

# Bounds of Neglect Benevolence in Input Timing for Human Interaction with Robotic Swarms

Sasanka Nagavalli  
Robotics Institute  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA, 15213, USA  
snagaval@andrew.cmu.edu

Shih-Yi Chien  
Information Sciences  
University of Pittsburgh  
135 North Bellefield Avenue  
Pittsburgh, PA, 15260, USA  
shc56@pitt.edu

Michael Lewis  
Information Sciences  
University of Pittsburgh  
135 North Bellefield Avenue  
Pittsburgh, PA, 15260, USA  
ml@sis.pitt.edu

Nilanjan Chakraborty  
Mechanical Engineering  
Stony Brook University  
Stony Brook, NY, 11794, USA  
nilanjan.chakraborty  
@stonybrook.edu

Katia Sycara  
Robotics Institute  
Carnegie Mellon University  
5000 Forbes Avenue  
Pittsburgh, PA, 15213, USA  
katia@cs.cmu.edu

## ABSTRACT

Robotic swarms are distributed systems whose members interact via local control laws to achieve a variety of behaviors, such as flocking. In many practical applications, human operators may need to change the current behavior of a swarm from the goal that the swarm was going towards into a new goal due to dynamic changes in mission objectives. There are two related but distinct capabilities needed to supervise a robotic swarm. The first is *comprehension of the swarm's state* and the second is *prediction of the effects of human inputs on the swarm's behavior*. Both of them are very challenging. Prior work in the literature has shown that inserting the human input as soon as possible to divert the swarm from its original goal towards the new goal does not always result in optimal performance (measured by some criterion such as the total time required by the swarm to reach the second goal). This phenomenon has been called Neglect Benevolence, conveying the idea that in many cases it is preferable to neglect the swarm for some time before inserting human input. In this paper, we study how humans can develop an understanding of swarm dynamics so they can predict the effects of the timing of their input on the state and performance of the swarm. We developed the swarm configuration shape-changing Neglect Benevolence Task as a Human Swarm Interaction (HSI) reference task allowing comparison between human and optimal input timing performance in control of swarms. Our results show that humans can learn to approximate optimal timing and that displays which make consensus variables perceptually accessible can enhance performance.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).  
HRI'15, March 2–5, 2015, Portland, Oregon, USA.  
Copyright © 2015 ACM 978-1-4503-2883-8/15/03 ...\$15.00.  
<http://dx.doi.org/10.1145/2696454.2696470>.

## Categories and Subject Descriptors

I.2.9 [Robotics]: Operator Interfaces; I.2.11 [Distributed Artificial Intelligence]: Multiagent Systems; H.1.2 [User/Machine Systems]: Human Factors; H.5.2 [User Interfaces]: Benchmarking

## General Terms

Algorithms; Design; Experimentation; Human Factors; Performance; Theory

## Keywords

Human-Robot Interaction; Robotic Swarms

## 1. INTRODUCTION

A robotic swarm is a multi-robot system where individual swarm members use control laws based on local spatial information about the environment and/or other members of the swarm within their spatial neighborhood to achieve various emergent collective behaviors, such as foraging or rendezvous. The key advantage of swarms is that the behaviors generated are robust to individual robot failures. Applications of swarm robotics are envisioned in environmental exploration, large scale emergency response and search and rescue and environmental cleanup (e.g. cleaning oil spills). Such missions may require use of a sequence of different swarm behaviors that may need to be imparted to the swarm by a human operator.

Since the state of a swarm not yet at equilibrium is continuously evolving, the supervisor's task is to supply inputs that divert the swarm from its current trajectory to a new desired trajectory. This Trajectory<sub>1</sub>-Input-Trajectory<sub>2</sub> sequence therefore defines the basic task that must be addressed in the study of human-swarm interaction (HSI).

There are two related but distinct capabilities needed to supervise a robotic swarm. The first is *comprehension of the swarm's state* and the second is *prediction of the effects of human inputs on the swarm's behavior*. Comprehension of the swarm's state requires a human to perceptually extract relations and regularities in behavior from observable data such as robots' positions and velocities. This can be a challenging task since crucial aspects of the robots'

internal states, such as distances from active goals or the conditions controlling current behaviors, may not be directly observable. In [24], for example, researchers were moved to install LEDs on their robots in order to display their internal states to aid in debugging swarm behaviors. Fortunately, regularities in relations and dynamics of widely employed swarm algorithms such as flocking are readily perceived by human viewers. In fact, the original purpose of the flocking algorithm [31] was to generate computer graphics that humans would find similar to the behavior of flocks of birds or schools of fish.

The second capability, predicting effects of human inputs, requires the operator to develop an “internal model” of the swarm’s dynamics. This capability is related to the first, since, if it is possible for the human to discern continuity in the behavior and state evolution of the swarm, then the human would have a better idea of the system dynamics, would be able to predict to some extent what the swarm would do next, and hence be able to choose when to insert input to bring about desired behavior. Even for relatively intuitive relationships such as commanding a leader who is followed by a swarm employing a flocking algorithm, issues involving limitations in communications, robot speed or lags in response require the operator to develop a model of the swarm dynamics in order to effectively control the swarm. For other inputs, such as a command to switch between algorithms, delivered by broadcast or propagation, the relation between input and a desired effect on the swarm may be even less clear. Difficulties in developing internal models of dynamics have been found to be particularly acute when timing of inputs or lags in response are present [38] as is likely in controlling a swarm. A close counterpart lies in industrial processes where substantial lags between input and response have been abstracted to laboratory tasks such as Crossman’s water bath [17]. In the waterbath task, operators were asked to control a heater under a beaker to heat the water in the beaker to maintain a setpoint with a lag introduced by putting the thermometer inside a testtube. Later, computerized versions of process control tasks [26] have used multiple tanks, valves and pumps but preserved the lagged integrative response of the waterbath. Common findings in these experiments are that performance improves with practice, but instruction in principles underlying the system does not improve performance [27]. Negative correlations often found between performance and verbalizable knowledge [5] further suggest learning of dynamics may be largely implicit making it difficult to improve or correct human performance.

