An Analysis of the Effects of Lifetime Learning on Population Fitness and Diversity in an NK fitness landscape

Dara Curran[1], Colm O'Riordan[2], and Humphrey Sorensen[3]

[1] Dept. of Computer Science,
University College Cork,
Ireland
d.curran@cs.ucc.ie
[2] Dept. of Information Technology,
National University of Ireland,
Galway
colm.oriordan@nuigalway.ie
[3] Dept. of Computer Science,
University College Cork,
Ireland
sorensen@cs.ucc.ie

Abstract. This paper examines the effects of lifetime learning on the diversity and fitness of a population. Our experiments measure the phenotypic diversity of populations evolving by purely genetic means (population learning) and of others employing both population learning and lifetime learning. The results obtained show, as in previous work, that the addition of lifetime learning results in higher levels of fitness than population learning alone. More significantly, results from the diversity measure show that lifetime learning is capable of sustaining higher levels of diversity than population learning alone.

1 Introduction

Much artificial intelligence research has focused on the interaction between learning and evolution, where individuals within a population of artificial organisms are capable of evolving genetically (population learning) and also of acquiring knowledge during their lifetime (lifetime learning). Hinton and Nowlan[1] were among the first to show that learning could guide evolution and a number of other researchers have since conducted experiments which support this view[2–8].

Each experiment simulates both a population of organisms and their surrounding environment. Typically, environmental stimuli are mapped to a specific problem domain. For instance, a population may inhabit an environment where individuals must navigate a maze or forage for food.

The majority of these investigations are concerned with the impact that learning has on the fitness of the population, finding in general that the addition of lifetime learning has a positive effect. In addition, studies have shown that both the Baldwin and Hiding effects can be simulated in populations of artificial organisms[9].

However, there is comparatively little research on other effects of learning, in particular with regard to population diversity. It is generally accepted that the maintenance of diversity can prevent an evolutionary algorithm from becoming trapped in local maxima. A number of techniques exist to both monitor and stimulate diversity[10–13].

The aim of this paper is to examine whether diversity can be maintained by lifetime learning. We employ an abstract model based on Hinton & Nowlan's work and conduct a number of experiments. The remainder of this paper is organised as follows. Section 2 describes related research, including the learning models on which this work is based and a discussion on diversity. Section 3 presents the model employed for these experiments, including a detailed description of the diversity measure. Section 4 outlines the experiment results and Section 5 presents conclusions.

2 Related Work

The following sections outline some background material, including learning models and diversity.

2.1 Learning Models

Population learning refers to the process whereby a population of organisms evolves, or learns, by genetic means through a Darwinian process of iterated selection and reproduction of fit individuals[14]. In this model, the learning process is strictly confined to each organism's genetic material: the organism itself does not contribute to its survival through any learning or adaptation process.

There exist species in nature that are capable of learning, or adapting to environmental changes and novel situations, at an individual level. Such learning, known as life-time learning, is often coupled with population-based learning, further enhancing the population's fitness through its adaptability and resistance to change. Much research explores the interactions between evolution and learning, showing that the addition of individual lifetime learning can improve a population's fitness[2, 3, 6–8].

Hinton & Nowlan Hinton and Nowlan employed a genetic algorithm to study the effects of lifetime learning on the performance of genetic evolution. In particular, they examined the Baldwin effect, which occurs when certain behaviours discovered through life-time learning become imprinted onto an individual's genetic material through the evolutionary processes of crossover and mutation[1, 15]. This should not be confused with Lamarckian theory, where an individual's lifetime experience is directly mapped to its genome and is directly inheritable[16].

Each agent in the model possesses a genome, represented by a string of characters which can be either *1*, *0* or *?*, mapping to a neural network structure. Each agent is allowed a number of rounds of lifetime learning where for each *?* in the genotype they 'guess' its value, assigning it either a *1* or a *0*. The *?* in the genotype represent the agent's phenotypic adaptability and any guessed values are lost at the end of the agent's life.

