Evolutions of communication

Alex Bell, Andrew Pace, and Raul Santos

May 12, 2009

Abstract

In this paper a experiment is presented in which two simulated robots evolved a form of communication to allow one of the robots to complete the task of finding a light placed in the world. The robots are constrained so that one has inputs about the world, seeing robot, and the other, blind robot, has no sensors besides an input from the seeing robot. The evolved communication allowed the seeing robot to develop an algorithm in which it told the blind robot to rotate until in was in front of the light and then the communication told the robot to move forward until it hits the light

1 Background

1.1 NEAT

Like most evolutionary robotics algorithms the NEAT algorithm has two parts, the evolution of the phenotype and search for weights between neurons [3]. The NEAT algorithm starts with the simplest configuration possible, input neurons directly connected to output neurons with random weights. This basic configuration is then modified with new nodes, connections, and deletions to grow the phenotype. There are three ways in which the NEAT algorithm differs from other evolutionary methods. The first is the manner in which crossover is conducted. Instead of random crossover the NEAT algorithm keeps track of all gene's creation points with an innovation number. During crossover only genes with the same innovation number are allowed to swap. Another problem of evolution is that complex individuals, which may produce better results after more development time, are eliminated too soon due to lower fitness scores against simpler individuals. To prevent this, the NEAT algorithm divides the population into different species based on topological similarity and only similar species are allowed to compete against each other for a certain number of generations. Lastly NEAT prevents overly large and poorly performing topologies through complexification. Complexification allows large topologies but only if they are better, by allowing evolution to add nodes but only topologies in which the added node is helpful are allowed to survive. Neat was chosen to evolve the neural network because it is able to generate its own phenotype and concurrently adjust the neural networks weights.

1.2 The emergence of communication in evolutionary robots

"The emergence of communication in evolutionary robots" is a paper in which two parts of research in communication in robots are discussed [2]. In the first part evolving neural nets send basic signals as to the value of objects. In the second part of the research the idea is extended to examine components of language such as verbs. In regards to our project the first part of their paper was most interesting and similar. Their experiment consisted of a robotic arm with 6 DOF in a 3D world with simulated gravity. The arm had touch sensors and in the world were a sphere and a cube. The parents of each evolving neural net had two output nodes which fed into the inputs of the childrens neural net. The fitness of the child was calculated based on the amount of time spent touching the sphere vs touching the cube. The interesting part was that the parent received no fitness for the childrens action. Only the child received fitness. But because of the evolutionary process only children who outputted correct values to their children would continue after a few generations. In the paper the robots were able to evolve a form of communication. The interesting question posed and answered in the paper was whether or not the communities that evolved communication were better off. Some robots evolved simply to poke around and if they found the cube stay away from it. In the paper they found that indeed communication did mean a more fit population. In our case we did not consider this point because it is obvious as we have set it up that the blind robot will never find the light, other than randomly, without communication. This paper however posses the interesting idea of perhaps giving the blind robot light sensors and determining whether communication in fact creates a fitter population.

1.3 The emergence of language in an evolving population of neural networks

In the paper "The emergence of Language in an evolving population of neural networks," a world is created in which evolving neural nets roam among ten healthy and ten poisonous mushrooms [1]. Each net has the ability to output a one or zero to the other nets. The individual fitness is calculated for each net depending on the number of healthy and poisonous mushrooms eaten. The interesting point about this paper was that there was no evolutionary pressure for the individual robots to communicate. In that they did not receive any increase in fitness for correct outputs. However the paper notes that a language of identifying mushrooms as good or bad did evolve and more impressively it was able to increase the performance of the robots so that communicating improve the general fitness of the population. In regards to our robot we thought it would be required for us to give the same fitness for both robots depending on the eating results of the blind robot. This may not have been necessary. In fact it may have been better if in stead we had somehow correlated the fitness function of the blind robot with how well it adopted to the outputs from the seeing robot instead of how many lights it ate. In this way we might have isolated exactly what we wanted which was for the seeing robot to guide the blind robot to the light. The example in the paper was different in that their robots hadthe same brain and the same task, however the interesting points suggested ways in which we could change our approach to the problem.

2 Experimental Procedure

The world used in this experiment consisted of two robots and one light, as shown in Figure 1. One robot, referred to as the blind robot and colored blue, had only one input node and two output nodes. The input nodes was the communication link between the two robots. The output nodes where linked to the blind robot's left and right motors. The other robot, referred to as the seeing robot and colored red, had three inputs and one output. The three inputs consisted of, the distance from the blind robot to the light, the cosine of the angle from the blind robot to the light, and the sine of this same angle. The only output from the seeing robot was the communication link to the blind robot, whatever value the output node had was passed to the input node of the blind robot.

