Chapter 6

Conclusions and Future Work

In this chapter, we summarize the results and contributions presented in this dissertation. We also offer some ideas for future work that we believe can contribute to the further development of the swarm intelligence field.

6.1 Swarm Intelligence Systems and Interference

Swarm intelligence is the problem-solving behavior of large groups of simple entities capable of autonomous perception and action (called agents) that collectively are referred to as a *swarm*. The term swarm intelligence evokes a mental image in which a large group of insect-like entities congregates and exhibits a purposeful behavior without a central authority to supervise the actions of each individual or issue commands to govern the group's behavior. Despite its name, which makes us recall science fiction works, swarm intelligence exists in nature. Bees form swarms to collectively find and choose the best location to build a new home. Ant colonies, which can be composed of millions of ants, build complex nests, search and retrieve food, maintain the young, etc. In each case, a swarm intelligence system performs a particular task without any single individual supervising or directing the actions of other members of the swarm. Swarm intelligence can also be the product of engineering efforts. Powerful optimization techniques and control mechanisms for groups of mobile robots have been designed exploiting swarm intelligence principles.

Artificial swarm intelligence systems are composed of numerous agents that interact locally with one another and with their environment. Through different mechanisms, but predominantly through self-organization and decentralized control, these kinds of systems exhibit a collective intelligence that allows them to solve problems that their constituent agents cannot solve individually. As in any system whose constituent agents interact with each other, there are interactions among the agents that form a swarm that reduce the efficiency of the system. These interactions are collectively referred to as *interference*. One of the most visible effects of interference in a swarm intelligence system is the reduction of the system's efficiency; that is, the time required by the system to reach a desired state is increased. Interference increases with the size of the population of agents. Thus, interference is a major problem in swarm intelligence systems since many of them require large populations to perform their tasks satisfactorily. Interference is thus a fundamental problem inherent to systems composed of many agents because it negatively affects the viability of the swarm intelligence approach when solving important practical problems.

6.2 Incremental Social Learning as a Mechanism for Reducing the Effects of Interference

In this dissertation, an original framework called incremental social learning (ISL) was proposed in Chapter 3. This framework aims to reduce the negative effects of interference in swarm intelligence systems. Two components form the core of the ISL framework. The first component directly manipulates one of the factors that causes interference: the number of agents that compose a swarm. A swarm intelligence system under the control of ISL starts with a small population. Gradually, the population grows until the system performs as desired or a maximum number of agents is reached. The second component of ISL is social learning. ISL exploits the fact that learning socially is less costly, in terms of trial-and-error trials for an individual, than asocial learning. Through social learning, newly added agents acquire knowledge from agents that have been part of the swarm for some time. As a result of the combination of these two components, a growing population size, and social learning, the effects of interference is reduced. Consequently, the swarm reaches a desired state more rapidly than without using ISL.

6.2.1 Incremental Social Learning in Particle Swarms

We demonstrated the effectiveness of ISL through two case studies. In the first case study, presented in Chapter 4, we applied ISL to particle swarm optimization (PSO) algorithms. These algorithms are commonly used to tackle continuous optimization problems and are composed of a population of searching agents called *particles*. PSO algorithms with a constant population size exhibit a trade-off between solution quality and number of objective function evaluations amenable to the application of ISL. With a small population size, the solution quality improves rapidly during the first objective function evaluations until it reaches a stable value. With large populations, the same solution quality reached by a small population is reached after many more objective function evaluations. However, if more evaluations are allowed, a better solution quality may be reached. The hypothesis that supports the application of ISL to PSO algorithms is that the trade-off between solution quality and number of objective function evaluations is due, at least partially, to interference among particles. Interference in PSO algorithms is the result of particles being attracted toward the best solutions found by other particles. Interference is large in big swarms because at the beginning of the optimization process too much information flows through the network of particles. This phenomenon makes particles spend objective function evaluations in regions that do not contain the optimal solution.

As a result of the application of ISL to PSO algorithms, three new PSO variants were designed. The first one, which serves as a basis for the other two, is an incremental particle swarm optimization algorithm that we call IPSO. In IPSO, the population of particles grows over time until the optimization process returns a solution of acceptable quality or until a maximum population size is reached. The rate at which particles are added to the system is scheduled and controlled through a parameter. Each time a new particle is added, its position in the objective function's domain (usually a subset of $\mathbb{R}^n$) is generated through a rule that biases the placement of the new particle toward the best-so-far solution. Through a thorough experimental evaluation, we could show how IPSO, with the appropriate setting of the population growth, could return solutions that are comparable to those that would be returned if multiple PSO algorithms with different constant population sizes were run in parallel and only the best solution found by any of those algorithms was returned. The other two algorithms that result from the use of ISL on PSO algorithms, called IPSOLS and IPSOLS+, repeatedly call a local search procedure from a particle's best found position in order to intensify the search. Each call of the local search procedure could be seen as simulating the individual learning of a particle. IPSOLS works in the same way as IPSO with an added step that consists in calling a local search procedure from each particle's best found position. IPSOLS+ is a further refinement of IPSOLS in which the local search is called more frequently from the particle's position that represents the best-so-far solution and in which the PSO rules are modified. IPSOLS's performance is comparable with state-of-the-art PSO algorithms. IPSOLS+'s performance is comparable with state-of-the-art algorithms for large-scale continuous optimization problems.