The experiments compared the performance of a population endowed with learning to one without. Results showed that the non-learning population was not capable of finding optimal solutions to the problem task. In contrast, once learning was applied, the population converged on the problem solution, showing that individual learning is capable of guiding genetic evolution. A number of researchers have since replicated the Baldwin effect in populations of artificial organisms[17, 6, 18–20].

Turney The model employed by Turney examined the shift in bias in simulations similar to those undertaken by Hinton and Nowlan [18]. The definition of bias in this context is all factors excluding input data that influence the selection of a concept. Bias can include the algorithm that the learner employs to search through concept space and the criteria for determining whether a selected concept is compatible with the training data the learner has received.

An agent's genome is a combination of bias direction genes and bias strength genes, where bias direction genes are *0*s or *1*s while the bias strength genes are real number values in the range 0 to 1. In a 64 element string, 32 alleles represent bias direction and a further 32 represent bias strength. If bias is strong at a particular allele, the individual will be more likely to select the bias direction present in its genome. However, if bias is weak then the individual will be more likely to guess at a bias direction rather than employing the genetically encoded value.

In other words, while the *?*s in Hinton and Nowlan's experiments were replaced with *1*s or *0*s with equal probability (which did not change during the course of the experiments), Turney's model allows the evolutionary process to shift biases towards particular guesses.

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[21–24].

Eriksson conducted an experiment designed to measure whether lifetime learning is capable of maintaining diversity [25]. The experiment employed Turney's model and measured the diversity of bias direction genes in the population. Experiments found that lifetime learning only produced higher diversity when bias strengths were explicitly initialised at low values. In all other cases, population learning produced higher diversity levels.

Previous work has examined the effects of cultural learning on both the phenotypic and genotypic diversity of populations of evolving neural network agents whose problem task was comprised of a number of turn-based games[26]. Cultural learning in this context is a form of lifetime learning where agents learn from other members of the populations. Results showed that the combination of cultural learning and population learning led to improvements in fitness and higher levels of genotypic diversity. However, the complexity of the model employed makes analysis difficult, highlighting the need for a simpler model of both agent and lifetime learning mechanism alike.

3 Model

Our model follows the structure of the original Hinton and Nowlan experiments. Each agent comprises a genecode of 20 bits and experiments employ a population of 1000 agents evolved for 400 generations. The experiments employ two populations: one using population learning alone (purely genetic evolution) and the other employing lifetime learning.

Where a population evolves solely by population learning, agent genomes consist of strings of 20 *1*s or *0*s randomly generated in the initial population with equal probability. Populations employing lifetime learning have genomes containing *1*s, *0*s or *?*s, where the *?*s represent the agent's phenotypic ability to guess either *1* or *0* in a similar manner to the original experiments. In the initial population these are randomly selected with probability 0.25, 0.25 and 0.5 respectively, corresponding to the original probability set in the Hinton and Nowlan experiment.

The fitness landscape of the original Hinton and Nowlan experiments was a particularly challenging one for populations attempting to evolve genetically, consisting in effect of a "needle in the haystack" problem where no evolutionary feedback is given to partial solutions. An individual with no lifetime learning would have a fitness of 1 unless its genome consisted of exactly twenty *1*s, in which case its fitness would be 20. It is highly unlikely that an individual will be randomly generated with such a genome and, since there is no path for the evolutionary process to follow, it is very difficult for population learning alone to solve this problem.

Our model employs the NK fitness model as the fitness landscape for the experiments [27]. The NK fitness model provides some evolutionary feedback to populations evolving genetically and therefore allows a fairer comparison between evolutionary and lifetime learning.

3.1 Fitness calculation

To calculate the fitness of each individual in the population, an NK fitness landscape is generated and the maximum fitness is calculated. The fitness of an individual is normalised with respect to the maximum possible fitness, obtaining a fitness measure in the range $[0, 1]$.

Populations employing purely genetic evolution (population learning alone) possess genomes comprising of *1*s or *0*s (no *?*s) and therefore have no opportunity to alter the way they interact with their environment. The fitness of agents in such populations is measured directly from their genomes.