Both robots were controlled using genomes coevolved through the use of NEAT. The NEAT program used in this paper had been modified to allow for simultaneous coeveolution of two separate agents. The modification set the current generation to evolve with the best genome from the previous generation of the other brain. For example, if two robots A and B are coevolving, and robot A is on generation 25, then it uses the best genome from generation 24 of robot B. A hall of fame was not implemented in this experiment because for each generation, its evolution affects the communication. Even if one generation may have a more effective communication, if the two robots had not evolved to associate the same meaning to the numbers, the trial would fail. In order to begin the coevolution,

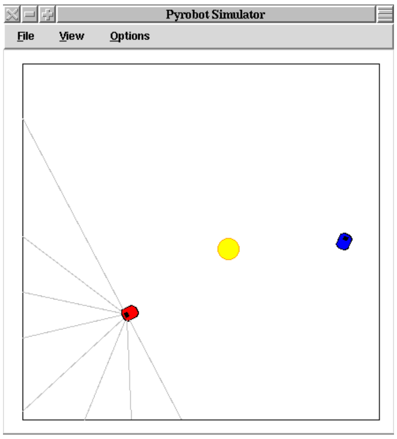

Figure 1: The world used throughout the experiment. The seeing robot is red and the blind robot is blue.

a random genome was initialized for both the seeing and the blind robots, so that in each generation 0 there was a genome for the other robot's brain.

In order to use NEAT, or any variant of a genetic algorithm, a fitness function needs to be defined. If the blind robot moved closer to the light in a given time step, then the score was increased by one. If the blind robot moved away, the genome's score was decreased by one. The purpose of this was to encourage movement towards the light without requiring the blind robot to randomly pass over it. The key aspect of the fitness function was increasing the score of the genome by 1000 whenever the blind robot passed over the light.

A trial of a genome would run as follows. First, the program would determine the distance from the robot to the light. Then the program determines the angle the robot needs to orientate itself towards the light. The value of the distance, and the sine and cosine of the angle are then set as input to the seeing robot. The output from the seeing robot is then passed, by the program, to the sole input of the blind robot. The output from the blind robot's neural net was then used to control the left and the right motors. If the blind robot ran over the light, the light was then randomly positioned elsewhere in the room.

3 Results

After the evolution was finished, four trials were made using the best chromosome in each generation. The fitness values of those trials are recorded in Figure 2. The last column indicates the average of the four trials. There number of trials was decided taking into consideration the amount of time that it took to run them, and the usefulness to normalize random behavior. The data indicates that generation 29 was able to achieve the highest average. Figure 3 plots the trial averages of the 40 generations.

4 Discussion

This section explores how the language evolved between the two robots over time. The example outputs presented below are for only one trial however it is illustrative of the majority of trials.

At generation 2 a trial is mapped in Figure 4 with the inputs to the seeing robot, which translate into the blind robots movements, over the length of the trial. In generation two the seeing robot

