METHODS OF COMMUNICATION

Gene T. Price
Swarthmore College
May 11, 2012

ABSTRACT

This paper presents the results of an experiment, which investigated different modes of communication by evolving a swarm of communicating robots to solve a collective searching task. Two types of information sending were evaluated: signals and movement knowledge. The first set of robots evolved communication based on sending abstract signals between members of the swarm. This set of robots was intended to demonstrate how robots could develop a grounded communication system based on highly abstract sensory inputs and outputs, but instead, these robots optimized their individual strategies and failed to develop a valid communication system. The second set of robots sent their motor values to communicate. This group illustrated how agents can interpret the motor states of their peers and extract pertinent social and environmental information. This group evolved robots with a heavy dependence on communication, and this group never managed to optimize its individual strategy. Lastly, a group of robots with both types of communication was evolved to illustrate how different data representations can be used in a communication system. Out of this group, there were individuals that either focused on communication or individual strategy, but no individuals focused on both. Unfortunately, due to the idiosyncrasies of the test environment and the evolution runs, the control group, without communication, outperformed the communicative robots.

I. INTRODUCTION

Inter-robotic communication and swarm communication are a fascinating area of developing research. Several papers have focused on how communication in a swarm of agents can develop grounded communication systems using very simple communication frameworks, but few have compared a variety of frameworks. What types of information should be communicated between robots to create the fittest individuals? Might it be better to communicate with an already grounded set of data, like motor values, and allow the robots to abstract this information into a communication system?

This paper divides communication into two somewhat overlapping classes: abstract and concrete. Concrete data is grounded and definite. It has an absolute meaning within the context of its environment. Abstract data on the other hand contains a meaning deeper than its concrete representation. For example, figure 1 is a set of concrete data corresponding to the amplitudes of a series of sounds, but what this representation does not hint at is the abstract representation behind the data. If a human, rather than a computer, had interpreted the data, they would have likely heard, “I am a robot”. Agents that have evolved communication take these concrete data sets and abstract them. In terms of a robotics framework, concrete data would be sensorimotor data, which needs to be interpreted into behavior cues to be useful, and abstract data would be a specific command from a controller, which only has meaning in terms of the robot’s interpretation of the command.

Another distinction between data is the degree of abstraction and concreteness. Most abstract data is a highly abstracted interpretation of concrete data; human speech for example can be broken down into a highly concrete data set of frequencies. Furthermore, some forms of data ride the line between abstract and concrete. Comparing human speech to human gesturing, speech is much more highly abstracted because it more often carries meaning independent of its raw

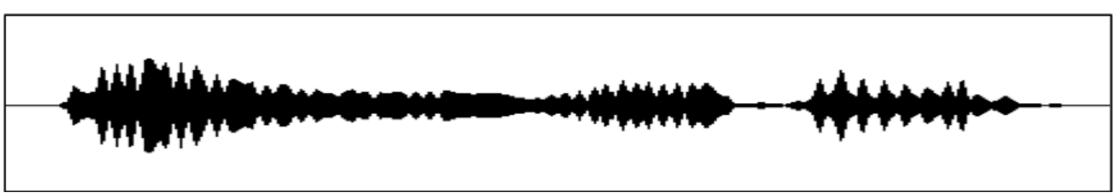

Figure 0: An amplitude chart of the author saying, “I am a robot”. This demonstrates the necessity of data abstraction in communication; the concrete representation of the data obscures the abstract meaning. To effectively communicate, agents must be able to interpret the raw data they receive.

data. Gestures, though, may or may not be intended for communication. If a person's hand moves towards an object, that person could be either gesturing at the object or moving to interact with it. In the first case, communication occurs and an observer can infer an abstract meaning, but in the latter case, it is an action to effect change on the environment, independent of observers. While both methods of communication can be abstract, gesturing is more concrete, because the communication is based on a grounded method of interacting with the environment.

Another important area of study comes from the combination of concrete and abstracted data in the same communication system. For example, body language and speech form a highly nuanced communication system when combined.

The goal of this paper is to observe how communicating with different types of data effects the behavior of a swarm of robots. This goal will be achieved in three ways: robots will be evolved to 1) show how robots can ground abstract data in social and environmental cues, 2) form an abstract interpretation of concrete data, and 3) integrate these different data modes into the same communication system.

