: Elsevier Academic Publishers.

Dorigo, M., Tuci, E., Gross, R., Trianni, V., Labella, H., Nouyan, S., et al. (2004). The swarm-bots project. *Proceedings of the Swarm Robotics: SAB 2004 International Workshop* (Vol. 3342, pp. 31-32). Santa Monica, CA: Springer Verlag.

Edmonds, J. (1965). Path, trees, and flowers. *Canadian J. Math.*, *17*(1), 449-467.

Fogel, D., Hays, T., & Johnson, D. (2004). A platform for evolving characters in competitive games. *Proceedings of 2004 congress on evolutionary computation* (pp. 1420-1426). Piscataway, USA: IEEE Press.

Funes, P., Orme, B., & Bonabeau, E. (2003). Evolving emergent group behaviors for simple humans agents. *Proceedings of the 7th European Conference on Artificial Life* (pp. 76-89). Berlin, Germany: Springer-Verlag.

Gautrais, J., Theraulaz, G., Deneubourg, J., & Anderson, C. (2002). Emergent polyethism as a consequence of increased colony size in insect societies. *Journal of Theoretical Biology*, *215*(1), 363-373.

Gomez, F. (1997). *Robust non-linear control through neuro-evolution. PhD thesis.* Austin, TX: Department of Computer Science, University of Austin, Texas.

Grefenstette, J. (1995). Credit assignment in rule discovery systems. *Machine Learning*, *3*(3), 225-246.

Hawthorne, D. (2001). Genetic linkage of ecological specialization and reproductive isolation in pea aphids. *Nature*, *412*(1), 904-907.

Haynes, T., & Sen, S. (1996a). Co-adaptation in a team. *International Journal of Computational Intelligence and Organizations*, *1*(4), 1-20.

Haynes, T., & Sen, S. (1996b). Evolving behavioral strategies in predators and prey. In G. Weiss & S. Sen (Eds.), *Adaptation and learning in multi-agent systems: Lecture notes in computer science* (pp. 113-126). Berlin, Germany: Springer-Verlag.

Hong, J., & Cho, S. (2004). Evolution of emergent behaviors for shooting game characters in RoboCode. *Proceedings of the 2004 Congress on Evolutionary Computation* (Vol. 1, pp. 634-638). Piscataway, USA: IEEE Press.

Hsu, W., & Gustafson, S. (2001). Layered learning in genetic programming for a cooperative robot soccer problem. *Proceedings of the 4th European Conference on Genetic Programming* (pp. 291-301). Como, Italy: Springer-Verlag.

Hsu,W., & Gustafson, S. (2002). Genetic programming and multi-agent layered learning by reinforcements. *Proceedings of the Genetic and Evolutionary Computation Conference* (pp. 764-771). New York: Morgan Kaufmann.

Ijspeert, A., Martinoli, A., Billard, A., & Gambardella, L. (2001). Collaboration through the exploitation of local interactions in autonomous collective robotics: The stick pulling experiment. *Autonomous Robots*, *11*(2), 149-171.

Kitano, H., & Asada, M. (2000). The RoboCup humanoid challenge as the millennium challenge for advanced robotics. *Advanced Robotics*, *13*(1), 723-736.

Komann, M., & Fey, D. (2007). Realizing emergent image pre-processing tasks in cellular-automaton-alike massively parallel hardware. *International Journal on Parallel, Emergent and Distributed Systems*, *22*(1), 79-89.

Kreiger, M., & Billeter, J. (2000). The call of duty: Self-organized task allocation in a population of up to twelve mobile robots. *Robotics and Autonomous Systems*, *30*, 65-84.

Kube, C., & Bonabeau, E. (1999). Cooperative transport by ants and robots. *Robotics and Autonomous Systems: Special Issue on Biomimetic Robots*, *1*(1), 20-29.

Laird, J., & vanLent, M. (2000). Human-level AI killer application: Interactive computer games. *Proceedings of the 17th National Conference on Artificial Intelligence and the 12th Annual Conference on Innovative Applications of Artificial Intelligence* (pp. 1171-1178). Menlo Park, USA: AAAI Press.

