

Human Factors issues for Interaction with Bio-Inspired Swarms

Michael Lewis*, Michael Goodrich**, Katia Sycara⁺, Mark Steinberg⁺⁺

* School of Information Science, University of Pittsburgh, Pittsburgh, PA. 15260

** Brigham Young University, 3361 TMCB, Provo, Utah, 84602

+ Robotics Institute, Carnegie Mellon University, 5000 Forbes Av., Pittsburgh, PA 15213

++ Office of Naval Research, 875 N Randolph St., Arlington, VA 22203

Robotic systems composed of a large number of entities, often called robot swarms, are envisioned to play an increasingly important role in applications such as search, rescue, surveillance, and reconnaissance operations. Nowadays, many mobile robots that are deployed for such applications are still tele-operated by a single or multiple operators. While these platforms are individually very capable, the development of cheaper hardware allows the consideration of swarm systems composed of many more robots but with each individual being far less powerful. Examples from biology indicate that such systems can be collectively more powerful than any individual robot within the team and also more than many larger, more sophisticated individual robots. Enabling a human to control such bio-inspired systems is a considerable challenge due to the limitations of each individual robot and the sheer number of robots that need to be coordinated to successfully complete a mission. Autonomous algorithms provide an opportunity to mitigate some of the complexity an operator faces in controlling such swarms, but it is not clear either (a) which tasks will ultimately need to be executed by the operator rather than by the swarm, or (b) what kinds of interactions would be needed.

Research Challenges

Michael Lewis

In computer science the notion of computational complexity, the time that must be used to solve a problem as a function of the size of its input, has proved fruitful for weeding out bad algorithms. Algorithms with high complexity may work for small problems, but fail or grow inefficient for even slightly larger ones. The task of controlling multiple robots is similar to an algorithm in that the operator must perform a repetitive sequence of decisions and actions to control a robot. If the robots are performing independent activities, the operator can devote the same attention to each in turn, resulting in a complexity of Order n , written $O(n)$, because each of the n robots requires the same set of actions and the total operator effort is proportional to the number of robots. Another benefit of independence is that more UVs can be controlled simply by adding more operators. A different form of control, such as designating a region to be searched by drawing it on a map, can command an arbitrary number of robots with a single act. Because the number of actions the operator must take are independent of the number of robots, control of this sort is $O(1)$ and has a constant effort. Dependent tasks such as box pushing, by contrast, can be arbitrarily difficult with command complexity, $O(>n)$, because dependencies among robots create cascading demands. When one robot pushes one corner of a box, for example, the operator must control the other robot to

push the other corner to straighten its path, after which the first robot needs attention again.

$O(1)$ tasks require substantial autonomy on the part of the robots but impose only a constant demand on the human operator. In general, $O(1)$ control is appropriate where a large number of UVs must be tightly coordinated with a relatively simple goal such as formation following or area search. $O(n)$ tasks, such as approving targets, or identifying victims, are robot-centric tasks that can be performed independently by one or more operators and impose a predictable additive demand. $O(>n)$ tasks, by contrast, cannot be specified simply and, depending on the task, could require arbitrarily large control effort on the part of the operator.

Human control of swarms is an $O(1)$ control problem of particular difficulty because there is no ready correspondence between human goals, swarm behaviors, and actions an operator might take to influence a swarm. Robots coordinated as “swarms” rely on simple control laws replicated across platforms which interact with each other to give rise to emergent organized behavior. Flocking behavior, for example, can be generated from three simple rules: 1) move away from any sensed robot closer than d_1 , 2) move toward any sensed robot further away than d_2 , 3) adjust heading to average heading of sensed robots. The balancing of attractive and repulsive forces and consensus on heading leads to a swarm that “sticks” together and moves in common, perhaps changing, directions.

Flocking is an example of biomimetic control because the control laws were chosen to mimic the

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE OCT 2012		2. REPORT TYPE		3. DATES COVERED 00-00-2012 to 00-00-2012	
4. TITLE AND SUBTITLE Human Factors issues for Interaction with Bio-Inspired Swarms				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S)				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Office of Naval Research, 875 N Randolph St, Arlington, VA, 22203				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES Presented at the HUMAN FACTORS and ERGONOMICS SOCIETY 56th ANNUAL MEETING, Boston, MA October 22-26, 2012					
14. ABSTRACT Robotic systems composed of a large number of entities, often called robot swarms, are envisioned to play an increasingly important role in applications such as search, rescue, surveillance, and reconnaissance operations. Nowadays, many mobile robots that are deployed for such applications are still tele-operated by a single or multiple operators. While these platforms are individually very capable, the development of cheaper hardware allows the consideration of swarm systems composed of many more robots but with each individual being far less powerful. Examples from biology indicate that such systems can be collectively more powerful than any individual robot within the team and also more than many larger, more sophisticated individual robots. Enabling a human to control such bio-inspired systems is a considerable challenge due to the limitations of each individual robot and the sheer number of robots that need to be coordinated to successfully complete a mission. Autonomous algorithms provide an opportunity to mitigate some of the complexity an operator faces in controlling such swarms, but it is not clear either (a) which tasks will ultimately need to be executed by the operator rather than by the swarm, or (b) what kinds of interactions would be needed.					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