6.2.2 Incremental Social Learning in Robot Swarms

In the second case study, presented in Chapter 5, we applied ISL to a collective decision-making mechanism for swarms of mobile robots. Two contributions are presented in that chapter. First, the collective decision-making itself, and second, the application of ISL to that mechanism. For this case study, we chose a foraging task that involves group transport in an arena that consists of two locations connected by two paths. During the execution of the task, which is to transport objects from one location to the other, a swarm of robots must choose one of the two paths. In one of the two locations, robots form teams to transport the objects, which cannot be transported by single robots. From that location, teams of robots transport objects to the other location. Robots have a preferred branch (encoded as an "opinion") and whenever they form a team, they advocate for their preferred path. The final team's decision is that of the local majority. Robots not only choose that path, but they also change their preference if it is different from the one they had before forming a team. After making a decision, a team moves from one location to the other using the chosen path. Once a team arrives at the target location, it disassembles and its component robots return as individuals using again the chosen path. Once they arrive at the initial location, robots can form new teams, repeating the process until the task is performed. The length of the paths induce a *latency* period during which robots can neither change opinion nor influence other robots. Thus, each opinion has a latency period whose duration depends on the length of the paths and on the number of robots in the environment. We showed through Monte Carlo and physics-based simulations that the dynamics of the system makes a swarm of robots reach a consensus. If the initial distribution of opinions in the swarm is such that half of the swarm prefers one opinion and the other half prefers the other opinion, the proposed collective decision-making mechanism makes the swarm reach consensus with high probability on the opinion associated with the shortest latency period. In the robotics setting described in Chapter 5, this means that a swarm reaches consensus on the opinion associated with the shortest path.

The aforementioned swarm robotics system shows a trade-off between performance and population size similar to the one observed in PSO algorithms. In this case, however, it is the population of "idle" robots, that is, those robots that are not engaged in the transportation task, that affects the system's performance. Our implementation of ISL manipulates this population. We start the process with only six robots (two teams). At each time step, we add a robot and let it copy the opinion of a randomly picked "idle" robot. If there are no robots to copy from, the opinion of the new robot is initialized at random. Because of the dynamics of the system, it is more likely for a new robot to copy the opinion associated with the shortest path. As a result, the population reaches a consensus on the opinion associated with the shortest path in fewer time steps than it would without ISL. The effectiveness of ISL, however, depends on the number of active robots in the environment. With more active teams, there are fewer "idle" robots, and thus, the effects of ISL diminish to the point at which there is practically no difference between the system that is using ISL and the system that is no using ISL.

6.2.3 Impact

One of the major challenges in swarm intelligence research is to design agent-level behaviors in order to obtain a certain desired behavior at the collective-level. Since a general methodology for achieving this goal has been elusive, most researchers in the field concentrate their efforts on specific applications. In doing so, a number of assumptions are made. One of these assumptions is that the size of a swarm of agents remains constant over time. In many cases, this assumption may not be well justified.

The framework proposed in this dissertation challenges the constant population size assumption. In the ISL framework, the population size changes over time and we have demonstrated that some benefits can be obtained with such an approach. As seen in Chapter 5, we are not the only ones to realize that an incremental deployment of agents (robots) can bring benefits and can even simplify the design of the agent-level behaviors. In fact, in many practical applications of swarm intelligence systems, in particular in swarm robotics and affine fields, such as sensor networks, it is actually more difficult to deploy hundreds of robots at once, than to deploy a few robots at different points in time. For example, consider the deployment of the 30 Galileo satellites (European Space Agency, 2010). It is not reasonable to assume that tens of satellites can be deployed at once. Rather, the deployment is painfully slow, with one or two satellites being deployed at a time. If these satellites were part of a swarm of satellites with specific tasks such as maintaining a formation in space, the rules needed to fulfill that task would be quite complex. Instead, with an incremental deployment, each satellite could take its position without disturbing the behavior of other satellites. In other words, the interference between satellites would be greatly reduced.

With our proposal, we hope that researchers in the swarm intelligence field will consider the possibility of an incremental deployment of agents in the design of new swarm intelligence systems.

