Genetic Evolution of Communication in Distributed Classifier Systems

5E−4

Adel T. Rahmani, Takuya Okamoto and Norihiko Ono
University of Tokushima

1 Introduction

Machine learning paradigms requiring a supervisor are not appropriate when faced with learning tasks for which no such supervision may be provided. For example, autonomous agents, such as robots cannot be given explicit information about the adaptive behavior corresponding to different situations they may encounter. Required knowledge must be inferred through interaction with the environment. Furthermore, since this knowledge is not available all at once, they must learn it incrementally. These two characteristics are the two most important features of a class of learning methods called incremental reinforcement learning.

Classifier systems, proposed by John Holland, are promising learning paradigm to fulfill these needs. They have been used to study the behavior of adaptive organisms and autonomous agents.

2 The Basic System Description

In standard classifier systems, there is only a single set of rules (classifiers) which interacts with the environment. Also in the parallel or distributed versions of classifier systems, generally a number of similar classifier systems are considered. Here we deal with a system having multiple copies of two different sets of classifier systems.

Basically, there are two learning mechanisms in classifier systems. The first is accomplished by the credit assignment algorithms by adjusting the associated strengths of a fixed set of rules. In Holland's model of classifier systems (Michigan approach) [1], Bucket Brigade Algorithm is responsible for this task. The second mechanism is creation of new rules by the rule discovery algorithms. Since our main concern has been to observe the evolution of a kind of communication protocol among a population of artificial animals, we have not used this mechanism. Instead all the CS's are initialized with a fixed set of rules and only their strengths are adjusted under Bucket Brigade Algorithm. This experiment also shows how effectively BBA can assign strength to classifiers under a delayed reward scheme and even form short classifier chains without classifiers being coupled to each other explicitly.

The basic execution cycle also differs from the standard one. The action of winner classifier from the first classifier system, after being placed into message list, instead of effecting the environment directly, is moved into the message list of a second classifier system. Then the action of winner classifier in this second classifier system will indicate how the environment will change.

3 The Problem Domain

The primary task domain we have used is mate finding in a population of artificial animals. This problem has been also studied by G. M. Werner and M. G. Dyer [3] using a recurrent neural network approach. We have tried to adhere to most of their assumptions and constraints so that the results obtained by the two approaches may be compared. In their implementation, they have used genetic operators to combine the animals' genome which encode the weights and biases of their neural networks, and no learning takes place in those networks. This almost corresponds to the strength adjustment of classifiers by BBA.

The environment is defined as a 50 by 50 grid where edges are continued toroidally. Each cell is either blank or occupied by an animal. We randomly place 50 females and 50 males into the environment, so the population density is 4%.

Our female animals have the capability to look at their surroundings and when detecting a male in a nearby position (within her visual field) then producing a sound which is represented as a signal pattern. They have a repertoire of eight such signals. When they do not detect any male, they just keep silent or the emitted signals do not have any significance since no male can hear them.

On the other hand, male animals can only hear the sound which has been emitted by the nearest female in their "auditory field". Upon receiving a signal they interpret it as one of four possible actions they can take (MOVE_FORWARD, TURN_RIGHT, TURN_LEFT, STAND_STILL). If no signal is heard then they just take an action as dictated by their rule set.

As there are two types of animal in our simulated environment, we consider two sets of classifier systems to represent each class. Each animal is associated with a classifier system of its own class which models the animal's behavior. For both systems we use classifiers with just one condition and one action.

At every iteration, at first all female animals scan their nearby positions in a specified order, and accordingly generate their signal patterns. Then male animals detecting these signals make a move and if a mating happens between a male and a female (i.e. the male gets to the female's position) then they both receive a payoff from environment. After the mating happens , they are placed into new random positions.

This problem falls within the category of "animat problem" suggested by Wilson [4] and may be considered as: incremental learning of multiple disjunctive concepts (though limited) under delayed payoff.