In view of these results, enabling humans to effectively exert supervisory control on swarms by understanding the underlying dynamics and giving control inputs at the right time to optimize performance objectives presents a huge challenge. In this paper, we study this challenge and explore ways to better understand whether and how humans comprehend swarm dynamics and also how to best aid them. In particular we want to answer the following questions: (a) Can humans develop a model of the dynamics expressed in moving from one swarm configuration to another, so that they approximate the optimal time to give their input to the swarm? (b) Can we improve human performance by providing an aid which makes the robots’ goals visible to the operator and provides a holistic guide to the swarm’s state? (c) More generally, can we extend human monitoring and control to distributed systems that do not conform to perceptual principles by engineering displays for humans rather than requiring swarms to follow inefficient algorithms to make their behavior humanly intelligible?

We conducted a study to investigate human performance at the Neglect Benevolence shape-changing HSI reference task to determine the degree to which human operators could approximate op-

imal performance, given experiences interacting with system dynamics of a swarm. We ran experiments in two conditions, one where the participants interacted via a usual display (unaided condition) and another where the participants had the use of an augmented display (aided condition). Our results showed that the participants indeed learned over time to improve their performance. Moreover, the success of the participants in the aided condition lends some support to our hypothesis that making swarm consensus variables perceptible to humans would increase the efficiency of human swarm interaction.

This paper makes the following contributions. First, it presents a framework characterizing Human-Swarm Interaction (HSI) as a sequence of human inputs diverting a swarm’s trajectory through state space that is reducible to a series of Trajectory<sub>1</sub>-Input-Trajectory<sub>2</sub> events and not restricted to behaviors involving only translation such as flocking, rendezvous, or deploy. Second, it introduces the swarm configuration shape-changing Neglect Benevolence Task as an HSI reference task allowing comparison between human and optimal input timing performance in control of swarms. Third, it introduces a Gestalt-based approach to characterizing the intelligibility of swarm behavior based on the perceptual saliency of the object of consensus. Fourth, it provides an initial test of human performance on the Neglect Benevolence reference task and a test of display augmentation based on characterization of intelligibility.

The paper is organized as follows. Section 2 presents related work. Section 3 briefly outlines the notion of Neglect Benevolence and optimal input time. Section 4 describes the configuration control and shape-changing task as a HSI reference task allowing comparison between human and optimal performance. Section 5 introduces an approach to characterizing the intelligibility of swarm behavior based on the Gestalt principle of common fate and evaluates common swarm algorithms including shape-changing. Sections 6 and 7 describe the experiment and results comparing human performance with conventional and aided displays at the Neglect Benevolence reference task. Section 8 presents discussion and Section 9 conclusions and future work.

## 2. RELATED WORK

Robotic swarm systems where a group of robots need to act through local interactions to collectively achieve a variety of behaviors have been extensively studied. Swarms can achieve behaviors such as flocking [31, 9, 33, 4], rendezvous [6], deployment [7], and foraging [13]. Formation control of robotic swarms has been proposed to adjust the spatial configuration of robotic systems so that they satisfy the requirements of deployment tasks or they can adapt to different environments while performing tasks [1, 36, 14]. However, schemes which include a human operator in the loop together with robotic swarms are required for some complicated tasks (e.g. complex surveillance and reconnaissance). For using swarm robotic systems in human-supervised missions, human swarm interaction (HSI) has been studied in [21, 8, 19, 15, 3] primarily at variants of foraging tasks. Study of human involvement in formation control includes transitions between flock and torus [3], leader choice and network configuration [12].

The concept of Neglect Tolerance has been introduced to capture the idea that a human operator can neglect robots that work independently of one another in a multi-robot system for a certain time before system performance degrades [29, 34]. One key issue in Neglect Tolerance is how to schedule the attention of the human operator among the multiple and independent robots so that the neglect time between servicing robots is minimized [25, 23]. In Neglect Tolerance, it is assumed that the robot system’s performance will degrade over time and the system will need human interaction

to restore performance to a desirable level. In contrast, the concept of Neglect Benevolence in human swarm interaction captures the idea that it may be beneficial for the human operator to wait for a certain length of time before applying input to the system while the swarm state evolves to stabilization (since the swarm performance may not degrade monotonically with time) [28].

### 3. NEGLECT BENEVOLENCE

In a previous experimental study using a foraging scenario [38], it was found that the performance of the human-swarm system was strongly affected by the time between two commands that the human applied to the robots. In particular, it was found that one group of subjects who performed well waited for some time after they issued a command before issuing another corrective command (when they wanted to change the direction in which the swarm was heading). The phenomenon was called *Neglect Benevolence*, since neglecting the swarm for some amount of time led to better performance of the system.

Further analysis of the results [38] found that when the swarm was in a transient state (i.e. moving towards one goal), applying another input (that changes the goal) could have different effects depending on the timing of the inputs.

To determine whether this was a spurious effect or not, [28] reported various simulations of swarm systems, starting at different configurations and performing rendezvous where the operator inputs changed the rendezvous point. The robots were moving under simple control laws given in [16] with a repulsive potential to avoid inter-robot collision and an attractive potential field to maintain swarm cohesion. It was observed that the simulation gave a variety of resulting outcomes depending on the times when the human input was given following the desired change in the swarm goal. More concretely, it was observed that giving the input immediately after the need arose to change the rendezvous point resulted in detaching several robots who subsequently never made rendezvous while some of the delayed inputs led to the whole swarm staying together and completing the rendezvous at the desired point.