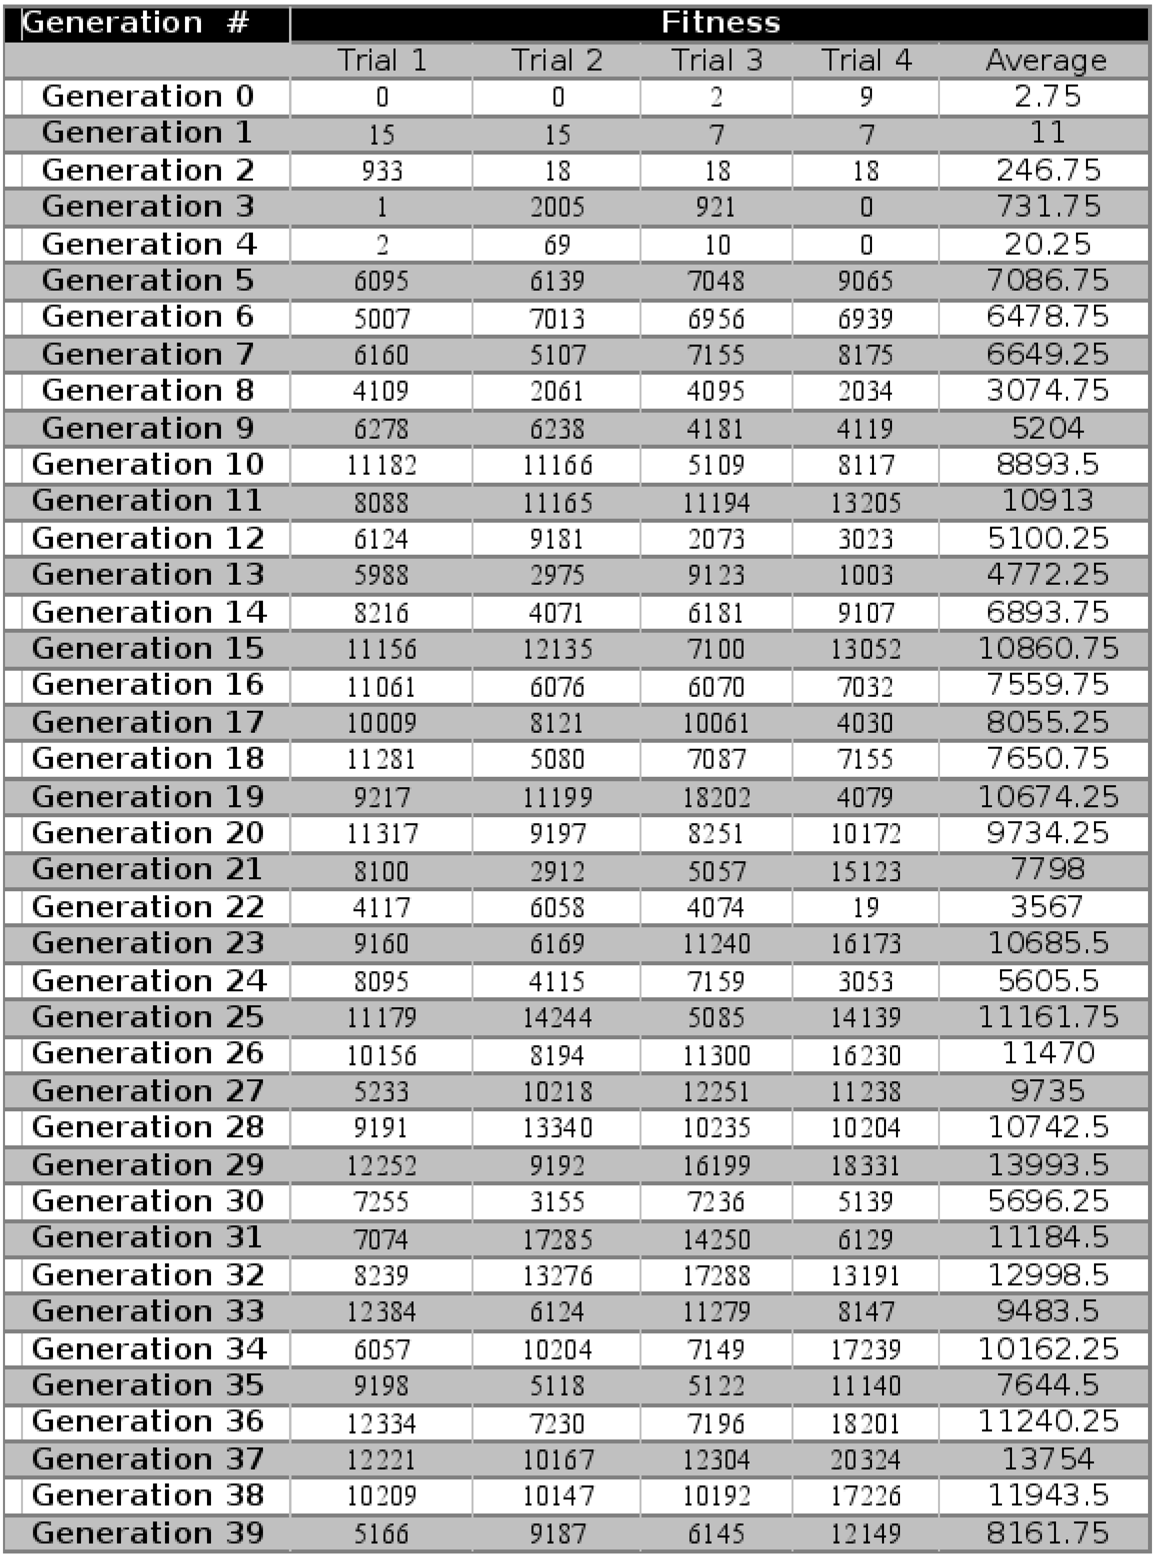

Generation #	Fitness				
	Trial 1	Trial 2	Trial 3	Trial 4	Average
Generation 0	0	0	2	9	2.75
Generation 1	15	15	7	7	11
Generation 2	933	18	18	18	246.75
Generation 3	1	2005	921	0	731.75
Generation 4	2	69	10	0	20.25
Generation 5	6095	6139	7048	9065	7086.75
Generation 6	5007	7013	6956	6939	6478.75
Generation 7	6160	5107	7155	8175	6649.25
Generation 8	4109	2061	4095	2034	3074.75
Generation 9	6278	6238	4181	4119	5204
Generation 10	11182	11166	5109	8117	8893.5
Generation 11	8088	11165	11194	13205	10913
Generation 12	6124	9181	2073	3023	5100.25
Generation 13	5988	2975	9123	1003	4772.25
Generation 14	8216	4071	6181	9107	6893.75
Generation 15	11156	12135	7100	13052	10860.75
Generation 16	11061	6076	6070	7032	7559.75
Generation 17	10009	8121	10061	4030	8055.25
Generation 18	11281	5080	7087	7155	7650.75
Generation 19	9217	11199	18202	4079	10674.25
Generation 20	11317	9197	8251	10172	9734.25
Generation 21	8100	2912	5057	15123	7798
Generation 22	4117	6058	4074	19	3567
Generation 23	9160	6169	11240	16173	10685.5
Generation 24	8095	4115	7159	3053	5605.5
Generation 25	11179	14244	5085	14139	11161.75
Generation 26	10156	8194	11300	16230	11470
Generation 27	5233	10218	12251	11238	9735
Generation 28	9191	13340	10235	10204	10742.5
Generation 29	12252	9192	16199	18331	13993.5
Generation 30	7255	3155	7236	5139	5696.25
Generation 31	7074	17285	14250	6129	11184.5
Generation 32	8239	13276	17288	13191	12998.5
Generation 33	12384	6124	11279	8147	9483.5
Generation 34	6057	10204	7149	17239	10162.25
Generation 35	9198	5118	5122	11140	7644.5
Generation 36	12334	7230	7196	18201	11240.25
Generation 37	12221	10167	12304	20324	13754
Generation 38	10209	10147	10192	17226	11943.5
Generation 39	5166	9187	6145	12149	8161.75