Li, L., Martinoli, A., & Mostafa, Y. (2002). Emergent specialization in swarm systems. In H. Yin, N. Allinson, R. Freeman, J. Keane, & S. Hubbard (Eds.), *Lecture notes in computer science: Vol. 2412. Intelligent data engineering and automated learning* (pp. 261-266). Berlin, Germany: Springer-Verlag.

Li, L., Martinoli, A., & Yaser, A. (2004). Learning and measuring specialization in collaborative swarm systems. *Adaptive Behavior*, *12*(3), 199-212.

Lipson, H., & Pollack, J. (2000). Automatic design and manufacture of robotic life forms. *Nature*, *406*(1), 974-978.

Luke, S., & Spector, L. (1996). Evolving teamwork and coordination with genetic programming. *Proceedings of the 1996 International Conference on Genetic Programming* (pp. 150-156). Stanford, USA: MIT Press.

Luke, S., Farris, C. H. 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, Germany: Springer-Verlag.

Martinoli, A., Zhang, Y., Prakash, P., Antonsson, E., & Olney, R. (2002). Towards evolutionary design of intelligent transportation systems. The *11th International Symposium on New Technologies for Advanced Driver Assistance Systems* (pp. 283-290). Siena, Italy: ATA Press.

Mataric, M. (1997). Behavior-based control: Examples from navigation, learning, and group behavior. *Journal of Experimental and Theoretical Artificial Intelligence*, *9*(1), 62-78.

Matsubara, H., Noda, I., & Hiraki, K. (1996). Learning of cooperative actions in multi-agent systems: A case study of pass and play in soccer. *Adaptation, Co-Evolution, and Learning in Multi-Agent Systems: Papers from the 1996 AAAI Spring Symposium* (pp. 63-67). Boston: AAAI Press.

Morley, R., & Ekberg, G. (1998). Cases in chaos: Complexity-based approaches to manufacturing. In *Embracing complexity: A colloquium on the application of complex adaptive systems to business.* (p. 97-702). Cambridge, USA: The Ernst and Young Center for Business for Business Innovation.

Müller, C., & Sick, B. (2006). Emergence in organic computing systems: Discussion of a controversial concept. *Proceedings of the 3rd International Conference on Autonomic and Trusted Computing, Volume 4158 Of Lecture Notes In Computer Science* (pp. 1-16). Berlin, Germany: Springer.

Murciano, A., & Millan, J. (1997). Learning signaling behaviors and specialization in cooperative agents. *Adaptive Behavior*, *5*(1), 5-28.

Murciano, A., Millan, J., & Zamora, J. (1997). Specialization in multi-agent systems through learning. *Biological Cybernetics*, *76*(1), 375-382.

Ng, Y., & Yang, X. (1997). Specialization, information, and growth: A sequential equilibrium analysis. *Rev Dev Econ*, *1*(1), 257-274.

Nishimura, S., & Takashi, I. (1997). Emergence of collective strategies in a predator-prey game model. *Artificial Life*, *3*(1), 243-260.

Nitschke, G. (2005). Emergence of cooperation: State of the art. *Artificial Life*, *11*(3), 367-396.

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.

Noirot, C., & Pasteels, J. (1987). Ontogenetic development and the evolution of the worker caste in termites. *Experientia.*, *43*(1), 851-860.

Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines.* Cambridge, MA: MIT Press.

Nolfi, S., & Parisi, D. (1997). Learning to adapt to changing environments in evolving neural networks. *Adaptive Behavior*, *1*(5), 75-98.

Nolfi, S., Baldassarre, G., & Parisi, D. (2003a). Evolution of collective behavior in a team of physically linked robots. In G. Guillot & J. Meyer (Eds.), *Applications of evolutionary computing* (pp. 581-592). Heidelberg, Germany: Springer Verlag.