When a population adds lifetime learning to population learning, its members are given the opportunity to replace each of the *?*s in their genomes with either *1*s or *0*s. Each agent guesses the value of its *?* with equal probability. The number of guesses an agent makes during its lifetime can be tuned to allow lifetime learning a greater or lesser influence on the evolutionary process. Each guess is evaluated using the fitness function and an agent's best guess is taken as its fitness value (see Fig. 1).

3.2 Diversity measure

The diversity measure examines the differences between members of a population and, for the purposes of this work, it is important that a diversity measure

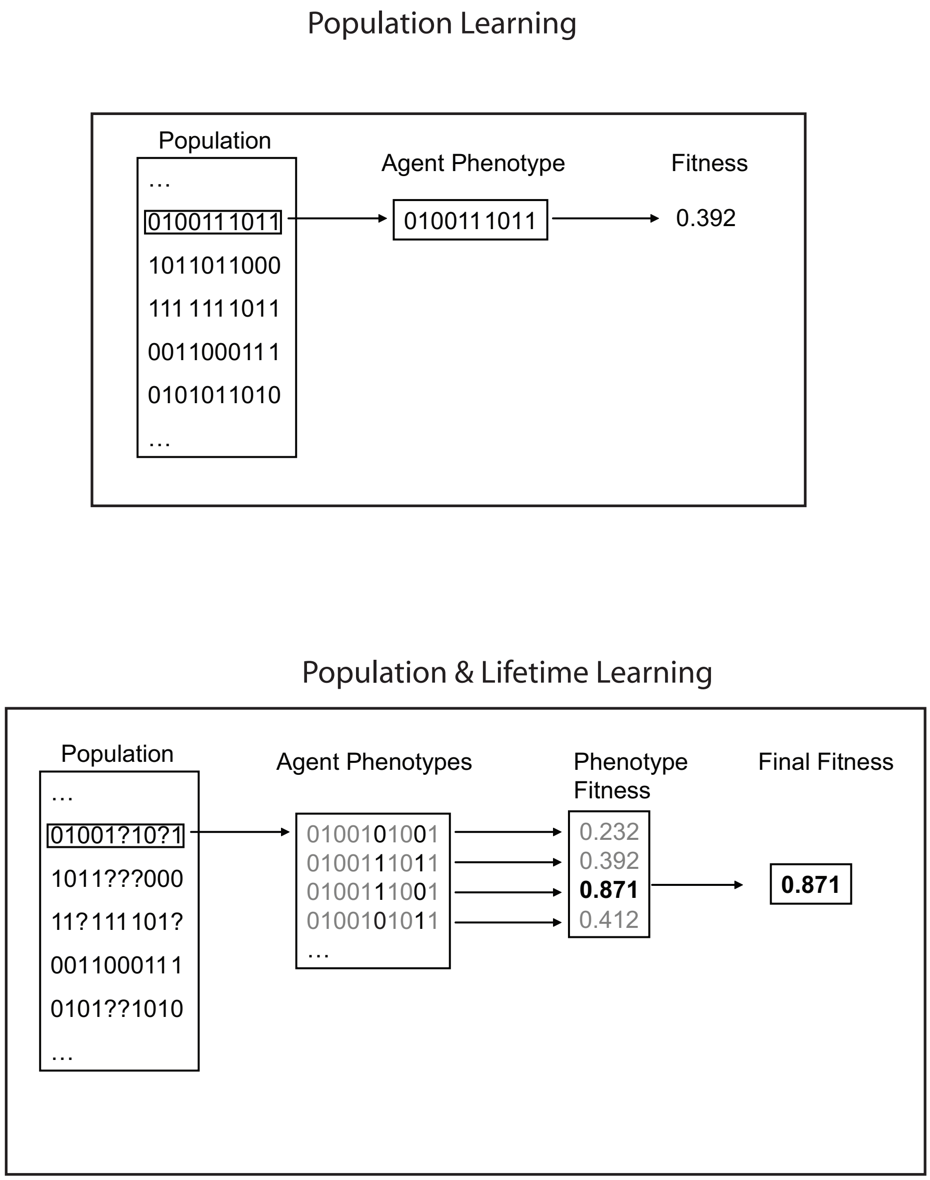

Fig. 1. *Fitness evaluation*

capture the differences between populations employing population learning alone and those employing both population and lifetime learning.

