Emergence of Communication in Competitive Multi-Agent Systems: A Pareto Multi-Objective Approach

Michelle McPartland
Artificial Life and Adaptive Robotics (A.L.A.R.) Lab
University of New South Wales
Australian Defence Force Academy
Campus Northcott Drive, Campbell, ACT 2600, Australia
m.mcpartland@adfa.edu.au

Stefano Nolfi
Institute of Cognitive Science and Technologies
CNR, Viale Marx 15
Rome, 00137, Italy
s.nolfi@istc.cnr.it

Hussein A. Abbass
Artificial Life and Adaptive Robotics (A.L.A.R.) Lab
University of New South Wales
Australian Defence Force Academy
Campus Northcott Drive, Campbell, ACT 2600, Australia
h.abbass@adfa.edu.au

ABSTRACT

In this paper we investigate the emergence of communication in competitive multi-agent systems. A competitive environment is created with two teams of agents competing in an exploration task; the quickest team to explore the largest area wins. One team uses indirect communication and is controlled by an artificial neural network evolved using a Pareto multi-objective approach. The second team uses direct communication and a fixed strategy for exploration. A comparison is made between agents with and without communication. Results show that as the fitness function vary differing exploration strategies emerge. Experiments with communication produced cooperative strategies; while the experiments without communication produced effective strategies but with individuals acting independently.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning - *Concept learning Connectionism and neural nets*

General Terms

Experimentation, Theory

Keywords

Communication, multi-agent systems, artificial life, evolutionary robotics and adaptive behavior

1. INTRODUCTION

The purpose of this paper is to investigate the emergence of communication strategies in a competitive multi-agent system. Evolutionary techniques are used to develop the communication strategy within a population of agents. The experiments encourage the emergence of communication strategies through a task where cooperation between individuals is conducive to the improvement of the group's behavior.

The task that has been selected consists of two populations of agents competing for exploration area of an environment. The first population (Blue) will evolve a competitive strategy through the use of indirect communication. This indirect communication is analogous to pheromone trails deposited by ants and only occurs among Blue agents. This encourages Blue agents to evolve an effective communication strategy which is cooperative and coordinated. Whether any strategies develop however, is left entirely to evolution. The opposing agents (Red) will use a fixed strategy involving direct communication. This strategy is based on complete knowledge of the environment and is used to compete against the evolving Blue population. Red agents will not be able to deposit or detect pheromone trails left by Blue agents. This situation provides an asymmetric problem space as Reds have complete information of the environment while Blues have only local information. However, both populations have the same goal: to increase the area explored by their team.

The following section reviews a small sample of the current literature on artificial life techniques relating to communication. Section 3 describes the experimental design and setup. Section 4 presents the empirical results. Section 5 provides an observational analysis of the agents' behavior that has evolved through communication. The paper concludes in section 6 with further directions this work might take.

2. BACKGROUND

The origins of communication systems apparent in sentient organisms are largely unclear. In the past decade researchers have been using artificial life techniques to study this phenomenon. Communication can be seen as the interaction between intelligent entities [11]. Even in insect societies this is evident in the use of pheromone trails by ants foraging for food and the waggle dance of bees to indicate the direction and distance of nectar.

Communication systems can be divided into two distinct categories: direct and indirect. In the former category the behavior of one individual directly affect the sensors (and potentially the behavior) of other individuals. In the latter category the behavior of one individual instead affects the external environment. The modification of the external environment might later affect other individual behaviors. Speech and emails are examples of direct communication. Pheromone trails are a biological example of indirect communication.

Evolutionary computation models of communication are represented by three different paradigms: signals, symbols and complex syntactical structure [3]. Significant work has been undertaken in understanding signal-based communication [9][10][11][12]. Signals generally occur in animal or insect societies and rely on the transmission of information through a sign or indication. In the work of [3] and [16] symbolic models have been evolved through object recognition tasks (physically and virtually respectively) which thereby develop lexicons. The third model of communication, complex syntactical structure, concerns the development of language. Refer to research by [3] and [6] for further information concerning complex syntactical structure. For an analysis of the adaptive factors that might favor the emergence of communication see [4][13][14].

Mataric uses simple broadcasting techniques to aid in the learning of multi-agent systems situated in complex, noisy environments [11]. These experiments illustrate the use of simple communication systems by robots to learn social skills. The robots broadcast information concerning their current behavior and reinforcement value (the value they recently received for performing a useful task). Robots in close proximity are encouraged to imitate this behavior thereby learning useful skills. The social behaviors learnt by the robots were broadcasting location of objects and yielding to other robots. These behaviors are non-trivial to learn as reinforcement encourages the individual to be selfish (i.e. they want to be the highest scoring robot in the population), which is not conducive to cooperative learning.

In the work by [10] we see the evolution of communication unaided by the designers. Agents were embodied with artificial neural networks (ANN) and were provided with the abilities to produce signals with varying intensities and to detect signals up to a certain distance. Evolved agents developed a communication system involving four different signals that allowed them to coordinate and cooperate in order to solve a collective navigation problem.

