Enabling Human Interaction with Bio-Inspired Robot Teams: Topologies, Leaders, Predators, and Stakeholders

BYU-HCMI Technical Report 2011-1

Michael A. Goodrich, P.B. Sujit, Sean Kerman, Brian Pendleton, and José Pinto

August 2, 2011

Abstract

This report formalizes the problem of human-interaction with bio-inspired robot teams (Hu-BIRT) and then systematically explores properties of various biological collectives that allow human interaction. The goal is to identify models, topologies, and control strategies that allow a decentralized agent collective to be appropriately influenced by a human. Evidence is provided that nearest neighbor topologies are more likely to be more cohesive than metric-based topologies but are less expressive. Importantly, randomly dynamic topologies appear to be as expressive as metric-based topologies and open up the possibility of probabilistically connected topologies, but this is not explored further in this report. Evidence is also provided that when nominal agents use switching-based control to respond to predators and leaders, the resulting behavior can be effectively influenced by a human and that this influence is more effective when using leaders than predators; however, the influence is obtained at the cost of changing the dynamic structure of the collective. Using non-switching controllers allows a human to influence the collective while retaining the collective's structure, but this requires more than one leader and risks fragmenting the collective. Finally, evidence is presented that groups of stakeholders, which seek both to influence and to be influenced by other agents, can sustain influence over certain types of collectives more effectively than cohorts of leaders.

1 Introduction

Two current research areas are receiving considerable attention in the recent literature: human-robot interaction (HRI) and bio-inspired robot teams (BIRT). Although HRI includes work on social inter-

action, this report emphasizes the design of HRI systems that help humans manage multiple remote robots [28, 21]. This type of HRI work expands decades-old work on supervisory control [14, 32].

Although bio-inspired robotics includes designing swimming robots or believable humanoids, this report emphasizes BIRT work on identifying principles and practices of biological societies [36]. After identifying those principles and practices, much BIRT research seeks to abstract and encode these principles in robots [31]. The resulting teams demonstrate so-called collective intelligence wherein simple individual robot behaviors produce colony-wide behaviors that appear collectively purposeful and goal-directed. Typical behaviors include swarming, flocking, foraging, and colony-building. This collective behavior tends to be robust to variations among robots even in the presence of low reliability.

Steinberg identified human-interaction with bio-inspired systems as an important research area for producing responsive, robust systems for complex surveillance and reconnaissance problems [35]. Research that combines elements of HRI with BIRT should yield robot teams that can be efficiently managed by humans but that retain robust qualities in the presence of unreliability. Such research has been called cooperative robotics and human-swarm interaction, but we will call the work in this paper human-BIRT (HuBIRT) to emphasize human-centered interaction with not only swarms but also more sophisticated bio-inspired collectives such as colonies.

2 Related Literature

Lewis et al. provide a useful framework for understanding systems with varying ratios of robots to humans, and therefore for several types of human-robot teams [40]. They identified three broad classes of systems, categorized by the complexity of a single human managing M robots. The first class of systems could be labeled as $O(1)$, meaning that a single human can manage a collective in constant time. This class is exemplified by centralized architectures wherein a human issues a command that is then executed by a team of robots. A simple example of this class is leading a fixed formation of robots wherein the human controls one robot and all other robots fall into an appropriate formation.

The second class of systems identified by Lewis could be labeled as $O(M)$, meaning that the complexity of managing M robots scales linearly with the number of robots. This class is often associated with the so-called fan-out metric [30, 10] which represents the number of independent, homogeneous robots that can be managed by a single human. A simple example of the $O(M)$ class is when a human schedules attention among M independent robots so that each robot gets input from the human in a timely way.

The third class of systems identified by Lewis could be labeled as $O(>M)$, meaning that the complexity of managing the system grows faster than linearly in the number of robots. This class includes tasks where robots must carefully coordinate execution, but where the human is responsible for managing the coordination.

In light of Lewis' overview of multi-robot systems managed by humans, it is useful to review the

state-of-the-art for real systems. Cummings et al. [11] provide an overview of several existing human-robot systems. They found that fan-out for existing systems ranges widely: from three or more humans per robot (e.g., Predator [12]) to systems of ten or twelve robots per person (e.g., Tomahawk missile systems [13]). Cummings reports that systems with high fan-out exhibit higher levels of autonomy, with a large jump in fan-out when robot autonomy allowed humans to use "management by exception" instead of "management by consent" [11]; this distinction in management policy splits Sheridan's levels of autonomy, with management by exception typified by robots who take initiative [33].

In terms of current research on human-robot team design, it is useful to begin with work that uses a centralized controller, which could be classified by Lewis as $O(1)$. Miller et al. advocate the use of playbook-style interactions [26], wherein a human calls plays that trigger predictable patterns of behavior. Simple plays, like grouping and searching, have been used to manage several large, simulated teams (50-200 robots) [17]. More complicated plays, like coordinated rendezvous or formation-following, have been applied to smaller teams [24]. The most obvious way to include a human in these approaches is for a human to control a leader agent who then influences other agents through conventional or bio-inspired means [15]. Managing patterns of behavior is another useful way to think about human interaction with bio-inspired teams. Recent work [19, 15, 41, 25] explores centralized methods for influencing these patterns so that a human can guide the team to accomplish effectively some human-specified mission.

In addition to these centralized methods for managing large robot teams, there are approaches that take a more decentralized approach. For example, Barnes et al. explored how to modify potential fields in response to directions from a human operator [3]. Similarly, simulations have been performed of robots herding animals [39, 22], producing an ad hoc collection of observations, such as the fact that herding a flock using a single sheepdog requires very different behaviors than using a team of sheepdogs.

Shifting from conventional robot teams to swarm-like teams, there are examples of so-called *human-swarm interaction* (HSI). Bashyal and Venayagamoorthy [4] presented an HSI approach that provided a human with a partial plan and global information, and then allowed the human to adjust the autonomy of a small subset of swarm members to influence swarm behavior. Although the paper presents a user study and an *ad hoc* set of desirable features, perhaps the most useful contribution of the paper is the observation that "The ideal man-machine interaction is ... one that functions autonomously while providing users with a method to inject knowledge and guidance so as to improve [system performance]." The authors take this design philosophy to an important extreme where a human's "control over the swarm is only as much as that of any other member in the swarm," meaning that the human may be able to control a single individual agent and thereby influence swarm behavior, but not perform any centralized control or exert any global influence on the entire swarm.

The GUARDIANS project seeks to use swarm robotic technology to support firefighters [1, 29]. These papers present use-cases for HSI as well as an artificial potential field-based implementation. The robots include the ability to autonomously respond to obstacles and to follow either a proximate human or a remotely controlled virtual avatar through an environment. This ability of a robot swarm

to track a human may be unique in the literature. Although these papers claim that a stability analysis has been performed, it seems more accurate to say that some simulations have been performed that demonstrate successful HSI.

It is useful to note that there are many models of how individual agent rules produce flocking and swarming behavior and that many of these models use some combination of switching controllers and additive dynamics used in this paper [7, 8, 20, 5]. Note, however, that at least one paper on wolf pack behavior does not use a discrete-time, dynamical system model but rather a finite probabilistic state machine model [23].

Since much of the literature reviewed in this section is limited to *designing robot behaviors*, it is useful to conclude this section with a note on technologies that *support human behavior* in managing multiple robots. Graphical user interfaces for controlling teams of independent robots have been studied and common design themes are emerging [17, 14]. Many of the design elements are similar to designs used in multi-agent video games, and this could be a fruitful area for designing HuBIRT systems [6].

3 Formalism

In this section, we present a formalism that is appropriate for capturing many types of flocking and swarming behaviors. We begin by reviewing some general principles of collective intelligence for swarms and other collectives, and then provide a formal model to encode many of these principles while simultaneously allowing human input. Although the model includes many existing swarm and multi-agent dynamics, we will impose some simplifying assumptions that we will evaluate in the remainder of this report.

3.1 Bio-Inspired Principles

Sumpter identified several principles that describe biological systems that exhibit collective behavior [36]. In this subsection, we briefly describe and label a subset of Sumpter's principles. In the next subsection, we will use these brief descriptions as the basis for a basic discrete-time state space dynamical HuBIRT model.

1. **Positive feedback** is described as "imitation or recruitment behaviour [which] continues [until] the number [of agents] performing an activity explodes exponentially. An isolated behaviour is quickly subsumed by a mass of similar behaviors .. [yielding a set of] collective patterns." To implement this, robots would require receptors that would allow signals from other robots to make it more likely to imitate a behavior.

2. **Negative feedback** is a counterbalance to positive feedback in that "positive feedback builds up a collective pattern ... [and] negative feedback ... stabilizes it. ... Negative feedback leads to

homeostasis, stable output in the face of varied input [via, for example] ... smooth adjustment [to the presence of changes]." To implement this, robots would require receptors that would allow environmental signals to discourage particular behaviors.

3. **Inhibition** is negative feedback that arises from other robots rather than from the environment in that members "of a group exhibiting one type of behaviour can inhibit the behaviour of others. When this inhibition is passive it is indistinguishable from negative feedback. ... Inhibition can also be active, whereby members of one group actively try to reduce another type of behaviour." Robots would require the ability to signal other robots and/or to receive signals from other robots that would make certain behaviors less likely or desirable.

4. **Individual integrity** means that each "... of the animals in a group is different, in terms of their ... previous experience [which can, for example, produce] ... different levels of response." Robots would require some ability to represent at least a partial history of their percepts and actions.

5. **Response threshold** make it so that agents can "... change their behavior in response to a stimulus reaching some threshold." Robots would need the ability to sense input from the environment and then switch behaviors based on a threshold.

6. **Leadership** means that "key individuals ... catalyze and organize the group [such as] ... the first [agent to find a resource] ... serving to inform [other agents] that a period of greater activity is about to begin. ... Often leadership is assumed by the possession of particular information." Leadership requires varying roles among robots; some robots must be capable of autonomously initiating a behavior while others must be capable of following.

Note that Sumpter also identifies *individual variability* as a key principle, but in the model below variability is a product of signal and feedback, not a property of its own. Similarly, Sumpter identifies a *redundancy* principle based on a surfeit of individuals so that there are multiple agents capable of performing any single needed function; redundancy will be useful in this report, but we will only use it to support HuBIRT by considering how a human may need to directly influence multiple leaders so that a redundant set of leaders sustains influence over the group. Sumpter continues and identifies a *synchronization* principle wherein agents adjust the frequency of a behavior so that coordination evolves in time rather than space; synchronization is beyond the scope of this report. Finally, Sumpter identifies a *selfishness* principle which means that "individuals should act to maximize their own fitness." We treat this as a form of bounded rationality, and restrict implementations of agents to simple form.

3.2 Mathematical Model

Sumpter's principles provide descriptive guidelines for what to look for in a system that exhibits collective intelligence, but they lack a formalism that could help HuBIRT design. However, Sumpter's

principle of *integrity* suggests a state-space dynamic model where an agent's state is an encoding for previous "experience." Denote agent i's state at time t as x_t^i, and the vector of $N > 1$ agent states as $\mathbf{x}_t$. The precise encoding of state varies from model to model, but typical models include some notion of position and velocity. There are several specific state-based dynamical systems models in the literature, so we will use Sumpter's principles to guide us to a general form of several of these models [7, 8, 20, 5].

Sumpter's principles of *positive feedback*, *inhibition*, and *response threshold* suggest how the states of other agents should affect agent i, and *negative feedback* and *response threshold* suggest how external signals should affect agent i. Recall that x_t^i denotes agent i's state, and let u denote an external signal; then we can write the state dynamics model as

$$\mathbf{x}_{t+1} = f(\mathbf{x}_t) + g(\mathbf{x}_t, u_t). \tag{1}$$

In terms of Sumpter's principles, Equation (1) encodes positive feedback and inhibition in $f(\mathbf{x}_t)$, and encodes negative feedback in $g(\mathbf{x}_t, u_t)$. In general, response thresholds can apply to either f or g, but this report will encode them only in f. These will be encoded using switching-based control in which the behavior of an agent can switch from one controller to another depending on whether other agents are within specific threshold distances.

A key element of bio-inspired teams is that *collective intelligence emerges as a function of local interactions.* This manifests itself in a set of sparseness properties on f and g, some of which are general and some of which are restrictions to the small set of HuBIRT interactions evaluated in this report. Consider agent i and let $\neg i$ denote all agents other than agent i. Extracting the dynamics from i^{th} agent in Equation (1) yields:

$$x_{t+1}^i = f^i(\mathbf{x}_t) + g^i(\mathbf{x}_t, u_t). \tag{2}$$

Partitioning $\mathbf{x}_t$ into x_t^i and $\mathbf{x}_t^{\neg i}$ yields,

$$x_{t+1}^i = f^i(x_t^i, \mathbf{x}_t^{\neg i}) + g^i(\mathbf{x}_t, u_t). \tag{3}$$

Since collective intelligence emerges from local interactions in BIRT, we will assume that the set of agents that affect agent i is small, that is, localness properties mean that only a handful of agents in $\mathbf{x}_t^{\neg i}$ affect x_{t+1}^i. We call this assumption the *locality* assumption. Let A_t^i represent the adjacency matrix of the graph induced by $f^i(x_t^i, \mathbf{x}_t^{\neg i})$. We call A_t the *cohesiveness matrix* because it represents spatio-temporal correlation of interagent dynamics that determine the way collective intelligence of the team changes over time. Note that the performance of both artificial and biological systems depend critically on the structure of A_t [38, 2]. Ballerini et al. [2] emphasize the importance of cohesion in real biological systems: "The main goal of the interaction among individuals is to maintain cohesion of the group." Furthermore, they add "Structure is the foremost effect of interaction." The cohesiveness matrix encodes such structure, created by the local interactions of agents.

In this paper, we will explore two different model types, one biomimetic (meaning biologically inspired) [8] and another physicomimetic (meaning inspired by physics-based interactions among particles) [34]. Neither biomimetic nor physicomimetic models explicitly include operator influence, but

we find it useful to include human influence in the same way that we treat environmental influence, namely from u through the g function. We call this assumption the *influence assumption* because it allows us to treat human influence similarly to environmental influence, and separate both of these from inter-agent influence. Thus, g is added to f.

Thus, the inputs to g includes the vector u_t, which denotes all external signals, both those that come from the environment and those that may be specified by an operator. Thus, we separate u_t into two components: a subvector from the operator, u_t^{op}, and a subvector from the environment, u_t^{env}. Thus, we note that $g^i(\mathbf{x}_t, u_t) = g^i(\mathbf{x}_t, u_t^{\text{op}}, u_t^{\text{env}})$. Future work will explore how teams can include autonomous responses to environmental signals, but in this report we assume that $g^i(\mathbf{x}_t, u_t^{\text{op}}, u_t^{\text{env}}) = g^i(\mathbf{x}_t, u_t^{\text{op}})$.

We next introduce an assumption that is useful for modeling useful swarming and flocking behaviors, but that appears to be too restrictive for more expressive collective behaviors. Assume that $g^i(\mathbf{x}_t, u_t^{\text{op}})$ is a function only of state x_t^i, that is, that external signals affect the next state of agent i only as a function of x_t^i and the external signals u_t. We call this the *autonomy assumption* because it says that agents can autonomously make decisions as a function only of their state and the signals that they receive independent of what other agents are doing. This yields, $g^i(\mathbf{x}_t, u_t^{\text{op}}) = g^i(x_t^i, u_t^{\text{op}})$.

In the next section, we will formally define another matrix, called the management matrix and denoted by B_t, that encodes the limited number of agents that a human can directly influence by controlling a leader agent. The idea behind the management matrix is that the human can influence a leader agent through g and then the leader agent influences the other agents through f; B_t encodes how the leader agent influences other agents through f.

Problems for which HuBIRT solutions are appropriate can have many different types of communication structures. We are most interested in problems where the human cannot receive information from all agents simultaneously, either because of communication limitations or because there are too many agents for a human to manage. Thus, we assume that the communication model restricts the number of agents that can be observed by the human. This can be formally encoded as

$$y_t = C_t \mathbf{x}_t,$$

where C is sparse. Although we will not do anything more with this equation in the rest of this report, this formalism will be useful in future work.

The remainder of the report will study variations of the following model

$$x_{t+1}^i = f^i(x_t^i, \mathbf{x}_t^{\neg i}) + g^i(x_t^i, u_t^{\text{op}}). \tag{4}$$

We also make two simplifying assumptions about the communications model, but note that future work should address these. First, we assume that local communication between agents is perfect. Second, we assume that the human can observe the locations of all agents all the time; i.e., we assume that C_t is a complete topology.

3.3 Measuring Collective Behavior: Performance, Cohesion, and Responsiveness

This report takes an empirical rather than a theoretical approach to evaluating HuBIRT performance. One purpose of this report is to analyze how certain (A_t, B_t) structures are correlated with cohesive and responsive collective behavior . Thus, we measure properties of (A_t, B_t) and correlate those properties to HuBIRT performance. One of the main results from decentralized control theory is that a fully connected group of agents often possesses strong collective properties like stability and cohesion [38]. Connectivity, in this context, refers to the connectivity encoded in the cohesiveness matrix and not in the management matrix. These results are important to robust HuBIRT performance because collective intelligence has been identified as the evolution of structure and cohesiveness as a time-invariant property of this structure.

Although connectivity in A_t appears to be necessary for cohesive behavior, it is not sufficient for HuBIRT. Cohesive and robust behavior must also be responsive to human input so that a human can direct collective behavior. Thus, HuBIRT performance, including responsiveness and cohesiveness, are encoded in (A_t, B_t). We now consider structures of these matrices and consider how these evolve as a function of different organization and management styles.

3.3.1 Inserting Human Influence: Leaders and Predators

For HuBIRT, we can use leadership principles to allow a human to influence team behavior. According to Sumpter, **leadership** means that "key individuals ... catalyze and organize the group." In bio-inspired robot teams, two leadership models have been studied by others: lead-by-attraction and lead-by-repulsion. Lead-by-repulsion is more commonly referred to as *predation* where the leader is a predator and agents are prey, so the leader influences the behavior of the agents by pursuing them. By contrast, lead-by-attraction is often associated with the colloquial use of the word leadership, meaning that a leader is one that gets ahead of a group and the group follows. For simplicity, we will call lead-by-repulsion models *predator models* and lead-by-attraction models *leader models*. Note that we will shortly introduce a third style of leadership that we call a *stakeholder* style. We describe this style when we introduce it.

The human influences a subset of agents in the collective via $\{g_t^i(x_t^i, u_t^{\text{op}}),\ i = 1, \ldots, N\}$, and then these agents influence other agents via f. For all but the stakeholder style, agents will either be influenced by the human or by other agents but not by both. Thus, agents will either be leader/predator agents or they will be nominal agents. For leader/predator agents, $f = 0$ and $g \neq 0$; for nominal agents, $f \neq 0$ and $g = 0$. Let B_t denote the adjacency matrix induced by the connectivity between leader/predator agents on one hand, and nominal agents on the other. We call B_t the *management matrix* because it encodes the graph that represents which agents can be directly influenced by the actions of the operator via leader/predator intermediaries. Note that B_t will often be sparse because,

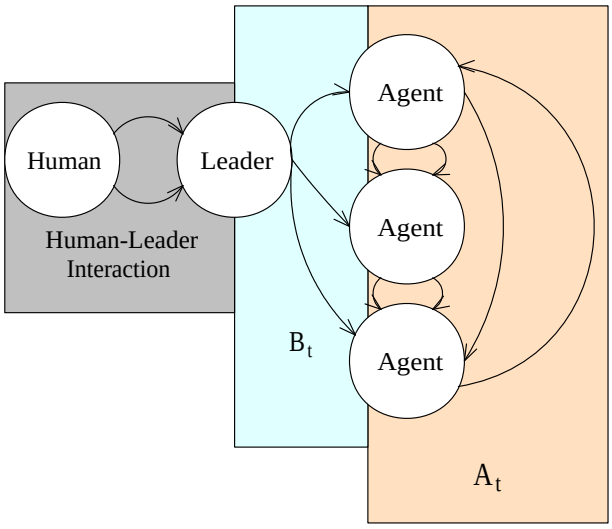

Figure 1: Leader-mediated HuBIRT

in keeping with limited inter-agent influence ranges, leaders/predators will only influence a subset of agents.

For the model in Equation (4), g encodes the relationship between human input and the set of leaders/predators, (f, A_t) captures interagent relationships, and B_t encodes the relationship between the set of leaders/predators and the set of nominal agents; see Figure 1.

As shown below, leader models have the desirable property that they tend to make the B_t matrix vary slowly in time; agents near the leader are attracted to the leader and tend to stay close enough to the leader to stay under the leader's influence. This allows the leader to sustain influence over the collective.

By contrast, predator models tend to make the B_t matrix vary quickly in time; agents near the predator are repelled by the predator and try to get away from the predator's sphere of influence. Thus, as the predator moves, many agents escape from the sphere of influence, making it difficult for a predator to sustain influence over the collective.

This suggests that HuBIRT organizations should be designed so that coherence is maintained, for example, through a fully connected A_t matrix and that influence is maintained, for example, through a slowly-varying B_t matrix.

Note, however, that strong spatio-temporal coherence may not always be desirable. Two examples will serve to illustrate this point. First, when a HuBIRT problem has multiple goals, it may be useful to split the group into two subgroups, each of whom then performs a different task. Predator models are likely to be more efficient than leader models for "decohering" the group (although a team of highly coordinated leaders may be able to do the same thing). For highly time-varying, task-rich problems, it

may be desirable to use multiple operators who may sometimes split a group to allow broader reach. Second, there are times when it is desirable to temporarily break local interactions among neighbors, such as when a flock or swarm needs to avoid a dangerous location as they travel to a goal. Predator models can be used to cause the swarm to continue to move forward but form a "safety pocket" around an undesirable region. Future work will explore decoherence.

3.3.2 Cohesiveness

Under HuBIRT, neighborhoods of any agent can be dynamic. We are interested in measuring how agents sustain influence amongst themselves even when neighborhoods are dynamic. To help measure sustained inter-agent influence, consider the time history of A_t over some temporal window,

$$\mathcal{A}_t = \sum_{k=0}^{T} A_{t+k}. \tag{5}$$

The i^{th} row of $\mathcal{A}_t(i)$ is a histogram of which of agent i's neighbors influence agent i during the time interval $[t, t+T]$.

$\mathcal{A}_t$ encodes a time-varying histogram of the cohesion matrix A_t. Cohesiveness manifests itself in $\mathcal{A}_t$ in two ways: First, if all agents interact with most other agents almost all the time, then $\mathcal{A}_t$ will be "tall and uniform". Second, if each A_t is sparse, indicating just local interactions, then $\mathcal{A}_t$ should be "lumpy", meaning that agents tend to influence the same neighbors for multiple time steps.

3.3.3 Manageability

To support HuBIRT, a subset of agents should be influenced by the human, and this subset should exert sustained influence over their neighbors, and so on until the entire collective is influenced. Since $\mathcal{A}_t$ is a function only of the cohesiveness matrix, A_t, it doesn't explicitly represent the responsiveness of a bio-inspired team; rather, it represents a *potential* to be collectively responsive. In order for this potential to be fulfilled, the collective must also be responsive to the human's input. Since responsiveness is a function of B_t, we can let

$$\mathcal{B}_t = \sum_{k=0}^{T} B_{t+k} \tag{6}$$

denote the histogram of what agents the human is influencing through leaders/predators. Sustained influence of a human over agents will be indicated by a "lumpy" or "pointy" visualization of $\mathcal{B}_t$. We will explore what properties of $\mathcal{B}_t$ will encode responsiveness.

3.4 A Bio-Inspired Task: Information Foraging

We will use *information foraging* to test HuBIRT performance. An information foraging problem is one where there are multiple tasks, represented as abstract resources, that appear at unknown locations in a spatial domain. Agents must discover the tasks and then assign a subset or subteam of the agents to perform the task. Each task takes time to complete, meaning that multiple agents must persist in the task for a satisfactory period of time before moving to another task.