The population employing population learning alone contains genomes comprising of *1*s or *0*s, while the population adding lifetime learning has an additional number of *?*s within its genomes. This difference between the two populations makes a direct genetic comparison difficult. Therefore, a phenotypic measure of diversity may be more suitable. In this model, the phenotype of an agent is defined as its response to its environment.

In the population employing population learning alone, an agent's response to its environment is fixed and directly influenced by its genome. Therefore, its genotype and phenotype are identical - the agent has no lifetime influence on its behaviour. However, in the population employing lifetime learning, an agent is allowed make a number of guesses which replace the *?*s in its genome with either *1*s or *0*s. Since the agent's fitness is calculated according to the best of these guesses, we chose to select an agent's best guess as its phenotype.

Employing this model, phenotypic diversity measurements of populations employing population learning alone and populations employing both population and lifetime learning are directly comparable, as both are measuring differences between strings of equal cardinality (containing either *1*s or *0*s).

The diversity measure is implemented in a manner similar to Eriksson's and is defined as the average of all individual Hamming distances between phenotypes of individuals x and y, $h(x, y)$, whose phenotypes are unique within the population. Thus, for a population of size n containing m unique phenotypic binary strings, the diversity measure can be calculated as:

$$\frac{2}{n(m-1)} \sum_{x=1}^{m} \sum_{y=x+1}^{m} h(x, y)$$

4 Experiments

The experiments employ three populations: one employing population learning alone, one employing both population and lifetime learning with 10 learning cycles and the final one employing both population and lifetime learning with 30 learning cycles. Both fitness and diversity are measured for each experiment. Populations of 1000 agents are allowed to evolve to 400 generations. Linear ranking selection is employed, 1-point crossover occurs with 0.6 probability and mutation occurs with 0.02. The NK landscape parameters are N=20, K=15. Results presented are averaged from 20 independent experiment runs.

4.1 Fitness

Fig. 2 illustrates the results obtained from the three populations over 400 generations. Each of the populations converges within 100 generations to high levels of fitness. However, it is clear that both populations employing population and lifetime learning have higher fitness levels than the population employing population learning alone.

Both lifetime learning populations begin the experiment with a higher fitness level than population learning alone and both gradually increase to convergence. By contrast, the non-learning population's increase in fitness is much sharper,

