Cultural Learning and Diversity in a Changing Environment

Dara Curran
Department of Information Technology, National University of Ireland, Galway
Ireland
dara.curran@nuigalway.ie

Colm O'Riordan
Department of Information Technology, National University of Ireland, Galway
Ireland
colmor@it.nuigalway.ie

Abstract- This paper examines the effect of cultural learning on a population of neural networks. We compare the genotypic and phenotypic diversity of populations employing only population learning and of populations using both population and cultural learning in a dynamic environment. We show that cultural learning is capable of achieving higher fitness levels and maintains a higher level of genotypic and phenotypic diversity.

1 Introduction

A number of learning models may be readily observed from nature and have been the focus of much study in artificial intelligence research. Population learning (i.e. learning which occurs at a population level through genetic material) is typically simulated using genetic algorithms. Life-time learning (i.e. learning which takes place during an organisms's life time through reactions with its environment) can be simulated in a variety of ways, typically employing neural networks or reinforcement learning models.

A relatively new field of study in artificial intelligence is synthetic ethology. The field is based on the premise that language and culture are too complex to be readily analysed in nature and that insight can be gained by simulating its emergence in populations of artificial organisms. While many studies have shown that lexical, syntactical and grammatical structures may spontaneously emerge from populations of artificial organisms, few discuss the impact such structures have on the relative fitness of individuals and of the entire population.

A robust multi–agent system should be able to withstand and adapt to environmental changes. This type of behaviour parallels that of the natural world where species capable of adaptation will have more chance of evolutionary success than ones that are rigid and incapable of such plasticity. At its most basic level, adaptation in nature takes the form of population learning. At a higher level, organisms capable of adapting their behaviour to suit a particular environment during their lifetimes will be more likely to survive in the long term.

The focus of this paper is to attempt to understand the effect of cultural learning on a population of artificial organisms subjected to a dynamic environment. This is accomplished by studying its effect on the population's fitness as well as its genotypic and phenotypic diversity. The remainder of this paper is arranged as follows. Section 2 introduces background research, including descriptions of diversity measures and cultural learning techniques that have been employed for this study. Section 3 describes the artificial life simulator employed for the experiments. Section 4 describes the experimental setup. Section 5 outlines the experiment results and Section 5 presents conclusions.

2 Background research

2.1 Cultural Learning

Culture can be succinctly described as a process of information transfer within a population that occurs without the use of genetic material. Culture can take many forms such as language, signals or artifactual materials. Such information exchange occurs during the lifetime of individuals in a population and can greatly enhance the behaviour of such species. Because these exchanges occur during an individual's lifetime, cultural learning can be considered a subset of lifetime learning.

An approach known as synthetic ethology [13, 19] argues that the study of language is too difficult to perform in real world situations and that more meaningful results could be produced by modeling organisms and their environment in an artificial manner. Artificial intelligence systems can create tightly controlled environments where the behaviour of artificial organisms can be readily observed and modified. Using genetic algorithms, the evolutionary approach inspired by Darwinian evolution, and the computing capacity of neural networks, artificial intelligence researchers have been able to achieve very interesting results.

In particular, experiments conducted by Hutchins and Hazlehurst [11] simulate cultural evolution through the use of a hidden layer within an individual neural network in the population. This in effect, simulates the presence of a Language Acquisition Device (LAD), the physiological component of the brain necessary for language development, the existence of which was first suggested by Chomsky [5]. The hidden layer acts as a verbal input/output layer and performs the task of feature extraction used to distinguish different physical inputs. It is responsible for both the perception and production of signals for the agent.

A number of approaches were considered for the implementation of cultural learning including fixed lexicons [21, 4], indexed memory [18], cultural artifacts [10, 3] and signal–situation tables [13]. The approach chosen was the teacher/ pupil scenario [1, 6, 4] where a number of highly fit agents are selected from the population to act as teachers for the next generation of agents, labelled pupils. Pupils learn from teachers by observing the teacher's verbal output and attempting to mimic it using their own verbal apparatus. As a result of these interactions, a lexicon of symbols evolves to describe situations within the population's environment.

