

Co-Evolution of Communication in a Competitive Setting

Ido Milstein
Department of Computer Science
Stanford University
ido@stanford.edu

March 9, 2000

Abstract

The following report summarizes the analysis of several experiments in the evolutionary development of communication in ANN based agents. A simple game is created as a model for the environment. The experiments in this model focus on the case where two (or more) clans may use communication to compete against each other. This competition involves both development of original communication by the clans, and exploitation of one clan's communication by the other.

1 Introduction

Communication appears in nature in many forms. Some of these forms take place within a specie, like bee communication by dancing, and ant use of pheromones. Other forms are cross-species communications like the honey-guide bird signals to the honey badger, where the communication helps both species [1]. In all of these cases a specie that can convey better information about food resources, predators in the area, mating locations and so forth will have better chances of survival and reproduction, and specifically the individuals within the specie, that make a better use of this communication will become dominant.

The issue of development and survival of species has been examined to some extent in a paper by David Ackley and Michael Littman [2], and will not be discussed in this paper. The focus here is on the development within a specie. We narrow the scope even further, to the case where two species observes the communication and/or behavior of each other, and fights for the same resources. This would correspond to a case where subspecies live in the same environment.

We would like to make a distinction between 2 types of communication that will be referred to throughout the paper. The first, and more obvious is communication in ways that are not inherently related to the communicated message. For example, when we say in English "food", the vocal sound "food" has no inherent connection to the food itself. This type of communication will be referred to as *voluntary* communication.

The second form of communication is *observational* information. For example when a zebra sees that the rest of its herd suddenly starts running it gets the information that there's a predator nearby. We call this a form of communication, because if the zebra didn't have information about what the herd was doing, it wouldn't flee (unless it spotted the predator).

2 The Model

In order to examine the above issues, we used a simple model. In the model we have two clans (and in some cases three), that contain individuals that are made of artificial binary neural networks (ANNs). Individuals from both clans are evaluated by competing against each other in a sequel of 2 by 2 games. After all individuals are evaluated reproduction occurs within each clan, using a genetic algorithm.

2.1 The Environment

A simple game serves here as a representative of the environment. In the game a group of 2 individuals from one clan plays against a group from the other clan. The players in each group are ordered. One of them will be called *the leader*, and the other will be called *the dummy*. Each player in a group gets to choose 1 of 2 moves, and the ordered pair of moves gives every group 3 possible moves A B C, and a non-valid move F. A wins over B, B wins over C and C wins over A. The scores are distributed as follow: a victory is worth 11 points, a tie 4 points, a loss 3 points, and a fault 0 points (and 11 for the winner). In the case of 2 faults both groups get 0.

This is a non zero-sum game, and it has some nice properties. First, some preliminary communication is required to avoid faults. This may have a drawback of specifically awarding communication, and so the motivation for communication is not altruistic in nature in some of our experiments (communication as altruism is explored in [2] and [3]). On the other hand, this facilitates the initialization of communication. Second, for each move that one group makes, the other group has the possibility of making a winner move, a tie, or a loss (and of coarse a fault). This makes interception of the other group's communication valuable. Third, it is beneficial for a clan not to get stuck in an infinite tie, and so a situation in which both clans stabilize on the same move, is not stable evolutionarily.

Whenever 2 groups play each other, they play a sequel of 64 games. This allows for more complex strategies, and more complex communication (because messages may depend on the previous self and opponent moves). This also makes the fitness measure more stable, as it is averaged over 64 games.

Table 2 gives a fair intuition to the possible fitnesses that groups may get.

2.2 Individuals

Individuals in a clan are represented by binary ANNs, where every neuron's output is a step function of the weighted sum of its inputs. The synapse weights get 8 bit signed values. Every network has 16 neurons, 9 of which are inputs, 2 are outputs and 5 are hidden, in the following order: [const, random, am I leader or dummy, ally said, ally

	Expected Fitness	
	group 1	group 2
(F,F)	5.5	5.5
(F,C)	4.5	7.5
(C,F)	7.5	4.5
(C,C)	6.0	6.0
T	4.0	4.0

Table 1: *Expected fitnesses for different strategies. F = Pick a random value including fault, C = Pick a correct random value (not including fault) T = Tie - Both populations always pick the same value.*

did, opponent leader said, opp. leader did, opp. dummy said, opp.dummy did, I say, I do, hidden 1, hidden 2, hidden 3, hidden 4, hidden 5].