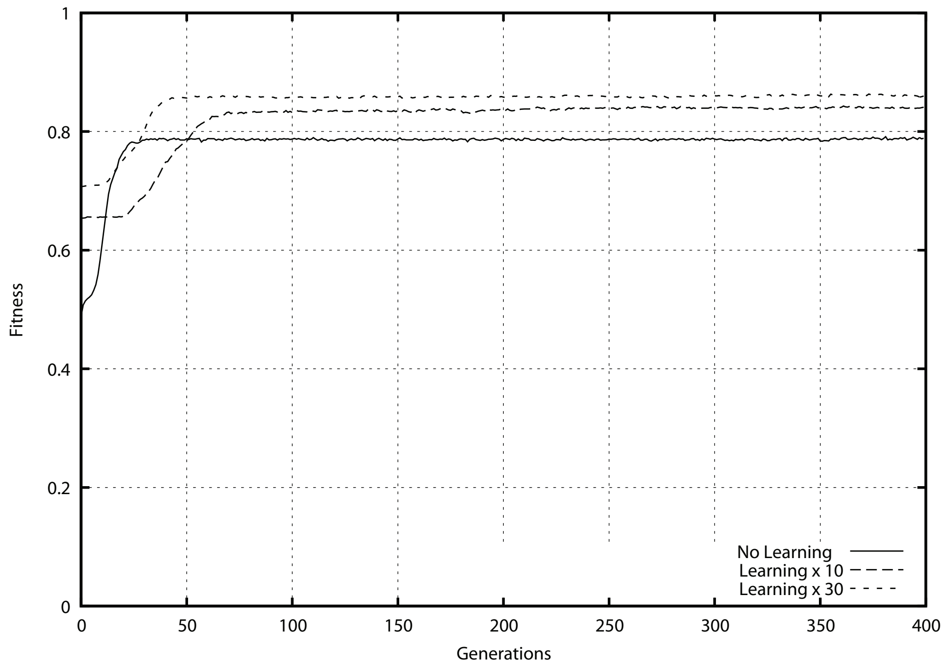

Fig. 2. *Average Fitness*

occurring within the first 10 generations. The more gradual fitness increase exhibited by the learning process is likely to be the result of a wider search space exploration, leading in turn to a higher level of fitness once the population converges.

4.2 Diversity

Fig. 3 shows the results obtained from the diversity measure for each of the three populations. Each of the three begin with a high level of diversity at the start of the experiment, but this quickly converges to a very low level. All populations become phenotypically similar by generation 100. However, the level of diversity exhibited by the populations employing lifetime learning converges at a higher level than that of population learning alone.

In addition, the initial levels of diversity exhibited by the lifetime learning populations is significantly lower than that of population learning alone. Furthermore, there is a more gradual decrease in diversity in populations employing lifetime learning. The gradual descent in diversity is clearly tied to the gradual increase in fitness levels described in the previous section. As the population slowly converges towards its optimum level, so does its diversity gradually decrease.

It is also clear from these results that the level of diversity does not increase as the number of learning opportunities increases. The population employing 10 learning cycles has a level of diversity that is significantly higher than the population employing 30 learning cycles (after 20 generations).

Diversity levels for all populations converge by generation 75, indicating that the population has reached an optimum and that a particular set of individuals

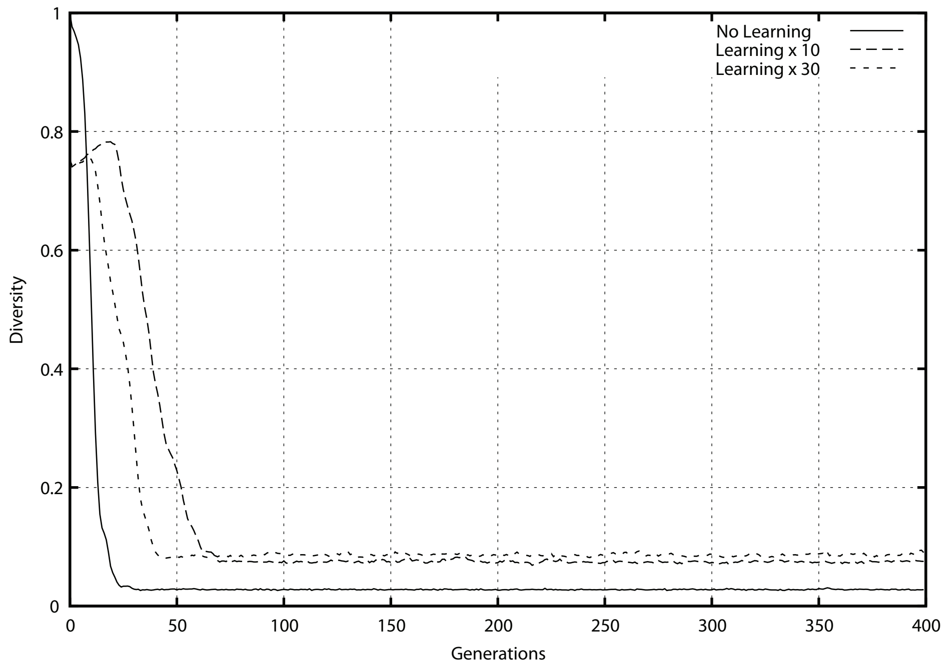

Fig. 3. *Average Diversity*

are dominant throughout the population. However, given that the fitness levels for populations employing both population and lifetime learning is significantly higher than the population employing population learning alone, it is clear that the initial maintenance of high diversity levels associated with learning allows those populations to more thoroughly search the problem space and converge on better individual solutions.