Here, we briefly discuss some results shown in [28] that serve as a useful background for the computations and task proposed in this paper. First, it was shown that Neglect Benevolence is present not only for nonlinear systems (e.g. swarms with collision avoidance) but also for linear systems. Second, a formal definition of Neglect Benevolence was introduced that allows performance to be defined by a variety of measures such as the time or energy used by the swarm to reach the human goal. Third, it was proven that any exponentially stable linear system exhibits Neglect Benevolence. Fourth, an algorithm that computes the optimal input time for insertion of human input to achieve least time to human goal for linear time invariant (LTI) systems was given. Fifth, it was shown that the dynamics matrix of a swarm that performs consensus for controlling swarm configuration control and shape-changing (the reference HSI task proposed in the current paper) falls within this class of (LTI) systems.

From the results reported in [28], it can be concluded that Neglect Benevolence is a nuanced and quantifiable notion that captures precise calculation of “too early” or “too late” in inserting the human input. In other words, if a system is neglected for a very long time, it may take a very long time to converge. On the other hand, if a system is neglected for too little time it may also take a very long time to converge. Therefore, in order to guarantee timely convergence, a bounded Neglect Benevolence must be determined, so that the neglect interval is neither too small nor too large. Hence, another way of looking at the study in this paper is to determine

whether humans have the ability to gauge the bounds of Neglect Benevolence.

We consider configuration control and swarm shape-changing as a base task within which to study Neglect Benevolence and optimal input timing in human swarm interaction. In the next section, we present more details about the underlying formal model and dynamics of this task.

### 4. CONFIGURATION CONTROL AS A BASE HUMAN SWARM INTERACTION TASK

In this section, we describe the dynamics of a swarm system using consensus for configuration control. Individual robots within the swarm communicate with their neighbors. The communication graph  $\mathcal{G} = (\mathcal{V}, \mathcal{E})$  captures the connectivity of robots within the swarm. An individual robot is represented as a vertex  $v_i \in \mathcal{V}$  of the graph and the edges  $\mathcal{E} \subset \mathcal{V} \times \mathcal{V}$  represent communication links such that each edge  $(v_i, v_j) \in \mathcal{E}$  indicates that robot  $v_i \in \mathcal{V}$  can communicate with robot  $v_j \in \mathcal{V}$ . Assume that all communication links are bidirectional. This means the communication graph is undirected (i.e.  $(v_i, v_j) \in \mathcal{E} \Rightarrow (v_j, v_i) \in \mathcal{E}$ ). The adjacency matrix  $\mathbf{A} = [a_{ij}]$  is a mathematical representation of the graph where where element  $a_{ii} = 0$ ,  $a_{ij} = 1$  if  $(v_i, v_j) \in \mathcal{E}$ , and  $a_{ij} = 0$  if  $(v_i, v_j) \notin \mathcal{E}$ . The Laplacian matrix  $\mathbf{L} = [l_{ij}]$  can be obtained from the adjacency matrix with each element  $l_{ii} = \sum_j a_{ij}$  and  $l_{ij} = -a_{ij}$ . Since the graph is undirected, the adjacency matrix and the Laplacian matrix are both symmetric. When the graph is connected, the Laplacian matrix is singular and has one eigenvalue equal to 0 and a corresponding eigenvector  $\mathbf{1}$ . Note that  $\mathbf{1}$  represents a column vector with all entries equal to 1.

The joint state of the robotic swarm is given by the state vector  $\mathbf{x}$ , which is composed of the concatenated states  $x_i$  of individual robots within the swarm. In continuous-time averaging consensus, the change in the state  $\dot{x}_i$  over time  $t$  of an individual robot  $v_i$  is based on averaging the states of neighboring robots. Two robots  $v_i$  and  $v_j$  are called neighbors if they can communicate with each other (i.e.  $(v_i, v_j) \in \mathcal{E}$ ). As presented in [30], the evolution of the joint state of a system performing averaging consensus can be described by the following equation, where  $\mathbf{L}$  is the Laplacian matrix of the system’s communication graph.

$$\dot{\mathbf{x}}(t) = -\mathbf{L}\mathbf{x}(t) \quad (1)$$

Note that Equation (1) describes a stable linear time-invariant (LTI) system. The Laplacian matrix  $\mathbf{L}$  has all non-negative real eigenvalues, so  $-\mathbf{L}$  has only non-positive real eigenvalues. Since all eigenvalues of the system are real and non-positive, the system with dynamics described by Equation (1) will always move monotonically from any initial state  $\mathbf{x}_0 = \mathbf{x}(0)$  towards a final state  $\mathbf{x}_f = \lim_{t \rightarrow \infty} \mathbf{x}(t)$  that lies within the null space of the system [28]. For a swarm with a connected communication graph, the final state is any state in which the robots all agree with each other. That is,  $\mathbf{x}_f = c\mathbf{1}$  where  $c \in \mathbb{R}$ .

When the robot states are their spatial positions, the averaging consensus dynamics in Equation (1) result in rendezvous at the centroid of the robots’ initial positions. In [28], this algorithm is extended to incorporate an input vector  $K$ .

$$\dot{\mathbf{x}}(t) = -\mathbf{L}(\mathbf{x}(t) - K) \quad (2)$$

It was shown in [28] that these dynamics cause the system to evolve towards a final state  $\mathbf{x}_f$  that (a) preserves the centroid of the initial positions of the robots and (b) preserves the differences in corresponding entries of the input vector  $K$ . Thus, if  $K$  specifies the robot positions for a formation in some arbitrary but consistent co-

ordinate frame, a robot swarm with the dynamics in Equation (2) will generate the specified configuration around the centroid of the initial positions. Thus, a robot swarm could be performing simple consensus and a human operator could broadcast an input vector  $K$  and with only local interactions, the robots would still generate the desired configuration.