4 Implementation

4.1 Representation

One of the most important steps in implementing a classifier system is deciding how to represent or encode the environmental states. One characteristic of a good representation is to reduce the number of rules which has to be learned by the system(solution set size) by allowing some generalization to be possible.

To make this argument clear, let us consider an example from our problem domain. There are 24 positions around a female which constitute her visual field. We have to consider an ordering for these positions to be scanned by the females. If we randomly or sequentially number those positions, then for the 24 positions and 4 directions, the classifier system has to learn about 96 rules to deal with all possible situations. Instead if we adopt the ordering shown in Figure 1 with an appropriate encoding of the four directions, we can reduce the number of rules which must be learned to 24 (reducing the solution set size by 75%!).

23	22	7	10	11
21	20	3	8	9
6	2	F	0	4
17	16	1	12	13
19	18	5	14	15

Figure 1: An appropriate ordering for female's visual field.

4.2 Bucket Brigade Algorithm

We use a standard Bucket Brigade Algorithm similar to the one described in [2]. We select the winner classifier by a probabilistic selection algorithm based on the effective bids of the competing classifiers. The effective bid is calculated as a power of each classifier's bid. We used a power of 2 to bias the selection toward classifiers with higher strength. We impose a biding tax of 5%. No head tax is necessary since we use a fixed set of rules.

We reduce the taxes from classifiers' strengths only if their strengths are more than the initial strength. This means that they have to pay these taxes from what they have gained, something like an income tax. This puts more pressure on the winning classifiers and if they cannot compensate for this tax then soon their strengths are reduced to the level of other competing classifiers. This simple strategy turned out to be very effective, specially in preventing looping. For example, when the males are trapped in *turning left and right successively or* just moving forward when they are out of visual field of any female.

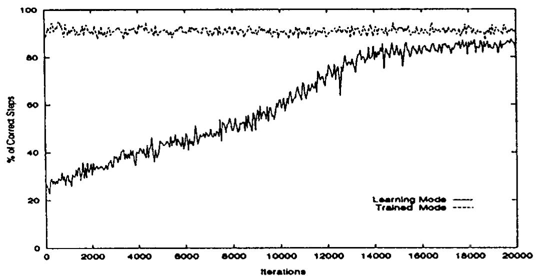

Figure 2: Performance of Learning Classifier System

4.3 Unification Mechanism

Apart from the bucket brigade algorithm, we use a unification operator to encourage coordinated behavior of these animals.

During the initial stages of our experiments, we noticed that if we let these animals behave freely, then almost always all the possible signals would be interpreted as MOVE_FORWARD by the male animals. This was not surprising for a number of reasons: first the rules with a MOVE_FORWARD action are the only possible candidates for receiving external payoff from the environment, second from the females' point of view, there are four situations from which a mating will happen by a male's MOVE_FORWARD. So even in the best case if we assume that a female will use only one signal in all those four cases, then without any coordination among females, we need about 50 (number of females) different signal patterns just to cover those situations.

To overcome this problem, we decided to let them share their experiences locally without making our animals superficially intelligent. We considered a unification-like operator. At some fixed intervals (say every 100 steps) a female is selected randomly and if she can find another female in a nearby position, then the strengths of all classifiers in their CS's are averaged pairwise. The same procedure is also performed for a pair of male animals. We call this a unification operator, because it will allow our animals to unify and coordinate their behavior with each other.

5 Results

To obtain a performance reference for comparing our results with, we performed a number of experiments with both CS's for males and females initialized with a perfect set of rules. That is we assigned one signal for each of the three useful moves (excluding STAND_STILL) and then initialized the females' set of rules with 25 classifiers; 24 classifiers for the 24 possible situations plus one for no male detection. Similarly, we initialized males' set of rules with 7 classifiers; 3 for when a signal is received from a female and 4 classifies for no signal case. So to see how our learning system works, we compare our results with the results we get from the above simulations.