Other successfully evolved communication systems demonstrate lessons concerning the origins of communication. Smith hypothesizes on the evolution of language through experiments involving natural selection and cultural transmission [15]. Learning biases, related to cultural transmission, of individuals in the population, were found to have a significant effect on the evolution of a reliable communication system. In some cases the learning abilities of the agents could cancel out the effect of natural selection. These results led to the implication that the evolution of language occurs in two distinct stages: firstly that a learning mechanism is adapted over a geological time scale and secondly that this learning mechanism enables the development of language over a historical time scale.

The need to investigate the evolution of communication is important for a multitude of reasons. Using artificial life techniques to evolve communication strategies may aid in the understanding of how communication (including language) may have evolved in animals, insects and humans alike [8][14].

This paper attempts to explore this issue by creating an environment which encourages cooperation among a team of agents through the use of simple indirect communication. A competitive environment is achieved by means of a second team who has access to complete information via direct communication. The use of communication is altered in the experiments to investigate the affect it has on the task at hand.

3. METHODOLOGY

The two populations of agents (Blue and Red) live in a rectangular environment of cells. Each agent occupies exactly one cell, and no two agents can occupy the same cell {see Figure 1}. Movement of the agents is restricted to the 8 surrounding cells in the environment, provided the cell is empty. All agents are moved synchronously through a random bidding system.

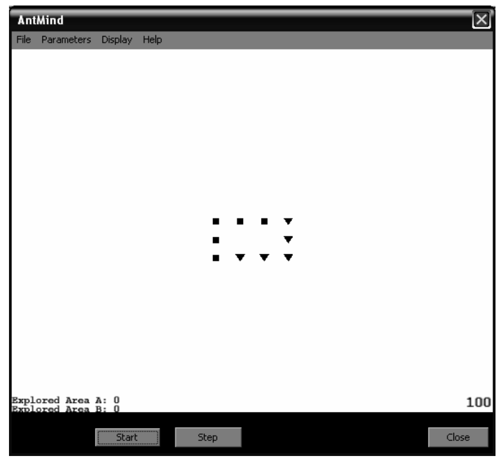

Figure 1. 20x20 environment with 5 Blue agents (squares) and 5 Red agents (triangles) in the center starting position.

The behavior of each Blue agent is controlled by 5 identical ANNs, analogous to the central hive brain of insects such as bees, wasps or ants (i.e. the team is homogeneous). Each ANN has 27 sensory neurons, 10 internal neurons (that receive connections from the sensory neurons), and 4 motor neurons (that receive connections from the sensory and motor neurons). Sensory neurons encode: (a) the intensity of pheromone trails in the 8 adjacent and the single current cell of the agent (9 neurons); (b) the presence of obstacles (agents or walls) in the surrounding cells; and, (c) the activation state of internal neurons at time t-1 {see Figure 2}. Agents can differentiate between a wall, a Blue agent and a Red agent. This is encoded through randomly generated numbers between the following different ranges for the four types of obstacles: 0.0 – 0.2 for no obstacle; 0.2 – 0.4 for a wall; 0.4 – 0.7 for Blue agents and 0.7 – 1.0 for Red agents. The motor neurons encode: (a) the direction of movement (3 neurons); and, (b) the amount of pheromone deposited in the current cell. The 3 output neurons were used to encode a binary direction of movement 0 to 7 where 000 represents moving north, 001 represents moving north east and so on in a clockwise fashion. This ensures the distance between each direction is the same (i.e. the amount of change needed to the weights of the ANN to switch from direction 1 and 2 is the same as for 7 and 0).

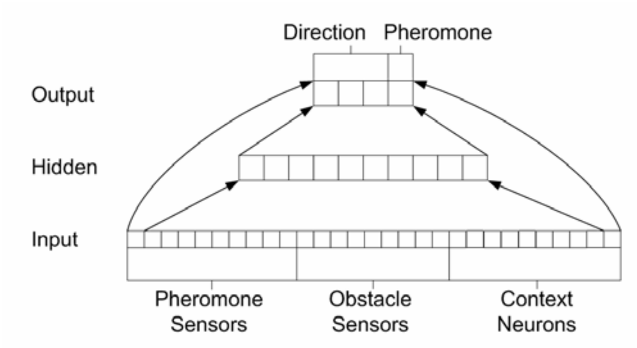

Figure 2. ANN topology.

The connection weights of the ANN controlling the Blue agents are evolved. Group fitness is evaluated at completion of the time steps by counting how many cells have been visited by Blue (and eventually Red) agents. Only cells visited for the first time are considered.

Each genotype encodes the connection weights of a corresponding ANN. A simple direct encoding was used in which each weight is represented with a real number. Each network is duplicated five times, embodied into five corresponding Blue agents, and tested in interaction with the Red agents.

