Evolving Swarm Communication with NEAT

Rex Chang   Sedinam Worlanyo
Swarthmore College
cchang2@swarthmore.edu, sworlan1@swarthmore.edu

Fall 2015

Abstract

We believe that communication in swarm robotics can be crucial and offer significant improvements in efficiency of certain tasks. Influenced by the paper *Comparative Evolutions of Swarm Communications*[2] written by previous cs81 students Amy Jin and Murphy Austin, we designed a task where having communication would provide an efficient solution. Inspired by the behaviors of ants collaborating to search and transport food, we created a world where three swarm agents are to search for a food object and all gather around the food. It is important to note that the agents do not have any information about the food besides a sensor that indicates whether the agents have reached food or not. This way, solving this task without communication would mean that the three agents need to individually stumble upon the food. On the other hand, if the agents are capable of communicating, an ideal situation would be that one agent discovers the food and through communication signals the two other agents.

We will utilize neural networks and the genetic algorithm NeuroEvolution of Augmenting Topologies (NEAT) to evolve such communication between the agents. Our hypotheses are that not only will all three agents reach the food object, but also effective communication strategies will be evolved so that the agents complete this task efficiently. In addition, we believe that NEAT is powerful enough to make sense of the inputs that we provide and generate meaningful outputs. To test these hypotheses, we created three testing groups: agents that have free communication, fixed communication, and no communication. After gathering the results, we discovered that all three groups were able to come up with strategies that allow them to complete the task, however we did not observe any meaningful communication being evolved. Yet, on a more optimistic note, there was evidence that the agents did react differently when they reached the food object, so this indicated that on some level NEAT still made a difference, and it is our future work to further understand what evolved behaviors our agents possess.

1 Introduction

Swarm robotics is a field in artificial intelligence inspired by the behavior of social insects such as ants. It is a field that has been applicable to several real world scenarios such as rescue missions whereby swarm robots are trained to cooperate with one another to achieve a specific goal. Evolutionary algorithms such as NEAT have been successful in evolving robots to perform a variety of tasks. NEAT utilizes three main techniques: tracking genes, implementing speciation and developing topologies incrementally from simple initial structures. Our final project seeks to explore how communication strategies between swarm robots can be evolved using minimal hard-coding in NEAT, meaning that we do not want to *provide* extraneous information and instructions but rather have the agents *evolve* complex behaviors.

1.1 Related Work

Roboticists, Floreano, Mitri and Magnenat discuss the necessary conditions for evolution of communication between robots to take place in their work[1]. In their experiment, they conduct repeated trials among robots that generate and emit visual signals to provide information about food location. Their environment consists of one "food" object which robots are attracted to and a poison object which each robot has an aversion to. In their experiment, they discover that effective communication amongst the robots was evolved and the robots went on to mimick biological natural selection by recombining successful genomes.

Additionally, researchers Marocco and Nolfi further discuss the relevance of communicative intelligence amongst robots[3]. They implement a navigation task in their environment whereby controllers have to make meaning of frequency values in order to build an effective communication system. They implemented bidirectional communication where their robots could play a hearing and a speaking role. They conclude that their robots were able to evolve a meaningful communication system without hard-coding.

Finally, in research executed by Yong and Miikkulainen's *Co-evolution of Role-Based Cooperation in Multigent systems*, they conclude that evolving cooperation was more effective through stigmergy rather than communication[4]. In their experiment, they co-evolve agents in their world with neural networks in separate sub-populations and then go on to test them together in a prey-capture task. This further incited us to test the value of communication in evolving cooperation.

2 Experimental Setup

The tools that we used for our experiment included the Python Graphics Library for the implementation of our simulator where we ran our experiments, and the Python NEAT Library which we utilized to train our swarm robots. The essential components of our experiment involved assessing how powerful NEAT is by providing minimal hard coding and information. To represent communication in NEAT, we had NEAT take in *frequencies* from the 4 quadrants of an agent, decide what frequency the agent should emit and finally decide how to behave according to the frequency inputs. Frequency is represented as a floating point value with range [-1,1] according to the tanh(hyperbolic tangent) activation function we used for NEAT. With this setup, each agent emits a frequency determined by the NEAT brain at each time step while receiving four frequencies from its four quadrants. Finally, in total the agents were given 8 inputs: stall, a boolean value to determine whether the agent has stalled; foodReached, a boolean value to indicate whether an agent has reached the food; the four frequencies from an agent's four quadrants; coverage, a float indicating the percentage of the world an agent has covered; and lastly a bias of 1. Based on these inputs, each agent output its motors: translate and rotate, as well as its communication content represented as a frequency.