2.2 Diversity

Diversity measures typically quantify the differences between individuals in a population. It is commonly accepted that a population that is capable of maintaining diversity will avoid premature convergence and local maxima.

Diversity measures for populations of neural networks have been the focus of considerable research, focusing mainly on genotypic diversity [20, 17, 2]. Many methods exist for the calculation of genotypic diversity, many based on binary representations. For the purposes of this research however, many schemes are unsuitable due to the nature of the marker-based encoding scheme used to represent each neural network.

Our scheme examines each block of the encoding and compares it to blocks of similar length in other encodings. Each encoding block contains a single node and a number of links emanating from that node. It is therefore intuitive to propose that blocks of similar length (having a similar number of emanating links) are suitable for mutual comparison.

There is comparatively little research on phenotypic diversity in evolutionary computation. Typically, phenotypic diversity is measured at the fitness level [7]. However, this measure tends to compress the available diversity information resulting in a coarse grained measure not useful in all situations. The approach adopted in this work is to examine the components of the fitness value of each individual, i.e. an individual's response to each bit-parity stimulus. By comparing the difference between all responses (and not just the aggregate fitness function) a finer grained measure of phenotypic diversity can be obtained.

2.3 Changing Environments

Much research has focused on the tracking of changing environments with regard to multi–agent and artificial life systems [16, 9, 14] focusing on Latent Energy Environments and fitness functions which vary over time. What is generally sought is the ability of a population to adapt to a change within a reasonable length of time and to guide evolution toward a level of plasticity otherwise difficult to attain.

The focus of this paper is to ascertain the effect of a changing environment on population diversity as well as fitness levels. While our approach to changing environments is straight-forward, it has the advantage of clarity: agents evolve learning to distinguish between food and poison bit-patterns (representing the 5-bit parity problem) and at a certain generation, food and poison become reversed (i.e. the bit pattern representing food now represents poison and vice-versa).

The approach is partly based on work performed by Nolfi et al [16] who compared the performance of a robotic agent employing genetic evolution (population learning) and that of agents employing back–propagation (life–time learning) in a changing environment.

3 Simulator

The architecture of the artificial life simulator (Figure 1 can be seen as a hierarchical structure. At the top-level of the simulator is a command interpreter which allows users to define an experiment's variables including the number of networks, the number of generations to run the experiment, mutation and crossover rates and the actual problem set which the population will be attempting to solve. The

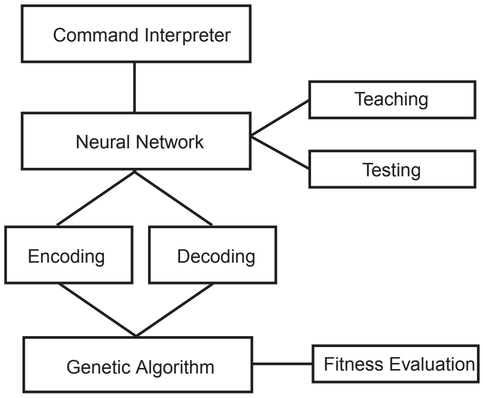

Figure 1: *Simulator Architecture*

neural network layer takes the variables set using the command interpreter and initialises a given number of neural networks. The layer then performs training and testing of the networks according to the parameters of the experiment. These network memory structures are then passed to the encoding layer which transforms them into genetic code structures for use in the genetic algorithm. The encoding mechanism used for this set of experiments is a modified version of marker based encoding.

The genetic algorithm layer uses the genetic codes and the data retrieved from the neural network layer's testing of the networks to perform its genetic operators on the population. A new population is produced in the form of genetic codes. These are passed to the decoding layer which transforms each code into a new neural network structure. These structures are then passed up to the neural network layer for a new experiment iteration. Once the required number of generations has been reached, the experiment finishes.

Two-point crossover is employed and weight mutation is employed which takes the weight value and increases/decreases the value according to a random percentage (200%). This approach was found, empirically, to be more successful and was adopted for this set of experiments.