The behavior of the Red agents is controlled by a fixed strategy involving complete information of the environment. The environment is projected onto a 3x3 resolution map where each of the nine segments store the number of Blue and Red agents. An attraction-repulsion system is used to decide on the movements of Red agents. The system uses a weighted sum algorithm where the segment with the maximum number of Blues and minimum number of Reds had the highest probability of being visited next. A random position in the calculated segment is selected for the agents heading {see Figure 3}. This calculation was repeated once the agent had reached its destination position. This strategy is designed so the Red agents traverse empty areas (of Red agents), while at the same time attempting to block the Blue agents from their exploration.

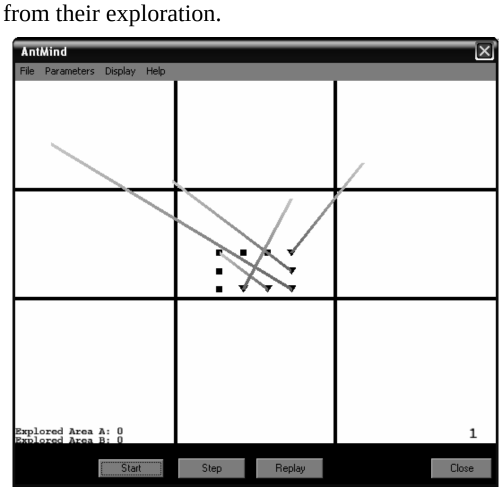

Figure 3. 3x3 resolution map and the selected segments represented by the directional lines from the Red agents.

We chose to use the self-adaptive Pareto artificial neural network (SPANN) algorithm proposed in [1][17]. The evolutionary algorithm involves the selection, mutation and reproduction of a population of ANNs over a number of generations. Each ANN has two fitness values associated with it, where the single objective case sets the second objective to a random number to promote diversity as suggested by [2]. The Pareto approach involves the preservation (elitism) of a Pareto set of ANNs. The ANNs are ranked according to layers, where the highest scoring rank is the Pareto front. This is made up of a minimum of three and maximum of ten ANNs. Therefore, two networks which have a high fitness in differing objectives are ranked equally. Reproduction involves three parents and the algorithm incorporates a self-adaptive crossover and mutation rate.

The Pareto approach has been shown to be more advanced than traditional evolutionary methods in relation to maintaining diversity in the population and efficiency in finding a good quality solution, though not necessarily the best, in a short time [17].

3.1 Experiments

In order to perceive an accurate view of the effect of communication on group cooperation, five experiments were necessary:

- T1. single objective with no communication;
- T2. single objective with random indirect communication;
- T3. single objective with indirect communication;
- T4. multi-objective with correlated objectives; and,
- T5. multi-objective with negatively correlated objectives.

The single objective used in the first three experiments involves the maximization of the Blue exploration area. This encourages Blue agents to explore but does not encourage strategies to block Red movement. No communication is included as a bench mark for the other experiments to reliably compare the effect of communication on the task. This simply involves the sensory neurons associated with the pheromones being set to 0. The second experiment introduces random pheromone trails. The justification for this was to 'shake' the weights of the ANN during evolution to produce a random communication strategy and to ensure that the pheromone trails used in the third experiment are not merely used as random sensors by the network. The third experiment institutes communication by enabling pheromone trails. The last two tasks use multiple objectives. The first objective in both tasks is the same as the single objective experiments, to maximize the Blue exploration area. The second objective in T4 is to minimize the exploration of Red. The last experiment, with negatively correlated objectives, attempts to maximize Blue and maximize Red. Due to the properties of the Pareto approach this will ensure the best case for the first objective (maximize Blue exploration) will be on one end of the Pareto front, while the best case for the second objective (maximize Red exploration) will be at the other end. This approach ensures greater diversity in the population because of the conflicting objectives.

The five experiments were repeated 30 times with different random seeds. The same parameters were used for all the experiments and can be seen in Table 1.

Table 1. Parameters for experiments

Parameter	Task
Task	T1…T5
Number of Blue Agents	5
Number of Red Agents	5
World Dimension	20x20
Pheromone Evaporation Rate	0.03
Generations	300
Number of ANNs	100
Hidden Neurons	10
Time Steps	100

4. RESULTS

This section is divided into three sub-sections. In section 4.1, the evolutionary results are discussed through observing the best network and average population over time for each of the aforementioned experiments. Section 4.2 looks at the workings of an evolved ANN through the correlation between the input and hidden neurons. The agent's behavior from each experiment is analyzed in section 4.3.

4.1 The Effect of Objective Functions

Five experiments were performed to explore the effect of varying objective functions. For each task the ANN that performed best in Blue exploration is displayed over time with the corresponding Red exploration area. The average of the best ANN over the 30 experimental runs was calculated and represented. Note the best Red exploration is not preserved over the generations in the single objective experiments. This is due to the single objective selection pressure only driving Blue exploration. As previously mentioned the second fitness value is set with random values, thereby the better Red strategies are not necessarily preserved.