behavior of flocking animals with the presumption that the animals, themselves, use some similar mechanism to coordinate their behavior. Swarm behavior can also be generated from analogs to physical laws by treating robots as point masses and using attractive/repulsive forces and artificial potential fields to produce emergent coordinated behavior. In this case the swarm is referred to as physicomemetic. In either case swarm behavior can only be influenced by altering the behavior of some members, altering the control laws, or altering the environment in which the swarm operates.

Our panelists take different positions on the ontological status of operator(s) in their efforts to influence swarm behavior. Katia Sycara will present an approach based on privileged operator communications that allow operators to alter swarm laws and/or parameters to influence their behavior. Michael Goodrich will present an alternate approach in which operators are limited to controlling or influencing a subset of members and must use unaltered control laws to propagate this influence to the remainder of the swarm.

Controlling Swarms through Manipulating Control Laws

Katia Sycara

In our research into human influence on swarms we have focused on human control over control laws and their parameters. Our work distinguishes two basic roles of the control laws: 1) maintaining swarm coherence and 2) producing behaviors that can be exploited to perform human desired tasks. Our research uses control laws from distributed robotics with known performance guarantees organized by (Bullo et al., 2009). Laws governing coherence are organized in the form of a “communications graph” that specifies the constraints on robot movement needed to maintain connectivity of the swarm. The severity of these constraints range from disc which requires that a robot never move to any position that takes it out of communication range of any robot it is currently in contact with to minimum spanning tree that allows a robot to move to any position that maintains at least one link to the swarm. Human desired tasks are then performed by requesting behaviors subject to these constraints achieved through algorithms for rendezvous (converge to consensus location) and deploy (distribute uniformly within area) with the addition of the primitive behaviors stop, random (all move randomly) and go to (heading) to help modulate behavior. Several initial experiments showed that exercising influence in this way allowed operators to

improve swarm behavior (on human desired tasks) in complex environments but not simple ones.

One difficulty with work in this area is that researchers have almost always made assumptions implicit in the way a swarm was displayed and commands issued that amount to operator omniscience. Rather than seeing accurate locations for swarm members and instantaneous communications most envisioned circumstances, such as undersea missions, where controlling robots as a swarm would be advantageous involve uncertainty, lags, and limited communications. We are currently focusing on problems where omniscience assumptions are violated.

One key challenge in the use of swarm robotic systems in human supervised tasks is to understand human swarm interaction in the presence of limited communication bandwidth, which is a constraint arising in many practical scenarios such as undersea missions or networks of limited capability robots. One of our human-subject experiments studied the effect of bandwidth limitations in human swarm interaction. We considered three levels of bandwidth availability in a swarm foraging task. The lowest bandwidth condition performed poorly, but the medium and high bandwidth condition both performed well. In the medium bandwidth condition, we display useful aggregated swarm information (like swarm centroid and spread) to compress the swarm state information. This corresponds to a situation in which there is reasonably high bandwidth between robots but limited bandwidth between the swarm and a remote human operator.

A second experiment investigated effects of locational error. In many operational settings, human operators are remotely located and the communication environment is harsh. Hence, there exists some latency in information (or control command) transfer between the human and the swarm. In this paper, we conduct experiments of human-swarm interaction to investigate the effects of communication latency on the performance of a human-swarm system in a swarm foraging task. We developed and investigated the concept of neglect benevolence, where a human operator allows the swarm to evolve on its own and stabilize before giving new commands. Our experimental results indicate that operators exploited neglect benevolence in different ways to develop successful strategies in the foraging task. Furthermore, we showed experimentally that the use of a predictive display could help mitigate the adverse effects of communication latency.

Although our positive results on mitigating problems of limited bandwidth and latency are promising the experiments were not designed to reflect ranges of

values and difficulties likely to be encountered in target applications. We believe that a broader more application sensitive series of studies are needed to expose the extent of problems and possible solutions for human control of swarms under realistic conditions.

The Middle Way: Cooperative Control of Swarms

Michael Goodrich

The relation between humans and swarms is often characterized in extreme ways. For control theorists seeking to design control laws to produce a desired set of emergent behaviors the human is frequently considered a disturbance. They consider their “job” to be designing the system so robustly that no matter what whacky input some operator might inject, the system remains conservative enough to stay with the bounds of safety. While designing to this criterion can preserve the coherence of the swarm and guarantee certain properties such as convergence to consensus it may so attenuate the operator’s influence that the swarm becomes incredibly sluggish and continues to behave autonomously even as the human seeks to control it. At the other extreme ceding full authority to a human operator can likewise cause undesirable results such as accelerating members at rates that lead to loss of coherence and break-up of the swarm or highly inefficient state transitions that might be achieved much more smoothly if done in consonance with ongoing behaviors.