There will be three different types of robots, one for each sub-goal. The first swarm of robots will be evolved to develop an abstract method of communication. This method of communication would be most similar to speaking, where the abstract data is ungrounded, so the swarm must evolve both the values and the interpretations of the signals. Since these robots use signals to communicate, they shall hitherto be referred to as signal senders. The second set will communicate using grounded, concrete information. In the implementation, these robots communicate by sending the previous time step's motor states. Since this method uses the movement of the robots as a form of communication, they will be referred to as movement type robots. The third will be evolved to communicate with both abstract and concrete data and will be evaluated based on the different ways in which it incorporates the data into its communication system. This robot will be a combination of the two previous robots, so it will be referred to as the conglomeration type robot.

In the following section, a short review of related literature will be made. Section III, details the experimental procedure, including the task environment and an overview of the four types of robots that were evolved. Section IV, presents and analyzes the qualitative and quantitative results, while section V discusses the successes, limitations, and failures of the experimental method and results. Lastly, section VI concludes this paper with a brief evaluation of the experimental goals and possible edits and continuations of the experimental method.

II. Related Works

Quinn et al. evolved swarms of robots for effective flocking behavior, staying close together and aligning their movement axis. They were given only a proximity sensor, but individuals evolved a method of signaling based on movement. Robots, could communicate the roles of leader and follower based on the way that individuals approached each other. If one robot is approaches another robot and stays close for an extended period of time, it is signaling that it would take the role of follower. If a robot experiences a high activation of its proximity sensors, then it would invert its direction and assume leadership of the swarm. This simple bi-modal communication strategy allowed individuals to effectively communicate without any explicit communication sensors. These robots were taking concrete data from their proximity sensors and abstracting them into social cues for effective swarming behavior.

In Marocco et al. (2003), evolved a controller for an arm robot tasked with discerning the shape of objects through tactile information. These controllers took turns being "speakers" and "hearers". The speakers were given tactile information about the object being held and had to interpret this data into a communication vector, a vector of floating-point values, which was then sent to the hearer. The hearer then reinterprets the communication vector to discern which type of object is being touched. The goal of the task was to pick up spherical items and avoid cubic objects, and the controllers with communicative abilities were able to create abstract representations for the concrete data and use this representation to solve the task faster. This paper showed how robots could evolve a method of abstracting concrete data and reinterpreting the data into environmental features.

Nolfi and Marocco (2004) evolved a swarm of communicating robots with a homogeneous neural network controller to solve a collective navigation task. They explored how robots capable of sending signals could develop and use a system of communication to improve their effectiveness at finding target areas. This experiment focused on how communication and goal solving abilities co-evolve, and how with a simple and open-ended communication system, a complex method of communication can arise. Through evaluating the fittest individuals in a 'deprived' condition, in which they took away the ability to communicate, they proved that in addition to creating a communication strategy, the searching strategy of individuals was also optimized. Their robots developed multi-modal types of communication, while communicating with just a single floating-point value.

The experimental method in this paper simplifies the Nolfi and Marocco paper's method, but tests a variety of communication types. Experiments have been done to evaluate the effectiveness of using abstract and concrete data for communication, but this paper compares the two approaches in the same evolutionary context. Rather than focusing on the development of complex communication systems, I examine the behaviors of robots that communicate using different forms of communication.

III. Experimental Design

This experimental setup was designed to evolve neural network controllers for a collective searching task. Each member of a swarm of robots is given. The controllers were evolved using NeuroEvolution of Augmenting Topologies (NEAT), a genetic algorithm that gradually complexifies simple neural networks by adding nodes and connections. Neat was given the ability to evolve recurrent connections. All four robots have the same controller.

In addition to the four robots, the environment consists of a square arena, and a circular target area. The target starts in the middle of the top half of the arena, and the robots are initialized randomly in the bottom half with initial headings in the range of [-22.5°, 22.5°] from straight up.

The goal of the robots is to search the arena, find the target and converge on it.

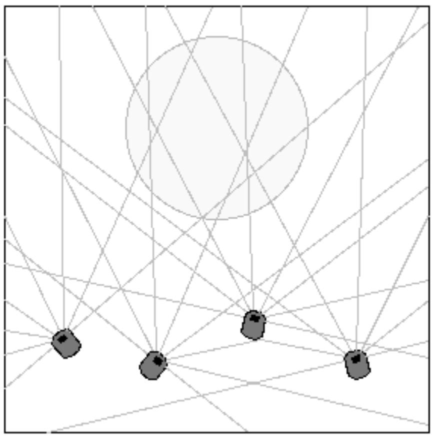

Figure 2: A diagram of the task environment and the initial starting points of the target and the robots. The outer boundary of the arena is demarked by the black square. The lighter lines are the rays of the sonar sensors. The circular area is the target area.