The performance measure we considered was the percentage of number of correct steps taken by males. We count the number of correct steps to mating when a male is within the visual field of a female and report the result every 50 iterations. Figure 2 shows the result. This is compared with the same measure obtained from experiments with perfect set of rules. The reason that even under a perfect set of rules the system does not reach a 100% correct performance is that there are many misleading interactions among these animals. For example, it is possible that more than one male be within the visual field of a female, but the female can only detect the nearest one and send signal appropriate for that one's position and orientation. But the same signal is received by the other male which may not be a good one in its case. This shows how noisy our environment may be.

This result just shows how these animals can learn to perform the designated task, but we also have to show how well they can communicate to aid performing this task. In the following table we summarize the females' responses in all 24 possible situations. We consider 8 different regions in which the females have to respond differently (Figure 3). Also the males' responses are shown in these tables. Actually, these two responses from females and males together can have any significance, otherwise just to illustrate one of them we can not interpret them correctly. Werner and Dyer in their paper [3] only presents the males' responses without showing how the females use those signals.

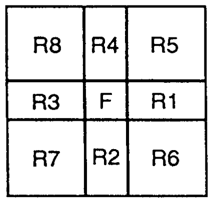

Figure 3: 8 different regions around a female

Table 1 shows the final result. By observing females' responses and males' interpretations, we can see a communication protocol has evolved among them. Only in a few cases there remains some misunderstandings, but the overall result is quite satisfactory!

REGION & DIR.	S#1	S#2	S#3	S#4	S#5	S#6	S#7
R1 N	0	0	0	0	0	50	0
R1 W	3	3	13	0	0	0	0
R1 E	0	4	3	8	7	17	11
R1 S	0	0	0	50	0	0	0
R2 N	0	6	14	0	0	0	0
R2 W	0	0	0	50	0	0	0
R2 E	0	0	0	0	0	50	0
R2 S	0	2	0	8	11	21	6
R3 N	0	0	0	50	0	0	0
R3 W	0	0	0	8	7	29	5
R3 E	5	20	1	0	0	0	0
R3 S	0	0	0	0	0	50	0
R4 N	0	0	4	17	2	18	9
R4 W	0	0	0	0	0	50	0
R4 E	0	0	0	50	0	0	0
R4 S	0	0	2	0	0	0	0
R5 N	0	3	3	1	7	26	7
R5 W	8	8	11	0	0	0	0
R5 E	3	6	5	22	6	3	1
R5 S	0	16	15	4	0	2	0
R6 N	6	11	16	0	0	1	0
R6 W	27	4	5	1	0	0	0
R6 E	1	3	1	1	17	18	6
R6 S	4	2	2	32	2	0	4
R7 N	4	3	40	0	0	0	0
R7 W	2	3	3	33	3	0	2
R7 E	23	0	5	0	0	0	1
R7 S	3	2	0	0	21	11	12
R8 N	1	0	7	34	2	4	0
R8 W	1	2	0	2	15	11	16
R8 E	0	17	27	0	0	0	0
R8 S	15	5	7	0	0	0	0
MOVE_FORWARD	50	50	50	0	0	0	0
TURN_RIGHT	0	0	0	50	0	0	1
TURN_LEFT	0	0	0	0	50	50	49
STAND_STILL	0	0	0	0	0	0	0

Table 1: Females' and Males' Responses At Step = 20000.

References

[1] Holland, J. H. (1986). "Escaping Brittleness: the possibilities of general-purpose learning algorithms applied to parallel rule-based systems", Machine Learning: an Artificial Intelligence Approach, Vol. 2.

[2] Riolo, R. L. (1987). "Bucket brigade performance: I. Long sequences of the classifier" Proceedings of the Second ICGA.

[3] Werner, G. M., Dyer M. G. (1991). "Evolution of Communication in Artificial Organisms", Artificial Life II.

[4] Wilson, S. W. (1987). "Classifier systems and the animat problem", Machine Learning, 2.