The resource size depletes at a rate based on the density of the number of assigned agents, that is, resource j depletes according to

$$S_j(t+1) = S_j(t) - \hat{N}s \tag{7}$$

where $\hat{N}$ represents the number of agents present less than r_s meters from the location of the resource and $s > 0$ represents the amount of resource to be reduced per agent. For the experiments, we used $s = 0.001$ and $r_s = 5$. New tasks can appear anywhere in the domain at any time. Bio-inspired agents are capable of performing some aspects of this task by themselves, but are generally inefficient at the task without having some kind of human input.

Although information foraging as encoded in Equation (7) is abstract, it represents several types of relevant tasks. The most obvious example is target prosecution that requires the persistent output of multiple agents to suppress, for example, enemy air defenses. IED discovery and de-mining are also similar to information foraging. Importantly, information foraging fits certain types of ISR[1] tasks such as maintaining persistent observation of a location of interest from multiple angles.

These examples illustrate some ecological validity of the information foraging task, but these examples are compatible with only certain types of bio-inspired behavior. Swarming and flocking behaviors lend themselves to collective foraging, but other collective structures are more appropriate for tasks like perimeter monitoring of the execution of multiple, spatially separate tasks. In the next section, we will identify a set of collective structures, and we will then explore cohesiveness and manageability of swarm and flock structures before discussing influence of other structures.

4 Cohesiveness and Expressiveness of Collective Structures

In this section, we explore what types of topological organizations and dynamic models yield what types of collective structures. We will explore both a biomimetic and a physicomimetic dynamic model[2].

[1] Intelligence, Surveillance, Reconnaissance.

[2] We recognize that the physicomimetic model is not bio-inspired and, consequently, that it doesn't technically deserve to be treated in a technical report on bio-inspired models. Nevertheless, since we are interested in a wide range of bio-inspired models, it is useful to include a different model in our evaluation. Moreover, the physicomimetic model is very similar to the bio-mimetic model assuming that agents are not constantly moving, it is also similar to another bio-inspired model recently appearing in the literature [27], and it appears often in literature on swarm behavior.

We will also compare metric-based and nearest-neighbor topologies. A metric-based topology is constructed by connecting to vertex x all other vertices y for which the Euclidean distance between x and y is less than some threshold in meters. A nearest-neighbor topology is constructed by connecting to vertex x the k closest vertices, regardless of their Euclidean distance. Ballerini et al. [2] gathered evidence that organizations based on topological distances better explain the highly cohesive behaviors of flocks of sparrows than organizations based on metric distances. This type of observation motivates the need to compare the different types of topologies.

Note that probabilistic topologies have also been explored based on their relevance to coordination in biological collectives [5]. Future work will explore this topology too.

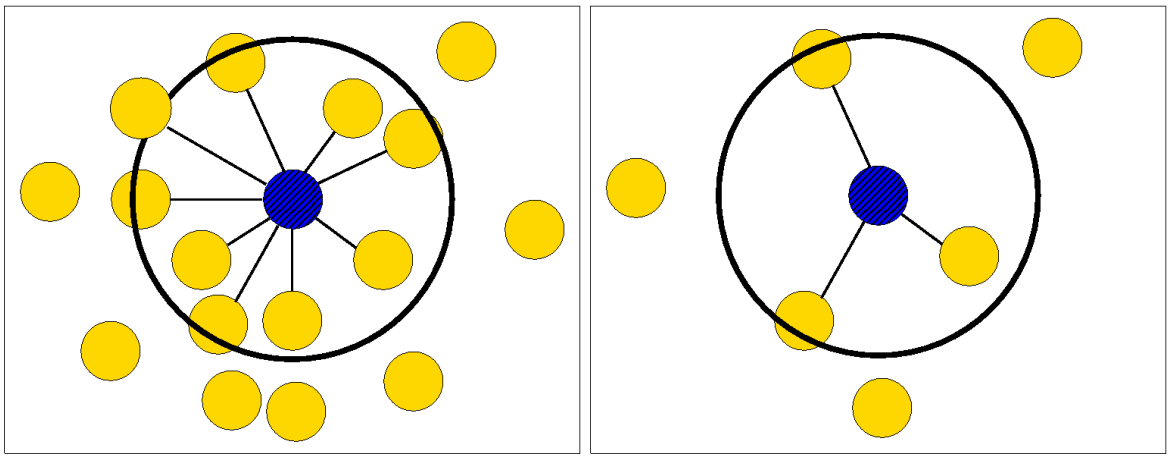

(a) (b)

Figure 2: For a metric-based topology, the blue agent is influenced by all yellow agents within a fixed radius for (a) densely clustered agents and (b) sparsely clustered agents.

Figure 2 illustrates how a metric-based topology is formed. An agent is influenced by all other agents within a fixed Euclidean radius. Note that if all agents have the same influence radius then this topology is symmetric. By contrast, Figure 3 illustrates how a nearest-neighbor topology is formed. In this topology, an agent is influenced by the k nearest neighbors, regardless of the actual physical distance. Note that this topology may not be symmetric, meaning that agent i may be influenced by agent j but not vice versa.

It is important to note that these two types of topologies correspond to different inter-agent communication models. *Metric-based topologies are power-limited*, meaning that the ability to sense and respond to other agents depends on how easy it is to sense other agents. Assuming that sensing capability is a function of the power of an agent's communication signal, there are natural distances beyond which agents cannot sense each other. By contrast, *nearest neighbor topologies are bandwidth-limited*,

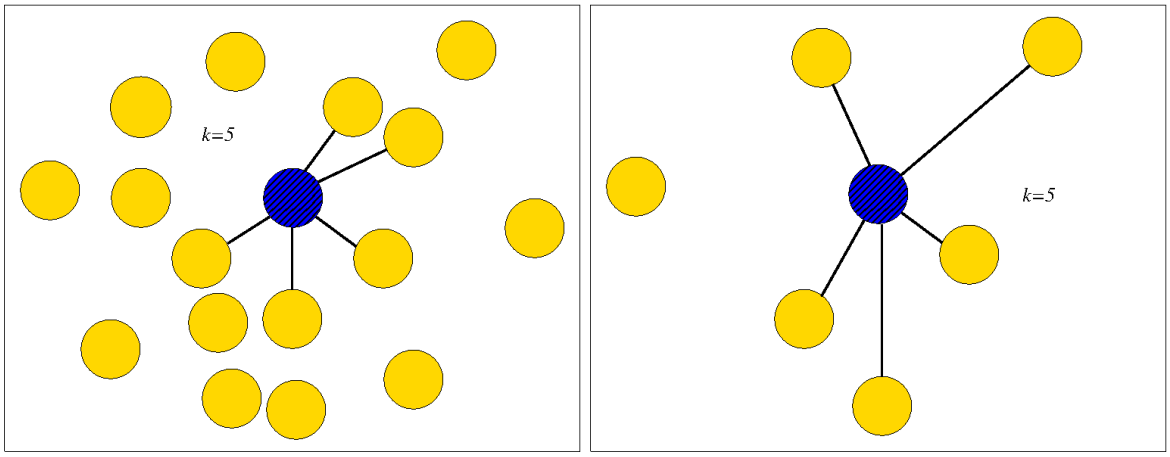

(a) (b)

Figure 3: For a nearest neighbor topology, the blue agent is influenced by the $N = 5$ nearest yellow agents for (a) densely clustered agents and (b) sparsely clustered agents.

meaning that the ability to sense and respond to other agents depends on a partitioning of a communication channel into separate bands, and an agent can only sense a fixed number of agents because only so much information can cross this band.

4.1 Cohesiveness under Static Connectivity

In this section, we want to know which is more likely to be connected, a metric-based or nearest neighbor topology. We evaluate this by randomly placing 100 agents in a 100 by 100 grid and then creating topologies using either nearest neighbors or metric neighborhoods. We then compare the likelihood that randomly placed agents form a connected topology for the two topology-construction methods.

Question 1: When is a nearest-neighbor topology connected? The first experiment randomly distributed 100 agents in a 100 by 100 grid and then created edges between the agents by finding the n nearest neighbors. We repeated this multiple times, varied the number of neighbors, and then determined what percentage of graphs are connected.

Results for $N = 100$ agents randomly placed on a 100×100 grid are shown in Figure 4. Note that since nearest-neighborhoods do not depend on the scale of the area, the results shown in Figure 4 are scale-invariant, that is, they apply to not only the 100×100 grid but any $\ell \times \ell$ grid for $\ell \in \Re^+$.

These results indicate that once an agent is capable of sensing and responding to seven or eight neighbors, the majority of the resulting graphs are connected, even for randomly placed agents.

Question 2: When is a metric-based topology connected? We repeated the above simulation but construct the edges using various distance thresholds. We evaluated the probability that

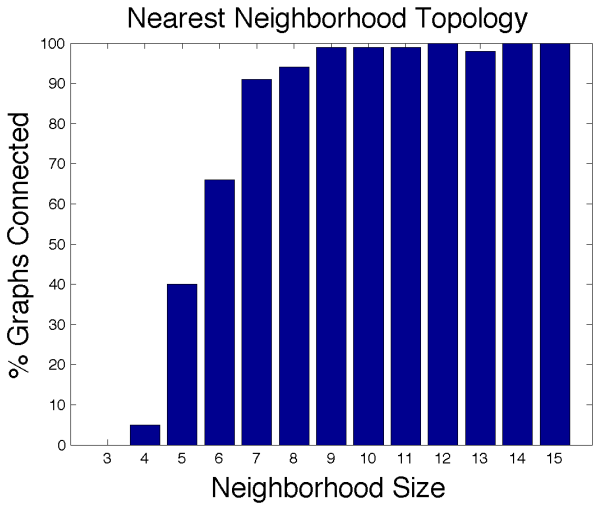

Figure 4: Percentage of randomly created graphs with edges produced using nearest-neighbors that are connected.

randomly-placed agents produced a connected graph as a function of the distance thresholds. We also measured the average valence of each node in the graph, because this gave us an estimate of the average neighborhood size for each agent. It also helped us to compare the results from the metric-based method with the results from the nearest neighbor method.

Results for $N = 100$ agents randomly placed on a 100×100 grid are shown in Figure 5. To make the results scale-invariant, the distance thresholds were set to a percentage of the total scale, so $\rho = 15$ implies that all agents within $d = \rho \times \text{SCALE}$ are connected to each other. The results shown in Figure 5 are shown for SCALE= 100.

Question 3: Are metric-based or nearest-neighbor topologies more likely to be connected? Insight can be gained about the two topology structures by identifying a way to compare them. Nearest-neighborhood topologies are regular, that is, all vertices have the same valence. By contrast, metric-based topologies are not regular and the valence of any given vertex depends on the random placement of other vertices in the space. For each distance threshold percentage, ρ, computing the median valence from the set of vertices yields Figure 6(b). Thus, Figure 6(b) gives an idea of the percentage of graphs that are connected as a function of the average neighborhood size for metric-based topologies. Figure 6(a), replicated from above, shows the percentage of graphs that are connected as a function of neighborhood size; comparing the two indicates that metric-based topologies are less likely to be connected than nearest neighbor topologies.

The results suggest that valence is correlated with connectivity, but that the valence variability induced using the metric-based method causes topologies that have a smaller probability of being

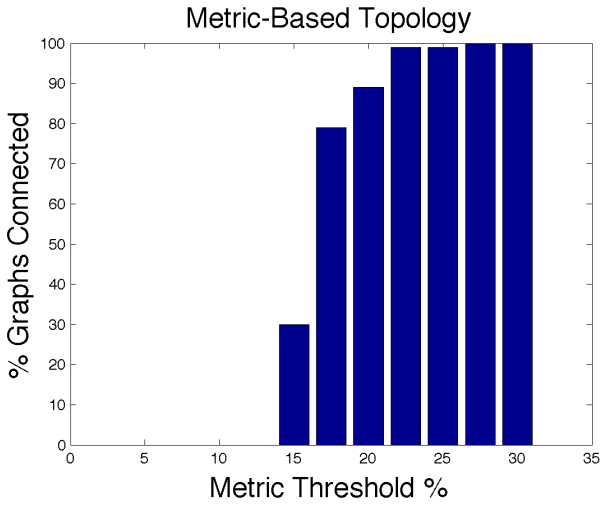

Figure 5: Percentage of randomly created graphs with edges produced using metric thresholds that are connected.

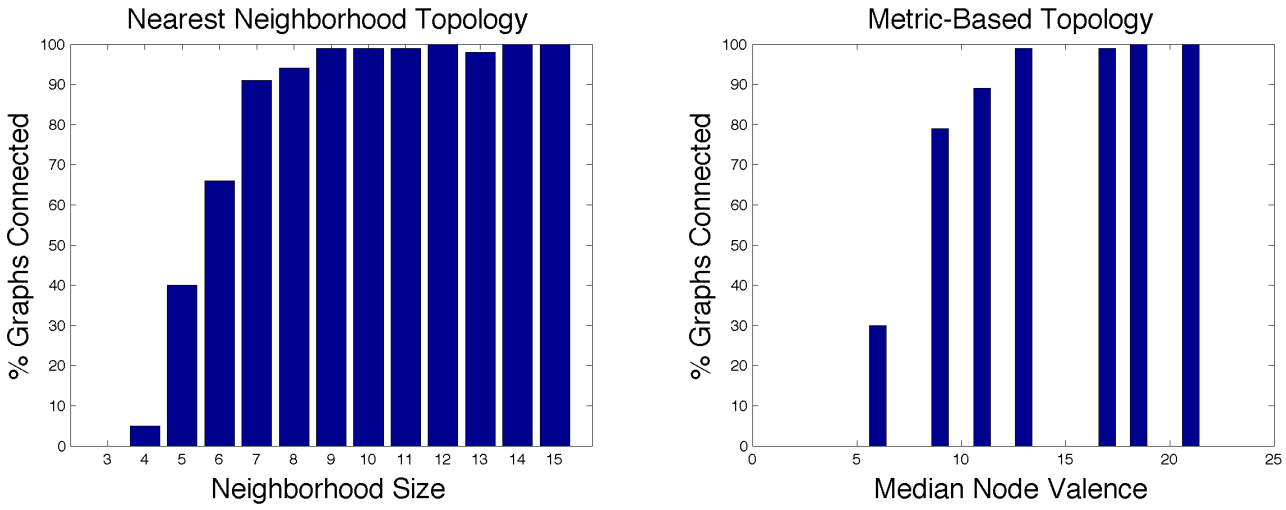

(a) (b)

Figure 6: Percentage of randomly created graphs as a function of valence for edges produced using (a) nearest-neighbor topologies and (b) metric-based topologies (results shown for median valence).

connected. Simply put, nearest neighborhood topologies are more likely to be connected than metric-based topologies. This will be important to HuBIRT when we explore how easy it is for a human to manage a team of agents that communicate through metric-based or nearest neighborhood topologies. We will show that it is easier for a human to manage a flock when the flock uses nearest neighborhood topologies.

4.2 Cohesiveness and Expressiveness of a Biomimetic Model

Evaluating randomly created, static graphs gives us a sense that nearest neighborhood topologies will be more likely to stay connected when the agents move. In this section we perform a series of simulation experiments to test this hypothesis. We begin with Couzin's model of fish-schooling behavior [8]. We will then move beyond this hypothesis and ask whether nearest neighbor topologies and metric-based topologies can perform the same set of tasks using the two topology types. More specifically, we will explore whether both methods of constructing topologies exhibit the same collective structures. We will explore five structures identified in the literature: swarms, toroids, dynamic parallel groups, highly parallel groups, and packs. For the number of agents that we consider, packs will emerge very rarely and will be unlikely to stay cohesive, so we will not explore them much.

4.2.1 Description of Model

Since Couzin's model is a metric-based model, we will generalize the model to allow both methods of forming the topologies: nearest neighbor and metric-based. Couzin's model defines three zones: repulsion, orientation, and attraction. If any other agent is in the zone of repulsion, other forces are ignored. If there are no other agents in the zone of repulsion, then attraction and orientation forces are applied. These zones, along with the notations used for the zones, are illustrated in Figure 7

In addition to the three zones, R^{rep}, R^{ori}, and R^{att}, Couzin's model has four other critical parameters: noise, blindspot angle, speed, and maximum turning rate. We will denote the blindspot angle by ϕ, the speed by s, and the maximum turning velocity by ω. The full model is given in Appendix A. To summarize the model, an agent computes a desired direction by being repelled from all agents within the zone of repulsion (or aligned to and attracted to all agents within the zones or orientation/attraction) and then moves toward that desired direction by turning no more than ω radians per second.

The simulations use a sampling rate of $\Delta t = 0.2$ samples per second and a noise value of ± 0.02 radians that is added to angular velocity each second. Following Couzin's example, all spatial units are defined with respect to R^{rep} which is set to unity. Speed is defined with respect to this unit as well, that is, $s = 1$ means that an agent moves one repulsion radius per second. In Couzin's model, forward speed is constant for all agents.

Only the blindspot needs further explanation. Figure 8 illustrates how an agent traveling in direction θ can only sense agents that are within $\pm\phi$ radians from the direction of travel.

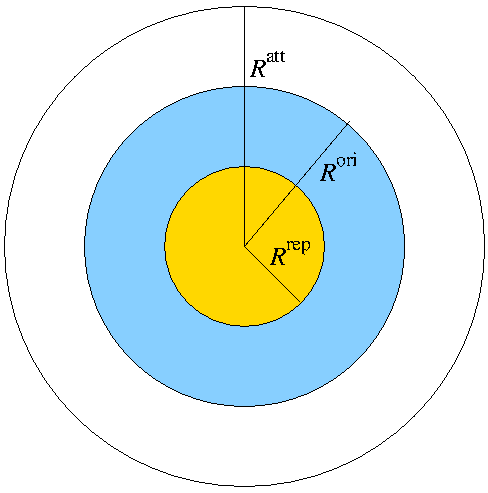

Figure 7: Couzin's zones of repulsion, R^{rep}, orientation, R^{ori}, and attraction, R^{att}.

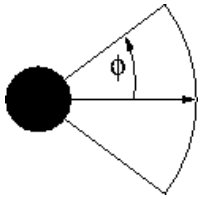

Figure 8: The blindspot is any area outside of $\pm\phi$ radians from the direction of travel.

Couzin's model is a type of metric-based topology since influence of neighborhoods is governed by R^{rep}, R^{ori}, and R^{att}. This is metric-based since all agents within the repelling radius affect the agent, and all agents within the orientation and attraction radius affect the agent (provided that no agents are within the repelling radius). All decisions are based on metric thresholds as in Figure 2. Note, however, that this is not a straightforward metric-based topology since the influence of those agents found within R^{ori} or R^{att} depends on whether or not there is an agent within the repelling neighborhood; if so, then all agents outside of the repelling are ignored, and if not then any agents within the other neighborhoods have influence on agent i. Thus, the controller for each agent switches between repulsion and attraction/alignment, depending on the metric thresholds.

4.2.2 Relating Couzin's Model to Equation (4)

It is useful to explicitly state how Couzin's model relates to Equation (4) so that we can compare the structure of this model to the physicomimetics model, with and without human influence. Before doing so, note that Couzin's model does not include human input, so at this stage of the report we will only

pay attention to the interagent dynamics function encoded in f^i.

The states of Couzin's model are the position vector, $\mathbf{c}_t^i$ and the direction of travel, θ_t^i. Thus,

$$x_t^i = \begin{bmatrix} \mathbf{c}_t^i \\ \theta_t^i \end{bmatrix}.$$

The position vector evolves over time as

$$c_{t+1}^i = c_t^i + \Delta t \begin{bmatrix} s\cos\theta_t^i \\ s\sin\theta_t^i \end{bmatrix},$$

where s is the speed of the agent. The direction vector evolves by first identifying and then tracking the desired direction of travel, $\hat{\theta}_{t+1}^i$. The desired direction of travel is obtained using the switching control law specified in Appendix A and illustrated in Figure 7 subject to the blindspot constraints illustrated in Figure 8. Tracking is then performed using

$$\theta_{t+1}^i = \begin{cases} \hat{\theta}_{t+1}^i & \text{if } |\theta_t^i - \hat{\theta}_{t+1}^i| > \omega\Delta t \\ \text{sign}(\theta_t^i - \hat{\theta}_{t+1}^i)(\theta_t^i - \omega\Delta t) & \text{otherwise} \end{cases}.$$

Naturally, this equation assumes appropriate modulo arithmetic in angle space.

4.2.3 Cohesiveness of Couzin's Model

In this section, we evaluate the connectivity properties of the metric-based implementation of Couzin's model. More precisely, there are a number of parameters that affect the behavior of the agents under Couzin's model, and we evaluate which parameters are compatible with topologies that stay connected. In doing so, we replicate the four collective structures, called phases, that Couzin observed.

Couzin used group angular momentum and group polarization, denoted m_{group} and p_{group}, respectively, to measure and classify the collective structures that can emerge when parameters are changed. Quoting from Couzin, "Group polarization increases as the degree of alignment among individuals increases" [8]. Simply put, if all agents are going the same direction, then polarization approaches unity. Again quoting from Couzin, "[Group] angular momentum measures the degree of rotation of the group about the group center." Simply put, if all agents are following a circle in the same direction with the same speeds, then group angular momentum approaches unity.

Figure 9 illustrates the four phases identified by Couzin, plus a phase identified by Woods and Ackland [42]. This figure is visual evidence that we successfully replicated the results from Couzin. The figure illustrates agents as circles with a line emanating from them in their direction of travel. Each subfigure in Figure 9 is a snapshot at one time instant. The swarm phase displays a highly dynamic collective that tends to stay stationary and exhibits short bursts of agent alignments interspersed with apparently random cycling about the swarms center. The toroid phase has agents that circle in the

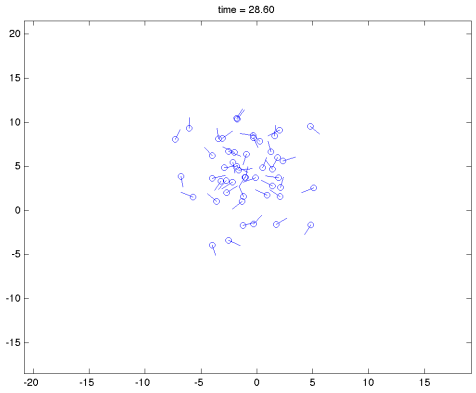

(a) Swarm

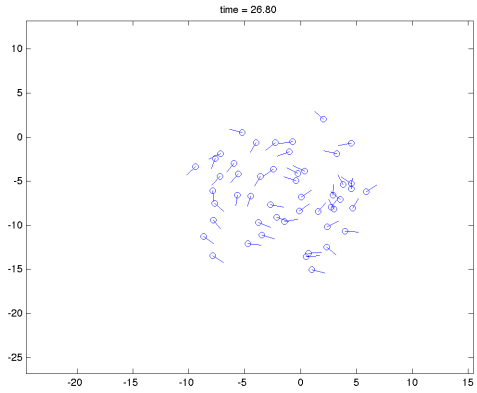

(b) Torus

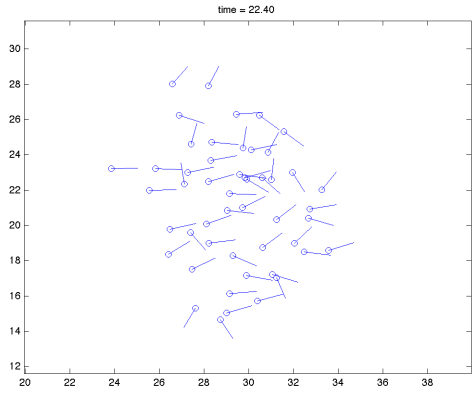

(c) Dynamic Parallel

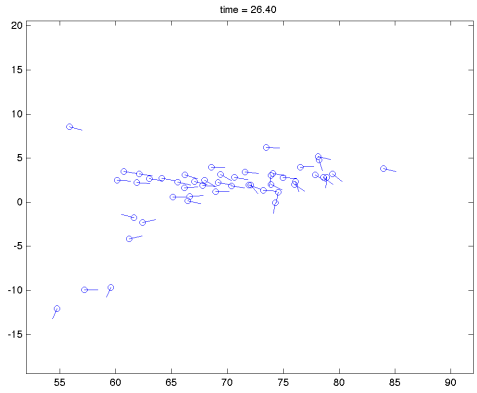

(d) Highly Parallel

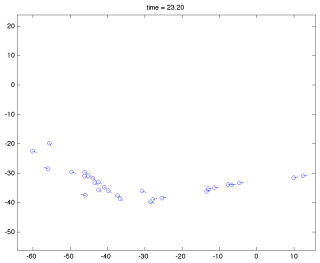

(e) Pack

Figure 9: Snapshots of various phases from Couzin's model: (a) Swarm. (b) Torus. (c) Dynamic parallel group. (d) Highly parallel group. (e) A possible pack in the style of Woods and Ackland.

same direction around a relatively stationary "hole" in the middle of the structure. The dynamic and highly parallel structures both exhibit collective movement in a unified direction, with the highly parallel structure exhibiting greater alignment and greater group velocity. The pack phase was observed by Ackland and Woods [42], but only for teams of agents with 80 or more agents. Although unclear from their description, we believe that Figure 9(e) illustrates a pack wherein agents tend to line up behind a lead agent and then follow; we were unable to reliably reproduce this collective structure and will only briefly touch on the pack structure in the remainder of this report.