The network is fully recurrent, so every non input neuron gets inputs from all neurons (including itself). This results in $(2 + 5) * 16 = 112$ synapses per network. The recurrency enables processing of information from previous games. The random neuron enables non deterministic strategies.

The ANNs are encoded in the genotype in a direct encoding, meaning that every synapse in the network corresponds to one allele in the chromosome. Note that because the networks are binary, the search space has many plateaus in it, as networks with similar weights may be functionally identical. But on the other hand, it simplifies the “brain” model, and allows for some analysis of it.

2.3 Reproduction

We do not attempt to represent the possibility of extinction of a clan, and so allow for reproduction only within a clan (different inter clan mating schemes are discussed in [2]). Individuals are selected for reproduction according to a non linear function of their rank within their clan with respect to the average score that they got in the games they played. This sometimes slows the evolution process at the beginning, but doesn’t let the GA get stuck when many individuals have the same fitness.

We used the conventional genetic operators of crossover and mutation (with probabilities of 0.5 and 0.001 respectively ¹), and an additional operator of inversion (probability of 0.1), to allow a greater stability of functional groups of genes.

Table 2 summarizes the problem, and the model.

3 Experiments and Results

3.1 Analysis Methods

As with the analysis of most evolutionarily developed objects, exact analysis of the evolved clans in this setting is practically impossible. In order to overcome this prob-

¹See further discussion on section 3.6

Objective:	Explore the development of communication in the presence of competition
Representation scheme:	Structure=Binary ANN Allele range (K) = [-128,127] Number of alleles (L) = 112 Mapping = Each allele maps to a unique synapse weight
Fitness cases:	Scores of 64 games
Fitness:	The average of the 64 scores
Parameters:	$P_C = 0.5, P_\mu = 0.001, P_I = 0.1$ Population size varies between experiments (20 to 600)
Termination criteria:	A fixed number of generations (usually 2000)

Table 2: *Summary of the genetic representation*

lem and extract some information about the developed species, and about the developed communication we used several method. These divide into two general classes of methods - methods that strive to understand the strategies of the developed agents, and methods that extract information about the existence and importance of communication.

The first, and more difficult of the two classes, is made of the following methods: Looking at the neural nets' values throughout a sequel of games. Looking at the weights' matrices, at cross-correlations between neuron values with and without the presence of input noise, and looking at activation correlations between neurons after some neurons are clamped to specific values and the others are selected randomly. Using these methods (and especially the last one) lets us understand some things about the evolved species and communication in a few cases, and understand almost nothing, in most other cases. In some of the cases that we examined, extremely complex strategies evolved. These used inter clan communication combined with different strategies for Leaders and for Dummies, and with interception of the other clan's behavior. Because of their complexity and minor attribution to the main issue, these will not be further discussed in this paper.

The second class of methods was used to extract simple features of the communications. In all the methods in this class we inject noise to some of the inputs of creatures from one group, and look at the resulting difference in fitness. These methods include noising the auditory input neuron within a group (neuron 4), noising both communication inputs between individuals within a group (neurons 4, 5, will be referred to as COM (communication) in following figures), and noising all 6 communication inputs in one group (neurons 4-9, will be referred to as INT (interception)). These measures are not accurate, because when one group changes its behavior due to noise it may cause a change in the other group's behavior, and so on, but are a good approximation to the correct values.

3.2 A Sample Case

Let's first look at one of the interesting and unusual evolutionary cases, that were seen. In this case both populations were of size 20. Figure 1 shows both the fitness of the evolved species, and the developed communication as described in the previous section using the COM and INT methods. The interesting part of this evolution occurs between generations 700-1500. Between generations 700-1000 both clans achieve the very high average fitness of about 7. As the middle figure shows, clan 2 clearly has strong communications at this point. Surprisingly, even though clan 1 is better off when the communication within group 2 is off (COM 2 in the top figure), it does slightly **worse** when clan 2 has no communication at all (INT 2 in the top figure, especially between generations 900-1000). So it's actually beneficial for clan 1 to let clan 2 receive its signals. What actually happens here is an emergence of collaboration between the clans, that allows both of them to achieve a high, but suboptimal fitness. In the games themselves clan 1 wins half the rounds, and clan 2 wins half the rounds, so the resulting fitness for any individual is on average $\frac{1}{2}(Win + Loss) = 7$. This is the highest fitness that both clans can achieve at the same time.