This introduces the important question of when a human operator should apply a particular input. Specifically, if the human operator has already specified an input  $K_{old}$  that generates a particular configuration, then when should the operator specify a new input  $K_{new}$  to generate a new configuration (i.e. change the goal of the swarm) as quickly as possible? Human intuition would indicate that the operator should specify  $K_{new}$  as soon as possible if they want to obtain the new configuration as quickly as possible. However, as shown in [28], this intuition is provably wrong in the general case! For any LTI system, given an input  $K$ , it is always possible to find a state  $\mathbf{x}$  in which it is beneficial to delay the application of the input. Since this is non-intuitive to the human operator, this motivates the need for appropriately aiding the operator in finding or approximating the optimal time to apply the input, maximizing the performance of the robot swarm.

## 5. INTELLIGIBILITY OF SWARM BEHAVIOR

The emergent behavior of a swarm results from many local interactions among its members. Most commonly, some variant of consensus (averaging values among neighbors) is used for the coordination giving rise to swarm behaviors. Perceiving an emergent swarm behavior requires the observer to recognize and associate similarities among the many local interactions among robots. Consequently, human perceptual access to swarm behavior depends on these local interactions/agreement being expressed in a way in which salient similarities and relations (consensus) can be recognized. Recognition of these regularities and the swarm's trajectory is a prerequisite for monitoring and decision making. The extraordinary ability of humans to perceive biological motion has been well established since Johansson's [18] experiments in which human actors with point-lights attached to major joints were recorded and later viewed as images showing only a few moving dots. Viewers easily identified the human figures and their activities, such as walking, even in the presence of substantial "masking" by other dots [11]. Even hidden characteristics of motion such as the weight of a box [32] or distance of a thrown object [20] were discernible to viewers of these point-light displays. Inverting the imagery (upside down human walkers), however, eliminates almost all of the effect [2] leaving observers unable to discriminate between scenes containing a walker and those with only a "mask". These results suggest a strong ecological and/or learned component to human abilities to perceive biological motion that may not fully extend to behaviors of a swarm. This suggests that the more general Gestalt principles of *proximity*, *similarity*, *closure*, *symmetry*, *continuity*, *common fate*, *pragnanz* and *past experience* may provide the best guidance for identifying representations likely or unlikely to lead to the perception of swarm behaviors. Of these principles, only *common fate* applies specifically to motion, while the rest largely prescribe how parts come to be perceived as wholes and could come into play as robots approach well defined formations.

The most common display of a swarm comprises spatially related icons on a plane or in a 3D space corresponding to robot positions. Swarm behavior following standard algorithms such as *flocking* or those collected in [6] such as *rendezvous*, *deploy* or *boundary following* is displayed through movements of these icons. Fortu-

nately, these displays (moving dots perceived as patterns) closely resemble those studied by psychologists and vision researchers so there is already research into the factors supporting the perception of group behavior in this context. *Common fate* refers to the grouping factor that leads a previously invisible grouping of dots to pop out of a noisy background when moved and has been found to be the most important factor in perceiving coordinated behavior although *proximity* [37], *similarity* [10] and *pragnanz* in the form of grouping features such as collinearity [37] have all been shown to play important roles and to interact [37] in strengthening the effects of *common fate*. Humans are very sensitive to perceiving *common fate* groupings with success at signal-to-noise ratios as low as 1-2% [35] and the effect is relatively robust to divergence among trajectories as long as some coordinated global direction is maintained [39]. There are suggestions, however, that *common fate* may be limited to a single group at a time [22]. For example, of two overlapping groups only one would appear to move coherently. Effects have been found strongest for linear or circular trajectories [40]. While spatial characteristics predominate in perception of moving groups other features such as common color can also contribute [10].

By this analysis, even without the benefits of biological motion, flocking behavior in which agents adopt the average heading and velocities of their neighbors seems ideally suited for perceptual recognition. In either its linear or torus form [3], the interaction among agents is precisely that giving rise to *common fate* while the regulation of distance (*proximity* and *pragnanz*) enhances this effect. For the rendezvous and deploy algorithms, the consensus process should be less strongly conveyed because while the global movement inward or outward and reflected in local interactions should be perceivable [40] it lacks the matching of velocities and headings, which are the dominant contributors to the perception of *common fate* [37].

By a similar analysis we would predict that the behavior of a swarm following the shape-changing algorithm would be extremely difficult for a human observer to follow because in moving from one configuration to another neither headings nor velocities are matched (*common fate*) nor constant distances nor boundaries maintained (*proximity* and *pragnanz*). That is to say that the movement of the swarm toward consensus as perceptually manifested by a centroid is masked by the formation. Since the consensus process is not perceptually accessible from robot behaviors, the human monitor would be left with the impression of robots that are moving about independently and be unable to either monitor their progress toward a goal or compose an input to alter that progress. As the target formation is approached, however, the entire set of Gestalt principles governing perception of form (*proximity*, *similarity*, *closure*, *symmetry*, *continuity*) would come into play making the goal and progress toward it increasingly apparent in this second phase.

These analyses suggest that intelligibility of swarm behavior could be improved for algorithms in which progress toward consensus is not perceptually available by augmenting the display in a way that makes this information available. Since the perceptually uncoordinated movements of robots to their positions in the target formation cannot be eliminated, the display can only be augmented in ways that make the robots' goals and progress toward their goals more apparent. A potential augmentation of this sort would be to draw lines from the robots' current positions to their predicted locations in the new formation. As the swarm evolves toward its initial goal shape these lines would shorten and lengthen until the supervisor chooses to divert the swarm to the desired shape at the point he/she judges to be the closest (shortest lines). In this augmentation, the line between each robot's current position and its final position is

a visual proxy for the distance in the swarm’s state space between current and goal state. The swarm’s state space has very high dimensionality and cannot be easily visualized, so this proxy is used as an aid. The next section describes an experiment comparing a display augmented in this way with a conventional display.

