

Discrete Dynamical Genetic Programming in XCS

Richard Preen & Larry Bull

Technical Report UWELCSG09-001
Learning Classifier Systems Group
UWE, Bristol, UK

ABSTRACT

A number of representation schemes have been presented for use within Learning Classifier Systems, ranging from binary encodings to neural networks. This paper presents results from an investigation into using a discrete dynamical system representation within the XCS Learning Classifier System. In particular, asynchronous random Boolean networks are used to represent the traditional condition-action production system rules. It is shown possible to use self-adaptive, open-ended evolution to design an ensemble of such discrete dynamical systems within XCS to solve a number of well-known test problems.

Categories and Subject Descriptors

I.2.6 [Artificial Intelligence]: Learning – *knowledge acquisition, parameter learning.*

General Terms

Experimentation.

Keywords

Learning Classifier Systems, XCS, Random Boolean Networks, Reinforcement Learning, Self-Adaptation.

1. INTRODUCTION

Traditionally, Learning Classifier Systems (LCS) [15] use a ternary encoding to generalize over the environmental inputs and to associate appropriate actions. A number of representations have previously been presented beyond this scheme however, including real numbers [35], fuzzy logic [31] and neural networks [4]. To date, no temporally dynamic representation schemes have been used in LCS, a potentially important approach since temporal behaviour of such kinds is viewed as a significant aspect of cognition in general.

In this paper we explore the use of a dynamical system

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

Conference '09, July 8-12, 2009, Montréal, Canada.

Copyright 2009 ACM 1-58113-000-0/00/0004...\$5.00.

representation within XCS [33] – what is herein termed “dynamical genetic programming” (DGP). Traditional tree-based genetic programming (GP) [19] has been used within LCS both to calculate the action [1] and to represent the condition [20]. DGP uses a graph-based representation, each node of which is constantly updated with asynchronous parallelism, and evolved using an open-ended, self-adaptive scheme. In the discrete case, each node is a Boolean function and therefore equivalent to a form of random Boolean network (RBN) (e.g., [18]). We show that XCS is able to solve a number of well-known immediate and delayed reward tasks using this temporally dynamic knowledge representation scheme.

2. RELATED WORK

A number of representations have been presented by which to enable the evolution of computer programs, the most common being tree-based LISP S-expressions [19]. Other forms of GP include the use of machine code instructions (e.g., [3]) and finite state machines (e.g., [11]). Most relevant to the form of GP used in this paper is the small amount of prior work on graph-based representations. Teller and Veloso’s [30] “neural programming” uses a directed graph of connected nodes, each with functionality defined in the standard GP way, with recursive connections included. Significantly, each node is executed with synchronous parallelism for some number of cycles before an output node’s value is taken. Poli (e.g., [26]) presented a very similar scheme wherein the graph is placed over a two-dimensional grid and executes its nodes synchronously in parallel. Other examples of graph-based GP typically contain sequentially updating nodes (e.g., [23]). Schmidt and Lipson [27] have recently demonstrated a number of benefits from graph encodings over traditional trees, such as reduced bloat and increased computational efficiency.

As noted above, tree-based S-expressions have been used within LCS. Recently, Wilson [36] has explored the use of a form of Gene Expression Programming (GEP) [10] within LCS. Here the rules are represented as expression trees which are evaluated by assigning the environmental inputs to the tree’s terminals, evaluating the tree, and then comparing the result with a predetermined threshold. Whenever the threshold value is exceeded, the rule is added to the match set.

The most common form of discrete dynamical system is the Cellular Automaton (CA) [32] which consists of an array of cells (lattice of nodes) where the cells exist in states from a finite set and update their states with synchronous parallelism in discrete

time. Traditionally, each cell calculates its next state depending upon its current state and the states of its closest neighbours. That is, CAs may be seen as a graph with a (typically) restricted topology. Packard [25] was the first to use evolutionary computing techniques to design CAs such that they exhibit a given emergent global behaviour. Following Packard, Mitchell et al. (e.g., [24]) have investigated the use of a Genetic Algorithm (GA) [14] to learn the rules of uniform binary CAs. As in Packard’s work, the GA produces the entries in the update table used by each cell, candidate solutions being evaluated with regard to their degree of success for the given task. Andre et al. [2] repeated Mitchell et al.’s work whilst using traditional GP to evolve the update rules. They report similar results. Sipper (e.g., [28]) presented a non-uniform, or heterogeneous, approach to evolving CAs. Each cell of a one- or two-dimensional CA is also viewed as a GA population member, mating only with its lattice neighbours and receiving an individual fitness. He shows an increase in performance over Mitchell et al.’s work by exploiting the potential for spatial heterogeneity in the tasks. Sipper and Ruppin [29] extended this approach to enable heterogeneity in the node connectivity, along with the node function; they evolved a form of random Boolean networks.

3. RANDOM BOOLEAN NETWORKS

The discrete dynamical systems known as random Boolean networks were originally introduced by Kauffman (see [18]) to explore aspects of biological genetic regulatory networks. Since then they have been used as a tool in a wide range of areas such as self-organisation (e.g., [18]) and computation (e.g., [22]). An RBN typically consists of a network of N nodes, each performing a Boolean function with K inputs from other nodes in the network, all updating synchronously. As such, RBN may be viewed as a generalization of binary CAs. Since they have a finite number of possible states and they are deterministic, the dynamics of RBN eventually fall into a basin of attraction. It is well-established that the value of K affects the emergent behaviour of RBN wherein attractors typically contain an increasing number of states with increasing K . Three phases of behaviour are suggested: ordered when $K=1$, with attractors consisting of one or a few states; chaotic when $K>3$, with a very large number of states per attractor; and, a critical regime around $K=2$, where similar states lie on trajectories that tend to neither diverge nor converge and 5-15% of nodes change state per attractor cycle (see [18] for discussions of this critical regime, e.g., with respect to perturbations). Analytical methods have been presented by which to determine the typical time taken to reach a basin of attraction and the number of states within such basins for a given degree of connectivity K (see [18]).