In generations 1000-1200 clan 2 breaks the collaboration with a sharp increase of fitness from communication within the clan. Later on, in generations 1200-1340 an increase in the communication in clan 1 and change of strategy makes both groups have equal fitnesses again. In the short period between generations 1340-1360 clan 2 suddenly understands the communication in clan 1 and takes advantage of it (can be seen by the great **increase** in fitness in clan 1 when its own communication is off (COM 1 in the top figure, and INT 2 in the middle figure)), that is quickly compensated in generations 1360-1410 when clan 1 regains a high fitness by **reducing** its own communication. This communication battle temporarily ends when in generations 1410-1500 both reach equilibrium by reducing their communication to 0. They actually both play A all the time and so the common score at this time is 4 (a tie).

3.3 Communicating in the presence of Noise

In order to understand the dependency of the development of *voluntary* communication on hostile intercepts, we tried two kinds of experiments. In both kinds we introduced noise in the auditory reception channels of all individuals (only in auditory channels that hold information about what the opponents said, neurons 6, 8), in regular intervals. The noise durations that we tried were 0, 25, 50, 75, 90 and 100 noise generations (will be called OFF times) out of every 100 generations. In the first experiments we added noise throughout all the experiment also to the observational channels of all players (neurons 5, 7, 9). In the later experiments we also tried this without the additional noise. In all cases the populations included 400 individuals, trained 2000 epochs.

The first experiments showed that when the amount of noise was small (0, 25) attempts to start communicating die out quickly, even though these attempts sometimes shift the populations to new strategic fixed points. What happens is that as one group develops communication the other group starts intercepting it, and using it for its own good. This can be seen well in Figure 2. In the 25 OFF case, we also saw fitness fluctuations every 100 generations after communication begins, which probably result from the ON/OFF cycle time. A sharp eye may notice that on Figure 3.




Figure 1: *A sample case of the evolution of communication. Top - Fitness difference for clan 1 due to injected noise. Middle - Fitness difference in clan 2 due to injected noise. Bottom - Fitnesses of both groups. For example, in the top figure, COM X lines show the difference in fitness for group 1 as a result of blocking the communication within group X. INT X lines show the difference in fitness when group X is not allowed to receive any communication at all.*

When the amount of noise is larger, we see fluctuations in fitness in decreasing magnitude, and great fluctuations in communication (at a constant magnitude), but interception is hardly ever observed. The fitness of both clans actually converges to 6 which is the average for clans that use random non fault moves. This probably happens because when you cannot know what the other group is planning, the best thing to do is pick a random non fault move.

The fluctuations in communication levels actually increase in frequency with time, maybe because the populations become more homogeneous, and this allows for changes in individuals to have a greater effect because the probability of communicating with another individual that “doesn’t speak your language” decreases, and so good individuals will be good no matter which partners they get. Figure 3 illustrates the difference in development of communication between the high noise cases and the low noise cases.

Another phenomenon that was observed is that in the presence of large OFF times (90, 100) there was a slower start of communication than the middle ones (50,75). It is not clear to us why this happened. This however means that letting the clans listen to each other’s communication for small periods, actually promotes the development of communication within the clans themselves.

In the later experiments we removed the observational noise, and applied noise with the same patterns to the auditory channels of the opponents (just like in the first ex-




Figure 2: *When the auditory channels are open to all, a few generations after group 1 starts to develop communication, group 2 starts to intercept it, and then the communication dies. However, the fitness of group 1 at the end is better than at the beginning.*

periments). Many of the experiments showed common developmental patterns, where at the first few hundreds of generations communication evolved rapidly, then died out for several hundreds of generations, and returned till the end of the run. Together with this evolved communication, there was a significant importance to interception of opponent signals. In almost all these cases all the communication was carried out as an *observational* communication. That is also the reason why there were no significant implications to the added ON/OFF noise to the auditory channels. It should be noted that in this experiment too, the cases where *voluntary* communication was seen, was when the clans could listen to each other for 50 or less generations (of each 100 generations).

3.4 Complicating the Environment

Does a complicated environment hinder the development of *voluntary* communication because of a difficulty to convey meaningful information, or does it make it more likely, as it becomes more important for survival?