1. Robot Topology

Four different types of robotic controllers were evolved, and each type, excepting the control, examines a different type of data representation. Since NEAT changes the topology of the networks as they evolve, the illustrated networks are only the initial starting points, and do not accurately represent the connectivity and hidden layers of the evolved topologies.

1. Control Group (no communication): this controller has no form of communication. Its neural network has three input nodes and two output nodes. The first two input nodes are left and right sonar sensors, and the third input node is a target sensor that returns a 0 if the robot is outside the target area and a 1 if it is in the target. The two output nodes are the left and right motors.

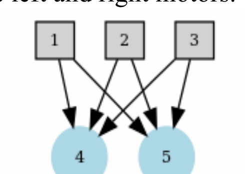

Figure 3: The neural network of the control type robots. Nodes 1 and 2 correspond to left and right sensors. Node 3 corresponds to a target sensor. Nodes 4 and 5 are the motor outputs.

2. Signal Type (abstract data): This type of robot was designed to evolve a method to ground arbitrary signal values in environmental and behavioral cues. These robots communicate with an arbitrary signal value, so the controller must evolve both what value signals to send and the meanings behind those values. This mode of communication represents the use of abstract knowledge because there is no the meaning of the signal is based not on the raw data of the signal value, but on the robot's interpretation of the value. In addition to the motor output nodes, this robot has a signal actuator that sends out a floating-point value from 0 to 1. This signal can be 'heard' by the additional sensor node that inputs the signal value of the closest neighbor robot. The last input node, the direction sensor, gives the direction of the closest robot relative to the heading of the robot.

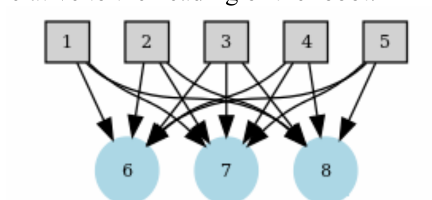

Figure 4: The neural network of the signal type robots. Nodes 1, 2, and 3 are the sonar and target sensors. Node 4 is the signal sensor. Node 5 is the direction sensor. Node 6 and 7 are the motor outputs, and node 8 is the signal actuator.

3. Movement Type (concrete data): This type was designed to evaluate how robots can develop abstract interpretations of concrete data. These robots communicate by sending their motor states, and unlike the signal senders, a concrete meaning grounds the data being transmitted. To effectively complete a task, this controller must be able to communicate and interpret important environmental features from movement, similar to how a bee signals to the swarm using a 'dance'. In addition to the sonar and target sensors, this robot has two motor sensors that return the previous time step's motor values for the nearest robot and one more input node, which gives the direction of the nearest robot.

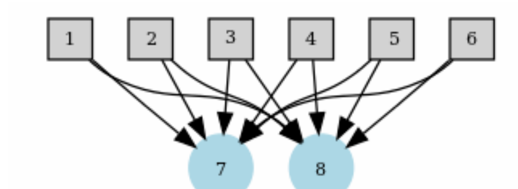

Figure 5: The neural network of the movement type robots. Nodes 1, 2, and 3 are the sonar and target sensors. Nodes 4 and 5 are the movement sensors, which input the motor values of the nearest robot. Node 6 is the direction sensor. Nodes 7 and 8 are the motor outputs.

4. Conglomerate Type (combination of concrete and abstract data): this controller combines the previous two methods, and can give insight as to how concrete and abstract data can interact to form more nuanced communication modes.

This controller combines all the nodes and outputs of the previous three controllers.

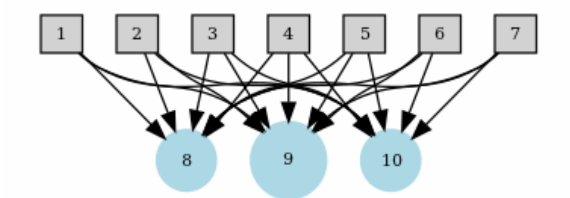

Figure 6: The neural network of the conglomerate type robots. Nodes 1, 2, and 3 are the sonar and target sensors. Node 4 is the signal actuator. Nodes 5 and 6 are the movement sensors. Node 7 is the direction sensor. Nodes 8 and 9 are the motor outputs, and node 10 is the signal actuator.

All the controllers with the ability to communicate were also given a direction sensor that gives the direction in degrees of the nearest neighbor relative to the robot's heading. This sensor was necessary for high-level communication because communication is much more useful in a collective searching task if you know where the communication is coming from.