The first experiment provided no means of communication for the Blue agents. This was included to observe the effect communication has on the emerged behavior in the system. The highest scoring ANN produced a fitness of 241 in Blue exploration; however, the Red exploration remains at a constant rate {see Figure 4}. Results confirmed that a reasonable strategy that maximized exploration of Blue can still emerge. However, as we will show in section 4.3, agents behaved independently, with no significant interaction between or within teams. As Red uses a fixed strategy, and as Blue can only detect Red in the immediate neighborhood, a no communication strategy is suited for this task. Only when Blue agents detect a Red in the immediate neighborhood, a Blue agent may have a better chance to behave in a non-random fashion. We use these results for benchmarking the other experiments.

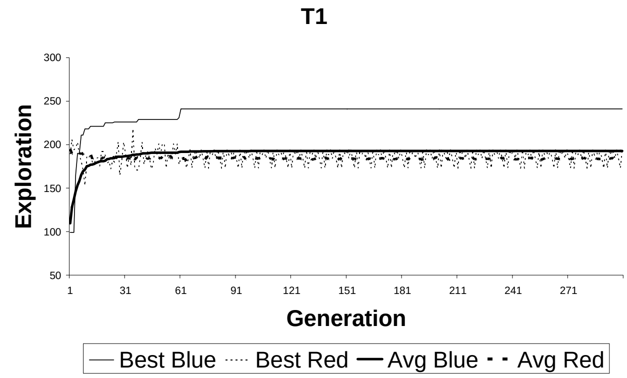

Figure 4. The best and average ANN(s) in T1.

T2 introduces the random communication system through the pheromone sensors being set to random values. Results show a sharp decrease in Blue exploration. Our initial hypothesis that by feeding these sensors with random values may 'shake' the network and produce a better result than T1 was incorrect. The random sensors introduced too much noise to the agents and the network was unable to filter this large amount of noise arriving simultaneously to the 9 sensors. The top scoring network produced a fitness of 208 in Blue exploration and 198 in Red {see Figure 5}. Of particular interest, the average of Blue and Red have a difference of one point in fitness {see Table 2}. This shows that the strategy for Red and Blue are evenly matched. We can also see the variance in the Blue population of ANNs is small by looking at the difference between the best ANN and the average of ANNs. We expect that the larger the difference the more variance in the set of solutions. This experiment is also a benchmark for the other three experiments because it gives us confidence that the network in the remaining experiments is using the pheromone in a sensible way and not just in a random fashion.

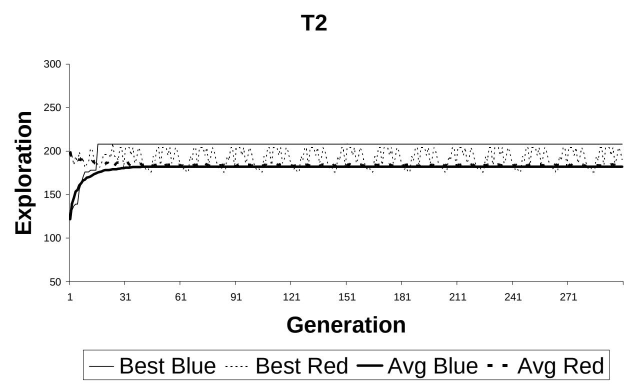

Figure 5. The best and average ANN(s) in T2.

T3 is the first experiment where we allow agents to sense the pheromone level they deposit; therefore, communication is allowed. The results show the best scoring fitness in Blue exploration at 269. This is the best fitness produced for Blue exploration throughout all the experiments. Red exploration is similar to both T1 and T2 {see Table 2} and follows a trend of a slight decrease from generation one to ten, then oscillates regularly for the remaining generations. The increase in Blue exploration has had next to no effect on Red exploration. The average of Red fitness reaches 185, compared to 183 and 184 of

T1 and T2 respectively {see Table 2}. A significant increase in variance of the Blue population is apparent from the distance between the best and average Blue compared with the distances in T1 and T2.

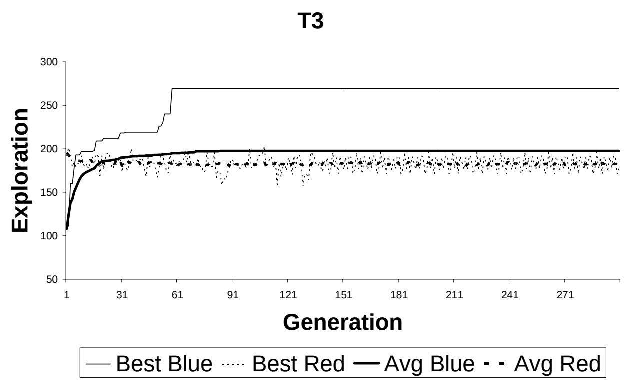

Figure 6. The best and average ANN(s) in T3.