Using $\Delta t = 0.1$ and $R^{\text{rep}} = 1$ we initialize the agents into a connected topology with all agents uniformly distributed around the center of the world. Ranges of other parameter values are shown in Table 1. We performed 10 trials under each parameter setting to evaluate how many of the teams stayed connected over a 30 second simulation. Initial orientations were uniformly chosen from the range $[-180°, 180°]$, and noise was added to the velocity sampled uniformly from

$$[-0.2, 0.2] rad/s \approx [-11.5, 11.5]°/sec.$$

For each parameter set, we recorded m_{group} and p_{group}

Parameter	R^{ori}	R^{att}	s	ω	N	ϕ
Units	unit	unit	unit/s	$°/s$	integer	$°$
Values	$R^{\text{rep}} + \{1, 5, 10\}$	$R^{\text{ori}} + \{0, 3, 8, 15\}$	$\{1, 3, 5\}$	$\{40, 70, 100\}$	$\{5, 10, 20, 50\}$	$\{180, 150, 120, 90\}$

Table 1: Parameter values used to evaluate when Couzin's model creates sustainably connected topologies.

We begin by noting a set of *ad hoc* observations that tell us something about the emergence of collective structures:

- Speed, s, does not impact the types of cohesion, that is, no matter what speed is used the different phases can emerge. The probability of an arbitrarily initialized team reaching one of those phases, however, does depend on speed.
- Small teams, $N \leq 10$ only reliably exhibit one type of cohesion, the dynamic parallel phase. Many of the other phases are seen during 30 second simulations, but they tend to only be transients and either break up or turn into a parallel group phase.
- No teams with viewing angle $\phi \leq 120°$ are reliably cohesive.
- Attraction and orientation have the strongest affect on the phase.
  - Low attraction tends to either be unstable or exhibit parallel group behaviors.

  - High orientation strongly tends to promote parallel group behaviors with some swarming behaviors.
  - Low orientation seems to be required for toroids.

- It is difficult to find a set of parameter values that reliably produce cohesive and highly dynamic parallel teams. This suggests a tradeoff between the speed of the team's centroid and the cohesion of the team. Fast teams may be less cohesive.

A more interesting observation is that we were able to replicate the clusters and phases that Couzin identified. We went a step beyond simply clustering parameter values into phase-based clusters, and also verified that the following parameter sets reliably produced the phases identified by Couzin; see Table 2 and Figure 9.

Phase	N	R^{ori}	ΔR^{att}	s	ω (°/sec)	ϕ (°)
Swarm	50	1	10	3	100	150
Torus	50	3	12	5	100	150
Dynamic Parallel Group	50	7	10	3	100	150
Highly Parallel Group	50	11	15	5	70	120
Pack	80	4	19	5	28	120

Table 2: Parameter values used to evaluate when model parameters produce the phases identified by Couzin. $R^{\text{att}} = R^{\text{ori}} + \Delta R^{\text{att}}$, that is, attraction radius was defined as an offset from the orientation radius.

In terms of group polarization and group angular momentum, Figure 10 is a useful way of representing clusters and their associated phases. Toroids spin and stay in place, so they have high angular momentum and low polarization; individual parallel groups tend to align and travel together, so they have low angular momentum and high polarization.

Before continuing, note that these results apply for a small but constant level of noise. Future work should evaluate how much noise affects which phases are likely to arise. Although this future work is important, we find it more interesting to explore how these different phases can be influenced by an external influence, namely, human input. We will return to this issue shortly.

4.2.4 What makes a Phase Cohesive?

Even if we eliminate the worst case, which occurs when $\phi = 90°$, only 34.72% of the simulations produce a cohesive group (one which is connected) at the end of 30 seconds of simulation. This suggests that parameters need to be selected so that they are likely to produce one of the cohesive phases. The parameter values that are most likely to produce connected topologies from the experiments in the

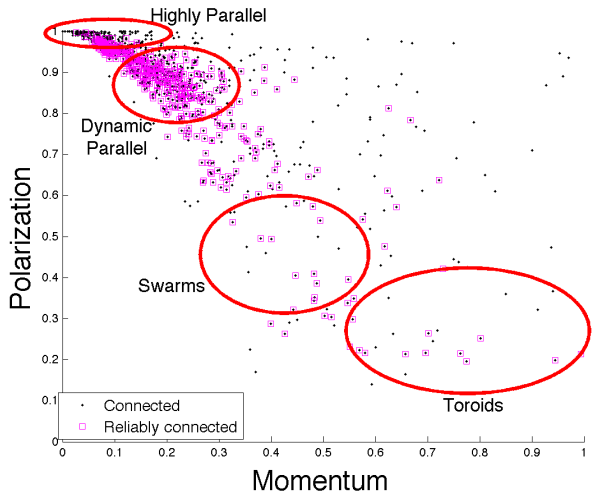

Figure 10: The clumping of phases suggested by Couzin with results of all teams with acceptable $(p_{\text{group}}, m_{\text{group}})$ values (solid dots) and results of simulation settings where 80% or more of the ten trials produced a connected topology.

previous section use $\omega = 100°/sec$ and $\phi = 150°$[3]. The remaining parameter values that were likely to produce a connected topology are shown in Table 3.

R^{ori}	ΔR^{att}	s
$\{1, 5, 10\}$	$\{8, 15\}$	$\{3, 5\}$

Table 3: Parameter values correlated with connected topologies.

We evaluated those parameter combinations from Table 3 that reliably produced connected groups. As before, we performed ten trials for each parameter combination and then evaluated which parameter combination was most likely to yield a connected graphs. The data indicated that a high value of ΔR^{att} is the strongest predictor of whether a topology is connected. This is illustrated in Figure 11, which displays histograms of the number of connected graphs out of ten trials for $\omega = 100°/sec$, $\phi = 150°$, and the values shown in Table 3. Also shown in the figure are results for lower values of ΔR^{att} for comparison.

These results indicate that, if we restrict attention to those parameters that are most likely to yield one of the four phases, then we often get collective behavior that is likely to be coherent. Stated

[3]Actually, $\phi = 180°$ was more likely to produce cohesive collectives, but this requires that agents be able to see everything around them, meaning that an agent doesn't have a blindspot. This seems overly optimistic for both biological systems and for robots with typical range sensors, so this value was not used in the simulations.

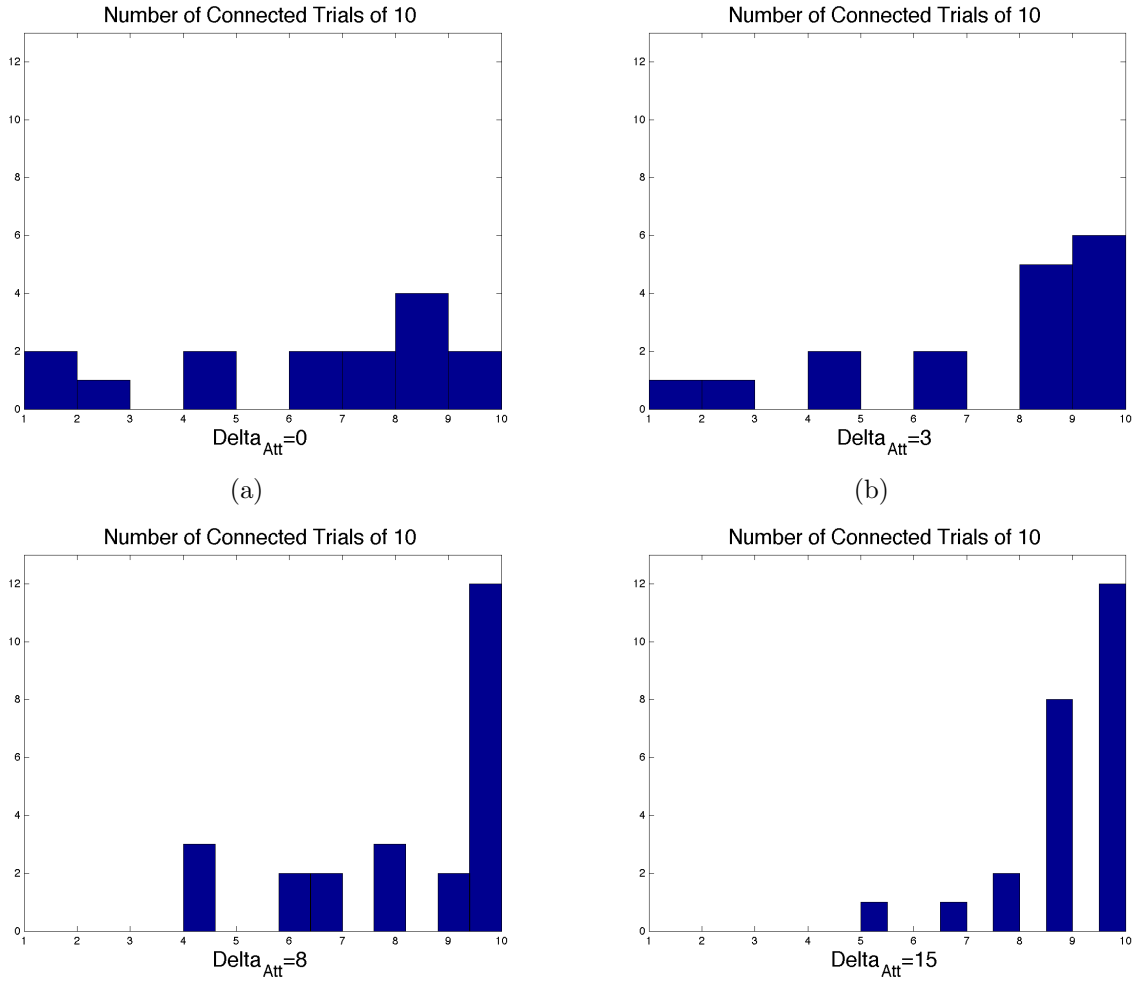

Figure 11: Attraction, ΔR^{att}, determines connection.

simply, a single parameter determines the likelihood of coherence, namely, ΔR^{att}, and high values are very likely to produce cohesive teams regardless of what happens with other parameters. The reason for this is that high radii of attraction mean that agents are rarely out of range of each other, so even very distributed individuals are likely to coalesce into one of Couzin's collective structures.

4.2.5 Couzin's Model with Nearest-Neighbor Topology

Consider two ways to define nearest-neighbor topologies. The first is to simply find some set of the N nearest neighbors as long as they are not too far away, that is, find the N neighbors that are both nearest to the agent and also within some large distance-based sensing threshold. The second way is to use the metric-based techniques of the previous section, but only allow an agent to process the N neighbors within the R^{ori} or R^{att} metric neighborhoods. For example, if there are 10 other agents within the R^{att} metric-neighborhood but $N = 5$ then only the 5 nearest neighbors affect orientation.

In this section, we use the former method. In essence, repulsion from other agents stays the same, but the controller for alignment and orientation is changed from one that uses fixed radii to one that uses a fixed number of nearest neighbors. In this section we explore how this strict counting of neighbors can affect team cohesion and behavior. We conducted ten trials each using the parameters shown in Table 4, and recorded various parameters including group angular momentum and group polarization. Results yielded the following ad hoc observations.

Parameter	NN^{ori}	ΔNN^{att}	s	ω	N	ϕ
Units	unit	unit	unit/s	$^\circ/s$	integer	$^\circ$
Values	$\{1, 7, 15\}$	$\{0, 3, 8, 15\}$	$\{1, 3, 5\}$	$\{40, 70, 100\}$	$\{5, 20, 50\}$	$\{180, 150, 120, 90\}$

Table 4: Parameter values used to evaluate when the nearest neighbor extension of Couzin's model creates sustainably connected topologies. NN^{ori} represents the number of nearest neighbors that affect how an agent orients. the number of nearest neighbors that attract an agent is given by $NN^{\text{att}} = NN^{\text{ori}} + \Delta NN^{\text{att}}$.

- Surprisingly, the number of neighbors doesn't make much of difference in terms of the percentage and types of phases, at least when orientation and attraction neighborhoods are the same ($NN^{\text{ori}} = NN^{\text{att}}$). See Figure 12.
- No teams with viewing angle $\phi = 90^\circ$ are reliably cohesive.
- Larger team sizes appear to exhibit a new phase $p_{\text{group}} \approx m_{\text{group}} \geq 0.75$. See Figure 13. If we look beyond the static polarization and angular momentum values to the dynamic structures, we discover that this is not a new phase but rather an artifact created when an individual splits from the group and loses contact with it because its field of view is low. This artificially biases p_{group} and m_{group}.

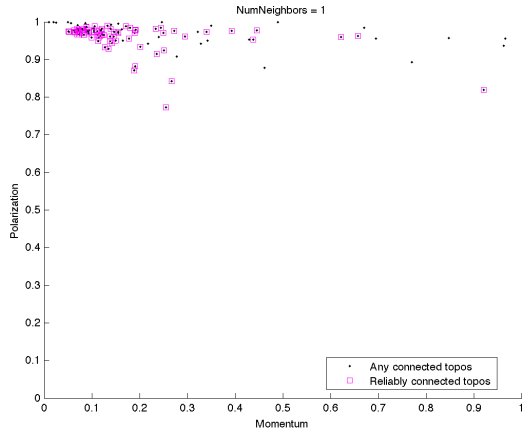

(a) 1 Neighbor

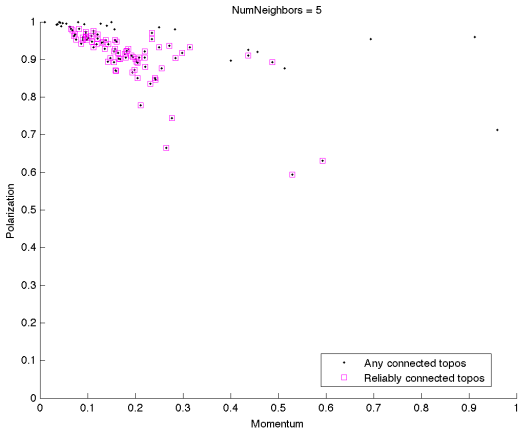

(b) 5 Neighbors

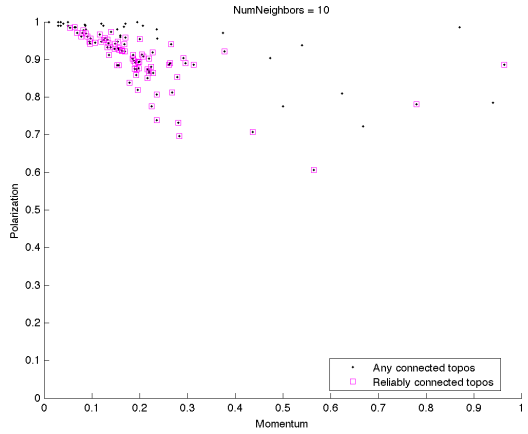

(c) 10 Neighbors

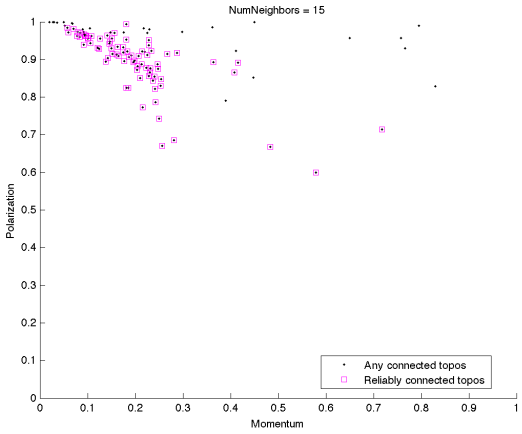

(d) 15Neighbors

Figure 12: Snapshots of various phases from a nearest-neighbor version of Couzin's model, where the number of attraction neighbors equals the number of orientation neighbors.

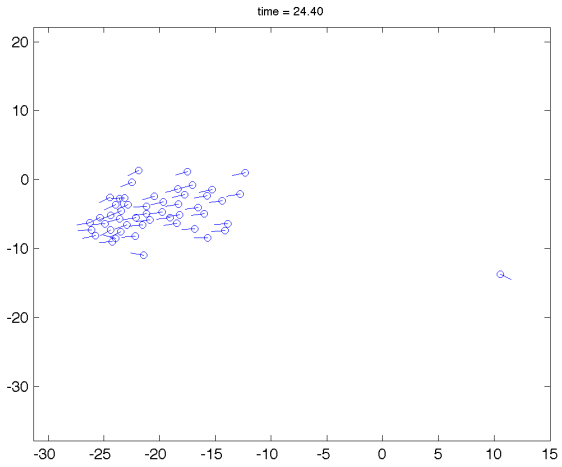

Figure 13: The artifactual new phase occurs when agents are very slow ($S = 1$) and they have limited fields of view $\phi = \pm 120°$. This phase is an illusion; it occurs when the group of agents is in a highly parallel group phase but where one agent has recently left the group and is traveling nearby but in the opposite direction to the group.

In addition to these ad hoc observations, we made two other observations that deserve special emphasis. First, toroids and swarms were not observed in the nearest neighbor topologies, even when there is a difference between the number of neighbors used for orientation and the number of neighbors used for attraction, with the former being smaller than the latter. Figure 14 illustrates that nearest neighbor topologies almost always produce dynamic parallel or highly parallel groups. Simply put, nearest neighbor topologies produce fewer collective structures than metric-based topologies.

Second, nearest neighbor topologies are much more likely to be cohesive than metric-based topologies. Even if we eliminate the worst case, which occurs when $\phi = 90°$, 79.34% of the simulations produce a cohesive group (one which is connected) at the end of 30 seconds of simulation. This is more than double the percentage for the metric-based topology (34.72%), suggesting that the nearest neighbor topologies are much less sensitive to specific parameter values.

Note that, to expedite simulations, we used an estimate for the nearest neighbor graph when we computed connectedness. This estimate said that two agents were connected if they were within 20 "sensing units" of each other. Thus, the 79.34% value above is slightly misleading since the estimate overestimated the number of neighbors in the graph. However, the estimate is much improved if we eliminate slow moving agents, because it is very unlikely for an uncohesive team to emerge simply because agents haven't had time to move more than 20 units apart yet. When we eliminate the slow agents, those for which $s = 1$, then nearest neighbor topologies produce a connected group 75.16% of the time, which is slightly less than the other measure. By contrast, eliminating $s = 1$ from the metric-

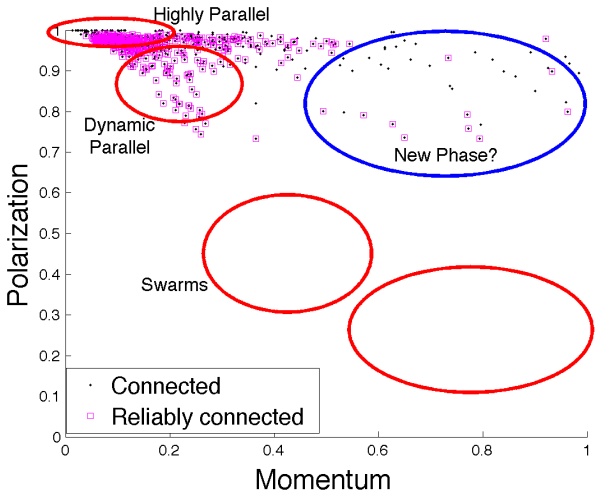

Figure 14: The clumping of phases suggested by Couzin with results of all nearest-neighbor-based teams with acceptable $(p_{\text{group}}, m_{\text{group}})$ values (solid dots) and results of simulation settings where 80% or more of the ten trials produced a connected topology. To make these results comparable with the metric-based results, we define "connected" using a metric-based connectedness criteria with $R^{\text{att}} = 30$. In other words, connectivity is measured on the graph produced by saying that agents share an edge if they are within 30 units of each other; since nearest neighbor topologies effectually induce unbounded metric neighborhoods, the 30-unit neighborhood can sometimes underestimate connectivity, at least for collective structures with bounded values of p_{group} and m_{group}.

based topologies reduces the number of connected groups to 25.81%. This means that even when we account for the slight overestimate in the nearest neighbor topologies, they are still much more likely to be connected than the metric-based topologies.

To summarize, nearest neighbor topologies are much more likely to produce cohesive teams than metric-based topologies, but metric-based topologies are capable of producing a wider range of behaviors.

4.2.6 Couzin's Model with a Randomly Dynamic Topology

Consider the bio-inspired alternative to Couzin's model from [5] where the topology is not purely metric-based and not purely neighborhood-based, but rather uses a probabilistic approach to determine an agent's neighbors. In this aproach, interaction neighbors are chosen with a probability inversely proportional to their distances from the agent. Once a neighbor is chosen, it influences the agent's behavior according to the rules associated with Couzin's three metric zones.

In this section, we explore a simplified version of such randomized topologies by randomly assigning neighbors independent of distance. In this simplified version, when there are other agents within the zone of repulsion, then these agents repel the agent as in the previous topologies; this prevents agents from colliding. If there are no other agents within the zone of repulsion then we form a randomly dynamic topology at each time step as follows. Each agent is assigned n neighbors with a uniform probability distribution from the group of $N + M$ total agents, with N nominal agents and M leaders.

The presence of this randomly dynamic topology opens up some possibility for future work into the probabilistic connectivity of this topology. Probabilistic connectivity arises from the fact that, with high probability over a sufficiently large time interval T, large values of n yield a connected group adjacency matrix and a connected management matrix. In this context, a connected group adjacency matrix means that each agent will likely be connected to each other agent for at least one time step during the interval T. Similarly, a connected management matrix means that at least one leader leader will likely influence each nominal agent at some time. This opens up the possibility of influence propagating throughout the group, even if there are temporarily isolated subgroups within the collective.

Simulations were run using the parameters listed in Table 5.

Parameter	N	n	R^{rep}	R^{ori}	R^{att}	ϕ	ω	s
Units	integer	integer	unit	unit	unit	degrees	degrees/s	unit/s
Values	50,80	1,5,15,30,50	1	1,3,10,50	∞	180	40,70,100	1,3,5

Table 5: Table of explored parameters for Random Neighbour Topology

The most important result from this section is that the random neighborhoods are capable of expressing all of the phases identified in Couzin's original model. In general, larger values of n tended to lead to more expressive groups, but many groups were able to form with as little as 1 neighbor.

Toroids tended to form when $R^{\text{ori}} = 3$, $n \geq 15$ and $s \geq 3$. On rare occasion, dynamic parallel groups formed under these parameter sets, which was more common for high values of ω and lower group sizes, N. In general, as n increased, the toroids formed more quickly and had higher group angular momentums. Since large group angular momentums reliably indicate the formation of the toroid group, we present Figure 15 to illustrate that toroids are more likely to form for large neighborhood sizes.