2. Fitness

At each time step (*i*) fitness was evaluated in three different ways:

1. To encourage effective searching behavior in the beginning of evolution, the controller was awarded a small amount of fitness (α) for each robot not stalling (A).
2. The controller was awarded the value of the cumulative distance that all the robots travelled towards the target from their initial starting points (B) multiplied by a constant (β).
3. For each robot in the target area (X) the controller was awarded a small amount of fitness (χ).

The controller was also awarded a large amount of fitness (E*)* for the number of times (ε) that all four robots reached and moved the target. These four values result in the following fitness function:

$$F = \varepsilon \mathrm{E} + \sum_{i=0}^{n} \alpha \mathrm{A} + \beta \mathrm{B} + \chi \mathrm{X}$$

In the experiments $\varepsilon = 300$, $\alpha = 0.1$, $\beta = 1$, and $\chi = 1$.

3. Rationale Behind the Task

The task was crafted to make sure that the movement type robots could be competitive with the signal sending robots. Signal sending robots have a highly flexible framework because their signals have no inherent meaning and can be tailored to any particular task. Furthermore these robots can communicate independently of their movement, unlike movement type robots. To make movement communication a viable strategy, changes in the environment needed to correspond to obvious changes in robot behavior. This allows movement type robots to communicate salient features of their position without compromising their movement strategy. To this effect, the environment is highly discretized into target area and non-target area, and the movement strategies in each of these areas should be completely different. In the non-target area, movement should be fast and wandering, to maximize the search area, but in the target area, movement should be kept to a minimum to stay inside the target area. This correlation between environmental states and motor states makes it feasible to communicate using motor states.

4. Evolution Runs:

Controllers were evolved for 300 generations with a population of 100 individuals. Each controller was given 4 trials with different random starting points and had 100 time steps to acquire fitness. An individual's fitness was the sum of each of the trials.

IV. Results:

In this section robots are evaluated qualitatively and quantitatively. Special interest is paid to the qualitative data because it is the best method of examining the behavior of the robots. None of the communicative controllers contained all of the different communication modalities exhibited by the entirety of the population, so rather than picking apart the communication strategy of one robot, there is a broad overview of the all the basic communication strategies that were developed.

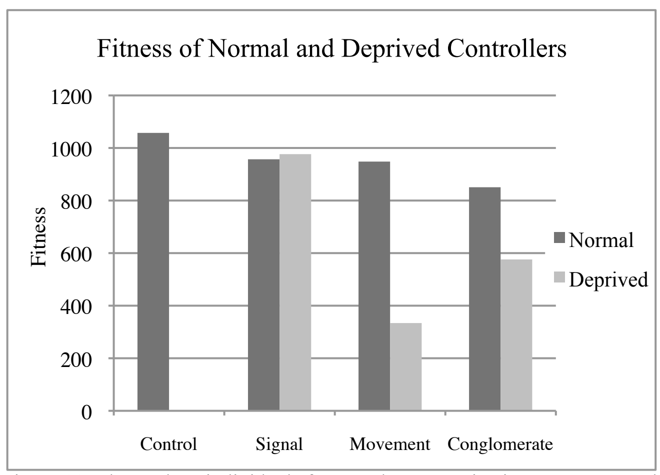

Figure 11: The ten best individuals from each communication type were each tested for 100 trials of 100 times steps, and the displayed fitness represents the average fitness of the controllers per trial. The same controllers, excepting the control group, were each tested with their communication values set to 0 for 100 trials as well. The 'Deprived' data represents the average of those trials.

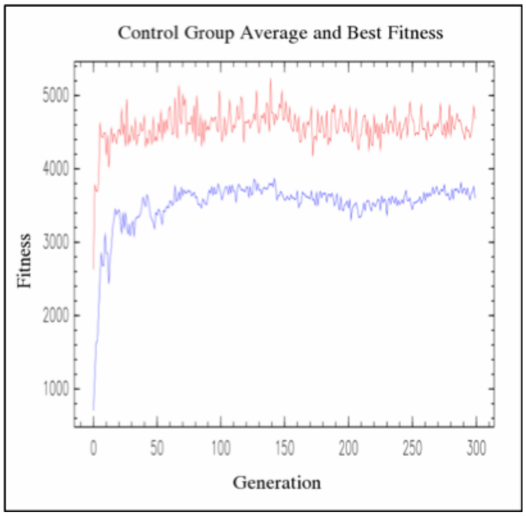

Figure 7: This graph illustrates the average (blue) and best (red) fitness at each generation for the control group. The control group had the highest average fitness, and it created many of the fittest individuals out of all the trials. (Note: the values on the y-axis are different than the other fitness graphs)