3.1 Encoding Scheme

An encoding scheme is necessary to map each agent's neural network structure to a genetic code. Many schemes were considered in preparation of these experiments, prioritising flexibility, scalability, difficulty and efficiency. The scheme chosen is based on Marker Based Encoding which allows any number of nodes and interconnecting links for each network giving a large number of possible neural network permutations.

Marker based encoding represents neural network elements (nodes and links) in a binary string. Each element is separated by a marker to allow the decoding mechanism to distinguish between the different types of element and therefore deduce interconnections[12, 15]. A gene code produced using this scheme is treated as a circular entity. Thus, the code parsing mechanism reading the end of the gene code will begin reading the start of the gene code once the end is reached until all available information is correctly retrieved.

3.2 Simulating Cultural Evolution

In order to perform experiments related to cultural evolution, it was necessary to adapt the existing simulator architecture to allow agents to communicate with one another. Agents communicate directly with each other and not through intermediary artifacts. This was implemented using an extended version of the approach adopted by Hutchins and Hazlehurst. The last hidden layer of each agent's neural network functions as a verbal input/output layer (figure 2).

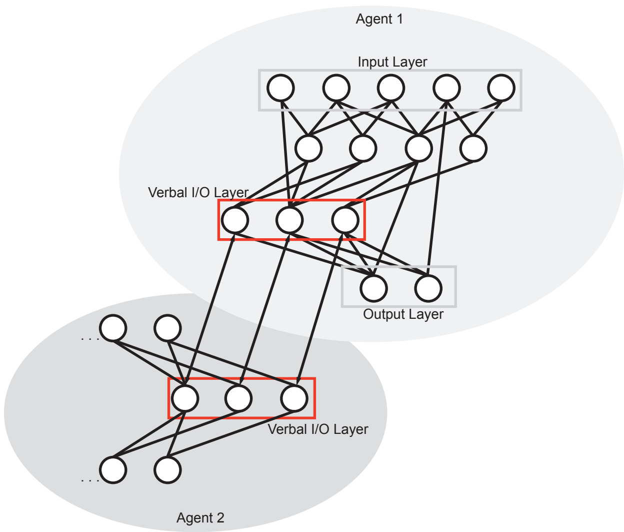

Figure 2: *Agent Communication Architecture*

At end of each generation, a percentage of the population's fittest networks are selected and are allowed to become teachers for the next generation. The teaching process takes place as follows: a teacher is stochastically assigned n pupils from the population where $n = \frac{N_{pop}}{N_{teachers}}$, where N_{pop} is the population size and $N_{teachers}$ is the number of teachers. Each pupil follows the teacher in its environment and observes the teacher's verbal output as it interacts with its environment. Verbal output is produces through the verbal I/O layer of the teacher's neural network as a result of stimulus received at the input layer. A teaching cycle occurs when the pupil attempts to emulate its teacher's verbal output, using its own verbal apparatus, via back-propagation. Once the number of required teaching cycles is completed, the teacher networks die and new teachers are selected from the new generation.

Unlike previous implementations, the number of verbal input/output nodes is not fixed and is allowed to evolve with the population, making the system more adaptable to potential changes in environment. In addition, this method does not make any assumptions as to the number of verbal nodes (and thus the complexity of the emerging lexicon) that is required to effectively communicate.

4 Experimental Setup

The following set of experiments each employs two populations. One population is allowed to evolve through population learning (by genetic algorithm), while the other employs both population and cultural learning.

Cultural learning is implemented based on a scheme developed by Hutchins and Hazlehurst [10] and further explored by Denaro [6] where the last hidden layer (or in Denaro's case, the output layer) of a neural network functions as a verbal input/output layer. At the end of each generation, a percentage of the best individuals in the population is selected to instruct the next. Pupil networks observe teacher networks as they interact with their environment and at each stimuli, teacher networks produce an utterance through their verbal I/O layer. The pupil responds to the utterance with its own, which is then corrected by back-propagation to approximate the teacher's. After the required number of these interactions (teaching cycles) have been completed, the teachers are removed from the population and the pupils continue to interact with their environment.