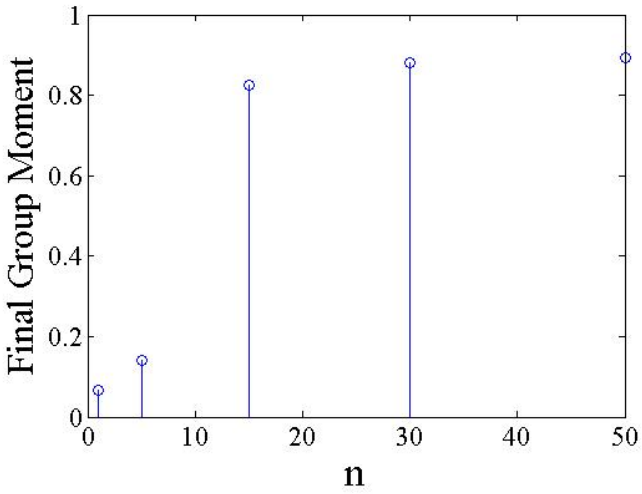

Figure 15: Plots showing group angular momentum in response to n for $N = 50$ $R^{\text{ori}} = 3$ $\omega = 40$ $s = 5$.

Dynamic parallel groups tended to form when the orientation radius was large. Thus, when $R^{\text{ori}} = 10$ dynamic parallel groups formed in almost all cases. On rare occasions, a torus formed, which happened most frequently when ω was low, and N, n and s were high. *Swarms* tended to form when the orientation radius was ignored because it is the same as the repulsion radius; when $R^{\text{ori}} = 1$, swarms were observed for all parameter sets. *Highly parallel groups* formed whenever $R^{\text{ori}} = 50$.

To summarize, simple randomly changing topologies are as expressive as the metric-based topologies.

4.3 Expressiveness of the Physicomimetics Model

Physicomimetics is a distributed control law that uses physics-based forces to control a swarm [34]. All agents are treated as point masses and inter-agent influence is modeled using a metric-based topology. There are two distance thresholds in physicomimetics: a radius of attraction, R^{att} and a radius of repulsion, R^{rep}. The zones of attraction and repulsion are illustrated in Figure 16. An agent is repelled by all agents within the zone of repulsion and attracted to all agents that are outside the zone of repulsion but within the zone of attraction. These forces endow all agents with simple autonomy from which collective patterns of agent behavior can emerge.

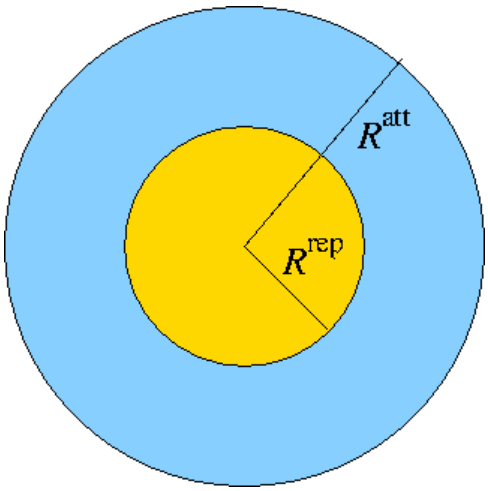

Figure 16: Physicomimetic zones of repulsion, R^{rep}, and attraction, R^{att}.

The total force is $F_i = \sum_{j=1}^{N} F_{ij}$, where N is the swarm size where interagent repulsion is given by $F_{ij} = -G(m_i m_j)/(r_{i,j}^d)$, and where interagent attraction is given by $F_{ij} = Gm_i m_j/(r_{i,j}^d)$; in gravitational physics, $d = 2$ representing the inverse square law, but physicomimetics allows us to vary this parameter. The velocity update of the particle at time $t+1$ is given as $v(t+1) = \gamma(v(t)+\Delta v)$, where $\Delta v = F\Delta t/m$ and where $\gamma \in (0, 1)$ is a coefficient of friction that stabilizes the system. Even if all agents begin in the same location, the repulsive forces cause agents to redistribute, enforcing a topology where each agent is influenced by only a handful of other agents.

Before continuing, it is helpful to summarize the differences between the physicomimetic model and Couzin's biomimetic model. We begin by relating the physicomimetic model to the parameters in Equation (4). Since the forward speed of agents in the physicomimetic model can change (unlike the fixed forward speed used in Couzin's model), the state of agents in this model is given by

$$x_t^i = \begin{bmatrix} \mathbf{c}_t^i \\ \mathbf{v}_t^i \end{bmatrix}$$

where $\mathbf{c}_t^i$ is the position vector and $\mathbf{v}_t^i$ is the velocity vector of agent i. We can identify f^i by identifying the dynamics of the various portions of the state vector. The dynamics of the position vector are given by

$$\mathbf{c}_{t+1}^i = \mathbf{c}_t^i + \Delta t \mathbf{v}_t^i,$$

the dynamics of the velocity vector are given by

$$\mathbf{v}_{t+1}^i = \gamma(\mathbf{v}_t^i + \Delta \mathbf{v}^i)$$

where $\Delta\mathbf{v}$ is given by

$$\Delta\mathbf{v}^i = \frac{\left(\sum_{j=1}^{N} \mathbf{F}_{ij}\right)\Delta t}{m_i},$$

and where $\mathbf{F}_{ij}$ is given by the switching control law wherein agent i is forced away from agent j if agent j is within the zone of repulsion, is forced toward agent j if agent j is within the zone of attraction, and is not affected by agent j if agent j is outside the zone of attraction. Unlike the biomimetic model, the physicomimetic model:

- Does not use orientation.
- Does not use the blindspot.
- Does not have a turn speed limitation.
- Includes friction.
- Controls all of velocity, both speed and direction, of each agent; agents can stop moving.

The collective structures produced by the physicomimetic model are much different than those produced by the biomimetic model. We conducted a series of experiments, varying the parameters given in Table 6.

Parameter	R^{rep}	ΔR^{att}	γ	d
Interpretation	Repulsion Radius	Attraction Radius $R^{\text{att}} = R^{\text{rep}} + \Delta R^{\text{att}}$	Friction Coefficient	Inverse d-Law
Value Range	$\{1, 6, 11, 16\}$	$\{0, 4, 8, 12\}$	$\{0.5, 0.7, 0.9\}$	$\{1, 2, 3\}$

Table 6: Physicomimetic model parameters varied to discover phases produced by the model.

Three qualitatively different phases emerge, as illustrated in Figure 17: isotropic, anisotropic, and structured. Each figure is actually a collection of subplots for various parameter sets. The text and captions are too small to see, but the figures are nevertheless included to demonstrate the different structures. Figure 17(a) shows how the different agents can be distributed more or less uniformly across the world, after being initialized in a random cluster. Notice the regular spacing among agents, especially for the bottom two subplots. Figure 17(b) illustrates how sometimes the distribution of agents takes on a cluster of clusters organization. The plot in the bottom left most clearly illustrates this phase. Figure 17(c) illustrates a phase where agents tend to distribute themselves around some type of central location. This phase surprises us, and future work should examine why it emerges.

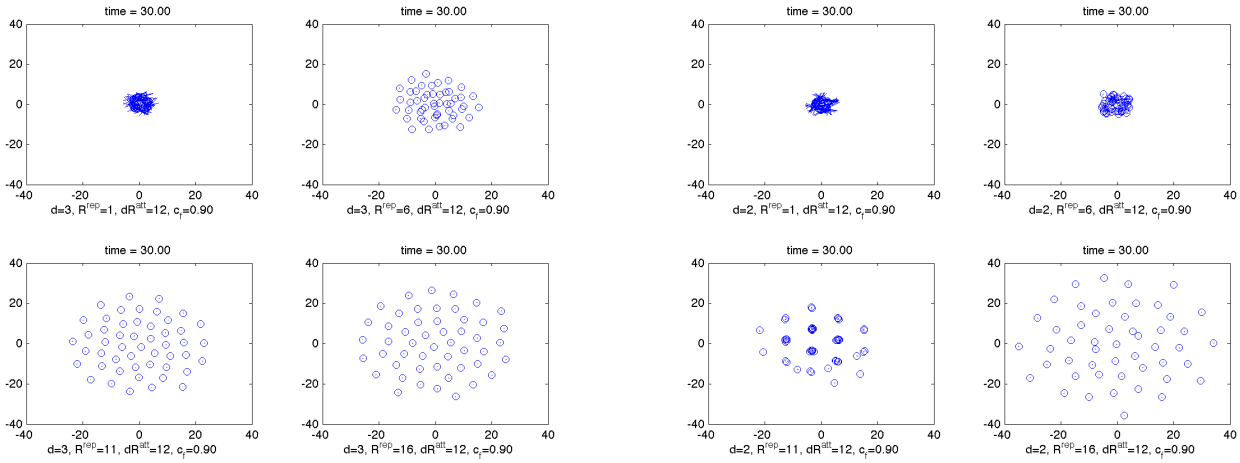

(a) Isotropic

(b) Anisotropic

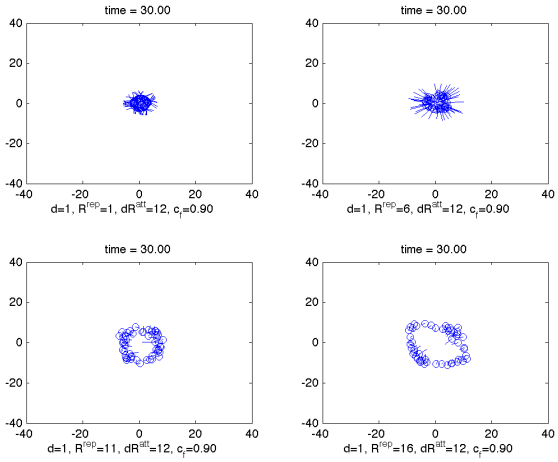

(c) Structured

Figure 17: Snapshots of various phases from the physicomimetics model.

Future work should also examine the expressivity of the physicomimetics model when modified to use nearest neighbor topologies rather than a metric-based topology. Other future work should include using the parameters of a radial basis distribution or pairwise correlation distribution to characterize the expressiveness of the three different physicomimetic phases.

5 Managing Parallel and Isotropic Phases Using Switching Controllers: Leaders and Predators

Results from the previous section do not include the potential for perturbations in the topology such as when leaders or predators seek to influence the movement of the group. This section explores how external perturbations affect connectivity, for both leaders and predators. More importantly, it explores what type of external perturbation is best able to sustain influence over various phases. In this section, we explore two types of "informed" perturbations:

- Predators: These special agents influence nominal agents within the collective only through repulsion forces.
- Leaders: These special agents influence nominal agents within the collective through attractive forces.

In the next section, we include another type of "informed" perturbation:

- Stake-holders: These special agents influence nominal agents by joining the collective, that is, by attracting, repelling, and aligning with other agents. What makes these agents special is the fact that they are also attracted to a given location in the environment, that is, they influence and are influenced by other agents but they also have a stake in where the collective moves.

The experiments presented in this section only consider a subset of the possible collective phases identified in the previous section. More precisely, this section considers how parallel groups (not swarms or toroids, biomimetics) and isotropic phases (not anisotropic of structured phases, physicomimetics) can be augmented with a new switching controller to respond to leaders and predators. The key insights from this section come from the case studies with human managers managing parallel groups and isotropic phases.

5.1 Physicomimetics

We begin with physicomimetics and explore three different types of controllers, two based on leaders and one based on predation. In the experiments, 100 agents occupy a $1000m \times 1000m$ grid in which

a new task appears with a probability of 0.0033 every second. Tasks are depleted at a rate given by Equation (7). To allow more rapid data collection, we set 10 of the 100 agents as leaders with the remaining 90 agents as nominal agents; this means that the human can control up to 10 leaders at any given time.

The implementation of each of these controllers distinguishes between two types of agents: leaders/predators, and nominal agents. Leaders and predators are controlled by the operator through $g^i(x_t^i, u_t^{\text{op}}) \neq 0$ without worrying about other agents; $f^i = 0$ if the agent is a leader or a predator. By contrast, nominal agents are not directly influenced by the operator, $g^i(x_t^i, u_t^{\text{op}}) = 0$, but are affected indirectly by leaders/predators through $f^i(x_t^i, \mathbf{x}_t^{\neg i}) \neq 0$.

This implementation requires two changes in the dynamics specified above, with both changes compatible with Equation (4). First, the state of the agent includes the type of agent. If we let $\alpha^i \in \{\text{Leader/ Predator, Nominal}\}$ denote the type of agent, then the agent's state becomes

$$x_t^i = \begin{bmatrix} \mathbf{c}_t^i \\ \mathbf{v}_t^i \\ \alpha^i \end{bmatrix}.$$

Second, the switching law for nominal agents encoded in f^i changes. If we let $\tilde{f}^i$ denote the new switching law and f^i denote the former law, then the new control law for nominal agents becomes

$$\tilde{f}^i = \begin{cases} f^i & \text{if no leaders/predators are within the leader/predator zone} \\ \text{attract to leader} & \text{if a leader is within the leader zone} \\ \text{repel from predator} & \text{if a predator is within the predator zone} \end{cases}. \quad (8)$$

In words, nominal agents behave like normal if no leaders or predators are nearby; otherwise, they move toward the leader or flee from the predator. This simply adds one more switch condition to the switching control law that governs nominal agent behavior. Note that experiment below uses only leaders or only predators (not both in the same experiment).

5.1.1 Virtual Leader Management (VLM)

In VLM, the operator deploys a virtual agent that attracts all agents within a radius of attraction. According to the switching controller in Equation (8), this radius of attraction supersedes the other zones (both interagent attraction and repulsion), meaning that the agents are only affected by the virtual agent. The virtual agent's radius of attraction is controllable by the operator but is nominally set to $75m$.

Once the agents are attracted to the virtual agent the operator drags the virtual agent to the resource location and, once the resource is depleted, the agents return to their nominal agent autonomy and re-distribute throughout the environment. Since the leader agent is virtual, the operator must

continually broadcast to the agents to maintain influence. As the number of tasks grows, operator workload grows quickly as the manager must attract agents to a virtual agent (selecting the team), move the virtual agent to the resource to draw the agents to the task, and then maintain agents in the proximity of the resource until it has vanished (as shown in Figure 18(b)).

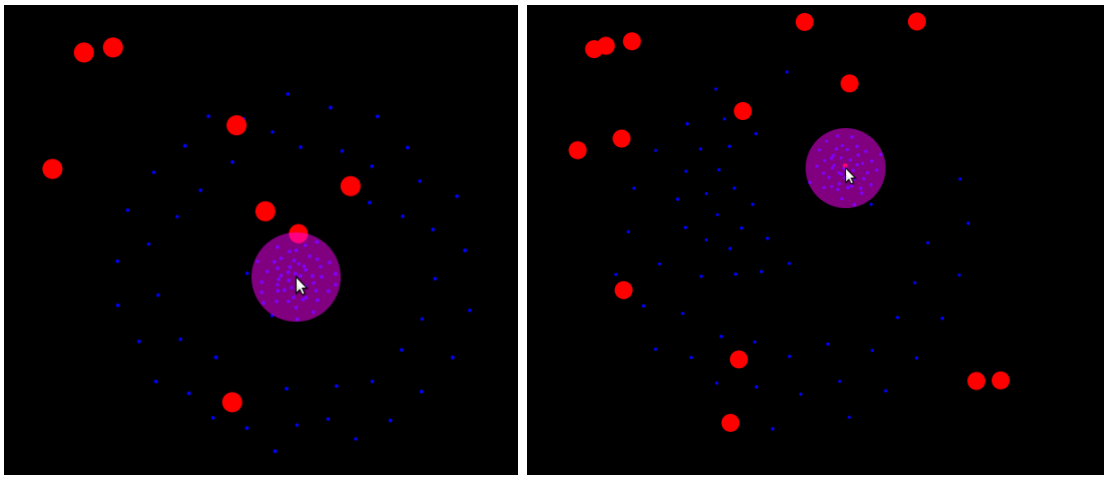

(a) (b)

Figure 18: (a) The agents are attracted to the circle in violet color which is the radius of attraction/influence for the virtual agent. (b) The circle is moved to the target and the resource is depleted.

5.1.2 Physical Leader Management (PLM)

In PLM, we assume that an operator can control a small set of leader agents. Each leader agent will recruit a number of agents and pull them to the resource location. The radius of attraction and the location of the resources is assigned to the leader by the operator. In compliance with Equation (8), when an agent is attracted to the leader it ignores interagent attraction and repulsion. The leader autonomously guides agents to the resources, and once the resource is consumed the agents redistribute. For the simulations we deployed a swarm of 100 agents. The parameters of the model were: $R^{\text{rep}} = 15m, R^{\text{att}} = 100m, g = 700$, and $F_{\text{max}} = 1$ (note that F_{max} is a saturation level; force magnitudes are limited to this value). The simulation runs for 20000 time steps. Every second, a new resource appears with probability of 0.0033.

The key difference between the PLM and VLM is the presence of delegation. As illustrated in Figure 20, PLM lets the human give commands to a leader (through$g$); all local interactions between the leader and the other agents are propagated through B_t. By contrast, in VLM, the human needs to remotely communicate with all agents in its sphere of influence through f. Delegation, in this context,

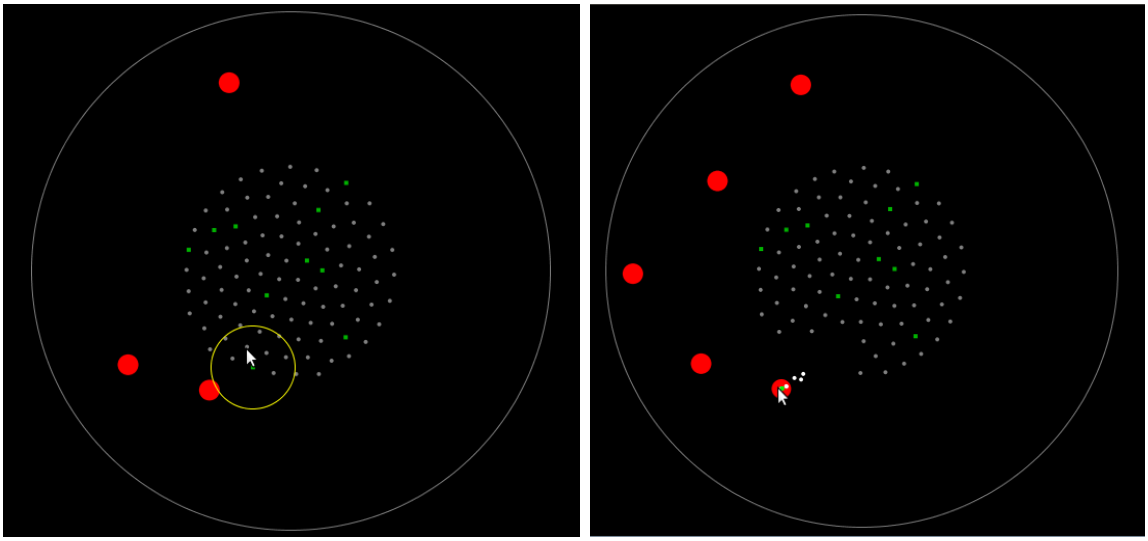

(a) (b)

Figure 19: (a) The radius of attraction/influence associated with a leader (green square) (b) The leader and its team at the assigned resource. Notice how the agents are in an isotropic phase.

is the ability of the human to control a small number of leaders but delegate the propagation of influence to the other agents through B_i; VLM does not include delegation because the human must constantly exert influence over agents through f. We will show that delegation is useful when local agent-to-agent communications are reliable but human-to-agent communication is spotty.

5.1.3 Virtual Predator Management (VPM)

VPM works similar to VLM but the virtual predator agent repels the agents present inside its radius of influence. In our experiments, the nominal radius of repulsion is subjectively set to $R^{\text{pred}} = 50m$ and the radius of attraction is $R^{\text{lead}} = 1.5R^{\text{pred}} = 75m$. By repulsion, the virutal predator can push the agents towards the targets. As shown below, when the agents move away from the virtual predator, the influence of the virtual predator is difficult to sustain making the performance of this model relatively poor; see below.

5.2 Biomimetics

In this section, we describe three types of management interfaces for an information foraging task using Couzin's biomimetic model. The scenario consists of 100 fish in a 120×120 area. Quantities of food, represented graphically as barrels in the simulation, are placed around the map to represent the

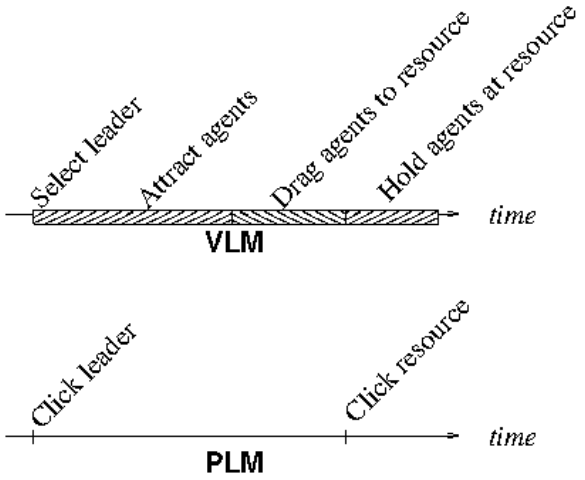

Figure 20: An illustration of required human input for VLM versus PLM.

information to be gathered. The "food" is depleted at 1 unit per second per fish whenever a fish is within range, that is, the initial quantity of food is measured in fish-seconds. For example, if N fish are within range of food for t seconds, then $N \times t$ food units are depleted.

With the exception of parameter-based control, which is a centralized method by which a human can influence the collective by broadcasting parameter changes to all agents simultaneously, the human will influence this collective using either a leader or predator. As with the physicomimetic model under human influence, described above, adding human influence to Couzin's model sets $g^i = 0$ for nominal agents and $f^i = 0$ for leaders/predators. The state of nominal agents is augmented to include the type of agent (nominal or leader/predator), but the switching control law for nominal agents is augmented slightly differently than it was for the physicomimetic model. If other agents are within the repulsion zone, then only the repulsion force influences agent behavior. If no agents are within the repulsion zone and if a leader (predator) is within the leader zone (predator zone) then the nominal agent ignores other agents and is attracted to (repelled by) the leader (predator). If no leader (predator) is present, then alignment and attraction forces apply. In the simulations, $R^{\text{rep}} < R^{\text{pred}}$ and $R^{\text{rep}} < R^{\text{lead}}$ so this slight modification from the physicomimetic model means that agents follow the leader or are repelled by the predator unless the agents are about to collide with each other.

5.2.1 Parameter-Based Control

A centralized version of control for this problem is parameter-based control, which serves as the baseline condition. Fish behavior is initiated by selecting parameters that cause them to spread out and keep a minimum distance from each other. The fish spread out over the map and consume resources. Over a

small time interval and over a small spatial scale (before the fish stop moving and adopt an isotropic structure) we see a new temporal phase, which we call "munch and march" that relies solely on repulsion and ignores alignment and attraction. For the simulations, the parameter values were subjectively optimized for agents to "munch and march", that is, consume information while they move apart from each other uniformly. Parameter-based control is illustrated in Figure 21 for a scenario where food resources are placed on a uniform grid throughout the world.

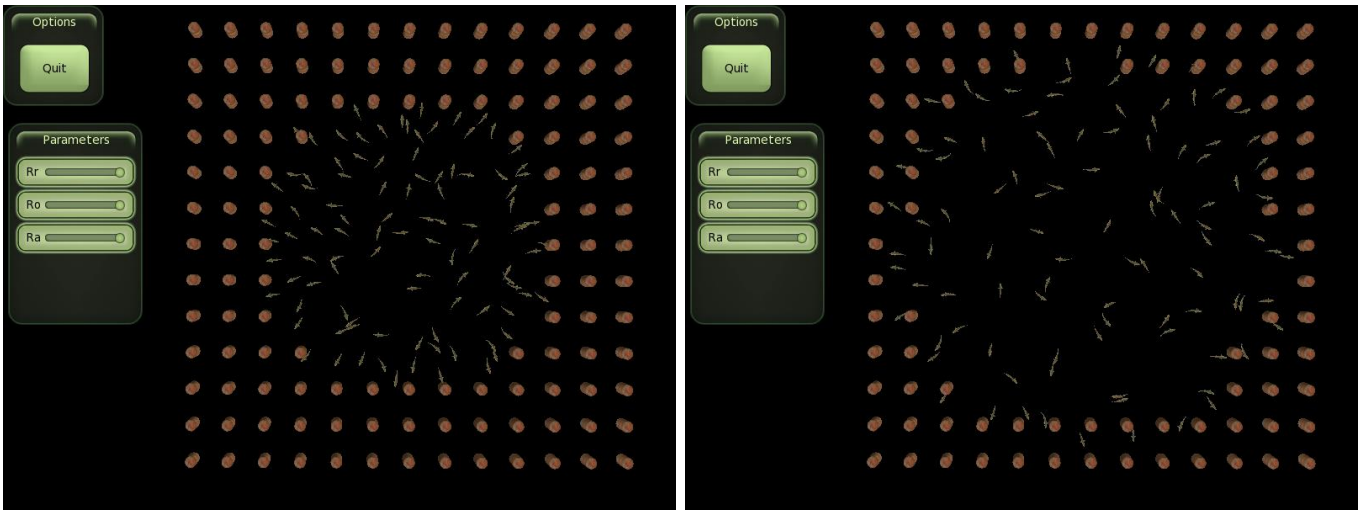

(a) (b)

Figure 21: Using Couzin's model with parameters set to "munch and march".