Nolfi, S., Deneubourg, J., Floreano, D., Gambardella, L., Mondada, F., & Dorigo, M. (2003b). Swarm-bots: Swarm of mobile robots able to self-assemble and self-organize. *Ecrim News*, *53*(1), 25-26.

O'Donnell, S. (1998). Effects of experimental forager removals on division of labour in the primitively eusocial wasp polistes instabilis. *Behaviour*, *135*(2), 173-193.

O'Riain, M., Jarvis, J., Alexander, R., Buffenstein, R., & Peeters, C. (2000). Morphological castes in a vertebrate. *Proceedings of the National Academy of Sciences of the United States of America*, *97*(24), 13194-13197.

Parker, G. (2000). Co-evolving model parameters for anytime learning in evolutionary robotics. *Robotics and Autonomous Systems*, *33*(1), 13-30.

Parker, G., & Nathan, P. (2006). Evolving sensor morphology on a legged robot in niche environments. *World Automation Congress 2006.* Budapest, Hungary: IEEE Press.

Perez-Uribe, A., Floreano, D., & Keller, L. (2003). Effects of group composition and level of selection in the evolution of cooperation in artificial ants. *Advances of Artificial Life: Proceedings of the 7th European Conference on Artificial Life* (pp. 128-137). Dortmund, Germany: Springer.

Pfeifer, R., Iida, F., & Gomez, G. (2006). Designing intelligent robots: On the implications of embodiment. *Journal of Robotics Society of Japan*, *24*(7), 9-16.

Potter, M., & DeJong, K. (2000). Cooperative co-evolution: An architecture for evolving co-adapted subcomponents. *Evolutionary Computation*, *8*(1), 1-29.

Potter, M., Meeden, L., & Schultz, A. (2001). Heterogeneity in the coevolved behaviors of mobile robots: The emergence of specialists. *Proceedings of the International Joint Conference on Artificial Intelligence* (pp. 1337-1343). Seattle, WA: AAAI Press.

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.

Resnick, M. (1997). *Turtles, termites, and traffic jams: Explorations in massively parallel micro worlds*. Cambridge, MA: MIT Press.

Revello, T., & McCartney, R. (2002). Generating war game strategies using a genetic algorithm. *Proceedings of the 2002 Congress on Evolutionary Computation* (pp. 1086-1091). Piscataway, USA: IEEE Press.

Reynolds, C. (1987). Flocks, herds, and schools: A distributed behavioral model. *Computer Graphics*, *21*(4), 25-36.

Robson, S., & Traniello, J. (1999). Key individuals and the organization of labor in ants. In J. Deneubourg & J. Pasteels (Eds.), *Information processing in social insects* (pp. 239-259). Basel, Switzerland: Springer Verlag.

Schultz, A., & Bugajska, M. (2000). Co-evolution of form and function in the design of autonomous agents: Micro air vehicles project. *Proceedings of the Workshop on Evolution of Sensors in Nature, Hardware, and Simulation* (pp. 154-166). Chicago, IL: AAAI Press.

Seligmann, H. (1999). Resource partition history and evolutionary specialization of subunits in complex systems. *Biosystems*, *51*(1), 31-39.

Sims, K. (2004). Evolving 3d morphology and behavior by competition. *Artificial Life IV: Proceedings of the 4th International Workshop on the Synthesis and Simulation of Living Systems* (pp. 28-39). Cambridge, MA: MIT Press.

Smith, A. (1904). *An inquiry into the nature and causes of the wealth of nations* (5th ed.), (E. Cannan, Ed.). London, United Kingdom: Methuen and Co., Ltd. First published: 1776.

Solow, D., & Szmerekovsky, J. (2004). Mathematical models for explaining the emergence of specialization in performing tasks. *Complexity*, *10*(1), 37-48.

Stanley, K., & Miikkulainen, R. (2002). Evolving neural networks through augmenting topologies. *Evolutionary Computation*, *10*(2), 99-127.