The problem domain for this set of experiments is the 5-bit parity problem which, while relatively simple in its structure, represents a reasonably complex classification problem[8]. Each network is exposed to bit patterns and must determine whether the pattern represents an odd or even number. Fitness is assigned according to the mean square error of a network.

The change in environment is implemented by reversing the food and poison representations such that the bit pattern representing food will represents poison and vice-versa. Only one environmental change is allowed to occur in each experiment.

Each experiment consists of a population of 50 neural networks evolving for 250 generations with crossover and mutation rates set at 0.6 and 0.02 respectively. The environmental change takes place at generation 125. The population employing cultural learning takes the fittest 10% of each generation as teachers which interact with pupils for five teaching cycles. An additional parameter, cultural mutation, adds noise to each interaction with probability 0.02. The results presented are averaged from 20 independent runs.

5 Experiment Results

The experimental results are divided into two sections. The first examines the relative performance of cultural learning and population learning through analysis of the error values for each population. The second section is concerned with genotypic and phenotypic diversity measures for each population.

5.1 Error

Each population's fitness is determined by its error value, where lower error is rewarded. Figure 3 shows the average error value for both populations over the experiment run. While both populations are equivalent up to generation 100, the population employing cultural learning begins to show improvement prior to the environment change.

At generation 125, both population errors increase dramatically (as could be expected). However, the cultural learning population's error value is significantly lower and its subsequent error reduction for the remainder of the experiment shows marked improvement over population learning alone.

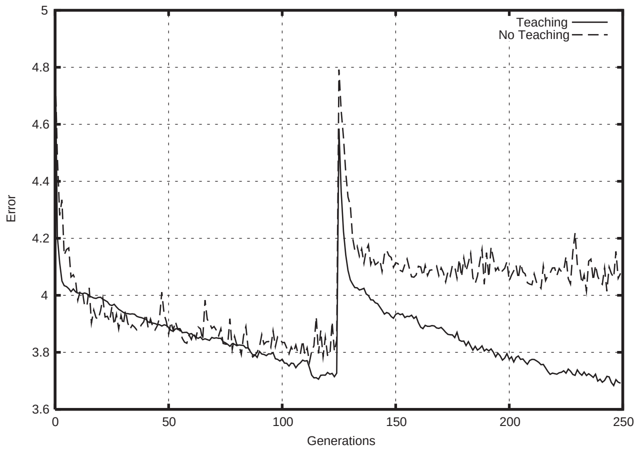

Figure 3: *Error Values*

The improvement delay experienced by the cultural learning population in the initial generations of the experiment can be explained by the fact that time is required to evolve and isolate potentially useful teacher individuals in the population. At the start of the experiment, such individuals are rare, and therefore cultural learning takes some time to achieve its performance boost.

To understand how the error reduction is achieved, both average maximum and minimum error values were obtained for each population. Figure 4 shows the results for maximum error values. The maximum error values for the population employing population learning are very unstable and remain persistently high throughout the experiment.

By contrast, the maximum error values for the cultural learning population are very stable and remain lower for the entire experiment run. This is an intuitive result, given that cultural learning will tend to reduce the population's error value as a whole through the teaching of poor individuals by superior teachers.

Interestingly, the environment change is only graphically apparent in the population employing cultural learning, suggesting that even the poorest individuals are being affected by the change. The worst individuals in the population employing population learning are equally inept in both environments, while even the worst individuals in the cultural learning population have learned some information about their environment and therefore react to the change.