5.2.2 Predator Management

The second control method involves using a single predator to steer groups of aligned fish. The original model is adapted by causing fish to be repelled by a predator if the predator is within a given radius. The predator moves slightly faster than the fish and can turn much more sharply. It has the desirable characteristics of providing direct operator input without directly manipulating the system parameters.

The parameters for this model induce a highly parallel group using

$$(R^{\text{rep}}, R^{\text{ori}}, R^{\text{att}}, \phi) = (1, 14, 14, 45^\circ)$$

and $R^{\text{pred}} = 7$. If a predator gets close then they are repelled by this predator. Since the radius of attraction is greater than the predator's radius, the fish tend to stay close together even when the predator starts to "chase" them. This is illustrated in Figure 22 which shows a predator (shark) steering a group of fish and splitting a group of fish; the scenario for this figure is for randomly placed food

(a) (b)

Figure 22: Using Couzin's model with parameters set to allow a predator (a) to steer or (b) to split a collective organized as a parallel phase.

resources. Note how the parameters in the upper left of the figures are set to support a parallel group phase; compare to Table 2.

5.2.3 Leader Management

The leader model is similar to the predator-based model above, but the fish are now attracted to the leader producing a tendency for fish to follow the leader. The radius of attraction for the leader is $R^{\text{lead}} = 7$, meaning that the predator-based and leader-based experiments use identical parameters, with the only difference being that predators repel and leaders attract nominal agents.

5.3 Performance and Robustness: Leaders and Predators

We begin by comparing the models in terms of task performance, measured as the time taken by a team to deplete the resource to zero, and of robustness, measured as the rate at which completion time increases as a function of the probability of communication loss or variations in resource distribution. We then explore characteristics of the cohesion and management matrices that correspond to good performance. We have designed the VLM and PLM physicomimetic models to allow us to compare robustness in the presence of communication loss, and the nominal and predator biomimetic models to allow us to compare robustness in the presence of resource distribution variation.

5.3.1 Physicomimetics

For the physicomimetic model, we conducted experiments that had two operators perform the tasks while we measured performance as a function of communication loss. In the first set of experiments, the probability of communication P is varied between $1, 0.5, 0, 1$ and 0.01; when $P = 1$ then all messages from the operator reach the agents, and when $p = 0.01$ then 99% of the messages sent from the operator to the agents are lost. The operator could control up to ten leader agents.

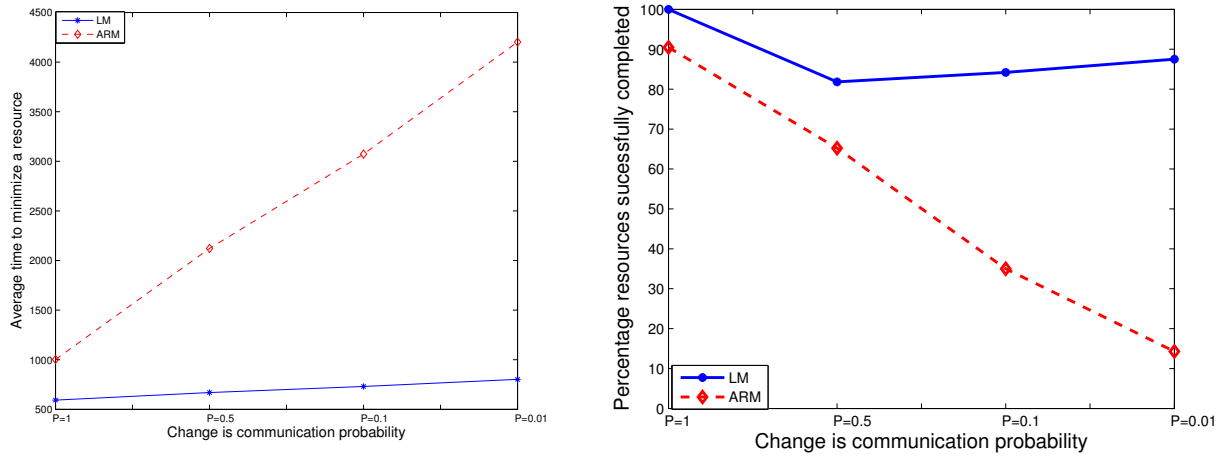

(a) (b)

Figure 23: (a) Average response time for different models with varying communication probabilities. (b) Percentage number of resources diminished to zero. Note that this plot has some old labels in it: ARM=VLM, and LM=PLM.

Figure 23(a) shows the performance for the VLM and PLM models when P is varied. For all values of P, PLM performs better than VLM. In terms of raw performance, PLM is superior because the operator can assign the radius of influence and the resource location and then attend to a different resource. By contrast, VLM requires that the operator continuously influence the agents. This requires a "sequential model" where the operator has to attend one resource, minimize it to zero and then attend to the next resource.

In terms of robustness, Figure 23(a) shows that the performance of the PLM model degrades slowly with decreases in P, indicating robust performance. This robustness occurs because once the leader gets the message through the noisy communication channel, it attracts nominal agents using the perfect agent-to-agent communication channel. By contrast, under VLM only a fraction of the operator's messages make it to the agents, so agents thrash between being attracted to the leader (when the message gets through) and being repelled by other agents that have drawn near the leader (when the

message doesn't get through). Since communication failures occur between the operator and the agents, it is not surprising that VLM is much less robust to communication failures that PLM.

To quantify robustness, we use the ratio of (a) the number of resources depleted to (b) the number of resources that appeared in the environment. Figure 23(b) shows that the ratio percentage decreases for both PLM and VLM, but that the decrease for PLM is much less then the decrease for VLM.

5.3.2 Biomimetics

For biomimetics, we conducted experiments to evaluate how variation in resource distribution affects performance. Prior to the experiment, we optimized parameter-based control to work best under a uniform distribution of small resource packets; this type of behavior will not work as well if the distribution varies because packets are not uniformly distributed or because larger packets require that the agents spend more time depleting them. Four different food distributions were evaluated. In the first two distributions, food was placed in a uniform grid, 10 units apart. The first distribution used containers containing one resource unit, and the second distribution again placed food in a uniform grid, but the size of the containers was increased to 10 resource units. In the third distribution, 10 containers of food were distributed randomly over the map following a uniform distribution on both the x-axis and y-axis. This scenario was designed to require fish to coordinate in schools, when the size of the food containers is large. To make the total amount of resource comparable to the second distribution, each food container held 100 resource units. The fourth distribution doubled the amount of resource in the barrels from the third distribution, and serves to illustrate how well the different methods scale as the need for coordination grows. Five runs were performed by a single operator for each distribution type and management strategy (parameter-based versus predator-based).

Note that these simulations are not intended to be a user study on the management interface, but rather a case study of the robustness and responsiveness of various management strategies. Task completion time is defined as the amount of time required for the fish to deplete all food sources. The world was constrained inside of a square; fish bounced off of walls when they reached the edge of the world. Consequently, food sources in corners don't get visited very often, so sources in the very corner were ignored in computing task-completion results.

Results are shown in Figures 24. The box and whiskers plots include the mean (red bar in the blue rectangle), the 25^{th}-75^{th} quartile range (the blue rectangle), and the entire spread of the distribution (the whiskers). Additionally, a thick, magenta line shows the trends of the average values.

In terms of raw performance, the parameter-based management completed the tasks more quickly for all simulations. Uniformly spreading the fish out in all directions by using parameters optimally chosen to do this produces fish that cover the area effectively. The predator-managed fish travel in highly dynamic groups (schools) and, therefore, take more time to cover the whole map. In the experiments, much of the completion time was spent guiding the fish from one resource to another. The predator-based approach naturally performed better near the end of the simulations when it was desirable to

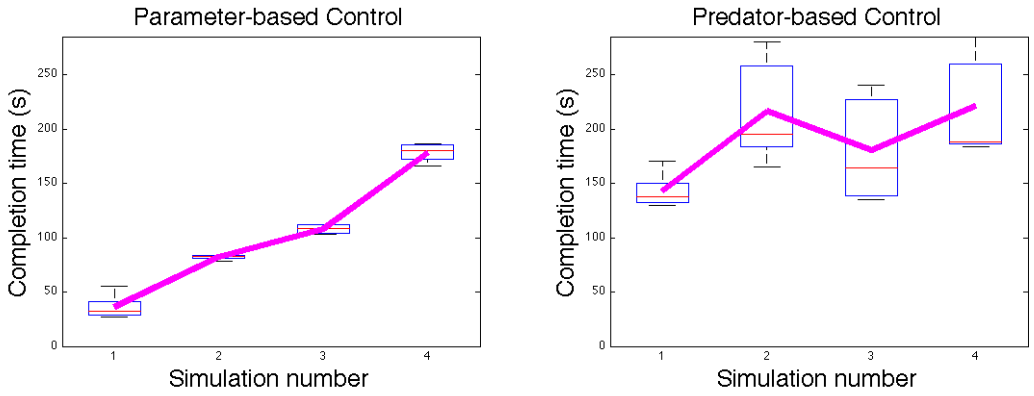

(a) (b)

Figure 24: Completion times for (a) parameter-based management and (b) predator-based management. Each simulation number corresponds to a different distribution type, in the order described in the body of text.

guide the fish towards the last few barrels, but this advantage occurred late enough in the task to be of little benefit in terms of task completion time.

Robustness tells a different story, namely that the predator-based control showed performance trends that were much more robust to variations in the distribution of food packets. Simply put, performance of parameter-based control rapidly deteriorates but performance for predator-based control stays relatively constant. Note the trends between the second through fourth distribution types. The predator-based model stayed fairly constant but the parameter-based model increased. This is because the predator-based model allowed a school of fish to focus on a concentrated resource for a long period of time, whereas the parameter-based model required the fish to continue to move about randomly, being repelled by each other on occasion or when they came near to walls. The predator approach seems to be potentially more robust to variations in the concentrations and distributions of the resources.

Before concluding this section, is is interesting that effective control strategies were learned by the operator over time using a sheep-herding metaphor. For example, the school could be controlled well by keeping the predator behind the school while they were moving forward and then circling them to make small adjustments to their direction. This motivates future work on autonomous predation.

5.4 Cohesiveness

For the biomimetics model, we ran experiments for two influence models, leader and predator, and two topologies, nearest neighbor and bio-mimetic. For each influence model and topology pair, we had a human control the leader/predator and we had the leader/predator perform a prescribed, autonomous

zig-zag path. We ran each experiments for two minutes and recorded the A_t and B_t matrices at one second intervals. The results in the previous section used $R^{\text{pred}} = R^{\text{lead}} = 7$, but the results in this section use $R^{\text{pred}} = R^{\text{lead}} = 30$, meaning that the human had quite a bit more spatial influence over the behavior the group.

In terms of performance under human influence, the fish were able to consume more resources under leader-based influence than predator-based influence. Subjectively, the leader-based control was also much easier for the human to use than predator-based influence. Although these results are interesting, it is more interesting to evaluate why leader-based control works better.

To evaluate this, we first consider whether the difference in cohesiveness can explain the difference in performance. We computed $\mathcal{A}$ as the sum of the A_t matrices over the entire two minutes of the simulation for all combinations of influence model and topology pair. Notice from Figures 25-27, that the $\mathcal{A}$ matrix, which was computed over 120 time steps, is approximately 120 units tall. This indicates that, regardless of topology and influence style, the agents tended to form and maintain a parallel group phase in which they stay close enough to each other that most agents were influenced by almost every other agent.

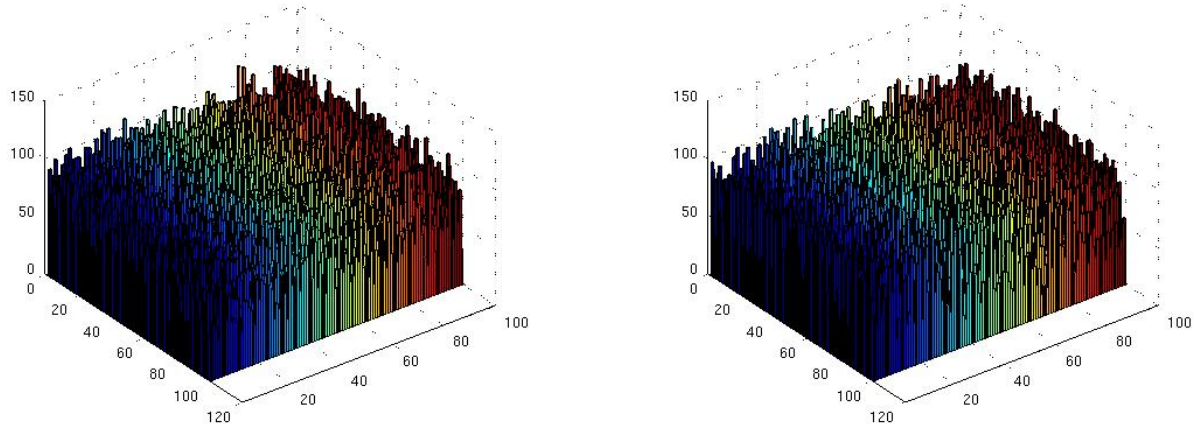

(a) Leader

(b) Predator

Figure 25: $\mathcal{A}$ using human influence under the metric-based topology for (a) leader model and (b) predator model

Note that the A_t coherence matrix described earlier is defined from the f matrix assuming homogeneous agents. The implementation in this section adds a new switch that causes agents near a leader/predator to ignore other agents and just attract to/flee from the leader/predator. For the data in this section, we constructed the A_t matrix assuming that an agent was connected to any other within

attraction or repulsion range regardless of whether either agent was actually ignoring the other because of the presence of the leader/predator.

This is important because the predator can cause agents to separate from each other during periods when they are being "chased" by the predator, since they ignore each other when the predator is within range. Nevertheless, cohesiveness stays strong because the radius of attraction exceeds the radius of predation, meaning that once agents are out or range of the predator they stay connected to each other, even if they are on opposite sides of the predator. This makes cohesiveness differences between leader and predator models very small, which is important since it allows the human to influence the group as a whole rather than as individuals.

There is more discrepancy in the $\mathcal{A}$ histograms when agents use autonomous zig-zag control. The zig-zag controller is a type of worst case operator input, because the leader/predator moves in a zig-zag pattern that disrupts the nominal topology as much as possible. Results, shown in Figure 26,

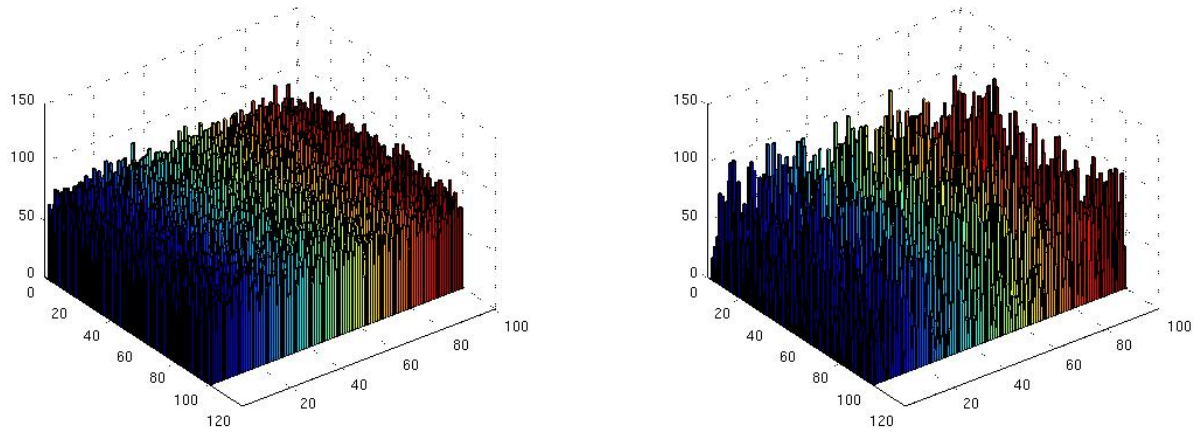

(a) Leader

(b) Predator

Figure 26: $\mathcal{A}$ using the zig-zag controller under the metric-based topology for (a) leader model and (b) predator model

indicate that leader-based influence maintains cohesion better than predator-based influence when the operator behaves in a somewhat worst-case manner. This illustrates that it is possible for a operator to "de-cohere" a group using predator-style interaction.

The experiments above were repeated using a nearest neighbor topology with the neighborhood size set to nine. Regarding the differences between nearest neighbor and metric-based topologies, nearest neighbor topologies were, subjectively, slightly easier to control and performed slightly higher, though remember that this is for a single human operator in a small number of trials. There appears to be more

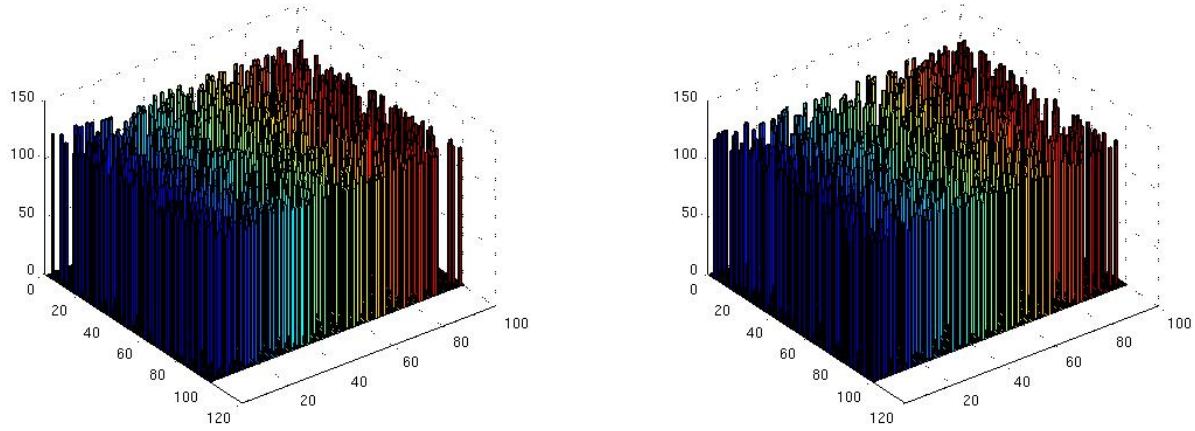

(a) Leader (b) Predator

Figure 27: $\mathcal{A}$ using human influence under the nearest neighbor topology for (a) leader model and (b) predator model

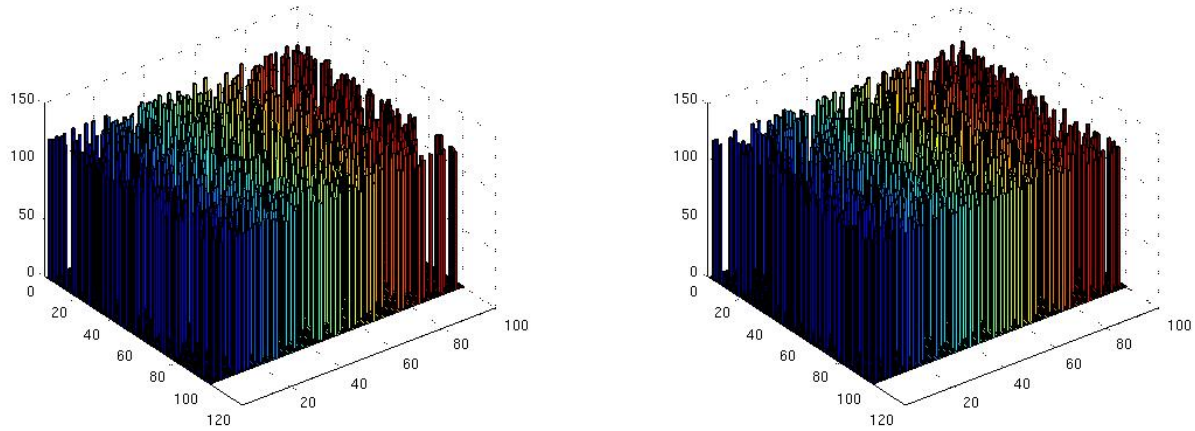

(a) Leader (b) Predator

Figure 28: $\mathcal{A}$ using the zig-zag controller under the nearest neighbor topology for (a) leader model and (b) predator model

uniformity in the nearest neighbor topologies under human control, as illustrated in Figure 27, than in the metric-based topologies in Figure 25, and there is almost certainly more uniformity with predator-style influence in the nearest neighbor topologies than metric-based control under zig-zag control; compare Figure 26 to 28. This is consistent with results from the previous section that indicate that nearest neighbor topologies are more likely to form connected topologies, even when agents are moving and are being perturbed by an external influence.

Results from physicomimetics show similar performance. Neither result suggests that cohesiveness is necessary for high performance, but since performance depends on having more than one agent in proximity to a food source it is unlikely that agents that are not part of a collective would perform as well as those that maintain a highly parallel group.

5.5 Manageability: Leaders and Predators

In this section, we evaluate whether manageability is the reason that leader-based influence outperforms predator-based influence. Manageability is reflected in how well humans maintain influence over agents. Under the physicomimetics model, we measured the $\mathcal{B}$ matrix by recording the number of time steps the human was able to influence the agents. Figure 29 shows that the agents using the PLM model have sustained human influence through the leader agents and hence their influence magnitude is higher compared to that of VPM model which has very low sustainable human influence through the virtual agent. The VLM model performs in between PLM and VPM models.

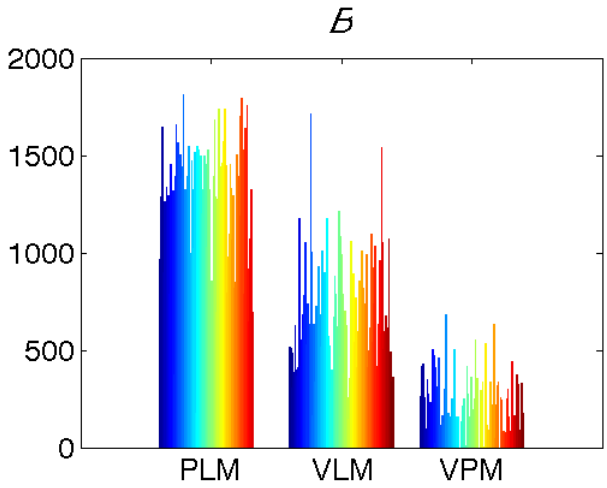

Figure 29: Human influence through leader or virtual agent for different models

The $\mathcal{B}$ matrices under biomimetics showed very similar trends. Using data from the experiments in

the previous section on biomimetics, we recorded the number of time steps that each agent was being influenced by the predator or leader. Results are shown in Figures 30-31. The important thing to notice in these figures is the scale on the y-axis. The scale on the leader matrices suggests that nominal agents were under the influence of the leader about 75% of the time for the metric-based topology, and higher than that for the nearest neighbor topology. By contrast, predators were able to sustain influence over the metric-based topology less than about 1/3 of the time and over the nearest neighbor topology even less than that. Simply put, leaders sustain influence longer than predators, and this is why it is easier to influence a cohesive swarm using a leader than it is with a predator.