Stanley, K., Bryant, B., & Miikkulainen, R. (2005a). Evolving neural network agents in the NERO video game. *Proceedings of the IEEE 2005 Symposium on Computational Intelligence And Games* (pp. 182-189). Piscataway, USA: IEEE Press.

Stanley, K., Bryant, B., & Miikkulainen, R. (2005b). Real-time neuro-evolution in the NERO video game. *Evolutionary Computation*, *9*(6), 653-668.

Stanley, K., Bryant, B., Karpov, I., & Miikkulainen, R. (2006). Real-time evolution of neural networks in the NERO video game. *Proceedings of the* 21st *National Conference on Artificial Intelligence* (pp. 1671-1674). Boston: AAAI Press.

Stone, P., & Veloso, M. (1998a). A layered approach to learning client behaviors in the RoboCup soccer server. *Applied Artificial Intelligence*, *12*(1), 165-188.

Stone, P., & Veloso, M. (1998b). Using decision tree confidence factors for multi-agent control. *Proceedings of the 2nd International Conference on Autonomous Agents* (pp. 110-116). Minneapolis, MN: ACM Press.

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.

Stone, P., & Veloso, M. (2002). Towards collaborative and adversarial learning: A case study in robotic soccer. *Evolution and learning in multi-agent systems*, *48*(1), 83-104.

Sutton, R., & Barto, A. (1998). *An introduction to reinforcement learning.* Cambridge, MA: John Wiley and Sons.

Tarapore, D., Floreano, D., & Keller, L. (2006). Influence of the level of polyandry and genetic architecture on division of labor. *The 10th International Conference on the Simulation and Synthesis of Living Systems (ALIFE X)* (pp. 358-364). Cambridge, USA: MIT Press.

Theraulaz, G., & Bonabeau, E. (1995). Coordination in distributed building. *Science*, *269*(1), 686-688.

Theraulaz, G., Bonabeau, E., & Deneubourg, J. (1998a). Fixed response thresholds and the regulation of division of labor in insect societies. *Bulletin of Mathematical Biology*, *60*(1), 753-807.

Theraulaz, G., Bonabeau, E., & Deneubourg, J. (1998b). Response threshold reinforcement and division of labor in insect societies. *Proceedings of the Royal Society of London B*, *265*(1), 327-332.

Theraulaz, G., Gervet, J., & Semenoff, S. (1991). Social regulation of foraging activities in polistes dominulus christ: A systemic approach to behavioural organization. *Behaviour*, *116*(1), 292-320.

Theraulaz, G., Goss, S., Gervet, J., & Deneubourg, J. (1991). Task differentiation in polistes wasp colonies: A model for self-organizing groups of robots. *The 1st International Conference on the Simulation of Adaptive Behavior* (pp. 346-355). Cambridge, USA: MIT Press.

Thompson, A., Harvey, I., & Husbands, P. (1996). Unconstrained evolution and hard consequences. *Towards Evolvable Hardware: The Evolutionary Engineering Approach* (Vol. 1062 LNCS, pp. 135-165). Berlin, Germany: Springer-Verlag.

Trumler, W., Klaus, R., & Ungerer, T. (2006). Self-configuration via cooperative social behavior. *Proceedings of the 3rd International Conference on Autonomic and Trusted Computing* (pp. 90-99). Wuhan, China: Springer.

Waibel, M., Floreano, D., Magnenat, S., & Keller, L. (2006). Division of labor and colony efficiency in social insects: Effects of interactions between genetic architecture, colony kin structure, and rate of perturbations. *Proceedings of the Royal Society B*, *273*(1), 1815-1823.

Watson, R., Ficici, S., & Pollack, J. (1999a). Embodied evolution: A response to challenges in evolutionary robotics. In J. Wyatt & J. Demiris (Eds.), The 8th *European Workshop on Learning Robots* (pp. 14-22). Lausanne, Switzerland: Springer Verlag.