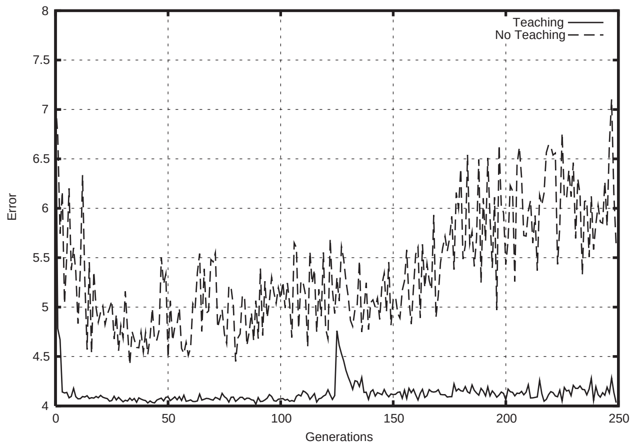

Figure 4: *Maximum Error Values*

Given these results, it could be suggested that cultural learning is achieving its higher average fitness simply through marginal improvements of poor individuals and is not infact guiding evolution by generating novel and superior individuals in the long run. This argument can be counteracted by examining the results for the average minimum error in both populations (Figure 5).

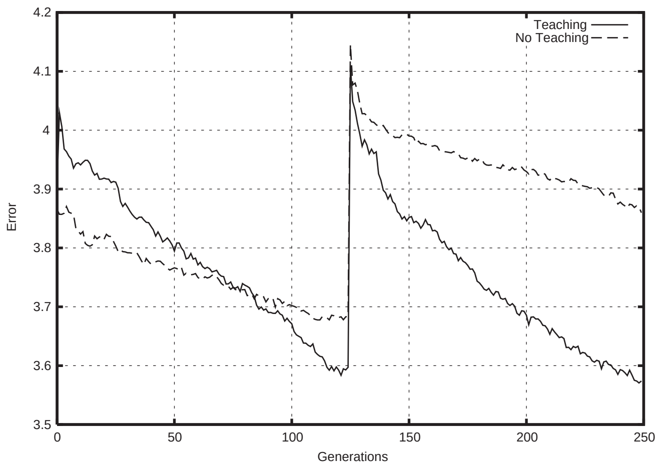

Figure 5: *Minimum Error Values*

It is clear from the minimum error data that cultural learning is producing better individual networks than population learning alone (at least after generation 100). As these results are averages of 20 runs, the possibility of a statistical anomaly to explain this phenomenon can be discounted. Rather, the cultural learning process is guiding evolutionary trends towards highly fit individuals.

5.2 Diversity

The results showing values for the first diversity measure, genotypic diversity, are illustrated in Figure 6. The overall trend for both populations shows an initial high value of genetic diversity and a subsequent drop. This is representative of an initial exploratory phase followed by convergence. However, there is a marked difference between each population. While the population employing population learning alone has a relatively short exploratory period and tends to converge very quickly, the cultural learning population spreads out is exploratory phase and converges around 50 generations later. Both populations have a roughly equivalent level of genetic diversity by the second half of the experiment.

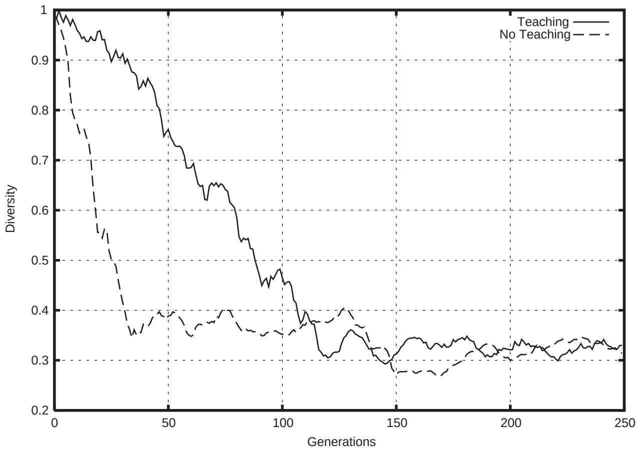

Figure 6: *Genotypic Diversity*

It is clear that cultural learning is not only affecting fitness levels, but that it also has a direct effect on the level of genetic diversity in the population. By maintaining a high level of diversity, the population contains quite disparate individuals, which are brought to some level of homogeneity by the cultural learning process. It can be speculated that by maintaining a high number of effectively redundant genes in each genome, the population is better suited to environment change.