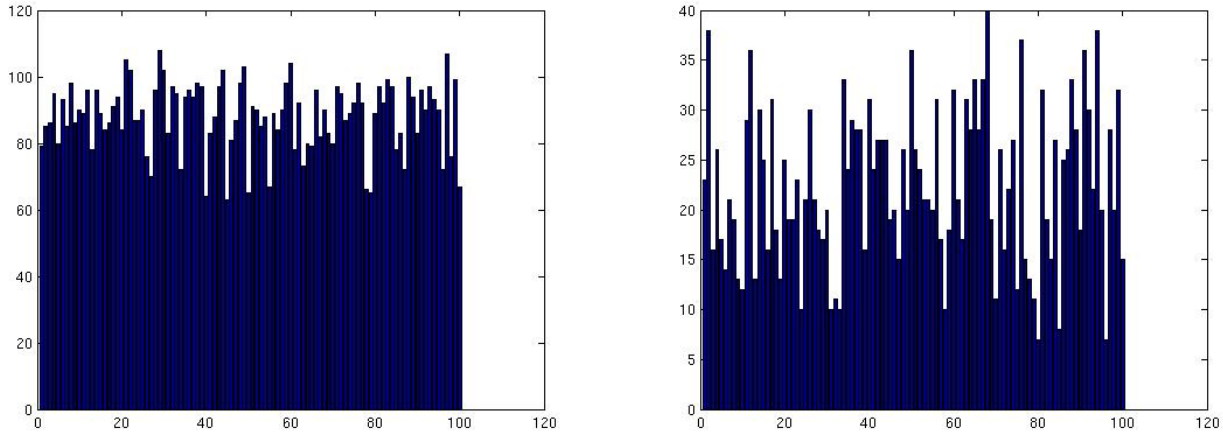

(a) (b)

Figure 30: $\mathcal{B}$ under the metric-based topology for (a) leader model and (b) predator model

To reinforce how leader models sustain influence we computed a power spectral density for $\mathcal{B}$ for various influence styles. For each agent, we created a time series of the number of new agents the leader interacted with plus the number of agents the leader no longer interacted with at each time index. We then computed the power spectral densities for these time series; see Figure 32. Note how the scales of the two plots are different, indicating that the number and frequency of changes for predator models is much higher than for leader models. Simply put, predators cause agents to scramble, making it more difficult to sustain influence.

5.6 Manageability: Metric-Based and Nearest Neighbor Topologies

Ballerini noted that some flocks of birds appear to have local connections that are based on a handful of their nearest-neighbors rather than all birds within a fixed distance [2]. To help understand this, the

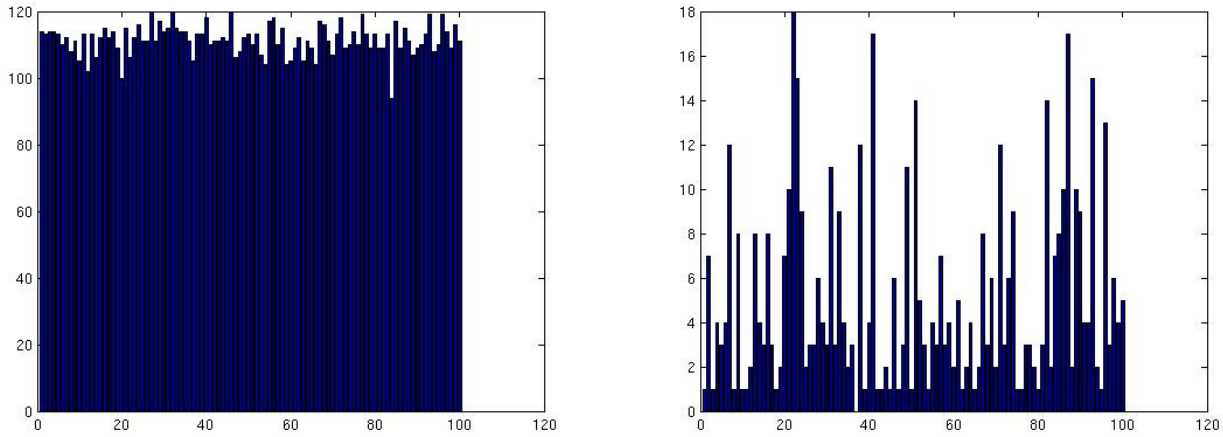

(a) (b)

Figure 31: $\mathcal{B}$ under the nearest neighbor topology for (a) leader model and (b) predator model

above simulations for biomimetic agents included both nearest-neighbor in addition to the metric-based topologies. The results for the bio-mimetic and physico-mimetic agents were similar, so we just present results for the former.

Figure 33 shows the power spectral densities for the nearest neighbor topologies. In comparing Figure 32 to Figure 33, note that the scale of the latter is much smaller. This indicates that neighborhoods change much less often for neighborhood-based topologies than metric-based. Simply put, neighborhood-based topologies are more cohesive, which is consistent with Ballerini's observation from nature [2]. More importantly, this is consistent with subjective observations that nearest neighbor topologies are easier for a human to manage than metric-based topologies.

6 Managing Other Phases Using Non-Switching Controllers: Leaders versus Stakeholders

From the previous sections, we can conclude that it is easier for a human to guide a team to a particular location through leadership than through predation, at least for highly dynamic groups (in Couzin's model) and for isotropic groups (in the physicomimetic model). Simply put, leaders act on a collective by attracting agents to them, so leaders tend to sustain influence over agents; predators act on a collective by repelling agents away from them, so predators tend to be unable to sustain influence. This is not to suggest that predators don't have a role, indeed a single predator can be used to split

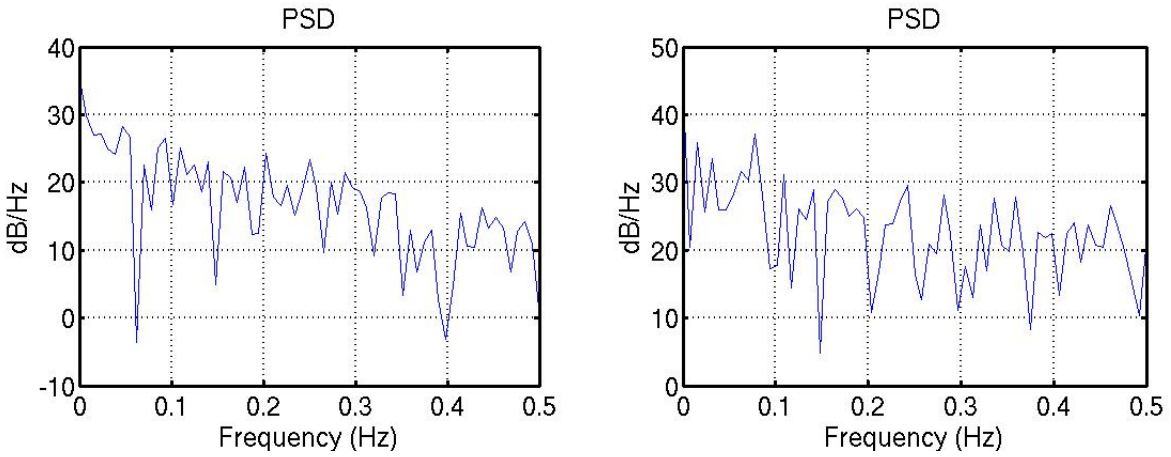

(a) Leader (b) Predator

Figure 32: Power spectral density of $\mathcal{B}$ in the metric-based topology for (a) leader model and (b) predator model .

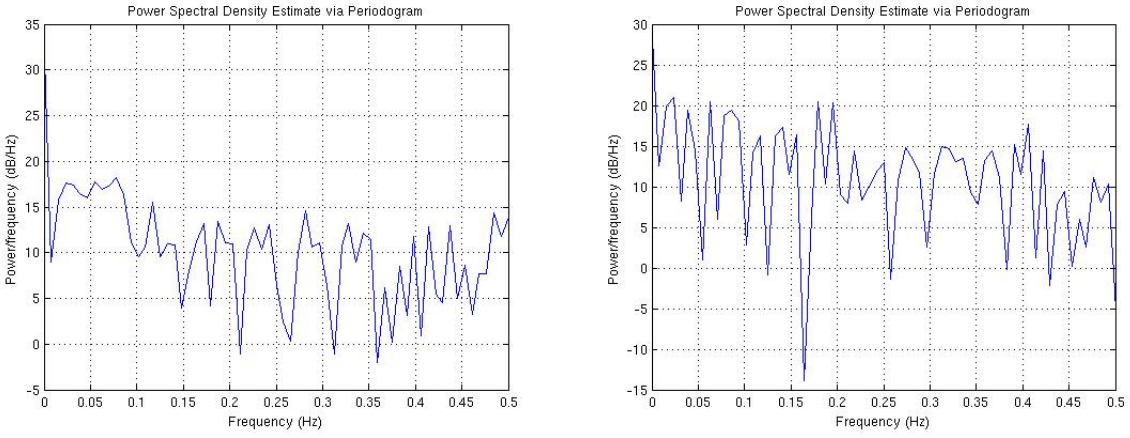

(a) Leader (b) Predator

Figure 33: Power spectral density for nearest-neighbor topologies of $\mathcal{B}$ for (a) leader model and (b) predator model .

a group or to cause a group to avoid an area of risk, and multiple predators can be used to guide collective behavior. Rather, it suggests that it will probably be easier for a human to lead a collective to a desired location than chase them there.

In this section, we address the problem that arises when we ask whether leaders can guide a swarm or a toroid phase. Subjective observations indicate that it is pretty hard to lead a toroid phase; the leader agent(s) tend to travel to a goal leaving agents in the toroid behind.

The results of the previous section indicate that the main reason that leaders can guide a collective better than predators is that leaders can sustain their influence for longer periods of time. This suggests a different type of agent influence: *stakeholder*. A stakeholder tries to sustain influence by remaining in close proximity to the collective, allowing itself to be influenced by the collective as well as to influence it. We implement this by determining both the desired direction using Couzin's equations as well as the direction to a fixed food source. The desired direction for using Couzin's equations produces $\hat{\theta}_i(t + \Delta t)$ computed from Equation (10). This is then added to $\hat{\theta}_i^{\text{food}}(t + \Delta t)$ to produce $\tilde{\theta}_i(t + \Delta t) = \hat{\theta}_i(t + \Delta t) + 0.8\hat{\theta}_i^{\text{food}}(t + \Delta t)$ where the 0.8 weighting is chosen subjectively to encourage the stakeholder to stay near the group more than going to the goal.

Qualitatively, a stake holder tends to occupy a location in a collective structure that is nearest the food. For example, the stakeholder joins the torus (when the group is in the torus phase) but tends to spend most of its time on the side of the torus nearest to the food. Similar behavior is observed for the swarm and parallel phases.

We conducted an experiment designed to evaluate if and when a stakeholder could influence the different collective phases of Couzin's model, and to compare a stakeholder's influence to that of a leader. To do this, we used the parameters given in Table 2. We then use zero or more stakeholders and zero or more leaders of various speeds to see how the different combinations influenced the position of the group.

6.1 No Leaders

As a point of reference, it is useful to show the final positions of the centroids of each phase after 45 seconds in the absence of any type of leader. This is shown in Figure 34(a). The legend indicates which symbols correspond to which phase. It is useful to give a short description of the collective movement patterns of each phase in the absence of any type of leader of stakeholder.

- Swarms tend to stay closely coupled together and rarely move far from the origin.
- Toroids also tend to stay closely coupled together and rarely move far from the origin. Occasionally, a toroid will decohere with a few agents spinning out of the toroid, causing the centroid of the toroid to be far from the origin.

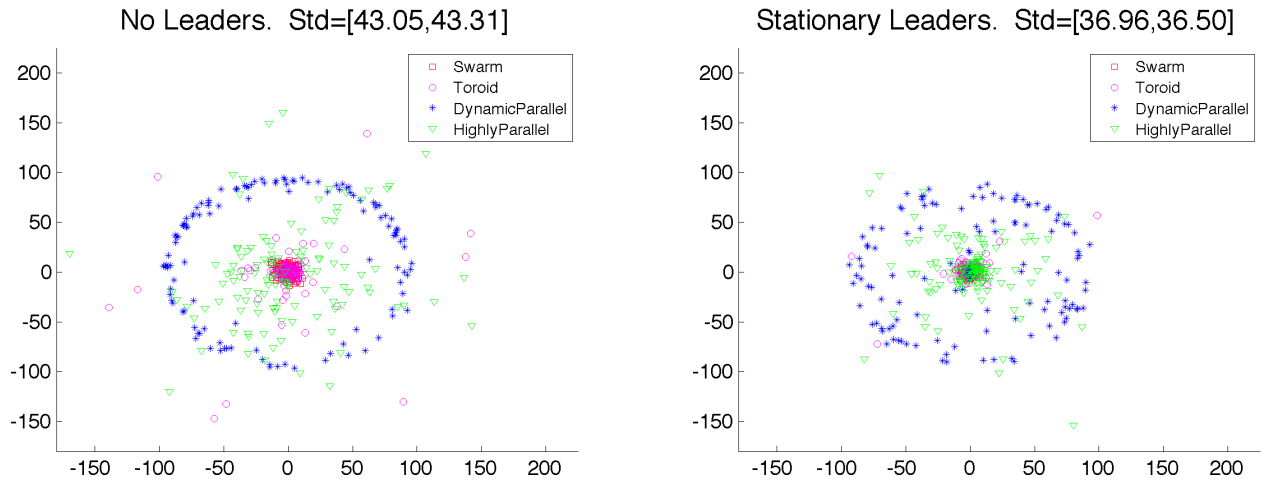

(a) (b)

Figure 34: Distribution of centroids for various phases with (a) no leaders and (b) stationary leaders.

- Dynamic parallel groups quickly cohere and then start traveling in a somewhat direct line in some random direction. This is indicated by the blue triangles clustered in a circle around the origin, suggesting that dynamic parallel groups all travel approximately the same straight line distance from the origin.
- Highly parallel groups quickly cohere, but they rarely travel in a straight line from the origin, usually making several turns causing a distribution that is on the inside of the "maximum distance" circle defined by the dynamic parallel groups.

6.2 Stationary "Leaders"

It is interesting to contrast the scatterplot in the absence of a leader to the scatterplot in the presence of agents that never move. Figure 34(b) shows the final positions of the centroids of each phase after 45 seconds when some agents do not move from their initial stationary positions for various numbers of stationary agents. These scatterplots are for the number of leaders shown in the first line of Table 7. By comparing Figure 34(a) to 34(b) we quickly conclude that leaders make a difference: toroids are more likely to stay coherent and highly parallel groups stay closer to the origin. Note that these observations can be observed in corresponding animations of the group dynamics, but the computation of the centroids is biased; for ease of computing, all agents including both stationary and non-stationary agents, were used to compute the centroid. This biases the locations of the centroids

so that all centroids are closer to the origin. Although this bias is undesirable, it will be important to allow it once we consider moving leaders and stakeholders.

Number of Leaders/StakeHolders	$\{0, 1, 2, 3, 5, 10, 15, 25\}$
Speed scaling factors	$\{0, \frac{1}{5}, \frac{1}{3}, \frac{1}{2}, 1, 2\}$

Table 7: Parameters used to create the scatterplots of centroids for each collective structure.

6.3 Leaders and Stakeholders

For moving leaders/stakeholders, the speed of the leader/stakeholder is determined by taking the speed of the nominal agents from Table 2 and multiplying by the speed scaling factors from the second line of Table 7. For example, the speed scaling factor of 0 means that the leaders/stakeholders are stationary, and the speed scaling factor of 1 means that the leaders/stakeholders are traveling at the same speed as the other agents.

The centroids for the different groups are computed, as in the previous section, by including the locations of the leaders and stakeholders. All leaders move toward a food source located at location $(30, 30)$ regardless of the behavior of other agents. All stakeholders are attracted toward a food source at the same location. Since the stakeholders are part of the collective, they contribute to the ultimate location of the centroid and should be used in computing the centroid of the collective at the end of the 40 second simulation. The leaders are not technically part of the collective since they ignore other agents, but we include their final location in the computation of the centroid so as to make sure that there is no unfair advantage given to the stakeholders in comparing scatterplots of stakeholders and leaders.

For both stakeholders and leaders, one and two leader/stakeholder agents had very little impact on the ultimate distribution of the phases; one or two such agents don't have much power to influence the behavior of a collective of 50 other agents. Three leaders/stakeholders in a collective does start influencing the collective toward the origin, producing an average location of the centroids across all phases that is shifted away from the origin toward the food source at location $(30, 30)$. There is not a substantial difference between how three stakeholders or three leaders influence the final locations of the collectives' centroids for three leaders.

But once the number of leaders is five or more, constituting 25% of the number of agents in the collective, we start to see a difference in how much influence stakeholders have in comparison to leaders, even given the fact that centroids for leaders should be biased more toward the food location since all leaders end up at that location. This is easily seen in Figures 35-37. This manifests itself in (a) centroid locations that are closer to the food source for the stakeholders than for the leaders and (b) smaller standard deviations of centroids over all the trials. To emphasize the differences in the means, two solid

lines intersect at the location of the food sources and two dashed lines intersect at the location of the mean of the centroids.

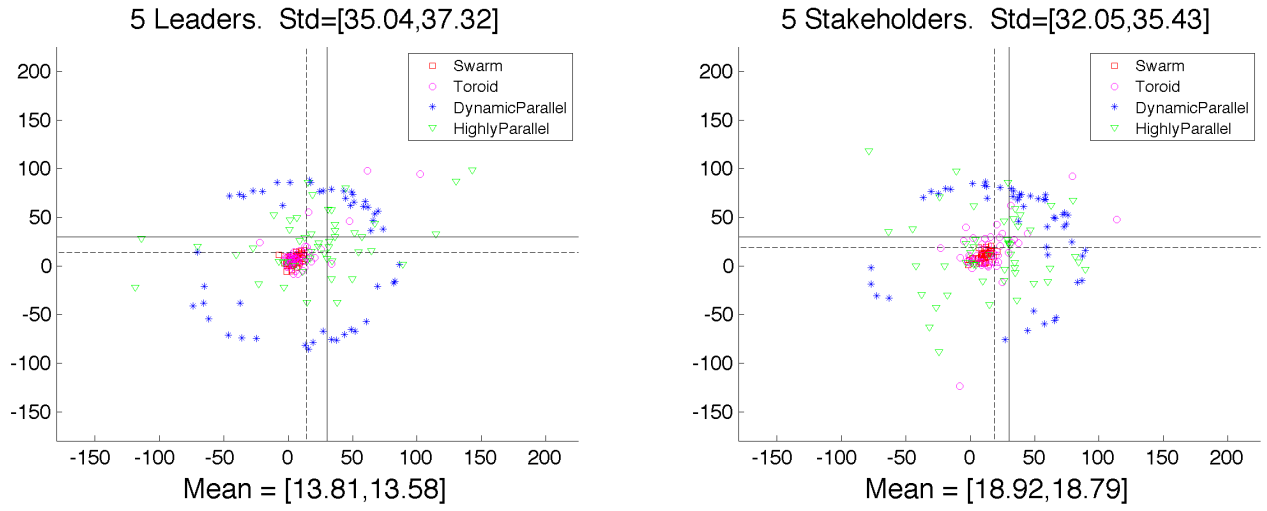

(a) Leader (b) Stakeholder

Figure 35: 20% of the agents try to influence the collective. (a) Five stakeholders influence collective behavior more than (b) five leaders.

If the conclusion from the previous section is that leaders sustain influence better than predators, and therefore have more influence over a parallel group phase, then the conclusion of this section is that stakeholders influence all groups better than leaders presumably because by being influenced by other agents a stakeholder can sustain its influence longer. This conclusion is consistent with claims from coactive design, "... coactive [design] is meant to convey the recriprocal and mutually constraining nature of actions and effects that are conditioned by coordination" [18]. Reciprocal and mutually constraining interactions allow influence to be sustained for longer, making coordination more likely. This seems to hold regardless of the phase of the collective, though some phases (e.g., toroid) require more stakeholders/leaders to induce movement.

Before ending this section, it is useful to make a couple of passing notes. First, note that some type of quorum signaling would go a long ways toward eliminating the problem of guiding swarms and toroids. If quorum signaling could be used, then the agents could switch to a highly parallel phase to move to a new location, and then resume their previous behavior. Second, note that the stakeholders seem appropriate for leading to a static point, but tracking a dynamic point will be problematic. In the simulation results, each simulation lasted for 40 seconds because otherwise littler progress is made toward the goal. We hypothesize that more expressive topologies (such as quorum signaling) or some sort of coordinated, multi-predator influence will be necessary to produce more highly responsive teams.

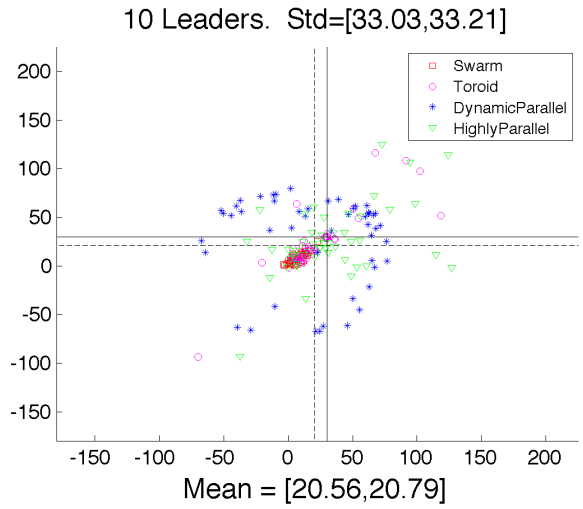

(a) Leader

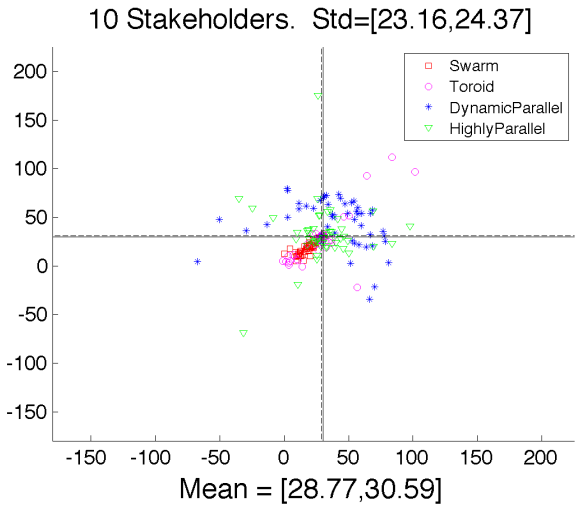

(b) Stakeholder

Figure 36: 40% of the agents try to influence the collective. (a) Ten stakeholders influence collective behavior more than (b) ten leaders.

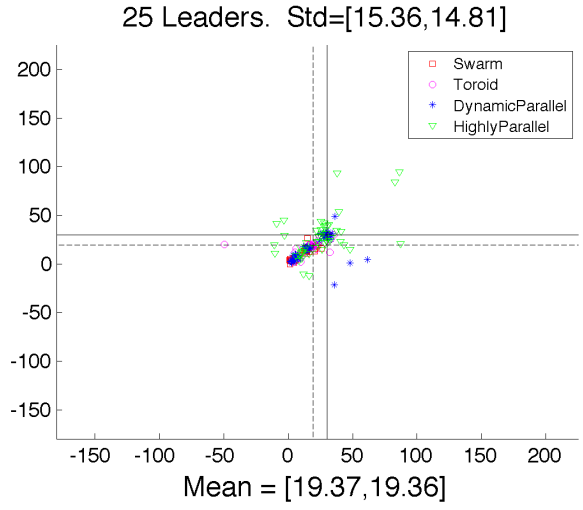

(a) Leader

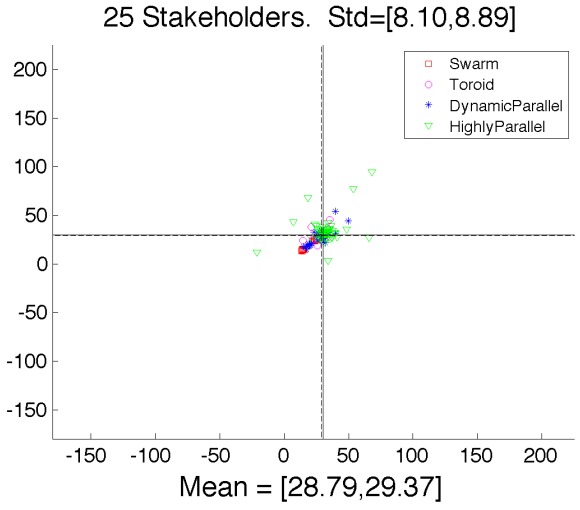

(b) Stakeholder

Figure 37: 50% of the agents try to influence the collective. (a) Twenty-five stakeholders influence collective behavior more than (b) twenty-five leaders.