These results show that learning has a clear effect on the diversity of populations – even small levels of learning produce an increase in population diversity. This is in contrast to the results obtained by Eriksson, who showed that diversity could only be increased through learning if initial parameters were carefully tweaked. However, future work should examine whether the impact of the initial proportion of *?* alleles contributes significantly to these results.

It is likely that the higher level of diversity comes from the hiding effect provided by the learning process. In other words, the learning process allows individuals without a genetic innate ability to survive by improving their fitness during their lifetime, thus hiding genetically weak individuals from the selection process.

An individual born in a non-learning population has no opportunity to influence its fitness because this is entirely derived from its genome and the selection process quickly moves to propagate highly genetically fit individuals throughout the population. Learning gives individuals the opportunity to search the problem space using their inherited genotypes as a starting point. Thus, if an individual is born with a mediocre genotypic fitness, it is possible it will significantly improve through the learning process and be selected for reproduction.

5 Conclusion

This paper examined both fitness and diversity levels in populations employing population learning alone and populations employing both population and lifetime learning. The model we employed is based on the Hinton & Nowlan model and was chosen for its relative simplicity and ease of analysis. The experiments that were conducted produced a number of interesting results.

Firstly, the results provide a further confirmation that combining lifetime learning with population learning leads to an increase in fitness. Unlike previous work based on the Hinton & Nowlan model, our fitness measure is not biased towards lifetime learning. The NK fitness model allows population learning to successfully develop solutions and highlights the fact that lifetime learning can provide a fitness improvement in a non-biased environment.

Secondly, we have demonstrated that diversity levels can be measured in an alternative manner to previous work based on modified versions of the Hinton & Nowlan model. Previous work examined the genetic differences between bias strength and direction, while our work is based on a phenotypic diversity measure examining the behaviour of the population in a much simpler setting.

Finally, our results show that lifetime learning maintains higher diversity levels than population learning alone. This is arguably a more intuitive result than that found in previous related work. Learning provides an opportunity for genetically weaker individuals to attain higher fitness levels. The high diversity levels in such populations attests to the fact that many behaviours are present in the population, leading to a larger, directed exploration of search space, which in turn leads to higher fitness levels.

Acknowledgements

The first author would like to acknowledge the support of the Irish Research Council for Science, Engineering and Technology.

References

1. Hinton, G.E., Nowlan, S.J.: How learning guides evolution. Complex Systems **1** (1987) 495–502
2. Nolfi, S., Parisi, D.: Learning to adapt to changing environments in evolving neural networks. Adaptive Behavior **5** (1996) 75–97
3. Floreano, D., Mondada, F.: Evolution of plastic neurocontrollers for situated agents. Animals to Animats **4** (1996)
4. Sasaki, T., Tokoro, M.: Adaptation toward changing environments: Why darwinian in nature? In Husbands, P., Harvey, I., eds.: Fourth European Conference on Artificial Life, Cambridge, MA, MIT Press (1997) 145–153
5. Pereira, F.B., Costa, E.: How learning improves the performance of evolutionary agents: A case study with an information retrieval system for a distributed environment. In: Proceedings of the International Symposium on Adaptive Systems: Evolutionary Computation and Probabilistic Graphical Models (ISAS 2001). (2001) 19–23
6. Watson, J., Wiles, J.: The rise and fall of learning: A neural network model of the genetic assimilation of acquired traits. In: Proceedings of the 2002 Congress on Evolutionary Computation (CEC 2002). (2002) 600–605