We tried to answer this question by making the game a little harder to play. We did this by inserting noise to the “am I a Leader/Dummy” input neuron of each player. One can get an idea of the way this complicates playing the game by thinking of one player playing as a partner of himself. In the absence of a random input neuron both players will always play the same, thus limiting them to only 2 possible moves (in our case one of these is the fault move). So, this practically makes a smart use of the

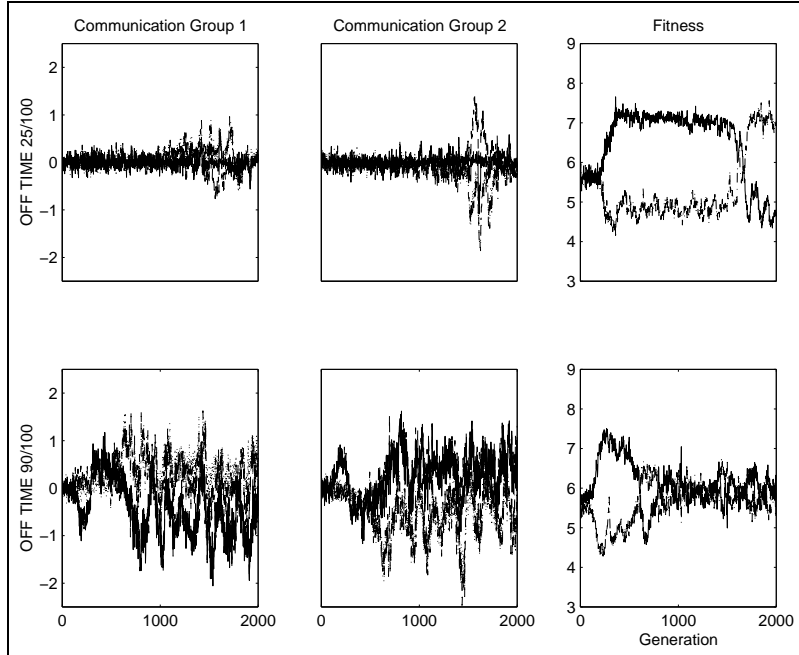


Figure 3: *The format of these figures is the same as that in figure 2. We can see that high OFF times result in communication fluctuations and a converging fitness, while low OFF times result in very brief and small communication periods.*

randomness necessary. The game setting that we used for these experiments included populations of 20/100/400 individuals with noise applied to the dummy neurons, and the auditory input neurons (in the same way as the later of the ON/OFF experiments, with 75 and 100 OFF times. See section 3.3).

The experiments that we carried out had two important results. First, development of communication was significantly hindered, and in many cases was not present at all. Second, when communication did appear it was either in the form of interception of behavioral information, or by using *voluntary* communication. *Behavioral* communication was not seen at all. Recall that the same setting but **with** Leader/Dummy information resulted in well developed *behavioral* communication and interception, and very low amounts of *voluntary* communication.

So the complicated environment resulted in a preference of *voluntary* communication over *behavioral* communication (maybe the players needed the “extra hand” that the audio channel provides while the “do” neuron is busy playing), but a reduced communication as a whole.

3.5 Changing the Population Size

We used the same setting of auditory input noise that was discussed in the last section, to examine the effects of changing the population size. Five population sizes were used - 20, 100, 200, 400, 600 (three of these population sizes (20, 100, 400) were also used in almost all the other experiments that we made, and had similar results). Small

populations have rapid and unstable changes in fitness and communication. Large populations develop slowly, and usually the changes that occur in them are kept for a long time. That is because it is harder for a new dialect of a language to invade a large population. Another interesting and less obvious result is that large populations tend to use more of the different communication and interception possibilities at one time. However, this difference between small and large populations was small, and further experiments are needed to examine and assert it.

Another point that has been mentioned earlier in the paper, is that the same global pattern of development appears when changing OFF times (in this setting), which is a fast development of communication and interception, followed by a “dead time”, and a slow rise later. This remains true even when changing the population size.

3.6 Changing the Genetic Operators Probabilities

In order to select values for the probabilities of the mutation and crossover operators we ran several experiments. These showed that changing the crossover rate in the range 0.3-0.9 doesn’t change much the evolution patterns, and that increasing the level of mutation results in very unstable fitnesses and no development of communication.