Figure 2: Fitness of the Best Individual in the Generation

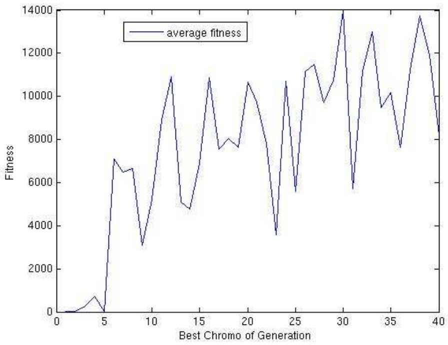

Figure 3: Fitness of the Best Individual in the Generation

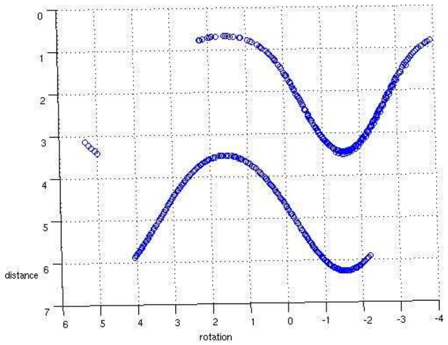

Figure 4: Generation 2

outputted a value of -1 for the entirety of the trial with no variation. The blind robot turned right in a constant radius circle for the entire length of the test. It is important to note however that the turn was to the right. This will be important in subsequent tests.

The graph of generation 5 is shown in Figure 5. While the graph looks very complicated it is simply a pattern which is repeated. The robots start in the ends of the straight section, for instance rotation 1.5 and distance 3.7. The robot rotates in place until a certain point, -2.5 rotation and 3.5 distance. At this point the robot travels straight towards the robot and the rotation angle appears to stay approximately the same with the distance decreasing. Throughout the trials the seeing robot is outputting one of three values, either a 1, -.5, or -1. In the case of of 1 the robot goes straight, -.5 the robot takes a sharp right, and -1 a medium right turn. However there are large sections of the graph which are hard to decipher due to the large number of points close together. The action the robots are taking as they near the light is causing the mass of points. The pattern which happens is the robot turns in place, -.5, until it is almost facing the light. Then it advances and periodically adjusts course by alternating between a medium right, -1, and straight,1. This alternating produces the zig zagging motion. However the motion is ineffective because it causes the robot to spend many steps zig zagging very close to the light.