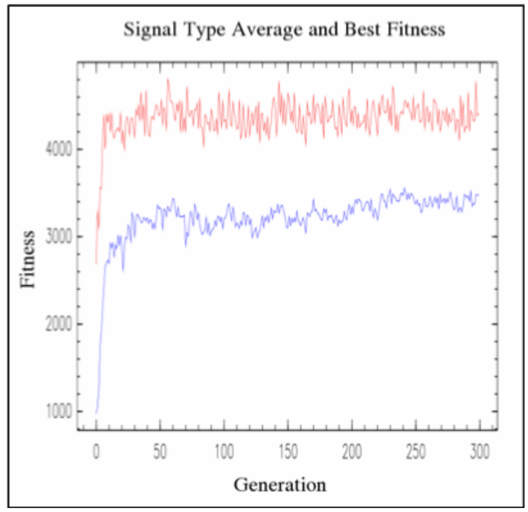

Figure 8: This graph illustrates the average (blue) and best (red) fitness at each generation for signal type robots. This evolution run developed very similarly to the control group, but it plateaued at a lower value, and the fitness is less consistent.

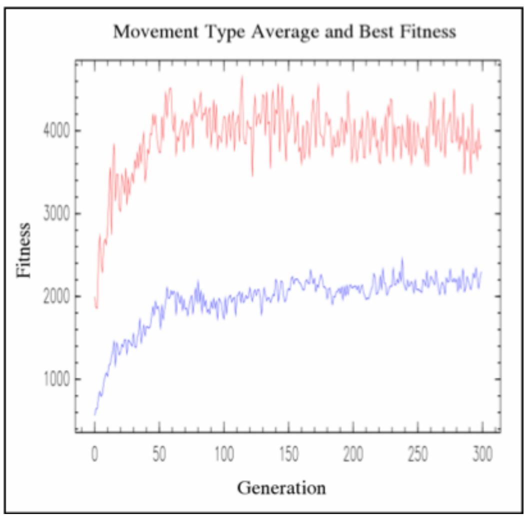

Figure 9: This graph illustrates the average (blue) and best (red) fitness at each generation for signal type robots. This evolution run developed much slower than the control or signal sending types. Even though its average fitness is much lower, it still managed to produce some individuals competitive with the signal sending type.

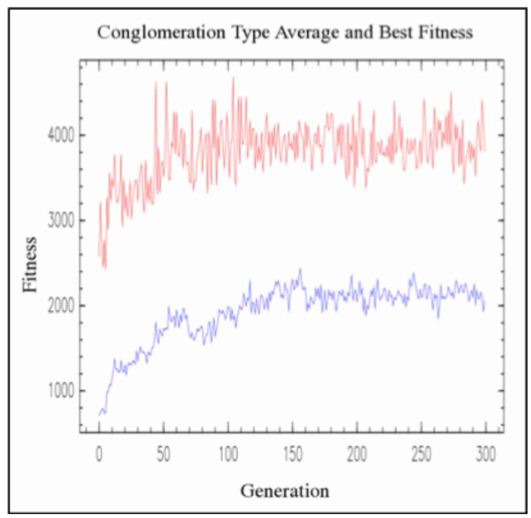

Figure 10: This graph illustrates the average (blue) and best (red) fitness at each generation for signal type robots. This evolution run developed very similarly to the movement type, just slower due to the higher dimensionality of the neural network. The fitness of the best individuals was highly erratic, belying the tremendous mix of strategies in contest.

1. Comparative Analysis:

The best 10 controllers from each evolution run were tested in the task environment with 100 times steps for 100 trials. To increase consistency, each controller was given the same set of 100 random starting points. Next, each of the 10 best communicative controllers was tested without being allowed to communicate. For these trials, all the inputs corresponding to communication nodes were set to zero. Controllers were tested in the deprived environment for 100 trials of 100 time steps each. The 'deprived' tests evaluated the controller's individual strategies; how robots move by themselves to optimize their ability to find the target area (figure 11). An individual strategy is the behavior that a robot uses when communication is non-effective or non-existent. This set of behavior includes any behavior independent of communication, including but not limited to effective wandering strategies, stopping inside the target area, and wall avoidance.

Much of the behavioral differences between controllers can be evaluated on how their fitness was split by independent and communicative strategies. Since the deprived tests evaluate how well a controller does without the ability to communicate, it provides a valid method of analyzing the individual strategies of each robot.