3.7 Experiments with Three Clans

In a two clan setting every clan gets to play only against one strategy - the one that is used by its opposing clan. This makes the intercepting the other clan’s strategy and communication easy. In a setting where each clan gets to play against opponents from several clans, we would expect interception to become significantly more difficult, and the development of good strategies also harder. We went on to examine this hypothesis in a series of the normal 2 against 2 player game, where the clans that play are selected randomly from a group of three clans. This was tried with, and without the noise ON/OFF signals that were described before. Almost all experiments ended with very little communication evolving. A nice case was seen in one of the experiments, where one clan developed an early communication, and got a high fitness. A little while later, the other clans began to intercept this dominant communication, and because in that clan the early communication resulted in a quick convergence of the population, it took it a very long while to overcome the interceptions, and it ended up having the lowest average fitness of all three clans. This is shown in Figure 4.

4 Summary

We have explored several settings for the evolution of communication in the presence of competition. These experiments have shown, among other things, that little competition actually speeds up the development of communication within a clan, and results in more complex communication schemes, while a large competition hinders the development of communication, especially that of the *voluntary* kind. We have also seen that complex and difficult environments hinder the development of communication as a whole, and biases the development of communication towards a *voluntary* type. An interesting follow up may be to explore the sophistication of strategies that individuals

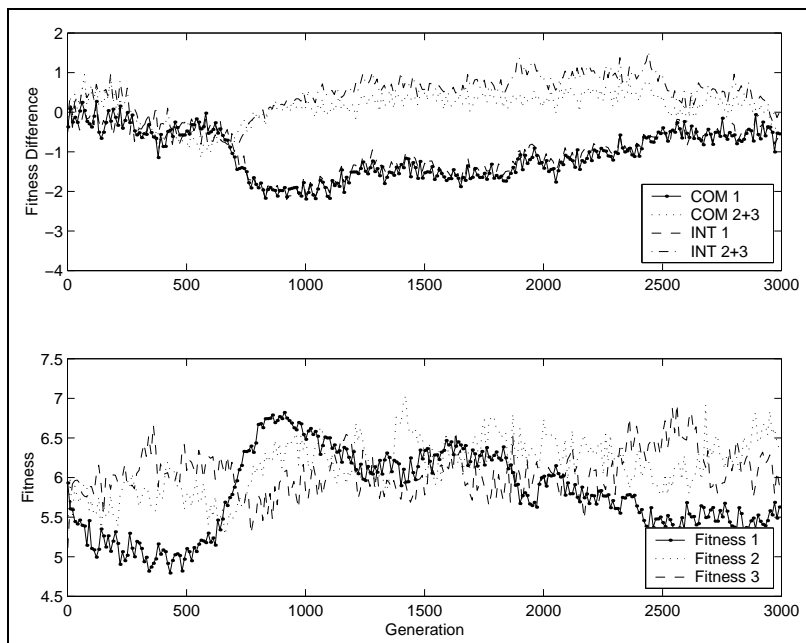


Figure 4: *Early development of communication in clan 1 causes premature convergence, and susceptibility to intercepts. The top figure shows the difference in fitness of clan 1, that results from noising communication inputs. COM 2+3 and INT 2+3 are differences that result from noising opponent inputs, so these are averages between COM 2/3 and between INT 2/3.*

in a clan may use after they develop communication. Could they use the language to “speak to themselves” when in the face of complicated tasks?

References

- [1] Michael Dyer, Toward synthesizing artificial neural networks that exhibit cooperative intelligent behavior: Some open issue in artificial life (*In Artificial Life*, Chris Langton, 1995).
- [2] David Ackley and Michael Littman, Altruism in the evolution of communication (*In Artificial Life V*, 1996).
- [3] J. Batali and P. Kitcher. Evolution of Altruism in Optional and Compulsory Games. *Journal of Theoretical Biology*. Vol. 175, pages 161-171. 1995.
- [4] J. Batali. Computational Simulations of the Emergence of Grammar. in *Approaches to the Evolution of Language: Social and Cognitive Bases*, J. R. Hurford, M. Studdert-Kennedy and C. Knight (eds). Cambridge University Press, pages 405-426, 1998.
- [5] M. Oliphant and J. Batali. Learning and the Emergence of Coordinated Communication. Submitted to *Cognitive Science*.