The graph of generation 25 is shown in Figure 6. In it is evident the same pattern of turning in place until a certain point and then advancing. However there is no zig zagging present in the robots motion. This is confirmed by examining the output of the seeing robot which has only two outputs -1,

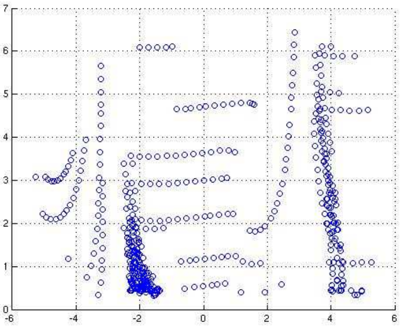

Figure 5: Generation 5

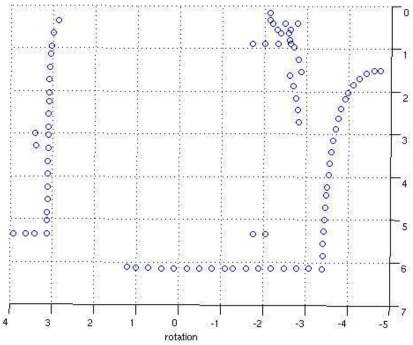

Figure 6: Generation 25

which is a hard right turn, and 1, which is travel forward. The results are much better than generation 5, the robot turns until is facing the robot and then travels straight right to it. There is no need for correction and none is taken. The robot reaches the target faster than does generation 5.

In the previous three graphs the language is shown evolving. At first there is no language designed to achieve the goal. However there is agreement or conjunction in action and signal because only one output is sent, -1, and only one action is taken, turn right. Interestingly by generation 5 when the robots are achieving the task they are building upon the agreement achieved in generation 2 that is, -1 means turn right. In generation 5 the robots have achieved the task however the method is much more complicated then it has to be and therefore does not work as well. Lastly the final example, generation 25, shows the language which is the last developed and is never replaced in subsequent generations.

The interesting characteristic of the language evolution is that the language appears as if a more complicated language is evolved first before the simpler language. From an adaptive robotics point it is expected that the evolution technique will find the simplest method possible. In addition in adaptive techniques care must be taken, through speciation, to give more complicated genotype a chance to evolve. Hence more complicated techniques are expected to evolve slower and require more time to reach good operation. However in this example we see a simple solution evolve after the more complex. We believe this is due to the communication required. The added necessity of communication we feel added another variable which prevented the simplest solution from being developed first. Because there was co-evolution which required two randomly evolving neural nets we believe this was enough to prevent the easiest solution from evolving first.

Looking at the data, we are able to discern that the robots were able to evolve a strategy to achieve the objective. Generation 5 shows fitnesses that were not due to random behavior. The highest possible fitness if no light is reached is 600. Generation 5 was able to achieve an average of 7000, which means that it reached 7 lights. Since there are only three trials per run, this means that it was able to reach about 2 lights per run. To do this, the seeing robot should be able to issue a signal depending on the location of the light in relation to the blind robot. The blind robot, in turn, should be able interpret the signal and move accordingly. As the generations progressed, their behavior got better and we were able to achieve fitness scores higher than 10000. Generation 29, the best generation, had an average of 14000, meaning that it was able to reach more than 4 lights per 200 steps.

The variation of the fitness in the trials is due to the random placement of the light. The light could potentially be very far or very close to the location of the robot, thus, taking it a variable number of steps to reach it. We can especially see this in generations two and three, where the light probably appeared close enough to the robot that it was able to reach it using its rudimentary strategy. This could have affected the evolutionary runs since an individual could have had many of the lights placed near it, and a better one might have had to travel further to reach its objectives. This would create the false sense that the first individual was better than the second and more of that species would be carried over to the next evolutionary run. We can infer this a couple times. Looking at the results of the best individual in generation 11, we can see that its average behavior was clearly better than that of the next few generations. Even though in three of the experimental trials it was able to achieve values higher than 10000, we can also observe that it had one trial with a fitness of 8000. If we look at the other generations, we can see that each of them was able to achieve a value higher than the lowest score for generation 11. Although this is a very possible scenario, a better strategy would be able to achieve higher values consistently, making it easier for it to survive the future generations.

References

[1] A. Cangelosi and D. Parisi. The emergence of 'language' in evoluing population of neural networks. *Connection Science*, 10:83–97.

[2] D. Marocco, A. Cangelosi, and S. Nolfi. The emergence of communication in evolutionary robots. *Philosophical Transactions of the Royal Society of London-A*, pages 2397–2421, 2003.

[3] Kenneth O. Stanley and Risto Miikkulainen. Competitive coevolution through evolutionary complexification. *Journal of Artificial Intelligence Research*, 2004.