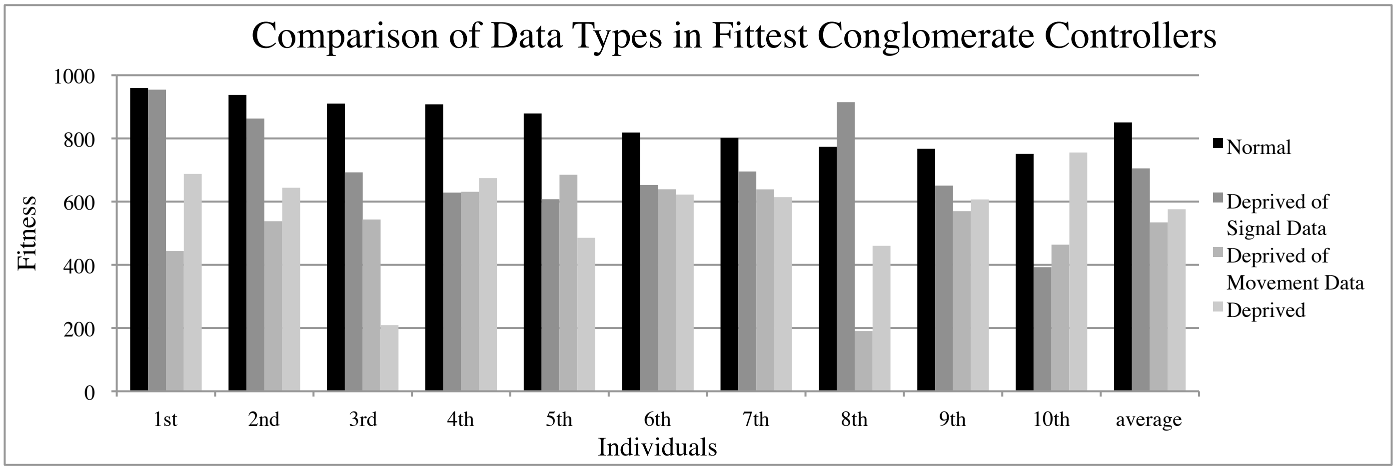

Figure 12: This graph represents the results of four experiments performed on the best 10 conglomerate type controllers. Each of the different types of test was run for 100 trials, and the averages of those trials are displayed. In the normal set, these robots were each tested in the normal task. The controllers are ordered by their level of fitness in the normal task with the 1st controller receiving the highest score, the second receiving the second highest score, and so on. For the deprived data set, the ability to communicate was taken away from the controllers. For the 'deprived of signal data' set, the conglomerate controllers were only able to communicate with movement knowledge. For the 'deprived of movement data' set, the conglomerate controllers were only able to communicate through signal sending. This graph represents the numerous ways in which the conglomerate type controller integrated signal and movement

Control Group:

This population had the highest average fitness and some of the fittest individuals (figure 7). It very quickly developed an effective searching strategy and then plateaued. One factor contributing to the higher average of its top ten controllers is because they all use a nearly identical, albeit very good, strategy. All three communicative types had individuals on par with the best of the control.

Signal Type:

This controller had the highest average fitness of the communicative robots (figure 8 and 11), but not a statistically significant amount more. It was actually superior (not by a significant amount though) in the deprived environment. It never found an effective way of grounding its signal in its behavior, so its communication sensor was just providing noise. Depriving this network of its ability to communicate produced such a small effect because it had already learned how to filter its communication sensor out.

Some of the best controllers from the early generations had the ability to communicate using simple feedback loops, but this ability was quickly lost as individuals with highly optimized individual strategies flooded the population.

To test where communication broke down, each of the ten best signal controllers was tested on a single robot placed alone in the task area, and the motor values were recorded for a range of possible signal values. There was no discernible change in motor values for 8 of the ten robots. When the signal values of the other two robots were observed while performing the standard task, their signal values never rose above 0.01. Depending on the controller, communication failed at either the front or back end. This type of controller produced robots that either refused to send or receive signals.

Movement Type:

These robots developed at a much slower rate than the signal type and control type, and also converged at a much lower fitness value (figure 9). One factor that slowed its growth is the codependence of communication and movement behavior; the ability of the movement type to communicate is dependent on its ability to move in the environment. Communication ability could not increase until the robot had some form of effective movement controller.