## 6. EXPERIMENTAL DESIGN

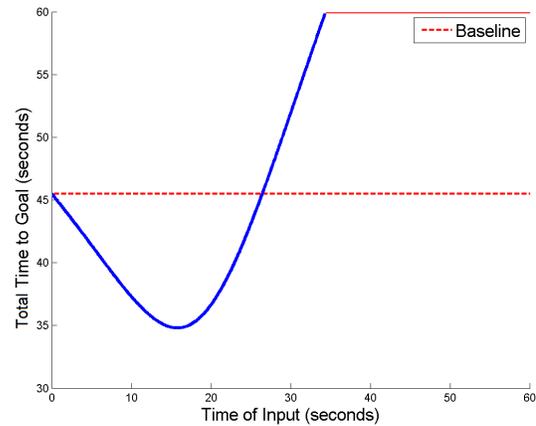
Our initial experiment investigating human performance at the Neglect Benevolence HSI reference task was to determine the degree to which human operators could approximate optimal performance given experience interacting with system dynamics (implicit learning). Comparison of human and optimal performance is a standard practice in the study of human-machine systems for benchmarking human performance and evaluating human competence within a work domain. While we would expect to assign easily automatable tasks such as input timing to the system, it may be important to allow the human to intervene in response to “out of band” information or other anomalies. In addition, the possibility of manual control may be an important element for preventing out-of-the-loop effects or complacency and enabling the operator to effectively monitor the system for failures or abnormal behavior. Possible outcomes included almost optimal performance, biased performance favoring late inputs (“finding the bottom” heuristic), idiosyncratic performance indicated by a wide dispersion between early and late inputs or some other explainable outcome. These are descriptive questions answerable by descriptive statistics.

We used a mixed-model design for our experiment, which divided subjects into two groups and had one group interact with the swarm through a conventional display showing only robot positions (unaided condition), while the other group interacted with the swarm through a display that included a Gestalt-based augmentation to the visualization of the swarm (aided condition). Interactions with the swarm were divided into discrete trials. Each group interacted with the swarm in the same number of training and test trials. The swarm dynamics and initial state remained consistent between corresponding trials in the unaided and aided condition. Thus, for the mixed-model design, the within-subject variable was the trial number and the between-subject variable was the presence or lack of the Gestalt-based visual aid. A regression against trials was used to measure practice effects. The comparison between the display with the Gestalt-based augmentation and the conventional display provides an initial test of our hypothesis that swarm behaviors can be made more intelligible and controllable by making progress toward consensus perceptually assessable.

### 6.1 Method

#### 6.1.1 Task

The experimental task involves human interaction with a simulated robotic swarm performing consensus on robot positions. The swarm includes the capacity to incorporate an input as described in Section 4. The task is for the subject to choose the best time at which to provide an input diverting the robotic swarm from the shape toward which it is initially moving (equilibrium point of the dynamical system) to another desired shape (new equilibrium point). This matches our earlier depiction of the canonical HSI task as redirecting the swarm — modelled as a dynamical system — from one “natural goal” state (Formation 1) to another “desired goal” state (Formation 2). If the system in a Neglect Benevolent state, it is beneficial to delay the application of the input to minimize the total time required to converge to the desired formation. Figure 1 shows an example of how delaying the application of the input could reduce the time required to reach the goal state.



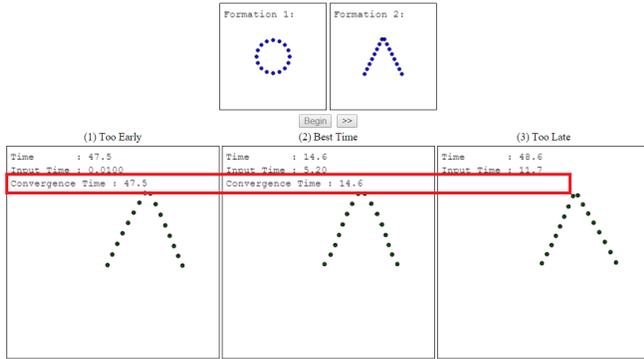
**Figure 1: This plot shows the relationship between convergence time (i.e. total time to goal) and input time for one of the trials in our experiment. The convergence time if the input is applied immediately at the start of the simulation is shown as the baseline time. Each trial has a maximum length of 60 seconds. For this trial, if the input is applied after approximately 33 seconds the swarm would not converge within the 60 second limit.**

Participants first read the standard instructions on how to use the interface to interact with the swarm robots and then proceed to complete each trial in the experiment. In the aided condition, there is a line between each robot’s current position and its final position in the desired formation. These lines help participants visualize the distance between the robots’ current positions and the desired formation (Formation 2). Since the swarm dynamics are LTI as described in Section 4, when the collective distance represented by all the lines is larger, more time is required to reach the desired formation. Formation 1 and Formation 2 are shown in thumbnails at the top of the screen following a summary of instructions.

The experiment has 12 training trials to familiarize participants with the system and allow them to observe the asymmetric effects of leading/lagging the optimal input time. As learning of system dynamics is presumed implicit, familiarization is needed to engender common strategies so performance might be compared. During the training trials, participants are informed of the effects of applying the input too early/late and compare them with the optimal time. The training trials are divided into two types: *passive*, where participants only observe, and *active*, where participants give input themselves. The first 6 trials are *passive*, where participants observe three simulation panels situated below thumbnails depicting Formation 1 and Formation 2. All three simulations begin with the robots in the same random initial configuration. When a user clicks “Begin”, the swarms in all three panels begin moving toward Formation 1 until given an input (at a different time for each panel). Visually, upon receiving the input, the swarm robots change color from red to green. Each panel shows how the input redirects the configuration of the swarm towards Formation 2. Upon converging to Formation 2, the convergence time is shown. The left, middle and right panels each display the effects when the input is applied “too early”, at the optimal time, and “too late” respectively (see Figure 2). The next 6 trials are *active*, where participants have the chance to apply the input themselves and get feedback as to whether they gave the input too early, too late, or optimally (see Figure 3). Participants are again presented with thumbnails showing Formation 1 and Formation 2, but followed by only one simulation panel. The panel shows the robot swarm in a random initial

configuration. After the participant clicks “Begin”, the swarm starts moving toward Formation 1. The participant then has the chance to give an input by clicking “Activate New Formation” at any time before the end of the simulation. Once the swarm has converged (or the simulation has reached the 60 second time limit), the convergence time resulting from the input and feedback as to whether their input was too early or too late (see Figure 3) is displayed.