We can find a middle ground by acknowledging that a swarm, particularly under conditions of noise/error and limited bandwidth may have more and deeper knowledge of its situation than a remote human operator. The human operator by contrast has a much more definite knowledge of his own goals but not necessarily of how the swarm may be able to achieve them. A solution is to require the operator to “work within the system” by injecting control through a small number of agents and allowing the system to adjust to these inputs over time. The resulting interaction uses the system’s inertia to protect it from disturbance like control “jolts” while remaining responsive to operator goals. The resulting style of interaction is one of persistent influence in which to bring the system to a new course the operator must persist in issuing a command over an extended time while persistently monitoring state information that may be noisy and lagged.

In a series of experiments we have investigated several approaches to modest injections of influence. In the first of these approaches we control agents called

leaders or predators that stand in a special relationship with other agents who are attracted and follow leaders and are repelled and flee predators. Operators influence the system by selecting agents that they can now control and designating those agents as leaders/predators to the remainder of the swarm. These experiments showed that use of leaders and attraction was a more effective approach under most conditions. More recently we have been investigating variants we are calling stakeholders and pacesetters that are operator “influenced” rather than controlled and exert influence over the swarm as peers rather than special types.

Another line of research is addressing the problems of overcoming inertia and sluggishness to influence naturally behaving swarms without excessive control effort. The solution is similar to a sprung garage door. By finely balancing the system between a closed state in which the spring exerts an upward force almost equal to gravity and an open state in which the spring has been unwound but winds as the door descends, a very small motor can move the door up and down. In a similar way we have found control laws that can shift a swarm’s behavior between flocking in which the swarm moves in linear direction to a torus in which the swarm moves in a circle. By choosing parameters carefully a swarm system can be created that can be shifted between the two attractors with relatively little exertion of human influence. We hope to find other such correspondences to help build a library of robust and humanly useful swarm behaviors with balanced transitions between them.

Futures and Challenges for Human-Swarm Interaction

Marc Steinberg

The past decade has seen substantial progress in moving from direct operator control to supervisory control of autonomous systems. This has included major field demonstrations with single and multiple air, sea, undersea, and ground systems. A primary goal of this research has often been to reduce the number of humans and/or the level of skill and training required to effectively manage a certain number of autonomous systems in a particular mission and environmental context. While this goal will remain important, future research may also focus increasingly on enabling a mixed human/robotic team to operate as (or more) effectively than a fully human one. However, given the challenges in this area, making progress may require advances across a much broader range of disciplines than was required to advance supervisory control. Progress in supervisory control has drawn particularly

on areas such as map and timeline-based displays, multi-modal interfaces, alert management, automated planning and payload management tools, rule and behavior-based systems, and autonomous navigation, guidance, and control. While this has advanced the field considerably, the current state of the art has significant limitations including brittleness of automation, inability to fully capture the human's intent in mission specification, and lack of flexibility to deal with situations the designer did not consider. The user may have a poor mental model of how the automation works, and the automation similarly may have little or no awareness of the user's status or perspective. Alerts must therefore often be very intrusive to ensure they are seen by the user. The automation also typically has only very limited ability to explain the rationale for its plans and behaviors, and the user may have poorly calibrated trust in the system.

There are particular challenges with enabling humans to interact and collaborate effectively with complex systems of multiple autonomous vehicles. Even for centralized and deterministic systems with much more understandable nominal behavior than many biologically inspired ones, it can sometimes be very challenging for operators to understand how their inputs will impact on the actions and behaviors of the system or to understand why the system is doing what it is doing. Trust also, can play a significant role in how operators choose to interact with these systems. For the most part, there has been limited thought about how users can effectively collaborate with these larger and more amorphous groups of unmanned systems. Further, a key reason for having a decentralized approach is communication limitations, and large asynchronous delays and system state uncertainty may greatly limit the potential role of the human in the loop. However, there are several approaches that suggest themselves as ways to move forward. First, it may be possible to abstract out the important characteristics of particular biologically-inspired systems and provide them in an understandable way to the user. This would involve some degree of "translation" of the system parameters to something much simpler and appropriate for a typical user to be able to understand. Though, this may require significant advances in analytical tools to be feasible, and it is not yet clear that this is even possible. In fact, it is currently not always possible even for the designers of such systems to ensure they will act in a desirable way or complete a particular objective. For those systems that have some type of leader or dominant animal, it may be possible to use the leader as a proxy for a human to control an entire group. This could be done with traditional leaders such as in the case of examples provided with wolves and lions or the more subtle

leaders of fish schools. Pheromone trails also suggest a way to support human interaction as has been explored to a limited extent.

Reference

Bullo, F. Cortes, J. & Martinez, S. (2009). *Distributed control of robotic networks*, ser. Applied Mathematics Series. Princeton University Press.