2.1 Learning Task

The learning task for our swarm robots is to learn how to cooperate and gather around food. Specifically, three agents will search in a fixed sized world for a food object, and the objective is for the agents to all reach the food as fast as possible. The learning goal is for our agents to evolve communication strategies that will enable them to solve the learning task efficiently and collaboratively.

Figure 1: This figure shows the environment with the three agents and food object

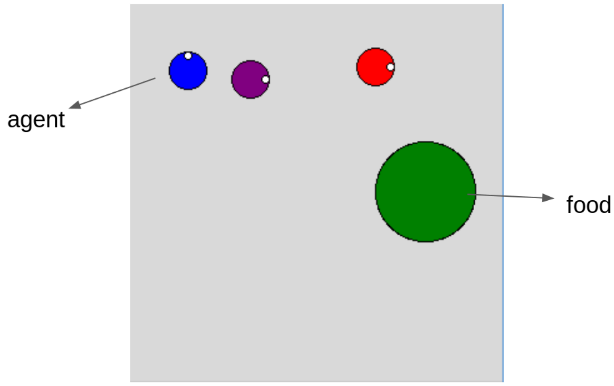

2.2 Environment

Our environment consisted of a 500 by 500 world that included 3 circular agents colored blue,red and purple of radius 20 units. It also included a food object of radius 70 units. In each run, both agents and the food were spawned in randomized positions. Figure 1 demonstrates such environment.

2.3 Communication

In order to allow each agent to receive frequencies from one another, we implemented a quadrant system. An agent was divided into four sections based on its heading. If another agent, say $a2$, lies within a quadrant, then the agent in question would receive the frequency of $a2$. On the other hand, if no agents lie within a quadrant, then the agent would receive a frequency value of 0 for that given quadrant. If there is more than one agent within the same quadrant, the agent only receives the frequency of the nearest agent. Figures 2, 3, and 4 further illustrate this in more detail. Quadrants are labeled from 0 to 3 beginning from the upper right in respect to the agent's heading going counter-clockwise and the signals the agents hear are represented in the same order.

2.4 NEAT

NEAT has demonstrated itself as an effective method in evolving neural networks. One of its distinct characteristics is its ability not only to search for fitting network weights for the neural network but also to incrementally complexify the entire structure, yielding more sophisticated and optimized solutions to the task.

2.4.1 Parameters

Table 1 presents the parameters that we used for our experiments. Some of these parameters were decided based on the experimental setup, such as having tanh as the activation function to allow for

Figure 2: Red Agent hears (0,0,f_p,f_b). Note: f_p and f_b represent the frequencies of the purple and blue agent

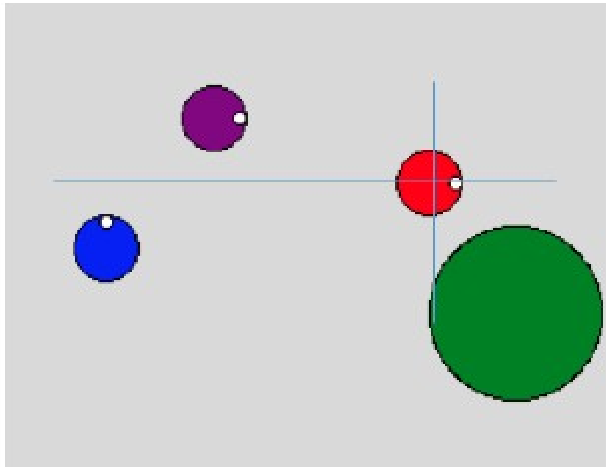

Figure 3: Blue agent hears (f_p,0,0,0).

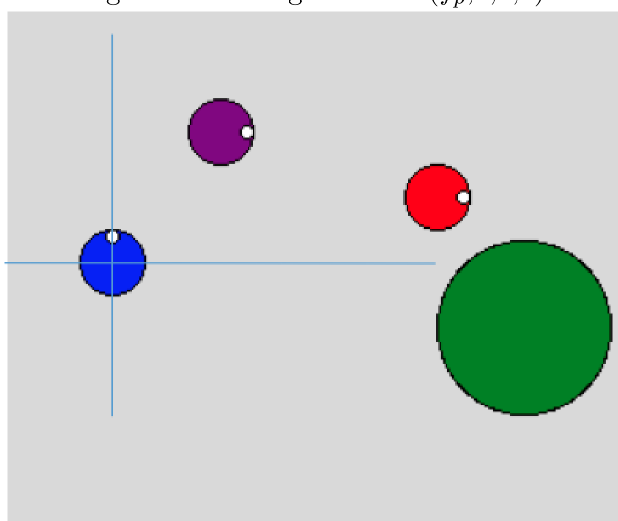

negative values and the sizes of the input and output nodes, whereas other parameters are results of trial and error- for example, we discovered that evolving agents over more than 25 generations did not produce novel behaviors.