This controller was the most dependent on its communication sensors (figure 11). It performed abysmally in the deprived condition. This controller depends on communication for its ability to stay in the target area

Conglomerate type:

This controller had the lowest average fitness, but produced the some of the fittest individuals. Furthermore, the signal type controller never found an effective use for communication, but this type had several individuals that made effective use of the signal sender. Unlike in the deprived conditions of the signal type robots, which were equally as good as the non-deprived individuals (figure 11), when the conglomerate type robot was deprived of its signal sending abilities, its fitness, on average decreased (figure 12). The higher complexity of the neural network for this controller made it slower to learn, but also made the population converge later.

This population acted much like an average of the values of the signal and movement type controllers. Its fitness in the deprived task (figure 11) is between the signal and movement type controllers, implying that it had an effective, but not optimized movement strategy.

2. Evaluation of Communication:

While comparative results are informative about the comparative strengths and weaknesses of differing types of communication, the real goal of this experiment is to analyze the behaviors of robots that developed communication. The fittest individuals were observed, and their unique approaches and modalities of communication were observed.

Movement Communication:

The most common use of communication in this type of controller was the use of feedback loops. This feature developed very early in this controller's evolution and stuck until the end. Dimensions were added to it as it progressed. In its simplest and weakest implementation, if one robot started going in circles, other robots receiving this signal would also start moving in circles. This helped weak controllers keep robots that had found the target inside of the target. By mirroring their peer's actions, they would, in effect, remind it to stay in the target area.

In later generations, this feedback loop was refined to be context specific. Rather than all listening robots duplicating an action, only robots in the correct environmental context will emulate the action. Once robots enter the target area, they start spinning, and if other robots enter the target area, they also start spinning. In the fittest movement controller, when there was more than one robot in the target area, the chances that a robot would exit the target became 0, and since they were continuously spinning, they also had no translational movement. In the deprived condition, this controller was highly ineffective because there was no feedback loop to remind the controller to stay in the target area.

The strongest strategy that developed was context independent unidirectional communication. Robots in the target area would move repeatedly in circles signaling the location of the target area encouraging robots to converge on its location. The one problem with this mode of communication is it does not work well with the feedback structure that earlier forms of communication use. When a robot converged on the target area with other robots in it, sometimes it would break the feedback reinforcement loops and cause another robot to move out of the target area. Since movement robots do not have effective individual controllers, as evidenced by their low deprived condition scores (figure 11), they are not well prepared to stay in the target area if they are not receiving feedback.

Conglomerate Communication:

This type of controller had the greatest variety of behaviors and communication styles, but rather than creating highly novel communication forms, it Number 10 behaved similarly to the non-communicative individuals in the signal type population. Rather than creating a strategy of communication, it found a way to block out communication and rely solely on an individual strategy, evidenced by the fact that the fully deprived controller is equally as good as the communicative controller. This population also includes individuals that closely mimicked the movement type controller. The eighth controller in particular took this mimicry to the extreme as it did better when it was deprived of its signal sending abilities. It was, in effect, just a movement type controller with the signal sending nodes only adding noise to its sensors.

A common use of both communication systems was in dual modality reinforcement. Rather than using different types of communication for different types of communication, this mode communicated with concrete and abstract data at the same time to reinforce each other. This strategy reinforces the social cues and provides a greater degree of specificity to communication. For example, most of the effective robots would spin in circles when they found the target area, but robots using dual modality communication would also send out a signal value of 1. This behavior figured very prominently in the individuals numbered 3, 5, 6, and 7 in figure 12. With only one of their communication modes active, they had a lower fitness than the normal, but a higher fitness than the deprived because they had the ability to communicate effectively with either type of communication. Furthermore, the behaviors of the partially deprived robots were almost identical because of the overlap in communication representations. Even though no new form of communication was created from the combination of the two types, controllers depended on both representations to reinforce their behavior. When the seventh individual was deprived of one of the communication modes, the robots had trouble staying inside the target area; if the sonar values of a robot in the target area got too low because another robot was approaching the target area, this controller would often back out of the target area to avoid stalling against it.

The tenth individual represents a particularly bizarre method of using communication, cancellation. In the pure deprived test, it performed just as well as the normal test, meaning it had a highly developed individual strategy, but give it only half of its communication abilities and it becomes a very poor controller. Rather than use the communication data to improve its strategy, it developed a very effective individual strategy and used the two methods of communication to cancel each other out.

V. Discussion

One of the reasons that the control group was so successful in this experiment is because it optimized its very simple, independent movement strategy. Starting in a low dimensional problem space, it quickly and effectively found a valid movement strategy. The communicative robots on the other hand never optimized their movement strategy, and this failure also prevented them from developing highly evolved form of communication. In many instances, communication is completely useless without an effective movement strategy. Take for instance a robot that has found the target area; the only way that it can effectively communicate the direction of the target area to its peers is if it knows to stay inside the target area.