After the subjects finish the 12 training trials, they are presented with 36 test trials in the same condition (aided or unaided). In the test trials, the subjects are presented with the same screen arrangement (see Figure 4) as in the *active* training trials, but although they are able to observe the convergence time, they are not given feedback as to whether their input was too early or too late. The participants are instructed to impose the new goal on the swarm at the time leading to the earliest convergence.



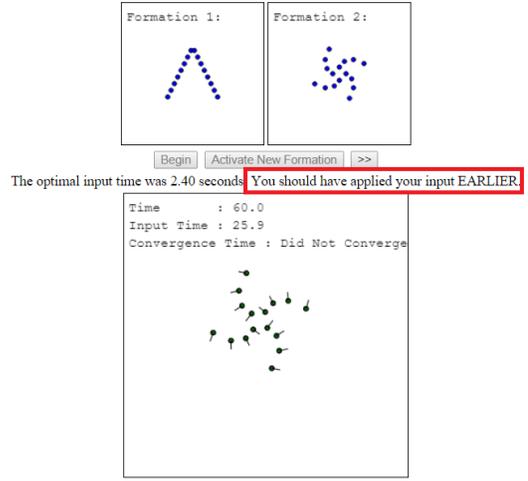
**Figure 2:** For each trial in the first (passive) training phase, participants were shown three simulations. In all simulations, robots started in the same initial positions, but the input telling them to move to Formation 2 was applied at different times: (1) too early, (2) at the best time, (3) too late. After the swarm converged to the final formation (see panel 1 and 2), participants were shown the time taken for the robots to converge.

### 6.1.2 Participants and Data Collection

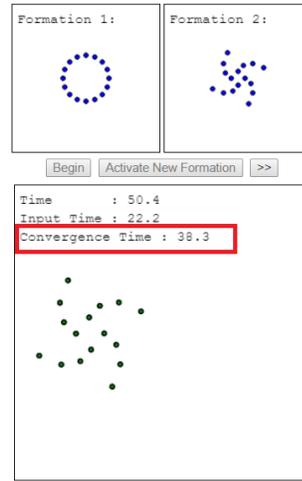
A total of 44 participants were recruited (22 for aided condition, 22 for unaided condition) from the University of Pittsburgh and Carnegie Mellon University communities through online solicitation. Data was collected through a HTML5 / JavaScript web application that recorded each subject’s trial data anonymously in a secured MySQL database. Participants were randomly assigned to either the control group (unaided condition) that interacted with the swarm through a conventional interface or the experimental group that interacted with the swarm through an interface with a Gestalt-based augmentation (aided condition). Participant input times were recorded for each trial. Optimal input times for each trial were computed offline using the algorithm in [28].

## 7. RESULTS

Input time deviation and convergence time served as dependent variables. Input time deviation ( $|t_{input} - t_{optimal}|$ ) was defined as the absolute difference between the time chosen by participants to divert the swarm and the optimal diversion time. Convergence time was defined as the interval from the beginning of the trial until the swarm converged to the desired formation (Formation 2). Input time measured the participants’ deviation from optimality directly while convergence time measured the effect of these suboptimal



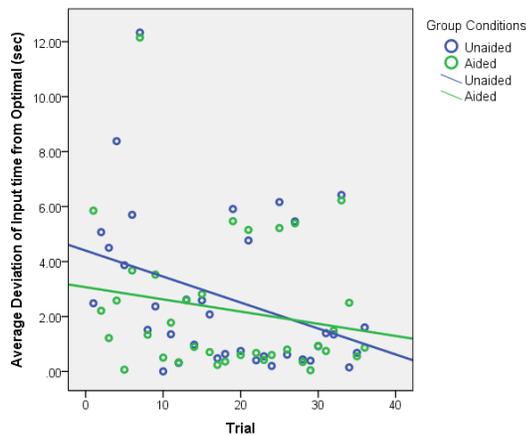
**Figure 3:** For each trial in the second (active) training phase, participants were asked to apply the input at whatever time they thought would minimize the total time required for the robots to converge to Formation 2. At the end of the trial, participants were told the optimal input time and whether they should have applied their input earlier or later.



**Figure 4:** For each trial in the test phase, participants were asked to apply the input at whatever time they thought would minimize the total time required for the robots to converge to Formation 2. They were not given any feedback.

inputs on performance. Since delays that were too long could lead to much longer convergence times than those that were too short, strategies favoring earlier inputs might be a rational response for participants having difficulty in approximating the optimal interval making convergence and input times slightly different measures. Independent variables were Display type (conventional or aided), Transition Type (circle-spiral, circle-V, spiral-circle, spiral-V, V-circle, V-spiral), and trial number.