2.4.2 Fitness Function

Our fitness function consisted of three main components: steps, distance and coverage. Steps was used to calculate how many steps it took for a robot to reach the food. Distance measured how close the agents were to food. Lastly, coverage was used to give the robots incentive to explore the world in order to maximize the chances of stumbling upon food.

Figure 4: Purple agent hears (f_r,0,0,f_b)

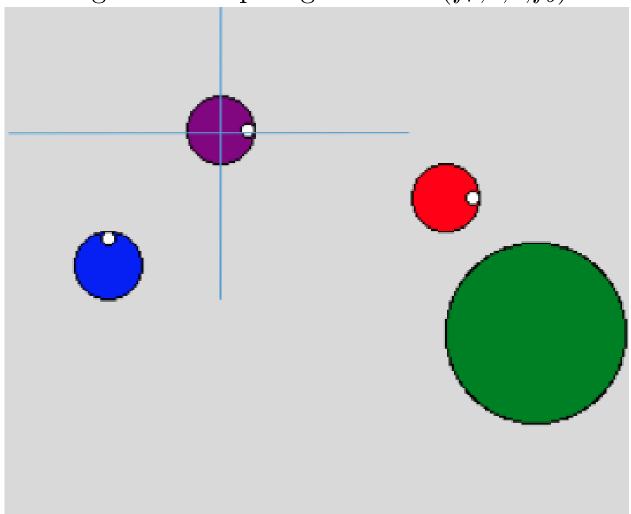

Parameter	Setting
Generations	25
Population size	50
Species size	10
Input nodes	8
Hidden nodes	0
Output nodes	3
Activation function	tanh
Prob. to add link	0.1
Prob. to add node	0.05

Table 1: NEAT parameter settings used in the experiments

Steps was calculated by the formula

$$f(s) = (max\ steps - steps)/max\ steps \tag{1}$$

Distance was calculated by the formula

$$f(d) = (max\ distance - distance\ away\ from\ food)/max\ distance \tag{2}$$

Finally, coverage was calculated by determining the percentage of the area that the agent covered in the world. We set max steps to 1500 steps and max distance to be the longest distance in the simulator world, which is $\sqrt{2} * 500$.

2.5 Test groups

We had 3 main test groups which we wanted to compare with our experiment:

1. Free communication: In free communication, we do not set any restrictions on how NEAT evolves the agents. In particular, we allow NEAT to determine what each agent should emit as their frequency. This is contrasted with the second test group, hard-coded communication.
2. Hard-Coded Communication: For hard-coded communication, an agent would constantly emit 0 as its frequency until it reaches the food, which then it would start emitting the frequency 1. We hope that this hard-coding scheme would simplify the process of evolving communication and instead allow NEAT to focus on reacting to the communication.
3. No Communication: Frequency was completely taken away so that the agents had absolutely no form of communication and this served as our control method.

2.6 Evolving and Procedure

It is important to note that we evolve each of our agents using the same NEAT brain. Each agent possesses its individual but identical NEAT brain and is evaluated based only on its performance. The final fitness function of one run then was composed of the average performance of the three agents. Below is the algorithm that we run to evolve our agents and determine the fitness score for one run:

```
//initialize
generate the three agents and the food on random positions
set the same NEAT brain to each of the agents

//evolve
for 1500 steps
    step the world for one step
    for each agent
        check if agent reaches food
    if all agents reached food
     end run

//get fitness
for each agent
    get its total fitness score
set overall fitness score to be the average total
```

Additionally, this is run three times for one population to once again account for the randomness of the initial positions.

3 Results

To verify our hypothesis, we evolved and evaluated the three test groups. Numerical results were collected to show how the three test groups performed while both qualitative observations and quantitative data were made to analyze the communication results.