Watson, R., Ficici, S., & Pollack, J. (1999b). Embodied evolution: Embodying an evolutionary algorithm in a population of robots. *1999 Congress on Evolutionary Computation* (pp. 335-342). Washington, DC, USA: IEEE Press.

Watson, R., Ficici, S., & Pollack, J. (2002). Embodied evolution: Distributing an evolutionary algorithm in a population of robots. *Robotics and Autonomous Systems*, *39*(1), 1-18.

Weis, T., Parzyjegla, H., Jaeger, M., & Mühl, G. (2006). Self-organizing and self-stabilizing role assignment in sensor/actuator networks. *Proceedings* *of the 8th International Symposium on Distributed Objects and Applications.* (pp. 1807-1824), Montpellier, France: Springer.

Wenseleers, T., Ratnieks, F., & Billen, J. (2003). Caste fate conflict in swarm-founding social hymenoptera: An inclusive fitness analysis. *Evolutionary Biology*, *16*(1), 647-658.

Whiteson, S., Kohl, N., Miikkulainen, R., & Stone, P. (2003). Evolving keep-away soccer players through task decomposition. *Proceeding of the Genetic and Evolutionary Computation Conference* (pp. 356-368). Chicago, IL: AAAI Press.

Wu, J., Di, Z., & Yang, Z. (2003). Division of labor as the result of phase transition. *Physica A*, *7*(1), 323-663.

Yannakakis, G., Levine, J., & Hallam, J. (2004). An evolutionary approach for interactive computer games. *Proceedings of the 2004 Congress on Evolutionary Computation* (pp. 986-993). Piscataway, USA: IEEE Press.

Zaera, N., Cliff, D., & Bruten, J. (1996). *(Not) Evolving collective behaviors in synthetic fish (tech. rep.)*. Bristol, England: Hewlett-Packard Laboratories.

Zhang, Y., Martinoli, A., & Antonsson, E. (2003). Evolutionary design of a collective sensory system. *The 2003 AAAI Spring Symposium on Computational Synthesis* (pp. 283-290). Stanford, USA: AAAI Press.

Endnotes

1 Examples of complex adaptive systems include social insect colonies, biological neural networks, traffic jams, economies of a nation, as well as industrial infrastructures such as energy and telecommunications networks (Resnick, 1997). We deem complex adaptive systems to be a subset of complex systems where autonomous software (simulated) or physically embodied (robots) agents operate in order to solve a given task.

2 The terms *task*, *activity*, *role*, and *caste* are defined as follows. Task: what has to be done; Activity: what is being done; Role: the task assigned to an individual within a set of responsibilities given to a group of individuals; Caste: a group of individuals specialized in the same role (Kreiger et al., 2000).

3 The terms *collective behavior system* and *artificial complex adaptive system* are used interchangeably throughout the chapter. Both refer to distributed

systems where specialization emerges as a property of a collective behavior dynamics.

4 We distinguish methodologies from methods. We assume the latter to be the actual algorithm, which is implemented for the purpose of solving a specific task. Where as, we assume the former to be a set of design principles for designing methods.

5 Various definitions for numerous types of specialization have been proposed across a broad range of disciplines. In *The Wealth of Nations*, (Smith, 1904) Adam Smith described economic specialization in terms of division of labor. Specifically stating that in industrialism, division of labor represents a qualitative increase in productivity, and regarded its emergence as the result of a dynamic engine of economic progress. Smith viewed specialization by workers as leading to greater skill and greater productivity for given tasks, which could not be achieved by non-specialized workers attempting to accomplish those same tasks.

6 The allocation of agent labor within a group of agents is analogous to resource allocation which derives from economic and game theory studies (Axelrod, 1984). Such studies attempt to derive methods that efficiently allocate a limited amount of resources so as to accomplish a given task with the highest degree of performance possible.

7 Heterogeneity, and hence behavioral diversity, was defined as the number of castes in the group, and specialization was the part of diversity that was required to increase task performance.

8 Such a case has been envisioned for swarm robotic systems (Beni, 2004).

Section IV

Social Science Perspectives