The first of the multi-objective experiments attempts to place some selection pressures on minimizing Red exploration. Figure 7 shows the best ANN reaching 245 in Blue fitness. The average of Blue is lower than T1 and T3; this is due to the introduced selection pressure for minimizing Red exploration. The interaction between Red and Blue was not encouraged in any of the previous three experiments. Once such an interaction is favored, the Blue team has to sacrifice some of its own fitness to block the path of Red. The reduction of 24 exploration units in Blue (when comparing T3 and T4) is balanced with a reduction of 34 units in the exploration of Red. The results show the first significant effect on the exploration of Red. Figure 7 shows the best fitness of Red oscillating sharply in the first 60 generations, and then tapers into a consistent low fitness of 143. The best ANN has evolved a reasonable strategy which increases Blue exploration while simultaneously decreases Red exploration.

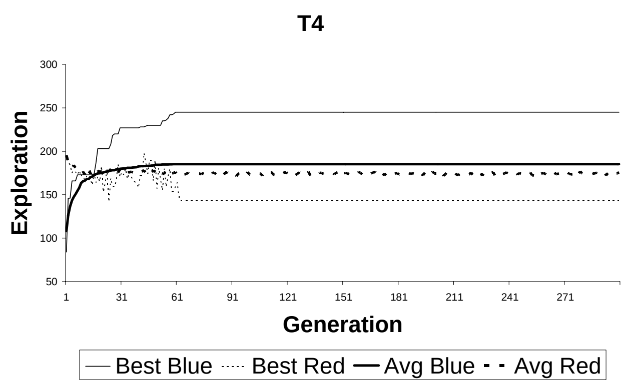

Figure 7. The best and average ANN(s) in T4.

The second multi-objective experiment involved negatively correlated objectives of maximizing Blue and Red exploration. It is evident in Figure 8 that this in fact occurred. Blue reached a high of 244 in fitness with a corresponding top scoring Red of 209. The average of the Red networks illustrates a slight increase of fitness over the first 100 generations. We hypothesize that given the Red strategy has not changed, Blue is actually evolving strategies that are avoiding Red to allow Red to explore more areas.

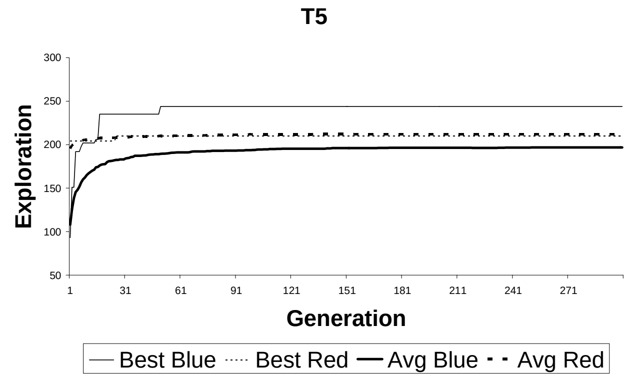

Figure 8. The best and average ANN(s) in T5.

The following table summarizes the empirical data displayed in the graphs.

Table 2. Summary of results for T1-T5

	T1	T2	T3	T4	T5
Best Blue	241	208	269	245	244
Best Red	186	192	182	148	209
Average Blue	192	182	197	185	196
Average Red	184	183	185	175	212

4.2 Activation States of ANN

In order to form a basic understanding of the neurons' behaviors we investigated the correlation of the activation states of the 9 pheromone neuron sensors, 10 hidden neurons and the one pheromone output neuron. T1 did not use the pheromone trails so no comparison was made. In T2, the correlation coefficients fall between -0.5 and 0.5 for many of the hidden neurons. This correlation explains the drop in the Blue population's performance because the random sensory input neurons were generating a high level of noise for the network.

The correlation coefficients between the sensory input neurons and the hidden neurons were very high for experiments T3-T5. T3 was found to have high correlation coefficients, ranging -1 to 1, with the higher correlations only periodically occurring throughout the experiments. T4 and T5 had similar correlation between the input, hidden and output pheromone neurons.

Figure 9 represents a 3-dimensional landscape of the correlation coefficients for the best network achieved in each of the 30 experiments for T5. The x-axis represents the hidden neurons (30 experiments x 10 hidden neurons in each experiment = 300 hidden neurons). The y-axis represents the 9 inputs (1-9) and the 1 output neuron (10). It is evident from Figure 9 that there exists a high correlation coefficient between the hidden neurons' activations and the pheromone values (both as input and output). This demonstrates that T3-T5, but more so in T4 and T5, use the pheromone inputs to evolve an exploration strategy.

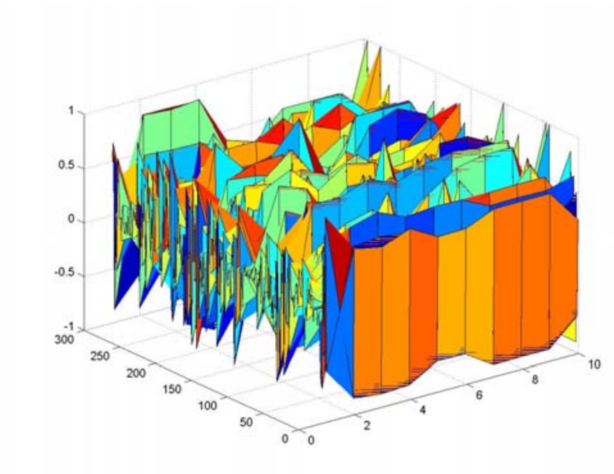

Figure 9. Activation correlation for input and hidden neurons for T5.