3.1 Fitness Comparison

Here we simply compare how the three test groups perform relative to each other based on the fitness function. Since each specie was evolved independently and may possess different strategies at completing the task, we decided to compute the fitness of each test group on each of their ten chromosomes. In addition, since the initial position of the food and agents were randomized, to account for this we evaluated each species three times and averaged their fitness. Finally, an average score for the ten chromosomes would represent the overall score of a test group.

Chromo	Average Fitness(free)	Average Fitness(hardcoded)	Average Fitness(no comm.)
0	0.29059	0.31120	0.54712
1	0.35578	0.33129	0.63062
2	0.40969	0.49447	0.44200
3	0.36733	0.40259	0.57396
4	0.31304	0.31826	0.53162
5	0.31421	0.31304	0.65641
6	0.40521	0.31421	0.57475
7	0.40478	0.40521	0.63000
8	0.25755	0.40478	0.63812
9	0.25755	0.25755	0.65514
Avg	0.31181	0.33213	0.58797

Table 2: Fitness Comparison Of All Three Test Groups

As we can see from Table 2, the no-communication test group actually performed the best, whereas the scores between the two communication test groups were trivial. This disproves our hypothesis, which was that the free communication group would perform the best because they would solve the task by collaboration, however the results demonstrated otherwise. Upon further investigation, these results in fact were reasonable. For one, the no-communication scheme had fewer input and output nodes, which made the network easier to evolve and simplified the amount of information the brain was taking in.

When observing the strategies that the no-communication group evolved, we discovered that the agents across all 10 species evolved identical approaches, that of travelling along the edges of the world in straight lines and angling slightly towards the center at each turn. This guaranteed that the food object would eventually be reached by all three agents by thoroughly covering the entire search space. As for the other two communication groups, the majority of the species did not evolve solutions to this task. For those that did however, a common trait was that they utilized variations of what we called "the battleship strategy", meaning that each agent would behave exactly the same way and try to spread out through the space to maximize the probability of reaching the food; this was quite similar in essence to the strategy evolved by the no-communication group, and by observation we did not notice any evidence that the agents were communicating.

3.2 Communication Observation

Besides assessing the performances of each test group based on their fitness scores, we observed how each group completed the task and made the conjecture that there were no signs of communication being evolved. Still, to investigate whether effective communication was evolved and to what extent did NEAT make sense of the inputs, we came up with two methods to examine the evolving of communication numerically. Note that we only performed these two tests on the free communication test group as the second method would not be applicable to the hard-coded communication group.

First, we wanted to see if the agents would change their behaviors when one agent reached food. Ideally, as one agent reaches food and communicates that information, the other two agents should somehow move towards where that agent is. In Table 3 we presented the time steps, whether the first agent(A1) to reach the food reached the food or not, and the motor outputs of the other two agents (A2, A3). The expected behavior is that as A1 reaches food, there would be noticeable changes in terms of the motor values of A2 and A3 because they would then shift towards A1. However, our results proved otherwise: in terms of speed, we can see that it stayed pretty consistent throughout. As for rotation, there was a pattern if we examine the values: each unique rotational value was about 0.005 away from the previous one and each rotational value would stay for about five time steps. The fact that A1 reaching food did not break these patterns demonstrated how communication did not affect the behaviors of the agents.

The second method that we used was to observe the frequencies of all three agents over time, and especially study whether reaching food would change the frequencies significantly. Figure 4 illustrates our results, and as we can see the frequencies of each agent stayed at about the same value until after a certain point spiked to another value and plateaued. The exciting thing about this result was that upon closer inspection, the time step at which the values spiked were exactly the one step after the agent reached food. In other words, the NEAT brain responded to the foodReached input and output a frequency that was significantly different.

Steps	**Food Reached A1**	**Speed M2**	**Rotation M2**	**Speed M3**	**Rotation M3**
260	False	-0.999993	0.129033	-0.999993	0.150571
261	False	-0.999993	0.134430	-0.999993	0.150571
262	False	-0.999993	0.134430	-0.999993	0.150571
263	False	-0.999993	0.134430	-0.999993	0.150571
264	True	-0.999993	0.134430	-0.999993	0.155934
265	True	-0.999993	0.139819	-0.999993	0.155934
266	True	-0.999993	0.145199	-0.999993	0.155934
267	True	-0.999993	0.145199	-0.999992	0.161287
268	True	-0.999993	0.145199	-0.999992	0.161287
269	True	-0.999993	0.150571	-0.999992	0.161287
270	True	-0.999993	0.150571	-0.999992	0.166632