Early convergence makes a population rigid in the face of changes. Cultural learning not only smoothes the change process, but also prepares the population against future environmental changes by including a high number of potentially useful genes in the gene pool. It is interesting to note that the change in environment is not apparent in the genotypic diversity results. Both populations have effectively converged by the time the environment change occurs and without external influence, it is clear that genotypic diversity cannot be refreshed through such changes.

Figure 7 shows results for phenotypic diversity for both populations over the course of the experiment. While both populations maintain a consistently high level of diversity, it is clear that the population employing cultural learning is producing more phenotypically diverse individuals. The environmental change is more significant in this diversity measure as can be seen by the falls in diversity levels in both population occurring at generation 125.

This drop is significantly less pronounced in the population employing cultural learning. Even after an environment change, there exist some able individuals which are capable of instructing the remainder of the populations, resulting in a less dramatic effect. The drop in diversity in both populations is interesting in that it shows that once the environment change occurs, each individual in the population performs

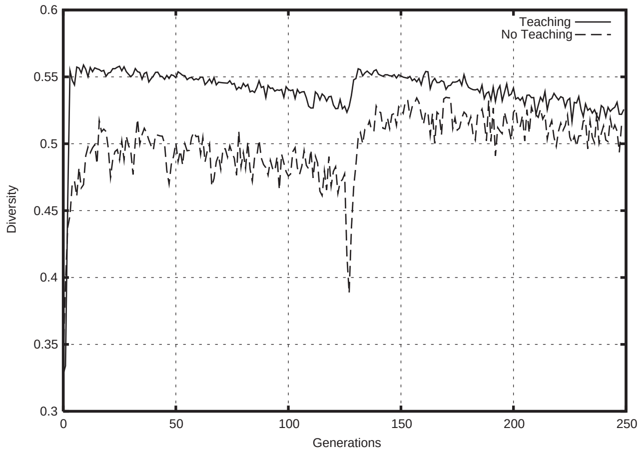

Figure 7: *Phenotypic Diversity*

on average in a very similar manner.

The fact that phenotypic diversity is so high in the population employing cultural learning is initially puzzling. One would expect intuitively that since a small minority of individuals are training each generation, the population would tend to become phenotypically similar and therefore that phenotypic diversity would be low. There are two factors which influence phenotypic diversity in a way which explains these results.

The first is that the teaching process does not result in clone pupils, rather each pupil is a rough approximation of the teacher's outputs. Furthermore, because each an individual's verbal input/output layer is a hidden layer and not its output layer, pupils do not approximate output values. Certainly, the effect on the hidden layer teaching has an effect on the output of the network (as exemplified by cultural learning's superior performance with regard to error), but it may produce rather different outputs from that of the teacher and indeed that of other individuals.

A second factor is genotypic diversity in both populations. The slow phenotypic diversity decline prior to the environmental change and the accelerated decline thereafter exhibited by the cultural learning population somewhat mirrors the corresponding convergence of genotypic diversity. Infact, by the end of the experiment, the phenotypic diversity of both populations is very similar. High genotypic diversity produces individuals which are also highly phenotypically different. Reducing genotypic diversity over time has a small but evident effect on the phenotypic diversity of the cultural learning population.

6 Conclusions

The results presented in this paper suggest some reasons as to why the addition of cultural learning often boosts a population's performance. In particular, the results highlight the ability of cultural learning to adapt to a dramatic change in environment. In addition, measures of genotypic and phenotypic diversity have been presented which shed some light on the evolutionary processes in each population. Cultural learning produces populations which tend to have a longer exploratory period, exhibited by increased genotypic diversity and later convergence. Furthermore, cultural learning produces higher phenotypic diversity which is less affected by changes in the environment. However, one must acknowledge that while cultural learning provides a performance enhancement, it is computationally expensive and may not be suitable for all problem domains. Future work will examine the performance of cultural learning and diversity with more complex problems and more extreme environment changes.

Acknowledgements

This research is funded by the Irish Research Council for Science, Engineering and Technology.

Bibliography

[1] A. Billard and G. Hayes. Learning to communicate through imitation in autonomous robots. In *7th International Conference on Artificial Neural Networks*, pages 763–738, 1997.