The other aspect of our proposal, the use of some form of social learning, can potentially have a bigger impact in the field. Social learning can be the mechanism that enables the appearance of a form of cumulative "culture" in a swarm that passes from one "generation" of agents to another. A continuous process of addition and elimination of agents can make this process possible as long as the knowledge acquired during the lifetime of one agent is not lost, but is instead transmitted to a new agent. This new agent in turn would have time to accumulate more knowledge to pass on to another agent, and so on. Perhaps the biggest impact of this idea will be in the field of swarm robotics, in which each robot has a lifetime determined by the capacity of its batteries. Before running out of power, a robot could pass on its knowledge to another fully charged robot, which will have more time to refine and accumulate more information.

6.3 Future Work

We believe that the work presented in this dissertation opens a number of potentially fruitful research avenues. In the remainder of this section, we will briefly describe some of them. Our presentation is divided in two parts. In the first, we describe future work that is directly related to the ISL framework. In the second part, future work derived from the two case studies presented in this dissertation is proposed.

6.3.1 Future Work Related to the Incremental Social Learning Framework

Theory

Interference has been identified by some authors, notably Matarić (1997); Helbing and Vicsek (1999) and Gershenson (2007), as an influence that we need to control in order to be able to design large multiagent and self-organizing systems. Unfortunately, very little theoretical work that could help us understand how to do that has been performed. Future work in this area, we believe, could significantly impact swarm intelligence, self-organizing systems, complex systems, and other related fields.

Throughout this dissertation, we have given empirical evidence of the effectiveness of the ISL framework. However, we have not determined analytically the conditions under which the ISL framework is guaranteed to reduce interference. Future work should be directed toward achieving this goal as this would increase the impact of the proposed approach.

More Applications

The performance of the optimization algorithms presented in Chapter 4 suggests that the ISL framework can improve the performance of other swarm intelligence-based optimization algorithms. In fact, in a recent paper, we explored the application of the ISL framework to an ant colony optimization algorithm for continuous optimization problems and obtained promising results (Liao et al., 2011). Another related and potentially fruitful research direction is the application of the ISL framework to evolutionary algorithms such as differential evolution (Storn and Price, 1997) or CMA-ES (Hansen et al., 1995). However, it should be noted that these algorithms' search dynamics are different from the search dynamics of swarm intelligence algorithms. Thus, even though it is straightforward to apply ISL to these algorithms, and that the two classes of algorithms share some common features, such as a population of candidate solutions, the results of the application of ISL to evolutionary algorithms may be different from the results obtained with swarm intelligence algorithms.

In swarm robotics, more studies about the possible use and benefits of using the ISL framework should be undertaken. In particular, it would be interesting to follow and build on the work of Winfield and Griffiths (2010) who are investigating how a "robotic culture" could emerge. The ISL framework could play the role of a knowledge transfer facilitator between "generations" of robots in those settings.

6.3.2 Future Work Related to the Case Studies

Tuning-in-the-loop Design of Optimization Algorithms

In Chapter 4, we described the redesign process of IPSOLS that led to IPSOLS+. This process relied on a parameter tuning tool, iterated F-Race, as a way to measure the impact of each important design decision. The result of this process was a highly competitive algorithm in the field of large-scale continuous optimization. We believe that a methodology that integrates parameter tuning tools as part of the optimization algorithm design process can have an important role in the emerging field of engineering stochastic local search algorithms (Stützle et al., 2007, 2009).

Collective Decision-Making Mechanisms based on Opinion Formation Models

The majority-rule opinion formation model which is at the basis of the collective decision-making mechanism introduced in Chapter 5 is only one of a large number of opinion-formation models that have been proposed in the statistical physics literature (Castellano et al., 2009). Considering the promising results that we were able to obtain, we believe that the swarm intelligence field could greatly benefit if more researchers consider using similar methods to address scenarios in which agents must choose among multiple choices. Also of interest is the study of the resulting systems' dynamics in changing environments. Domains in which the agents agree on a continuous quantity instead of on a discrete one should also be explored.

6.4 Concluding Statement

In this dissertation, we have introduced the incremental social learning framework. Its design is aimed at reducing interference in systems composed of many interacting agents. To show its potential, we instantiated the framework in the context of particle swarm optimization algorithms and swarm robotics. The results obtained represent evidence that the framework indeed reduces interference, which in turn makes the systems have a better performance.

We hope that these results motivate other researchers interested in multiagent systems, swarm intelligence, and other affine fields, to integrate the incremental social framework into a set of agent deployment strategies. Such a set of strategies can indeed simplify the agent design process because, as we demonstrated, by reducing the levels of interference, it is possible to simplify the rules that govern agent interactions.