Table 3: Change in motors as agent 1 reaches food

4 Discussion

To reiterate our results, we first concluded that the no-communication test group performed the best by the fitness score. Next, by observing how the agents solve the task, we conjectured that no meaningful communication was evolved in either of the two communication test groups. To further verify that, we examined the relationship between one agent's foodReached input with the motor

Figure 5: Frequencies of the three agents in free communication

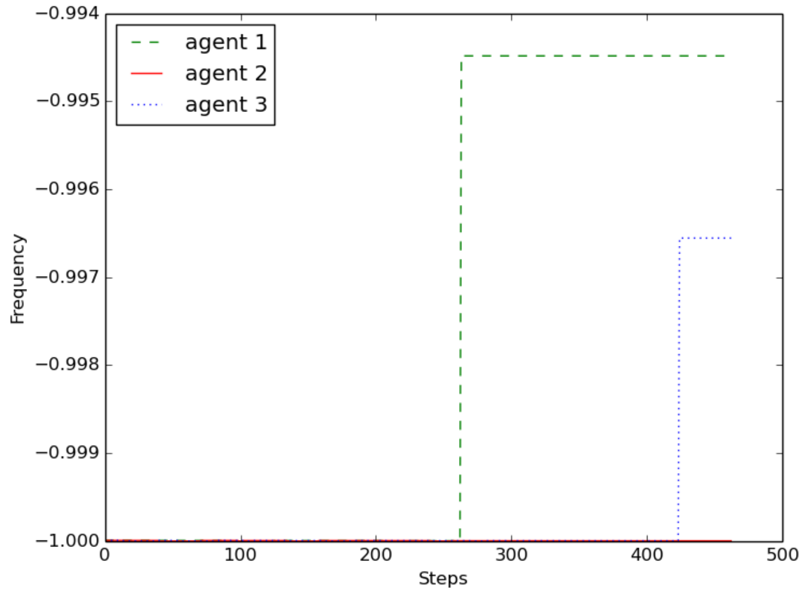

outputs of the two other agents to see if the other two agents would react differently as one agents reaches food, and our examination confirmed our conjecture that in fact whether one agent reaches food or not does not make any impact to the behaviors of the other two agents. However, when we looked at the frequencies that each agent emitted in relation to when they reached the food, we did find out that each individual agent responded to their own foodReached input by outputting an apparently different frequency than before.

Reflecting upon the experimental procedure and results, there were perhaps several explanations as to why communication was not being evolved. For one, we believe that the amount of randomness used in this experiment significantly made a negative impact on the evolution. A new position of the agents and the food made it a lot harder for the agents to learn about their tasks and come up with meaningful strategies since an approach that worked for one configuration may not work for another, and it is very possible that a promising strategy could easily be discarded because a new configuration makes that strategy unfavorable. We tried to account for this randomness by running each population three times and averaging out the fitness score, but still we realized that this wasn't sufficient as we would still get varying scores for the same species. For another, as we showed in the results, each individual agent was able to learn the correlation between whether it reached food or not and what frequency it should emit. However, it was the communication among the agents that was not effective.

Each agent had four input nodes that represented the four communication content from its four quadrants, yet it seemed that this way of representing communication did not work out. It is possible that having four inputs for communication were just too difficult and complex to make sense of. Perhaps a further improvement is to investigate alternative communication schemes that would be just as explicit but easier and simpler to evolve. Still, even with this failed attempt of evolving communication, we believe that communication still plays a vital role in swarm robotics, particularly in tasks where the agents do not have much knowledge of the world or information of others. Even though our results showed that the no-communication test group performed the best and the most consistently by far, imagine if the size of the world is significantly larger. It is then obvious that this brute-forced method that the no-communication test group evolved would not be feasible and that more clever strategies using communication are necessary.

References

[1] Dario Floreano, Sara Mitri, Stphane Magnenat, and Laurent Keller. Evolutionary conditions for the emergence of communication in robots. *Current Biology*, 17, 2007.

[2] Amy Jin and Murphy Austin. Comparative evolutions of swarm communication. 2014.

[3] Davide Marocco and Stefano Nolfi. Emergence of communication in embodied agents evolved for the ability to solve a collective navigation problem. *Connection Science*, 19, 2007.

[4] Chern Han Yong and Risto Miikkulainen. Coevolution of role-based cooperation in multiagent systems. *IEEE Transactions on*, 19(2), 2011.