[2] G. Brown. *Diversity in Neural Network Ensembles*. PhD thesis, University of Birmingham, 2003.

[3] A. Cangelosi. Evolution of communication using combination of grounded symbols in populations of neural networks. In *Proceedings of IJCNN99 International Joint Conference on Neural Networks (vol. 6)*, pages 4365–4368, Washington, DC, 1999. IEEE Press.

[4] A. Cangelosi and D. Parisi. The emergence of a language in an evolving population of neural networks. *Technical Report NSAL–96004, National Research Council, Rome*, 1996.

[5] N. Chomsky. On the nature of language. In *Origins and evolution of language and speech*, pages 46–57. Annals of the New York Academy of Science, New York. Vol 280, 1976.

[6] D. Denaro and D. Parisi. Cultural evolution in a population of neural networks. In *M.Marinaro and R.Tagliaferri (eds), Neural Nets Wirn-96.New York: Springer*, pages 100–111, 1996.

[7] G. K. E. K. Burke, S. Gustafson. Diversity in genetic programming: an analysis of measures and correlation with fitness. In *IEEE Trans. Evolutionary Computation*, volume 8(1), pages 47–62, 2004.

[8] S. E. Fahlman. Faster-learning variations on back propagation: An empirical study. In *Proceedings of the 1988 Connectionist Models Summer School*, 1988.

[9] J. J. Grefenstette. Genetic algorithms for dynamic environments. In R. Maenner and B. Manderick, editors, *Parallel Problem Solving from Nature 2*, pages 137–144, 1992.

[10] E. Hutchins and B. Hazlehurst. Learning in the cultural process. In *Artificial Life II, ed. C. Langton et al.*, pages 689–706. MIT Press, 1991.

[11] E. Hutchins and B. Hazlehurst. How to invent a lexicon: The development of shared symbols in interaction. In N. Gilbert and R. Conte, editors, *Artificial Societies: The Computer Simulation of Social Life*, pages 157–189. UCL Press: London, 1995.

[12] H. Kitano. Designing neural networks using genetic algorithm with graph generation system. In *Complex Systems, 4, 461-476*, 1990.

[13] B. MacLennan and G. Burghardt. Synthetic ethology and the evolution of cooperative communication. In *Adaptive Behavior 2(2)*, pages 161–188, 1993.

[14] F. Menczer. Changing latent energy environments: A case for the evolution of plasticity. In *Technical Report CS94-336*, 1994.

[15] G. F. Miller, P. M. Todd, and S. U. Hedge. Designing neural networks using genetic algorithms. In *Proceedings of the Third International Conference on Genetic Algorithms and Their Applications*, pages 379–384, 1989.

[16] S. Nolfi and D. Parisi. Learning to adapt to changing environments in evolving neural networks. In *Technical Report 95-15, Institute of Psychology, National Research Council, Rome, Italy.*, 1995.

[17] D. W. Opitz and J. W. Shavlik. Generating accurate and diverse members of a neural-network ensemble. In D. S. Touretzky, M. C. Mozer, and M. E. Hasselmo, editors, *Advances in Neural Information Processing Systems*, volume 8, pages 535–541. The MIT Press, 1996.

[18] L. Spector. Genetic programming and AI planning systems. In *Proceedings of Twelfth National Conference on Artificial Intelligence*, pages 1329–1334, Seattle, Washington, USA, 1994. AAAI Press/MIT Press.

[19] L. Steels. The synthetic modeling of language origins. In *Evolution of Communication*, pages 1–34, 1997.

[20] X. Y. Y. Liu and T. Higuchi. Evolutionary ensembles with negative correlation learning. In *IEEE Transactions on Evolutionary Computation*, volume 4(4), pages 380–387, 2000.

[21] H. Yanco and L. Stein. An adaptive communication protocol for cooperating mobile robots. In *From Animals to Animats 2. Proceedings of the second International Conference on Simulation of Adaptive Behavior*, pages 478–485. MIT Press, Cambridge Ma., 1993.