7. Curran, D., O'Riordan, C.: On the design of an artificial life simulator. In Palade, V., Howlett, R.J., Jain, L.C., eds.: Proceedings of the Seventh International Conference on Knowledge-Based Intelligent Information & Engineering Systems (KES 2003), University of Oxford, United Kingdom (2003) 549–555
8. Curran, D., O'Riordan, C.: Artificial life simulation using marker based encoding. In: Proceedings of the 2003 International Conference on Artificial Intelligence (IC-AI 2003). Volume II., Las Vegas, Nevada, USA (2003) 665–668
9. Mayley, G.: Guiding or hiding: Explorations into the effects of learning on the rate of evolution. In: Proceedings of the Fourth European Conference on Artificial Life, MIT Press (1997)
10. O'Reilly, U.M.: Using a distance metric on genetic programs to understand genetic operators. In: IEEEInternational Conference on Systems, Man, and Cybernetics. Volume 5. (1997) 4092–4097
11. De Jong, K.A.: Analysis of Behavior of a Class of Genetic Adaptive Systems. PhD thesis, The University of Michigan (1975)
12. Booker, L.B.: Improving the performance of genetic algorithms in classifier systems. In: Proc. of the International Conference on Genetic Algorithms and Their Applications, Pittsburgh, PA (1985) 80–92
13. Collins, R.J., Jefferson, D.R.: Selection in massively parallel genetic algorithms. In: ICGA. (1991) 249–256
14. Darwin, C.: The Origin of Species: By Means of Natural Selection or the Preservation of Favoured Races in the Struggle for Life. Bantam Classics (1859)
15. Baldwin, J.: A new factor in evolution. American Naturalist **30** (1896) 441–451
16. Lamarck, J.B.: Philosophie Zoologique. Chez Dentu, Paris (1809)
17. Boers, E., Borst, M., Sprinkhuizen-Kuyper, I.: Evolving Artificial Neural Networks using the "Baldwin Effect". Technical Report TR 95-14 (1995)
18. Turney, P.D.: How to shift bias: Lessons from the baldwin effect. Evolutionary Computation **4** (1996) 271–295
19. Suzuki, R., Arita, T.: The baldwin effect revisited: Three steps characterized by the quantitative evolution of phenotypic plasticity. In Banzhaf, W., Christaller, T., Dittrich, P., Kim, J.T., Ziegler, J., eds.: ECAL. Volume 2801 of Lecture Notes in Computer Science., Springer (2003) 395–404
20. Arita, T., Suzuki, R.: Interactions between learning and evolution: The outstanding strategy generated by the baldwin effect. In: Proceedings of Artificial Life VII, MIT Press (2000) 196–205
21. Liu, Y., Yao, X., Higuchi, T.: Evolutionary ensembles with negative correlation learning. In: IEEE Transactions on Evolutionary Computation. Volume 4(4). (2000) 380–387
22. Opitz, D.W., Shavlik, J.W.: Generating accurate and diverse members of a neural-network ensemble. In Touretzky, D.S., Mozer, M.C., Hasselmo, M.E., eds.: Advances in Neural Information Processing Systems. Volume 8., The MIT Press (1996) 535–541
23. Brown, G.: Diversity in Neural Network Ensembles. PhD thesis, University of Birmingham (2003)
24. Burke, E.K., Gustafson, S., Kendall, G.: Diversity in genetic programming: an analysis of measures and correlation with fitness. In: IEEE Trans. Evolutionary Computation. Volume 8(1). (2004) 47–62
25. Eriksson, R.I.: An initial analysis of the ability of learning to maintain diversity during incremental evolution. In Freitas, A.A., Hart, W., Krasnogor, N., Smith, J., eds.: Data Mining with Evolutionary Algorithms, Las Vegas, Nevada, USA (2000) 120–124
26. Curran, D., O'Riordan, C.: Increasing population diversity through cultural learning. Adaptive Behavior **14** (2006)
27. Kaufmann, S.A.: Adaptation on rugged fitness landscapes. Lectures in theSciences of Complexity **1** (1989) 527–618