A main problem with the experiment was the isolation of strategies possibly caused by the speciation method of NEAT. It separated members of the population that specialized in different aspects of the task. Some of the best members of each of the communicative populations developed very effective movement strategies but non-existent communication strategies. When these individuals were tested in normal and deprived conditions, they did slightly better in the deprived conditions. Other members of the population developed highly effective communication strategies, but had weak individual movement strategies. These individuals performed very poorly under deprived conditions, proving that effective communication can be an extremely powerful strategy.

Another precursor to effective communication is effective spatial organization of the robots. Since robots only hear their closest neighbors, the swarm of four robots could and often did break down into two pairs of two robots communicating with each other. In this state, even if one pair found the target area, the other pair would never receive a signal from the first pair. This is likely one of the reasons that communicative robots either had a sub-optimal communication or movement strategy. If a species of robots were evolving communication then they needed to evolve a method to effective orient themselves to prevent isolation. Spatial organization strategies prevented the development of effective individual strategies because rather than moving to maximize their individual ability to find the target, robots moved to maintain the ability to communicate, so that once a robot found the target, the rest could converge on its location.

Another explanation for the peculiar behavior is the fitness function and the incremental nature of the task. The fitness function gives more fitness for robots that are in or near the target area, and this feature was meant to encourage robots to stay in the target area and to move towards their peers in the target area. A problem though arises from the fact that robots could receive more fitness if there are three robots in the target area and only one robot outside of it. In this configuration, controllers would receive a sizeable amount of fitness at each time step, which could easily amount to a higher fitness than getting a large amount of fitness by moving the target and less fitness per time step.

Even if moving the target provides a large fitness increase it is still a risky strategy. Rather than finding the target and moving it immediately, it could benefit robots to find the target, stay close to it, and only move it near the end of the trial. Furthermore, moving the target and trying to find it again is a risky strategy; if the robots do not find it again within the allotted time steps, they would get a much lower fitness than they usually do. This inconsistency would make them more likely to be removed from the population if they have a particularly bad fitness run.

This combination of factors resulted in the premature convergence of the populations on sub-optimal strategies. It was too hard to consistently move several targets, so evolution favored controllers that stayed close to the target for extended period of time and only moved it once.

In a task that rewarded robots for staying in the target area and not moving it, it makes sense that non-communicative robots were better; if the target area is moved shortly after it is found, less fitness will be gained at each time step.

Even though the task had some flaws and the control group was equally as good as the communicative populations, this experiment still succeeded in its central goal of evolving and evaluating robots using different data types to communicate.

VI. Conclusion:

There are some changes that could be made to the experimental method to improve and change the results. The values used in the fitness function need to be tweaked to more clearly reward robots converging on the target area. The task could also benefit from an increased number of time steps and trials per fitness evaluation. In the Marocco paper (2006), an individual's fitness was based on 20 trials with 1000 time steps per trial. Due to the limited time of the experiments, this level of evaluation could not be implemented.

Despite these problems, robots were evolved that had the ability to communicate using concrete data. These robots found an abstract representation of this data that enabled them to connect it to important behavioral cues. Robots with the ability to communicate with abstract values were never able to ground the data with any relevant meaning. Some of the early, ineffective controllers used these values to communicate with feedback loops, but these weak communication strategies were quickly made obsolete by highly effective individual strategies. Communication based strategies never became fit enough to supplant the dominance of these non-communicative strategies. Robotic controllers that combined the different data representations also learned how to communicate effectively rather combining the different data representations to communicate in radically different ways, these controllers used both methods to communicate similar data. The duality of the data representation was a powerful tool for these robots in cueing behaviors.

VII. References

Marocco D., Nolfi S. (2006). "Emergence of Communication in Embodied Agents Evolved for the Ability to Solve a Collective Navigation Problem". Institute of Cognitive Science and Technologies.

Miikkulainen R., Stanley K. (2004). "Competitive Coevolution through Evolutionary Complexification". *Journal of Artificial Intelligence*: 63-100.

Marocco D., Cangelosi A. & Nolfi S. (2003). "The Emergence of Communication in Evolutionary Robots". *Philosophical Transactions of the Royal Society London - A*, 361: 2397-2421.

Quinn M., Smith L., Mayley G., Husbands P. (2003). "Evolving Controllers for a Homogeneous System of Physical Robots: Structured Cooperation with Minimal Sensors". *Philosophical Transactions of the Royal Society of London – A*, 361: 2321-2344.