A Blue agent was extracted from T4 to study its behavior in an experiment. Figure 10 shows the directions the agent moved over one hundred time steps. The direction is represented by the number corresponding to the binary encoding where north is 0, north east is 1, east is 2 and so forth in a clockwise direction. The extracted agent has a preference for moving in west, north westerly direction for the majority of the time. Periodically the agent moves east, south east and then south in different formations. Interestingly the agent never moves north east, and only moves north and east once. This agent may be representative of one of the Blue agents with the specialization behavior of attracting Red agents which will be seen in following sub section.

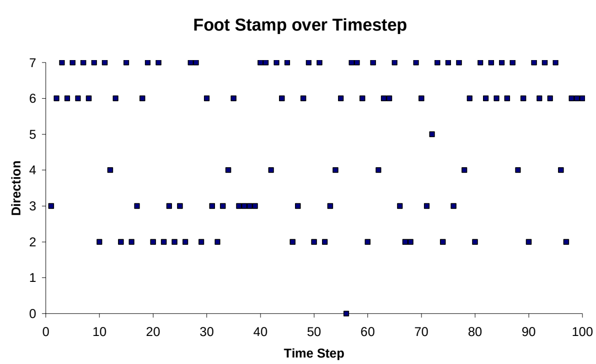

Figure 10. Foot stamp for an agent from T4.

4.3 Behavior of Agents

This section presents screenshots of the agent's exploration area for the best ANNs evolved for each of the objectives. Blue agents are represented by squares and Red by triangles. The cells explored by the Blue agents are shaded with dark grey, the Red with light grey.

Observation of the Blue agents' behavior showed some interesting results. The agents appear to move synchronously with each other in similar pattern formations. The pattern consists of medium sized, approximately eight-step, circular movements with tails. The results of this pattern can be seen in the checkerboard trail in white and dark grey space in Figure 11. These are the spaces avoided from the circular movements. The Blue agents are able to explore a wider space through the tails directing them on a new course. Although there are no pheromone trails to guide them, the Blue agents explore the environment using their circular pattern producing an effective solution.

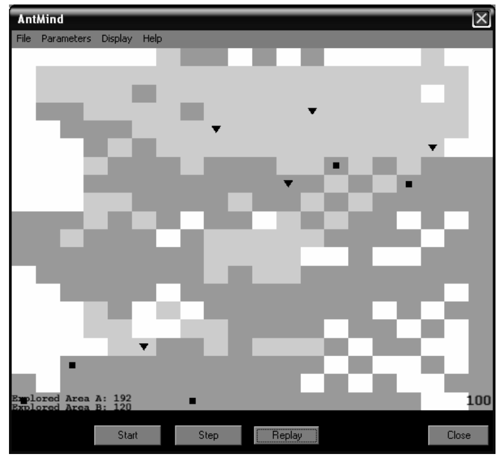

Figure 11. Screenshot from T1.

The empirical data from T2 showed the Blue and Red strategies to be equal. Observation of a number of solutions showed this to be true. Figure 12 represents a common theme occurring from the set of solutions from T2. From a glance it is evident both sides are closely matched. No pattern is discernible of the Blue agents; they genuinely appear to be moving in random directions.

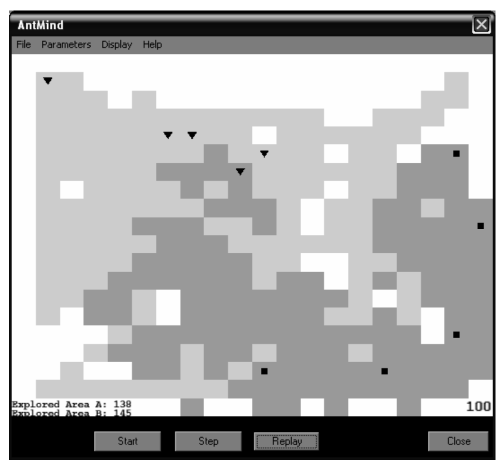

Figure 12. Screenshot from T2.

The single objective experiments produced the highest exploration area for Blue agents and the highest average. A common strategy that emerged from the set of solutions was distinct from the other experiments. This involved a systematic strategy of long straight movements. A typical agent would travel in a straight line until encountering an obstacle; this could be a wall or another agent. Once an obstacle was encountered the agent would move up or down for a few steps and then begin traveling in the opposite direction, although sometimes they would retrace their own steps. Every so often the agent would abandon the straight path behavior for a one-step diagonal path system which consisted of a typical pattern such as: one step north, one south east, one north, one south east and so forth. As seen in Figure 13, lines of white unexplored space are present near the dark grey of Blue exploration (see middle eastern and western portion of the environment). This systematic behavior of straight movement produces good coverage of the environment.