In a previous section, we explored how switching controllers could be used to cause nominal agents to be influenced exclusively on leaders or predators. This switching behavior was not a problem in dealing with parallel groups and isotropic phases because the agents quickly resumed a structure that mimicked the original collective structure. In this section, we treated the influence of leaders/predators as equal to that of other agents – there was no switch to attract to a leader of flee from a predator, just an extra attraction toward a leader or an extra repulsion from a predator. One of the problems in removing the switch is that the agents either tended to ignore the leader or they tended to fragment thereby destroying coherence. In the next section, we explore how agents can dynamically switch between different phases if we force them to stay coherent.

6.4 The Dynamics of Leadership

In the previous section, the radius of attraction was low enough that the collective structures sometimes fragmented, leaving groups of agents that were not within each others radii of influence. In this section, we explore what happens when we guarantee that agents stay coherent by making the radius of attraction infinite and by eliminating the blind spot. This will allow us to see what happens when the influence of leaders is sufficiently high to force at least some agents to follow. Interestingly, doing so produces emergent behavior that is sometimes effective in automatically causing collectives to switch from one structure to another. In the interest of space, we only present a sample of results focusing on results that illustrate potentially useful phenomena. Most observations are qualitative, but we include some quantitative data when possible.

6.4.1 Leading Toroids

Let $s_\ell = 1$ denote the speed of the leader and consider what happens when the speed of the nominal agents is faster than the leader. Set the parameters of Couzin's model to $R^{\text{rep}} = 1$, $R^{\text{ori}} = 3$, and $R^{\text{att}} = \infty$, and eliminate the blind spot by setting $\phi = 180°$. Consider a group of $N = 80$ agents, and consider groups of 30 and 50 leaders. In this set of simulations, leaders are added to the nominal group, so when a group of $N = 80$ has 30 leaders there are a total of 110 agents in the simulation. Thus, for $N = 80$, 30 leaders constitute approximately 27% of the total population and 50 leaders constitute approximately 38% of the total population. The leaders travel due east at constant speed, but travel more slowly than the centroid of the torus. Note that leaders are excluded from any calculations of centroid, moment, and polarization.

When agent speeds are moderately higher than the speed of the leaders ($s = 10$ and $s = 5$), the torus tended to follow the leaders but had difficulty keeping up with it (this is true for $\omega \in \{40, 70, 100\}°/sec$). When the torus got behind the leaders, it would break from the torus formation, temporarily form a dynamic parallel group, catch up with the leader, and then resume the torus.

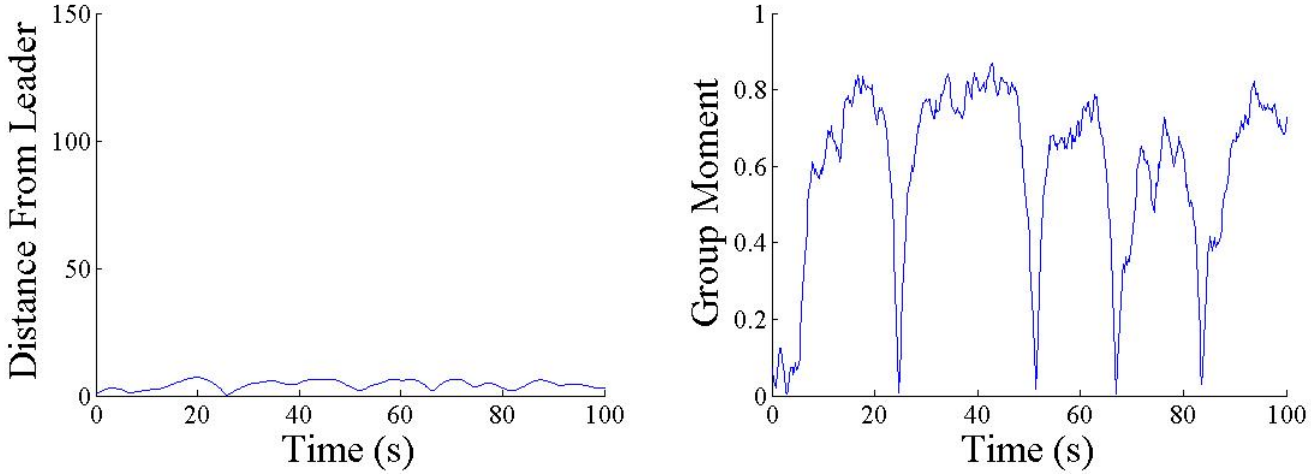

(a) Distance from Leader

(b) Angular Momentum

Figure 38: When the distance of the group gets large, the group changes from a torus to a parallel group.

Figure 38 illustrates this emergent switching behavior, where the switching occurs at the level of the collective rather than at the individual. We choose to illustrate this switching using group angular momentum, since angular momentum is very high for a torus and very low for a parallel group. Although the individual agents are not programmed to switch behaviors when the leaders get too far away, but the collective structure of the agents automatically changes when the leaders get too far away. The figure illustrates this in a series of changes in group angular momentum that correspond to large distances from the centroid of the toroid and the locations of the leaders.

This natural emergence of switches in collective structure could be very useful for HuBIRT design.

6.4.2 Leading Dynamic Parallel Groups

Let $s = 5$ denote a constant speed of each agent and set Couzin's parameters to form a dynamic parallel group: $R^{\text{rep}} = 1$, $R^{\text{ori}} = 10$, and $R^{\text{att}} = \infty$, and with no blind spot by setting $\phi = 180°$. Since dynamic parallel groups tend to choose a direction and go, it is interesting to evaluate what happens when the speed of the leaders is less than or equal to the speed of the agents, that is $s_\ell \in \{1, 3, 5\}$.

In general, when $s_\ell < s$ the group tended to follow random patterns around the leader; see Figure 39(a). However, for some parameter sets, orbits (Figure 39(b)) and half-orbits (Figure 39(c)) emerge. Following orbits could be useful in a perimeter monitoring task, and following half-orbits could be useful in flank, front, or rear protection. Interestingly, when $s = s_\ell$ the agents line up and follow the leaders, but gradually fall farther behind because the randomness in the agents cause them

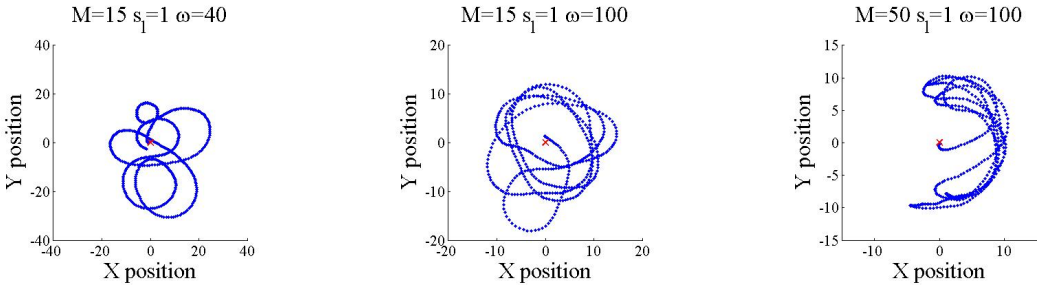

(a) Random (b) Orbits (c) Half Orbits

Figure 39: Patterns displayed by dynamic parallel groups as they follow a leader.

to fail to follow perfectly.

6.5 Leading Randomly DynamicTopologies

Since nearest neighbor topologies were not as expressive as metric-based topologies, it was not possible to explore how the different phases induced by the nearest neighbor approach were influenced by leaders. By contrast, since probabilistic topologies are as expressive as metric-based topologies, it is useful to see how leaders influence each collective structure. This section illustrates that leaders can guide each collective structure, assuming that there are enough leaders. As before, parallel phases are easier to lead than swarms and toroids. Moreover, swarms and toroids may change their structures under the influence of leadership, especially if the leaders move too fast. This latter result emphasizes that it may be necessary for leaders to be influenced by nominal agents if it is essential for the groups to maintain their structures, but that switching structures is possible if desired.

In keeping with the other models presented in this section, we let nominal agents be influenced by leader agents as if they were other nominal agents, that is, nominal agents do not switch their behaviors depending on whether a neighbor is another nominal agent or a leader agent. Simulations in this section were performed in the same manner as those in the previous section using a small subset of successful simulation parameters, and varying the number of leaders and the speed of those leaders.

6.5.1 Torus Groups with Random Neighbour Assignments

Simulations were performed using the parameters listed in Table 8; the notation s_ℓ denotes the speed of the leaders, all of whom travel with the same orientation due east. Note that the radii are set so that the nominal group structure is a torus. We say that a torus is successfully lead if the distance from the group to the leader was approximately constant (that is, the group does not fall too far behind the leader) and the group angular momentum remained above 0.5 after successfully rising above 0.75.

This latter criterion simply means that the group forms a clear torus at some point (indicated by the high angular momentum) and then retains the general characteristics of a torus for the remainder of the simulation.

Parameter	N	n	M	R^{rep}	R^{ori}	R^{att}	ϕ	ω	s	s_ℓ
Units	integer	integer	integer	unit	unit	unit	degrees	degrees/s	unit/s	unit/s
Values	80	1,5,15,30,50	50	1	3	∞	180	100	15	1

Table 8: Table of explored parameters for torus configuration with random neighbor assignments.

When neighborhoods were small, $n < 50$, the group always formed a swarm rather than a torus (Figure 40), though the group tended to become more aligned as n increased (Figure 41). In all trials, the group closely followed the leader. For $n = 50$, weak toruses formed that successfully followed the leader (Figure 42) in all trials. It is useful to note that unlike simulations without a leader, the torus forms more slowly in the presence of several leaders and tends to have less structure.

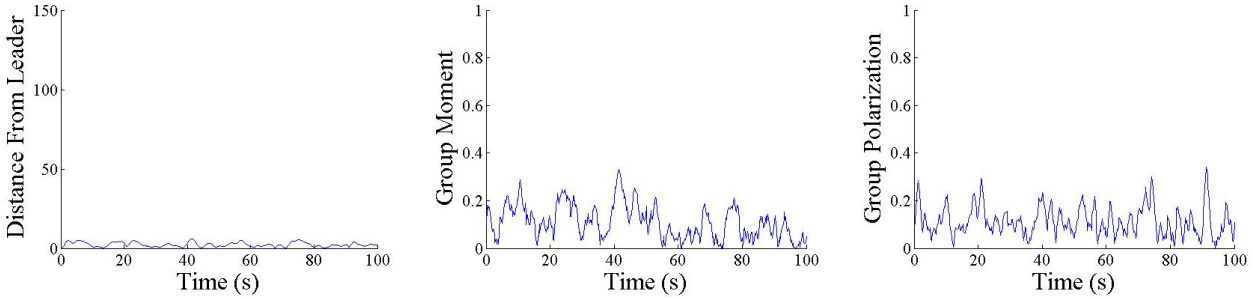

(a) Distance From Leader (b) Group Moment (c) Group Polarization

Figure 40: Plots showing the swarm like leader following behaviour of a group with torus parameters and $n = 15$ randomly assigned neighbors. (a) demonstrates that the group was able to successfully follow the leader. (b) and (c) show that the group formed a swarm instead of a torus.

In conclusion, a torus can be formed and lead using random neighbor assignment for relatively large values of n. However, the torus takes longer to form and has a lower group angular momentum than a torus formed using Couzin's model in the absence of leaders.

6.5.2 Dynamic Parallel Groups with Random Neighbours

Dynamic parallel group simulations were run with random neighbor assignments using the parameters listed in Table 9. Note how the radii are set so that the nominal group structure is a dynamic parallel group.

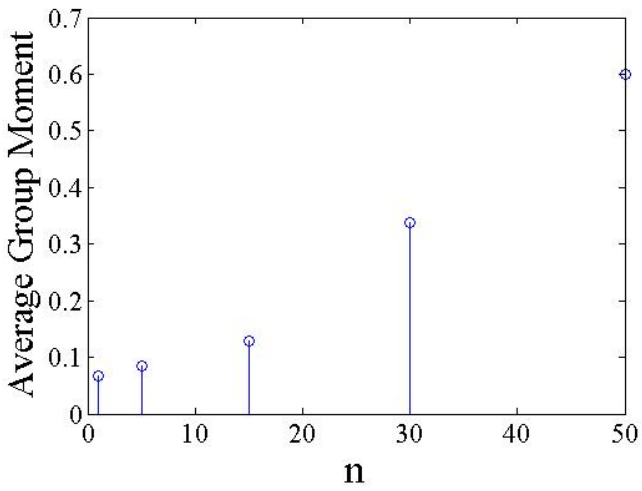

Figure 41: Average group angular momentum of a group with torus parameters for various values of n.

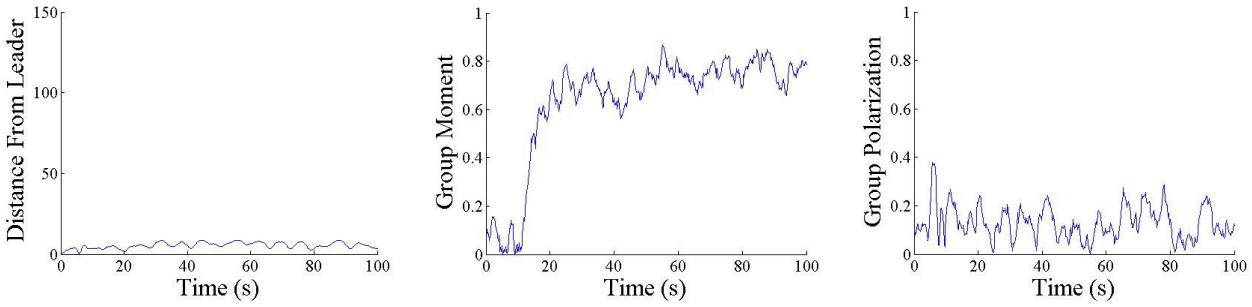

(a) Distance From Leader (b) Group Moment (c) Group Polarization

Figure 42: Group distance from the leader, moment, and polarization for a single trial of leading a torus with $n = 50$ randomly assigned neighbors. (a) demonstrates that the group was able to successfully follow the leader. (b) and (c) show that a weak, but stable, torus formed.

Parameter	N	n	M	R^{rep}	R^{ori}	R^{att}	ϕ	ω	s	s_ℓ
Units	integer	integer	integer	unit	unit	unit	degrees	degrees/s	unit/s	unit/s
Values	50	1,5,15,30	1,5,15,30,50	1	10	∞	180	40,70,100	5	1,3,5

Table 9: Table of explored parameters for a dynamic parallel configuration with random neighbor assignments.

In all cases, the group responded to the leader and tended to maintain a parallel structure. For $n = 1$, swarms tended to form instead of dynamic parallel groups. For all other values of n, dynamic parallel groups successfully formed. Interestingly, the dynamic patterns that the group displayed were very similar to those seen with Couzin's model, although the flight patterns tended to be a little more noisy; compare Figure 43, produced by random topologies, to those in Figure 39, produced by metric topologies.

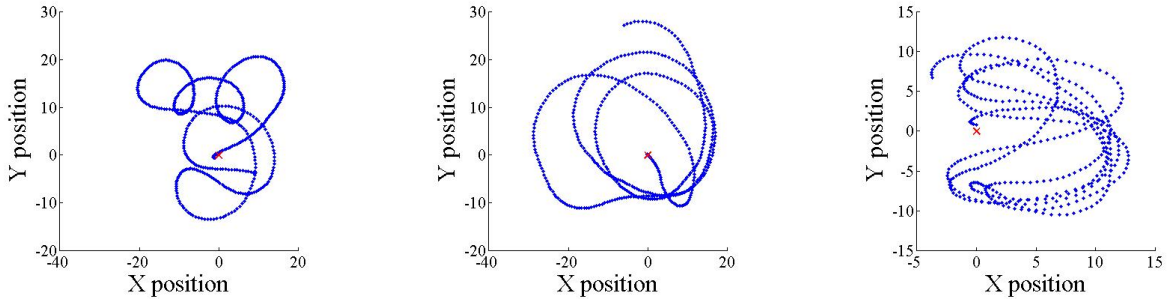

(a) $M = 15$ and $\omega = 40$ (b) $M = 5$ and $\omega = 100$ (c) $M = 50$ and $\omega = 100$

Figure 43: Plot of various flight patterns seen in dynamic parallel groups with random neighbor assignments with $n = 15$ and $s_\ell = 1$. Other varied parameters are shown in individual figure captions. The blue dots show the position of the swarm relative to the leader over a 100 second simulation. The red X in the figure represents the leader with a relative position of zero. (a) shows the random flight patterns exhibited by many parameter sets. (b) and (c) show an orbit and half orbit respectively. Compare to Figure **??**

In conclusion, dynamic parallel groups can be formed with relatively few randomly assigned neighbors.

6.5.3 Highly Parallel groups with Random Neighbours

Highly parallel group simulations were run with random neighbor assignments using the parameters listed in Table 10. As before, note how the radii are set so that highly parallel groups are likely to form. Simulations were performed in the same manner as those in the previous section.

Parameter	N	n	M	R^{rep}	R^{ori}	R^{att}	ϕ	ω	s	s_ℓ
Units	integer	integer	integer	unit	unit	unit	degrees	degrees/s	unit/s	unit/s
Values	50	1,5,15,30	1,5,15,30,50	1	50	∞	180	100	5	1,3,5

Table 10: Table of explored parameters for highly parallel groups with random neighbor assignments.

For all parameter sets, including when $n = 1$, the group successfully formed a highly parallel group. For all values of n, M, and s_ℓ the leader successfully maintained control of the group. Groups with higher n tended to respond more quickly to leaders than groups with lower n. For all values of $n \geq 5$ both standard and half orbit patterns were observed, although the groups with lower n had less well defined patterns; see Figure 44.

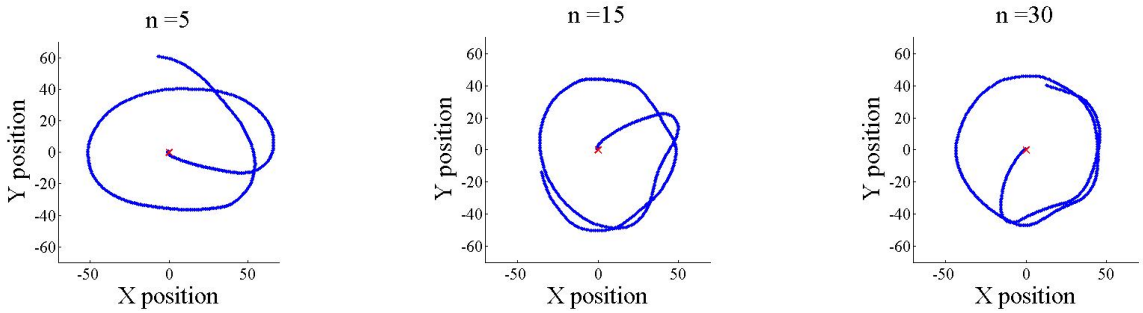

(a) (b) (c)

Figure 44: Plots for standard orbits for various values of n. The figure shows that fewer agents lead to a less defined pattern in response to the leader.

In summary, highly parallel groups can be controlled with relatively few randomly assigned neighbors, though models with more neighbors were more responsive than those with fewer neighbors.

7 Conclusion and Summary

In this report, we presented a simple model of human interaction with bio-inspired robot teams. This model had the form of

$$x^i_{t+1} = f^i(x^i_t, \mathbf{x}^{\neg i}_t) + d^i(x^i_t, u^{\text{op}}),$$

which describes the evolution of agent i's state, x^t_i over time as a combination of two parameters: interagent dynamics encoded in $f^i(x^i_t, \mathbf{x}^{\neg i}_t)$ and dynamics between operator input and agent i's state encoded in $d^i(x^i_t, u^{\text{op}})$. This model says that the next state of an agent is an additive combination of interagent dynamics and response to operator input.

The first set of experiments identified conditions under which a bio-inspired team was likely to exhibit cohesive behavior, and then replicated prior work that described the qualitative structures that emerge when behavior is cohesive. This included a comparison of metric-based and nearest neighbor topologies, demonstrating that nearest neighbor topologies were more likely to be cohesive but also less expressive than metric-based topologies. This section also showed that randomly dynamic topologies

could be as expressive as metric-based topologies, with a passing observation that when the number of neighbors is high enough that the group may display some probabilistic connectivity properties.

The next set of experiments was designed to understand when a human could be injected into the team to guide team behavior. This set of experiments never actually used the full model. Instead, we extended both physicomimetics and biomimetics in two ways. First, we introduced a leader or predator agent type for which $f^i = 0$ and $g^i \neq 0$. Second, we modified the nominal agent type, those for which $f^i \neq 0$ and $g^i = 0$, to respond differently to leaders/predators than they would to other nominal agents. Leaders and predators ignored all other agents and only paid attention to operator input, and nominal agents ignored all operator input and only paid attention to other agents.

The reason that these models allowed a human to effectively manage a team is that f^i could be designed so that the leader or predator agent could get the nominal agents to do what the operator wanted done. The way that this was done was through extending the switching control demonstrated in both Couzin's model and the physicomimetic model: an agent switches its control depending on what is happening around it. If other agents are in agent i's repulsion zone, all other agents are ignored until agent i is away from those agents. If the repulsion zone is clear, then agent i switches to a completely different type of control that is intended to attract and align to other agents. We added another switch. For example, under predator-based control a predator is within a special"predator repulsion zone" then agent i ignores all other agents and just flees from the predator. If no predator is present, then agent i switches to one of the control modes encoded in Couzin's or the physicomimetics model. Leader-based control performed similarly.

One consequence of this switching controller is that it destroys the phases precisely because agents switch away from behaviors that preserve the phases. This doesn't manifest itself in the first sections because we only really paid attention to highly parallel groups (isometric groups). In the presence of a leader, the group behaved almost the same, falling in behind the leader but still moving as a highly parallel group (or establishing an isometric distribution). In the presence of a predator, the group would deform when the predator was in range, but then resume it's highly parallel structure (isometric distribution).

Results from analyzing this approach to agent design demonstrated that leader-based influence performed better than predator-style influence, precisely because leaders could sustain influence more easily than predators. Moreover, nearest neighbor topologies were easier for a human to manage than metric-based topologies, but these differences were less significant than the differences between leaders and predators.

The next set of experiments removed the individual switching behavior and instead introduced a new type of agent, which we termed the stakeholder. Stakeholders seek to preserve the collective phase by allowing agents to preserve their nominal autonomy, but the way that the human influences the stakeholder includes both f^i and g^i. This special type of agent, the stakeholder, is influenced not only by the human but also by by the other agents. It is for this reason that stakeholders can influence the agents while preserving phase, and sustain their influence with fewer numbers of agents than required

when leaders are used.

The final set of simulations again sought to preserve the phases while injecting human input, that is, to allow a human to lead a team without causing the team to switch out of their phase. This was done by abandoning the switching control in the nominal agents. As before, leader agents only paid attention to the operator by setting $f^i = 0$ and $g^i \neq 0$. But the nominal agents treated leader agents as if they were other agents, being repelled from them if too close, and aligning or being attracted to them if the leaders were in an appropriate zone. This set of experiments provided information on how a human could influence a bio-inspired collective while preserving both the nominal agent autonomy and the collective's phase; these results held for both metric-based and randomly dynamic topologies. Agents may switch phases to follow a leader, but they can also preserve their phase because we have not programmed the phase out of them. Moreover, agent collectives can exhibit emergent behaviors that include switching between structures and also orbit and half-orbit behaviors.

Acknowledgments

This work was partially funded by the Army Robotics CTA program and ONR's Science of Autonomy Program. The opinions expressed in this report do not necessarily reflect the opinions of the funding organizations.