Data were analyzed both in an ANOVA design treating Display type and Transition type as factors and a regression incorporating trial numbers as well. An ANOVA found main effects for both Display types ( $F_{1,1572} = 6.214, p = .013$ ) and Transition type ( $F_{1,1572} = 33.402, p < .001$ ) for input time deviation. An ANOVA



**Figure 5: This plot shows the average absolute deviation of the participants' input time from the optimal input time for each trial in the unaided and aided condition.**

for convergence time also found main effects for both Display type ( $F_{1,1572}=3.387$ ,  $p=.006$ ) and Transition type  $F_{1,1572}=45.418$ ,  $p<.001$ ). Both groups on average preceded the optimal time with participants in the aided display condition responding nearly twice as early ( $M=-1.23$  sec,  $SD=5.11$ ) as those in the unaided condition ( $M=-.65$  sec,  $SD=7.49$ ) although this difference was not significant at the  $p<.05$  level. A Mann-Whitney U test of deviations from optimal input times, however, found that participants in the aided condition (mean rank=767.63) were significantly closer to the optimum ( $z=-2.164$ ,  $p=.03$ ) than those in the unaided condition (mean rank=817.37). A comparison of deviations from optimal convergence times found that participants in the aided condition (mean rank=762.55) were significantly closer to the optimum than those in the unaided condition (mean rank=822.45);  $z=-1.97$ ,  $p=.009$  on this measure as well. Regression of input time deviations against trials found a significant improvement ( $F_{1,34}=4.918$ ,  $p=.033$ ) in the unaided Display condition but not in the aided condition.

## 8. DISCUSSION

Our results as shown in Figure 5 indicate that humans did, in fact, learn to approximate the optimal input time over the course of the experiment. This suggests that both Display types provided sufficient cues for humans to learn to time their input in order to decrease convergence times. The quality of this approximation, however, began at a high level for the group using the aided Display but developed only slowly for those in the unaided condition. We believe this occurred because the lines drawn to terminal locations in Formation 2 on the augmented Display provided viewers with a reference for judging the swarm's evolution while the movement of robots in the unaided Display did not. Relying only on Gestalt influences over the perception of form as formations were approached, unaided participants did not have a basis for earlier responses. This conjecture is supported by the finding that those in the aided condition responded on average earlier than those in the unaided condition. There are several possible explanations for the improvement in performance of participants in the unaided condition. Participants might have acquired an implicit model of shape-changing dynamics that allowed them to better assess the swarm's state as it evolved. Alternately, they may have developed greater sensitivity in recognizing an evolving formation's similarity to their goal formation. A third alternative might be that par-

ticipants learned a timing heuristic for delaying input. These and similar questions can only be answered through additional properly designed experiments to test those explanations.

The success of the augmented condition in improving the performance of naive participants lends some support to our hypothesis that difficulties can arise in interaction with swarm systems because the variables on which the swarm is coordinating are not perceptually accessible to the human. Many distributed algorithms involving human interaction such as flocking, biology-mimicking control through leaders and predators, or physicomimetic direction via potential fields depend in some part on their intelligibility to human users and their programmers. If we could develop principled methods for augmenting the behavior of swarm algorithms with display of variables under consensus it could extend human control and use to a wider range of potentially more efficient algorithms.

## 9. CONCLUSION

Some of the main contributions of this work are methodological and lie in providing a framework for a more controlled study of human-swarm interaction. Characterizing HSI as the act of diverting a swarm's trajectory through state space from the current goal to an alternate goal provides a new basis for evaluating HSI in isolation from normally confounding factors. In particular, by focusing on progress toward goal states rather than simply measuring transient effects of behavior, it allows an experimenter to define a strict performance reference for human behavior unlike earlier efforts such as [21] which have relied on comparing performance to that of intelligent agents.

The shape-changing Neglect Benevolence Task provides an example of an HSI reference task of this sort that allows direct comparison between human and optimal performance in control of swarms. The methods and tools developed in this process provide a roadmap for other researchers wishing to develop distributed algorithms and displays that can be evaluated in these ways. We have presented and provided an initial test of a Gestalt-based approach to characterizing the intelligibility of swarm behavior based on the perceptual salience of the object of consensus. While results are preliminary they are promising in providing a closer linkage between human cognition, distributed algorithms and the way we display them.

## 10. ACKNOWLEDGMENTS

This research has been sponsored in part by ONR Grant N0001409-10680 and an NSERC PGS D scholarship.

## 11. REFERENCES

- [1] T. Balch and R. C. Arkin. Behavior-based formation control for multirobot teams. *Robotics and Automation, IEEE Transactions on*, 14(6):926–939, 1998.
- [2] B. I. Bertenthal and J. Pinto. Global processing of biological motions. *Psychological science*, 5(4):221–225, 1994.
- [3] D. S. Brown, S. C. Kerman, and M. A. Goodrich. Human-swarm interactions based on managing attractors. In *Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction*, pages 90–97. ACM, 2014.
- [4] D. J. Bruemmer, D. D. Dudenhoeffer, M. D. McKay, and M. O. Anderson. A robotic swarm for spill finding and perimeter formation. Technical report, DTIC Document, 2002.
- [5] A. Buchner, J. Funke, and D. C. Berry. Negative correlations between control performance and verbalizable knowledge: Indicators for implicit learning in process control tasks? *The*