In the previous three experiments, it seems that the two teams became more specialized on different parts of the environment, where the Blue mostly explore the bottom half while the Red explore the top half.

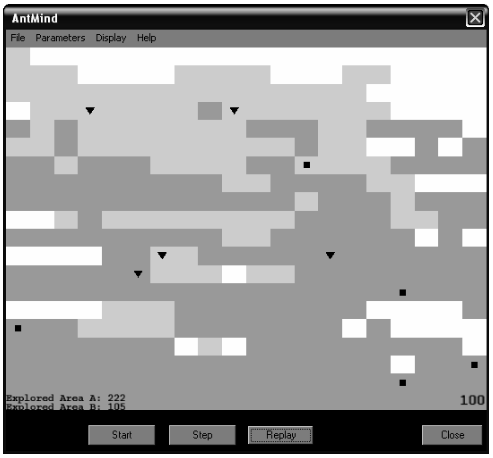

Figure 13. Screenshot from T3.

The first of the multi-objective tasks (T4) evolved solutions containing high Blue fitness and low Red fitness. Observation of the resultant behavior provided insight into a strategy the Blue agents used to minimize Red exploration. In the run represented in Figure 14, the Blue agents begin by splitting into two groups. The first group began by moving in a north westerly direction in jagged steps (e.g. north, west, north, north, west etc.). After a small amount of exploration in the north west corner of the map the agents appear to become stuck. The three agents' position in the north western corner of the environment in Figure 14, was held for the majority of the time steps.

To gain an insight into this seemingly unintelligent behavior we must look at the actions of the second group. One of the remaining two Blue agents began moving in a southerly direction and the other in an easterly direction. Both agents began to thoroughly explore their areas, occasionally bumping into each other and then moving away again. The area they are exploring has not been traversed by the Red agents as they are all hanging around in the north western area of the environment The other Blue 'stuck' agents are acting as a decoy by attracting the Red agents through the weighted sum strategy. The Red agents head towards the north western segment as it has the majority of Blue agents in it. Once they reach their destination, there are too many Red agents in the segment so the new calculated position heads them away from the corner. However, once they reach this new destination they head back to where all the Blues are hanging around as the Reds have all left. This means the Red agents keep traversing the same paths, back and forth, leaving a huge area for the two Blue agents to explore. We see signs of specialization within the group, even though all the agents are controlled by a single ANN.

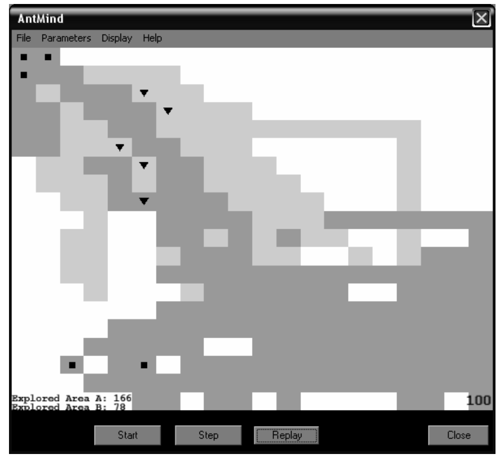

Figure 14. Screenshot from T4.

The second multi-objective experiment produced an exploration strategy which maximized Blue and Red exploration. The Blue agents' movement generally focused on the southern and eastern regions of the environment. The behavioral pattern of the Blue agents consisted of a diagonal movement strategy. A typical agent may follow a sequence similar to: north, south east, east, north, south east etc. It can be discerned from Figure 15 that the Blue agents generally avoid the areas the Red explore thus creating a cooperative environment. Even though cooperative behavior has evolved between the two teams, the Blue still has coordinated behavior within its team. The agents move out in different directions, singularly exploring the area using the same diagonal movement strategy. This eventuated in all the agents heading towards the north eastern most corner of the map. They all reached this point at the same time of time step 95.

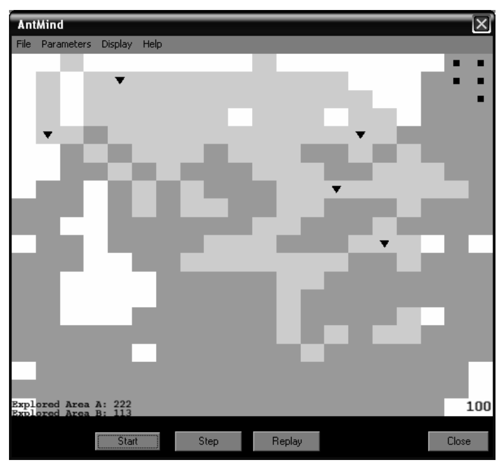

Figure 15. Screenshot from T5.