A Couzin's Model

In this section, we present Couzin's biomimetic model [8] and his notation. There are two in this notation from what we've used throughout the rest of this paper, namely that agent index is used as a subscript and time is an argument of the state. Using Couzin's notation, let $\mathbf{c}_i$ be the position vector of agent i and let $\mathbf{r}_{ij} == \frac{\mathbf{c}_j - \mathbf{c}_i}{|\mathbf{c}_j - \mathbf{c}_i|}$. For two dimensions, $\mathbf{c}_i = [x_i, y_i]^T$ where x and y denote the position of agent i in two dimensions. Similarly, let $\mathbf{v}_i$ denote the direction or velocity vector of agent i. This vector is completely described by the angle that the agent is traveling, denoted by θ_i. Thus, $\mathbf{v}_i = s[\cos\theta_i, \sin\theta_i]^T$, where the magnitude of the velocity, s, is constant. Finally, let $\hat{\theta}_i(t + \Delta t)$ denote the desired direction after time interval Δt; note that we use $t + 1$ to denote one time step in the main sections of this report, but use $t + \Delta t$ in this section since dynamics depend explicitly on the time interval.

Couzin's model can now be written as follows:

$$\hat{\theta}_i(t + \Delta t) = \begin{cases} -\sum_{j \neq i}^{N} \alpha_{ij}^{\text{rep}} \frac{\mathbf{r}_{ij}(t)}{|\mathbf{r}_{ij}(t)|} & \text{if } \exists j \text{ such that } \alpha_{ij}^{\text{rep}} = 1 \\ \frac{1}{2}\Big(\sum_{j \neq i}^{N} \alpha_{ij}^{\text{ori}} \frac{\mathbf{v}_{ij}(t)}{|\mathbf{v}_{ij}(t)|} + \sum_{j \neq i}^{N} \alpha_{ij}^{\text{att}} \frac{\mathbf{r}_{ij}(t)}{|\mathbf{r}_{ij}(t)|} \Big) & \text{otherwise} \end{cases} \tag{9}$$

where $\alpha_{ij}^{\text{rep}} = 1$ if agent j is within agent i's zone of repulsion and zero otherwise; $\alpha_{ij}^{\text{ori}} = 1$ if agent n is within agent i's zone of orientation and zero otherwise; and similarly for α_{ij}^{att}. Formally,

$$\alpha_{ij}^{G} = \begin{cases} 0 & |\mathbf{r}_{ij}(t)| \leq R^{G} \\ 1 & \text{otherwise} \end{cases} \quad \text{for } G \in \{\text{rep, ori, att}\}$$

Note that α_{ij}^{ori} is also zero if agent n is in agent i's blind spot.

To include the blindspot, we need only make one modification to α that sets α_{ij} to zero if agent j is in agent i's blindspot.

We can rewrite Equation (9) in a single equation as:

$$\hat{\theta}_i(t + \Delta t) = -\sum_{j \neq i}^{N} \alpha_{ij}^{\text{rep}} \frac{\mathbf{r}_{ij}(t)}{|\mathbf{r}_{ij}(t)|} + \frac{1}{2}(1 - \max_{j \neq i} \alpha_{ij}^{\text{rep}}) \Big(\sum_{j \neq i}^{N} \alpha_{ij}^{\text{ori}} \frac{\mathbf{v}_{ij}(t)}{|\mathbf{v}_{ij}(t)|} + \sum_{j \neq i}^{N} \alpha_{ij}^{\text{att}} \frac{\mathbf{r}_{ij}(t)}{|\mathbf{r}_{ij}(t)|} \Big). \tag{10}$$

Note that the radii of influence are set so that the zone of repulsion is smaller than the zone of orientation, which is smaller than or equal to the zone or attraction. In other words, $R^{\text{rep}} < R^{\text{ori}} \leq R^{\text{att}}$.

Equation (10) provides the desired orientation for agent i. If the agents have momentum, then it is unlikely that they will be able to instantaneously change from one orientation to another. Using a simple model, we suppose that the agent can turn no faster than some discrete angular velocity, denoted by ω. Let θ denote the angle of the current velocity vector and recall that $\hat{\theta}$ denotes the desired angle. Then,

$$\theta_i(t + \Delta t) = \begin{cases} \hat{\theta}_i(t + \Delta t) & \text{if } |\theta_i(t) - \hat{\theta}_i(t + \Delta t)| < \omega \Delta t \\ \text{sign}(\theta_i(t) - \hat{\theta}_i(t + \Delta t))(\theta_i(t) - \omega \Delta t) & \text{otherwise} \end{cases} \tag{11}$$

where Δt is the sampling interval, and where $\text{sign}(x)$ returns the sign of x, using appropriate modulo arithmetic. That is, the new angle equals the desired angle if this can be reached in a single time step given the angular velocity, and otherwise the new angle is the maximum allowable step in the direction of $\hat{\theta}(t + \Delta t)$.

B Task Complexity and Collective Expressiveness

The goals of the previous sections were (a) to characterize some of the ways existing collectives can express collective structure, and (b) to identify mechanisms by which a human could influence simple agent collectives like swarms and flocks. In this section, we outline a framework that will allow us to identify more sophisticated collectives and the types of tasks for which they may be well adapted.

David Sumpter recently called for work that combines mechanistic and functional explanations for animal interaction [37]. He says that "Mechanistic explanations look at how animals interact to produce

group level patterns" and "Functional explanations are based on arguments about why a behaviour has evolved through natural selection." A mechanistic explanation might be in the form of a set of rules that generate collective intelligence, such as the rules encoded in Couzin's model above, and a functional explanation might be in the form of the evolutionary forces that produce, for example, flocking behavior that evolves to avoid predation. He cites the excellent paper by Conradt et al. [7] in which they combine a functional description (how hungry a subset of fish are) with a mechanistic description (how assertive or speedy a fish is) for how small subgroups of fish can lead a larger school to an objective that differs from the goal of the majority.

B.1 Task Complexity

One way to relate the mechanistic approaches of Couzin's model or a physicomimetics model to a functional approach is to identify those elements of the environment that may influence survival. Based on a preliminary review of existing literature, there are at least three qualitatively different types of influences: location/type of resource, obstacles/threats, and homeostasis. Homeostasis refers to the ability to maintain a stable internal state; homeostasis emerges, for example, in insect colonies that must maintain temperature (as in a termite colony or in a beehive).

Clearly, there is going to be some interdependence between these three types of influences, but treating them independently allows us to represent each factor as a separate dimension in a taxonomy. Suppose that we encode the natural complexity of an environment as a probability distribution on the location/type of resource, presence and type of obstacles/threats, and complexity of regulation associated with homeostasis. Denote these three dimensions by ℓ, o, and s, respectively. The joint probability $p(\ell, o, s)$ represents the probability that a collective will need to perform a set of activities that will accomplish tasks encoded in ℓ in the presence of obstacles encoded in o while maintaining homeostasis as encoded in s. In short, $p(\ell, o, s)$ encodes how challenging an environment is, and therefore what capabilities a collective would need to survive and thrive. When we assume independence of the different factors, then we can factor the joint probability into the product as follows:

$$p(\ell, o, s) = p(\ell)p(o)p(s).$$

Although it is unlikely that these factors are independent, this assumption allows us to create an intuitive representation of the complexity of an environment, namely, through the information that would need to be extracted from the environment by a bio-inspired team. Let $H(\text{env})$ denote the entropy associated with the environment. In this context, the term *entropy* refers to the information theoretic properties of the environment, namely "the average uncertainty in the random variable [that describes the environment]" [9, page 5]. Stated another way, $H(\text{env})$ "turns out to be the minimum expected number of binary questions required to determine the *value* of [the environment]" [9, page 15, emphasis added]. For bio-inspired team, the "value of the environment" is, loosely speaking, the set of activities that must be performed for the collective to thrive.

An environment with high entropy is more complex because more "questions" must be "asked" by the collective to determine what will allow the collective to thrive. Above, we associated the environment with three random variables, env $= (L, O, S)$, so that when we apply the independence assumption we get

$$H(\text{env}) = H(L, O, S) = H(L) + H(O) + H(S). \tag{12}$$

Thus, under the simplifying assumption of independence, the entropy of the environment is the sum of entropy due to the location/type of resources, obstacles/threats, and need to maintain homeostasis.

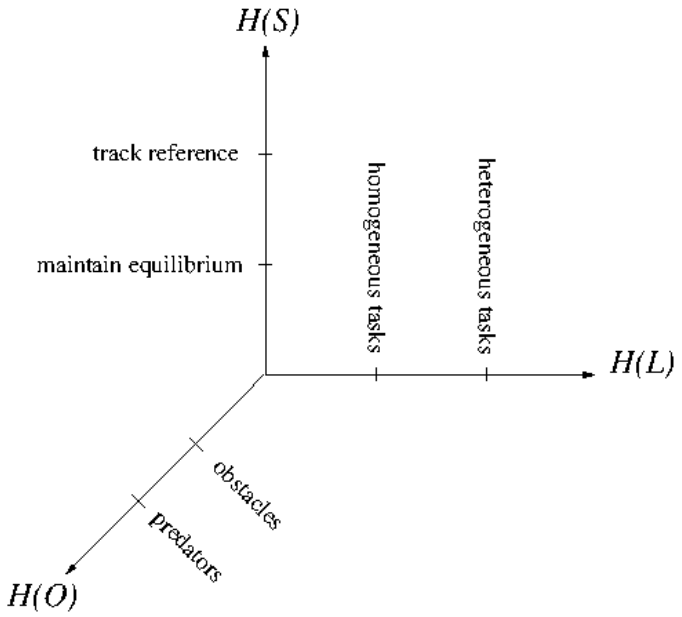

Figure 45: A visual representation of Equation (12).

Figure 45 represents Equation (12). Under the assumption that L, O, and S, are independent, we can treat each component of the environment as a different dimension. Figure 45 includes some labels to indicate some locations where qualitative switches can occur that affect the complexity of the environment. For $H(L)$, one level of entropy is reached when we introduce more than one homogeneous task, and this entropy can increase as the distribution of these tasks becomes less informed (i.e., when the distribution ceases to be clustered and is, instead, uniformly distributed). A second qualitative shift occurs when tasks are no longer homogeneous but heterogenous. This qualitative shift means that the collective must be able to perform more than one type of task, making the world more complex. Similarly, for $H(O)$ there is a qualitative shift in complexity when we move from a world with either static and dynamic passive obstacles to a world where predators are actively trying to disrupt and destroy the collective. Finally, for $H(S)$ there is a qualitative shift between a world that only requires

a collective to maintain, for example, the temperature of the nest or hive to a world that requires the collective to track a reference or adapt to some persistent but time-varying requirement, such as migrating to a new location when the amount of daylight changes.

B.2 Topological Capacity

A collective must be capable of handling the complexity in the environment. In other words, collectives must be capable of expressing spatio-temporal behavior that allows them to thrive. In this section, we identify attributes of collective expressiveness.

The term *collective expressiveness* refers to the ability of various topologies to perform various functions. For example, although nearest neighbor topologies tended to be more cohesive, they failed to express the full range of group behaviors seen among metric-based topologies. Simply put, the nearest neighbor topologies performed better in terms of cohesiveness, but had less capacity because parameters could not be adjusted to exhibit other collective structures (swarms and toroids) that might be useful given different task requirements.

Thus, some topologies perform very well under certain conditions but may not have as much capacity in terms of the breadth of tasks that they can adequately perform. This concept is illustrated in Figure 46. This figure introduces a performance threshold above which the group's performance is

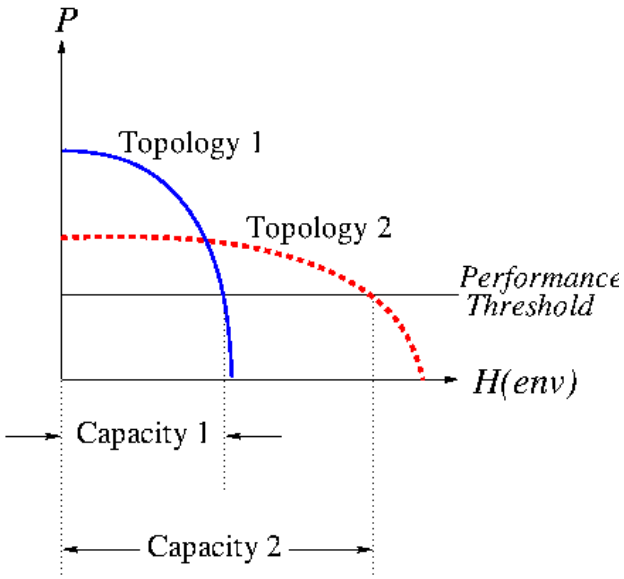

Figure 46: A notional illustration of how various topologies can perform very well in certain capacities but fail to perform adequately across a wide range or tasks.

considered adequate.

Consider an example to illustrate the concept of topological capacity. Under the biomimetic model, leaders can efficiently guide highly parallel groups in the information foraging task by quickly getting the group to where it needs to be do consume resources. We showed that, for this task, the nearest neighbor topology performed better than the metric-based topology. But the nearest neighbor topology appears to be unable to "express" swarming and toroidal structures, so if a collective is presented with a task that requires a toroid structure then a metric-based collective will be more capable of succeeding in this task than a nearest neighbor-based collective. In terms of Figure 46, the nearest neighbor topology assumes the role or "Topology 1", performing well on a small subset of tasks but being incapable of succeeding over a larger range of tasks. The metric-based topology assumes the role of "Topology 2", not performing as well in simple tasks but being capable of performing well enough over a larger range of tasks.

We have treated Figure 46 as a useful metaphor, but work in human-robot interaction has used similar notions do characterize efficient remote interactions between a human and a robot [10]. The human-robot interaction work required the definition of instantaneous and average performance and the specification of a performance threshold. Future HuBIRT work should consider following the human-robot interaction work and define better metrics for average, accumulated, and instantaneous performance of a collective.

C Definitions from Graph Theory

One way to better establish a theoretical basis is to explore algebraic graph theory properties of various team topologies. We hypothesize that there are a set of algebraic graph theoretic properties that correlate with or predict coherent, responsive teams. We further hypothesize that other algebraic graph theory properties will correlate with expressive responsive teams. In this section, we include a set of definitions that may be useful not only for this report, but that may also be useful in framing future work.

valence the number of edges of a node.

path "A path of length r from x to y in a graph is a sequence of $r+1$ distinct vertices starting with x and 4nding with y such that consecutive vertices are adjacent" [16, page 4].

connected "If there is a path between any two vertices of a graph X, then X is *connected*" [16, page 4].

distance "The *distance* $d_X(x, y)$ between two vertices x and y in a graph X is the length of the shortest path from x to y" [16, page 5].

diameter "The *diameter* of a graph is the maximum distance between two distinct vertices" [16, page 16].

regular "A graph in which every vertex has equal valency k is called *regular* of valency k or k-regular" [16, page 5].

We use Lemma 8.1.2 from [16] to compute when a graph is connected. This lemma says that the number of paths of length r from node i to node j is $[A^r]_{i,j}$. This notation means that we take A to the r^{th} power and then look at the $(i,j)^{\text{th}}$component. Thus, if we compute

$$\mathcal{A}^{[N]} = \sum_{k=1}^{N} A^k,$$

where N is the number of agents, and find that there are no zeros in $\mathcal{A}$ then it is possible to reach any node from any other node; in other words, the topology is connected.

References

[1] L. Alboul, J. Saez-Pons, and J. Penders. Mixed human-robot team navigation in the GUARDIANS project. In *Proccedings of the International Conference on Safety, Security, and Rescue Robotics*, Sendei, Japan, October 2008.

[2] M. Ballerini, N. Cabibbo, R. Candelier, A. Cavagna, E. Cisbani, I. Giardina, V. Lecomte, A. Orlandi, G. Parisi, A. Procaccini, M. Viale, and V. Zdravkovic. Interaction ruling animal collective behavior depends on topological rather than metric distance: Evidence from a field study. *Proceedings of the National Academy of Sciences of the United States of America*, 105(4):1232–1237, 2008.

[3] L. Barnes, M. Fields, and K. Valvanis. Umanned ground vehicle swarm formation control using potential fields. In *Proceedings of the 15th Mediterranean Conference on Control & Automation*, Athens, Greece, July 2007.

[4] S. Bashyal and G. K. Benayagamoorthy. Human swarm interaction for radiation source search and localization. In *IEEE Swarm Intelligence Symposium*, St. Lous, Missouri, USA, 2008.

[5] N. W. F. Bode, D. W. Franks, and A. J. Wood. Limited interactions in flocks: Relating model simulations to empirical data. *Journal of the Royal Society Interface*, 8, 2011.

[6] R. A. Chadwick, D. J. Gillan, D. Simon, and S. Pazuchanics. Cognitive analysis methods for control of multiple robots: Robotics on $5 a day. In *Proceedings of HFES*, 2004.

[7] L. Conradt, J. Krause, I. D. Couzin, and T. J. Roper. "Leading According to Need" in self-organizing groups. *The American Naturalist*, 173(3), March 2009.

[8] I. D. Couzin, J. Krause, R. James, G. D. Ruxton, and H. R. Franks. Collective memory and spatial sorting in animal groups. *Journal of Theoretical Biology*, 218(1), September 2002.

[9] T. M. Cover and J. A. Thomas. *Elements of Information Theory*. Wiley Interscience, 1991.

[10] J. W. Crandall, M. A. Goodrich, D. R. Olsen Jr., and C. W. Nielsen. Validating human-robot interaction schemes in multi-tasking environments. *IEEE Transactions on Systems, Man and Cybernetics — Part A: Systems and Humans*, 35(4), 2005.

[11] M. L. Cummings, S. Bruni, S. Mercier, and P. J. Mitchell. Decision support for networi-centric command and control. *The International C2 Journal*, 1(2):1–24, 2007.

[12] M. L. Cummings, A. S. Brzezinski, and J. D. Lee. The impact of intelligent aiding for multiple unmanned aerial vehicle schedule management. *IEEE Intelligent Systems: Special issue on Interacting with Autonomy*, 22(2):52–59, 2007.

[13] M. L. Cummings and S. Guerlain. The tactical tomahawk conundrum: Designing decision support systems for revolutionary domains. In *IEEE International Conference on Systems, Man and Cybernetics*, pages 1583–1588, 2003.

[14] M.L. Cummings, S. Bruni, S. Mercier, and P.J. Mitchell. Automation architecture for single operator, multiple UAV command and control. *The International C2 Journal*, 1(2):1–24, 2007.

[15] X. C. Ding, M. Powers, M. Egerstedt, S. Young, and T. Balch. Executive decision support: Single agent control of multiple uavs. *IEEE Robotics and Automation Magazine*, 2009.

[16] C. Godsil and G. Royle. *Algebraic Graph Theory*, volume 207 of *Graduate Texts in Mathematics*. Springer, 2000.

[17] B. Hardin and M. A. Goodrich. On using mixed-initiative control: A perspective for managing large-scale robotic teams. In *Proc. of ACM/IEEE Intl. Conf. on Human-Robot Interaction*, San Diego, March 2009.

[18] M. Johnson, J. M. Bradshaw, P. J. Feltovich, R. R. Hoffman, C. Jonker, B. van Riemsdijk, and M. Sierhuis. Beyond cooperative robotics: The central role of interdependence in coactive design. *IEEE Intelligent Systems*, May/June 2011.

[19] Z. Kira and M. A. Potter. Exerting human control over decentralized robot swarms. In *Proceedings of International Conference on Autonomous Robots and Agents*, pages 566–571, 2009.

[20] H. Levine, W-J. Rappel, and I. Cohen. Self-organization in systems of self-propelled particles. *Physical Review E: Statistical, Nonlinear, and Soft Matter Physics*, 63(1), 2000.

[21] M. Lewis, H. Wang, S. Chien, P. Scerri, P. Velagapudi, K. Sycara, and B. Kane. Teams organization and performance in multi-human/multi-robot teams. In *Proceedings of IEEE International Conference on Systems, Man, and Cybernetics*, Istanbul, October 2010.

[22] Jyh-Ming Lien, O. B. Bayazit, R. T. Sowell, Sm. Rodriguez, and N. M. Amato. Shepherding behaviors. In *Proceedings of the IEEE International Conference on Robotics and Automation*, New Orleans, Louisana, 2004.

[23] J. Madden, R. C. Arkin, and D. McNulty. Multi-robot system based on model of wolf hunting behavior to emulate wolf and elk interactions. In *Proceedings of the IEEE International Conference on Robotics and Biomimetics*, Tianjin, China, December 2010.

[24] T. McLain and R. W. Beard. Cooperative rendezvous of multiple unmanned air vehicles. In *AIAA Guidance, Navigation and Control Conference*, Denver, Colorado, USA, 2000.

[25] 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 *Proceedings of AAAI Spring Symposium*, Stanford, CA, USA, 2006.

[26] C. A. Miller, H. B. Funk, M. Dorneich, and S. D. Whitlow. A playbook interface for mixed initiative control of multiple unmanned vehicle teams. In *Proceedings of the 21st Digital Avionics Systems Conference*, volume 2, pages 7E4–1 – 7E4–13, November 2002.

[27] H. Min and Z. Wang. Group escape behavior of multiple mobile robot system by mimicking fish schools. In *Proceedings of the IEEE International Conference on Robotics and Biomimetics*, Tianjin, China, December 2010.

[28] R. R. Murphy, E. Steimle, C. Griffin, C. Cullins, M. Hall, and K. Pratt. Cooperative use of unmanned sea surface and micro aerial vehicles at Hurricane Wilma. *Journal of Field Robotics*, 25(3):164–180, March 2008.

[29] A. M. Naghsh, J. Gancet, A. Tanoto, and C. Roast. Analysis and design of human-robot swarm interaction in firefighting. In *Proceedings of the International Symposium on Robot and Human-Interactive Communication*, Muhich, Germany, August 2008.

[30] D. R. Olsen Jr. and S. B. Wood. Fan-out: Measuring human control of multiple robots. In *Proceedings of Human Factors in Computing systems*, 2004.

[31] E. Sahin. Swarm robotics: From sources of inspiration to domains of application. In *Lecture Notes in Computer Science*, volume 3342. SpringerLink, 2005.

[32] T. B. Sheridan. *Telerobotics, Automation, and Human Supervisory Control*. MIT Press, 1992.

[33] T. B. Sheridan and W. L. Verplank. Human and computer control of undersea teleoperators. Technical report, MIT Man-Machine Systems Laboratory, 1978.

[34] W. Spears, D. Spears, R. Heil, and W. Kerr. An overview of physiocomimetics. In E. Sahin and W. Spears, editors, *Swarm Robotics*, Lecture Notes in Computer Science, pages 84–97. Springer, Berlin, 2005.

[35] M. L. Steinberg. Biologically-inspired approaches for self-organization, adaptation, and collaboration of heterogeneous autonomous systems. In *Proceedings of SPIE Volume*, volume 8062, April 2011.

[36] D. J. T. Sumpter. The principles of collective animal behavior. *Philosophical Transactions of the Royal Society B*, 361:5–22, November 2006.

[37] D. J. T. Sumpter. Group behavior: Leadership by those in need. *Current Biology*, 19(8):R325–R327, April 2009.

[38] H. G. Tanner, A. Jadbabaie, and G. J. Pappas. Stable flocking of mobile agents part ii: Dynamic topology. In *Proceedings of the 42nd IEEE Conf. on Decision and Control*, Hawaii, December 2003.

[39] R. Vaughn, N. Sumpter, J. Henderson, A. Frost, and S. Cameron. Robot control of animal flocks. In *Proceedings of the International Symposium of Intelligent Control*, pages 277–282, Gaithersburg, Maryland, 1998.

[40] J. Wang and M. Lewis. Assessing coordination overhead in control of robot teams. In *Proceedings of IEEE International Conference on Systems, Man and Cybernetics*, pages 2645–2649, Montreal, Canada, 2007.

[41] J. Wang, M. Lewis, and P. Scerri. Cooperating robots for search and rescue. In *Proceedings of AAMAS '06*, Hakodate, Japan, May 8-12 2005.

[42] A. J. Wood and G. J. Ackland. Evolving the selfish herd: Emergence of distinct aggregating strategies in an individual-based model. *Proceedings of the Royal Society B*, 274(1618):1637–1642, 2007.