- Quarterly Journal of Experimental Psychology*, 48(1):166–187, 1995.
- [6] F. Bullo, J. Cortés, and S. Martinez. *Distributed control of robotic networks: a mathematical approach to motion coordination algorithms*. Princeton University Press, 2009.
- [7] H. Choset. Coverage for robotics—a survey of recent results. *Annals of mathematics and artificial intelligence*, 31(1-4):113–126, 2001.
- [8] G. Coppin and F. Legras. Autonomy spectrum and performance perception issues in swarm supervisory control. *Proceedings of the IEEE*, 100(3):590–603, 2012.
- [9] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and N. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of theoretical biology*, 218(1):1–11, 2002.
- [10] L. J. Croner and T. D. Albright. Image segmentation enhances discrimination of motion in visual noise. *Vision research*, 37(11):1415–1427, 1997.
- [11] J. E. Cutting, C. Moore, and R. Morrison. Masking the motions of human gait. *Perception & Psychophysics*, 44(4):339–347, 1988.
- [12] J.-P. de la Croix and M. Egerstedt. Controllability characterizations of leader-based swarm interactions. In *AAAI Fall Symposium: Human Control of Bioinspired Swarms*, 2012.
- [13] F. Ducatelle, G. A. Di Caro, and L. M. Gambardella. Cooperative self-organization in a heterogeneous swarm robotic system. In *Proceedings of the 12th annual conference on Genetic and evolutionary computation*, pages 87–94. ACM, 2010.
- [14] M. Egerstedt and X. Hu. Formation constrained multi-agent control. *Robotics and Automation, IEEE Transactions on*, 17(6):947–951, Dec 2001.
- [15] M.-A. Fields, E. Haas, S. Hill, C. Stachowiak, and L. Barnes. Effective robot team control methodologies for battlefield applications. In *Intelligent Robots and Systems, 2009. IROS 2009. IEEE/RSJ International Conference on*, pages 5862–5867. IEEE, 2009.
- [16] V. Gazi and K. M. Passino. Stability analysis of social foraging swarms. *Systems, Man, and Cybernetics, Part B: Cybernetics, IEEE Transactions on*, 34(1):539–557, 2004.
- [17] E. Grossman and E. Cooke. Manual control of slow-response systems. *Ergonomics: Major Writings*, page 281, 2005.
- [18] G. Johansson. Visual perception of biological motion and a model for its analysis. *Perception & psychophysics*, 14(2):201–211, 1973.
- [19] Z. Kira and M. A. Potter. Exerting human control over decentralized robot swarms. In *Autonomous Robots and Agents, 2009. ICARA 2009. 4th International Conference on*, pages 566–571. IEEE, 2009.
- [20] G. Knoblich and R. Flach. Predicting the effects of actions: Interactions of perception and action. *Psychological Science*, 12(6):467–472, 2001.
- [21] A. Kolling, S. Nunnally, and M. Lewis. Towards human control of robot swarms. In *Proceedings of the seventh annual ACM/IEEE international conference on human-robot interaction*, pages 89–96. ACM, 2012.
- [22] B. R. Levinthal and S. L. Franconeri. Common-fate grouping as feature selection. *Psychological science*, page 0956797611418346, 2011.
- [23] S. Mau and J. M. Dolan. Scheduling to minimize downtime in human-multirobot supervisory control. 2006.
- [24] J. McLurkin, J. Smith, J. Frankel, D. Sotkowitz, D. Blau, and B. Schmidt. Speaking swarmish: Human-robot interface design for large swarms of autonomous mobile robots. In *AAAI Spring Symposium: To Boldly Go Where No Human-Robot Team Has Gone Before*, pages 72–75, 2006.
- [25] P. Mitchell, M. Cummings, and T. Sheridan. Management of multiple dynamic human supervisory control tasks. In *10th International Command and Control Research and Technology Symposium*, 2005.
- [26] N. Moray, P. Lootsteen, and J. Pajak. Acquisition of process control skills. *Systems, Man and Cybernetics, IEEE Transactions on*, 16(4):497–504, 1986.
- [27] N. M. Morris and W. B. Rouse. The effects of type of knowledge upon human problem solving in a process control task. *Systems, Man and Cybernetics, IEEE Transactions on*, (6):698–707, 1985.
- [28] S. Nagavalli, L. Luo, N. Chakraborty, and K. Sycara. Neglect benevolence in human control of robotic swarms. In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 6047–6053, May 2014.
- [29] D. R. Olsen and M. A. Goodrich. Metrics for evaluating human-robot interactions. In *Proceedings of PERMIS*, volume 2003, page 4, 2003.
- [30] W. Ren, R. W. Beard, and E. M. Atkins. A survey of consensus problems in multi-agent coordination. In *American Control Conference, 2005. Proceedings of the 2005*, pages 1859–1864. IEEE, 2005.
- [31] C. W. Reynolds. Flocks, herds and schools: A distributed behavioral model. *ACM SIGGRAPH Computer Graphics*, 21(4):25–34, 1987.
- [32] S. Runeson and G. Frykholm. Visual perception of lifted weight. *Journal of Experimental Psychology: Human Perception and Performance*, 7(4):733, 1981.
- [33] W. M. Spears and D. F. Spears. *Physicomimetics: Physics-based swarm intelligence*. Springer, 2012.
- [34] A. Steinfeld, T. Fong, D. Kaber, M. Lewis, J. Scholtz, A. Schultz, and M. Goodrich. Common metrics for human-robot interaction. In *Proceedings of the 1st ACM SIGCHI/SIGART conference on Human-robot interaction*, pages 33–40. ACM, 2006.
- [35] F. Stürzel and L. Spillmann. Perceptual limits of common fate. *Vision research*, 44(13):1565–1573, 2004.
- [36] M. Turpin, N. Michael, and V. Kumar. Decentralized formation control with variable shapes for aerial robots. In *Robotics and Automation (ICRA), 2012 IEEE International Conference on*, pages 23–30. IEEE, 2012.
- [37] W. R. Uttal, L. Spillmann, F. Stürzel, and A. B. Sekuler. Motion and shape in common fate. *Vision Research*, 40(3):301–310, 2000.
- [38] P. Walker, S. Nunnally, M. Lewis, A. Kolling, N. Chakraborty, and K. Sycara. Neglect benevolence in human control of swarms in the presence of latency. In *Systems, Man, and Cybernetics (SMC), 2012 IEEE International Conference on*, pages 3009–3014. IEEE, 2012.
- [39] S. N. Watamaniuk. Ideal observer for discrimination of the global direction of dynamic random-dot stimuli. *JOSA A*, 10(1):16–28, 1993.
- [40] S. N. Watamaniuk, S. P. McKee, and N. M. Grzywacz. Detecting a trajectory embedded in random-direction motion noise. *Vision research*, 35(1):65–77, 1995.