It is evident from observations of the agents' behaviors that the differing fitness functions have produced significantly different strategies. T1 has evolved a reasonable solution to the problem; however, the agents act independently from each other. Observation of T2 supports the empirical data that the random use of pheromone trails creates too much noise for the ANN and thereby creates random behavior of the agents. T3 showed the best results in Blue exploration, as the fitness function aimed solely to increase this. The first three tasks produced diverse strategies for the exploration task; however, no cooperation between the agents is evident. T4 saw an interesting development in behavior by producing a blocking mechanism to stop the Red agents from exploring large areas. The Blue agents sacrificed their own exploration to do this. The last experiment, T5, demonstrated cooperation between the two teams. The Blue team explored the area quite well while avoiding the Red agent exploration areas.

5. CONCLUSION

The results of these experiments demonstrate how a group of evolving agents, that are able to modulate the amount of pheromone released in the environment and detect pheromone intensity, are able to exploit indirect communication to coordinate and cooperate. The effect of communication is apparent from the empirical comparison of the no communication experiment, the random communication experiment and the three communication experiments. Moreover, the results obtained show how multi-objective optimization methods can be successfully applied to a collection of interactive agents that cooperate in order to achieve conflicting collective goals. Observation of the agents' behaviors supported the empirical evidence in hypothesizing that the different fitness functions have a dramatic effect on the emergence of exploration strategies. There was a clear relationship between the fitness functions and the behaviors that emerged in evolution. For example, cooperation between the agents was seen in T5 through avoidance of the Red team, while the competitive nature of T4 forced Blue agents to specialize and maneuver with the Red team.

The experiments presented in this paper would extend well into the field of robotics. In order to do this it may be worthwhile changing the domain from discrete to continuous, in which different solutions may emerge. Other experiments could be performed with differing environmental conditions, for instance the effect of number of agents, world size or starting positions.

6. ACKNOWLEDGEMENTS

This research is supported by a CSIRO grant and the ARC Centre on Complex Systems grant number CEO0348249.

7. REFERENCES

[1] Abbass, H.A. Speeding up back-propagation using multi-objective evolutionary algorithms. Neural Computation, 15(11):2705-2726, 2003.

[2] Abbass, H.A. and Deb, K. Searching under multi-evolutionary pressures. In Zitzler et al., editor, Proceedings of the Fourth Conference on Evolutionary Multi-Criterion Optimization, Spain, 2003.

[3] Cangelosi, A. Evolution of Communication and Language Using Signals, Symbols, and Words. IEEE Transactions in Evolution Computation, 2001. **5**(2): p. 93-101.

[4] Cangelosi, A. and Parisi, D. The emergence of a 'language' in an evolving population of neural networks. Connection Science, 1998. 10: 83-97.

[5] Elman, J.L. Finding structure in time. Cognitive Science, 14(2):179-211, 1990.

[6] Gerkey, B.P and Mataric, M.J. Principled Communication for Dynamic Multi-Robot Task Allocation. Experimental Robotics VII, 2001: p. 353-362.

[7] Jones, C. and Mataric, M.J. Automatic Synthesis of Communication-Based Coordinated Multi-Robot Systems. in IEEE/RSJ International Conference on Intelligent Robots and Systems. 2004.

[8] Kirby, S. Natural Language from Artificial Life. Artificial Life, 2002. **8**(2): p. 185-215.

[9] MacLennan, B.J. The Emergence of Communication through Synthetic Evolution, in Advances in the Evolutionary Synthesis of Intelligent Agents, M. Patel, V. Honavar, and K. Balakrishnam, Editors. 2001, MIT Press.

[10] Marocco, D. and Nolfi, S. Emergence of Communication in Embodied Agents Evolved for the Ability to Solve a Collective Navigation Problem, (in press).

[11] Mataric, M.J. Using Communication to Reduce Locality in Distributed Multi-Agent Learning. in Proceedings AAAI-97. 1997. Providence, Rhode Island.

[12] McKennoch, S., McNew, J.M. and Bushnell, L.G. A Biologically-Inspired Platform for the Evolution of Communication in Multi-Agent Systems. in Proceedings of the 2003 IEEE International Symposium on Intelligent Control. 2003. Houston, Texas.

[13] Noble, J. Cooperation, Conflict and the Evolution of Communication. Adaptive Behavior, 1999. **7**(3/4): p. 349-370.

[14] Nolfi, S. and Marocco, D. Emergence of Communication in Embodied Agents: Co-Adapting Communicative and Non-Communicative Behaviours. In A. Cangelosi et al. (eds.), Modeling language, cognition and action: Proceedings of the 9th Neural Computation and Psychology Workshop. Singapore: World Scientific (in press).

[15] Smith, K. Natural selection and cultural selection in the evolution of communication. Adaptive Behavior, 2002. 10(1): p. 25-44.

[16] Steels, L. and Kaplan, F. AIBO's first words: The social learning of language and meaning. Evolution of Communication, 2001. 4: p. 3-32.

[17] Teo, J. and Abbass, H.A. Automatic Generation of Controllers for Embodied Legged Organisms: A Pareto Evolutionary Multi-Objective Approach. Evolutionary Computation, 2004. 12(3): p. 355-394.