Closely akin to the work described here, Kauffman [18] describes the use of simulated evolution to design RBN which must play a (mis)matching game wherein mutation is used to change connectivity, the Boolean functions, K and N . He reports the typical emergence of high fitness solutions with $K=2$ to 3, together with an increase in N over the initialised size.

As noted above, traditional RBN consist of N nodes updating synchronously in discrete time steps, but asynchronous versions have also been presented, after [13], leading to a classification of the space of possible forms of RBN [12]. Asynchronous forms of CA have also been explored (e.g., [17]) wherein it is often

suggested that asynchrony is a more realistic underlying assumption for many natural and artificial systems.

Asynchronous logic devices are known to have the potential to consume less power and dissipate less heat [33], which may be exploitable during efforts towards hardware implementations of such systems. Asynchronous logic is also known to have the potential for improved fault tolerance, particularly through delay insensitive schemes (e.g., [8]). This may also prove beneficial for hardware implementations.

Harvey and Bossomaier [13] showed that asynchronous RBN exhibit either point attractors, as seen in asynchronous CAs, or “loose” attractors where “the network passes indefinitely through a subset of its possible states” [ibid.] (as opposed to distinct cycles in the synchronous case). Thus the use of asynchrony represents another feature of RBN with the potential to significantly alter their underlying dynamics thereby offering another mechanism by which to aid the simulated evolutionary design process for a given task. Di Paolo [9] showed it is possible to evolve asynchronous RBN which exhibit rhythmic behaviour at equilibrium. Asynchronous CAs have also been evolved (e.g., [29]).

4. DISCRETE DGP-XCS

To use asynchronous RBN as the rules within XCS, the following scheme is adopted. Each of an initial randomly created rule’s nodes has K randomly assigned connections, here $1 \leq K \leq 5$. There are as many nodes N as input fields I for the given task and its outputs O , plus one other, as will be described, i.e., $N=I+O+1$. The first connection of each input node is set to the corresponding locus of the input message. The other connections are assigned at random within the RBN as usual. In this way, the current input state is always considered along with the current state of the RBN itself per network update cycle by such nodes (see Figure 9). Nodes are initialised randomly each time the network is run to determine [M], etc. The population is initially empty and covering is applied to generate rules as in the standard XCS approach.

Matching consists of executing each rule for T cycles based on the current input. The value of T is chosen to be a value typically within the basin of attraction of the RBN. Asynchrony is here implemented as a randomly chosen node being updated on a given cycle, with as many updates per overall network update cycle as there are nodes in the network before an equivalent cycle to one in the synchronous case is said to have occurred (see [12] for alternative schemes).

In this study, when well-known Boolean problems are explored there are only two possible actions, only one output node is required. Where well-known maze and woods problems are explored there are eight possible actions and accordingly three required output nodes. An extra “matching” node is also required to enable RBNs to (potentially) only match specific sets of inputs. If a given RBN has a logical ‘0’ on the match node, regardless of its output node’s state, the rule does not join [M] (see Figure 9). This scheme has also been exploited within neural LCS [4]. A ‘windowed approach’ is utilised where the output is decided by the most common state over the last W steps up to T . For example, if the last few states on a node updating prior to cycle T is 0101001 and $W=3$, then the ending node’s state would be ‘0’ and not ‘1’. In this paper, W is set to 3. Thereafter, match set and

action set processing proceeds as standard in XCS (the reader is referred to [7] for an algorithmic description of XCS).

When covering is necessitated, a randomly constructed RBN is created and then executed for T cycles to determine the status of the match and output nodes. This procedure is repeated until an RBN is created that matches the environment state.

Parameter self-adaptation was first explored in LCS by Bull et al. [5] wherein the mutation rate is a locally evolving entity in itself; each rule has its own mutation rate μ . Mutation only is used here and applied to the node's logic and connectivity map at rate μ . Node logic is represented by a binary string and its connectivity by a list of K integers in the range $[1, N]$. Since each rule node has a given fixed K value, each node maintains a binary string of length 2^K which forms the entries in the look-up table for each of the possible 2^K input states of that node, i.e., as in Packard's [25] aforementioned work on evolving CAs, for example. These strings are subjected to mutation on reproduction at the self-adapting rate μ for that rule. Hence, within the RBN representation, evolution can define different Boolean functions for each node within a given network rule, along with its connectivity map. Specifically, each rule has its own mutation rate stored as a real number and initially seeded uniform randomly in the range $[0.0, 1.0]$. This parameter is passed to its offspring. The offspring then applies its mutation rate to itself using a Gaussian distribution, i.e., $\mu = \mu e^{N(0,1)}$, before mutating the rest of the rule at the resulting rate.

Due to the need for a possible different number of nodes within the rules for a given task, the DGP scheme is also of variable length. Once the node logic and connections have been mutated, a new randomly connected node is either added or the last added node is removed with the same probability μ . The latter case only occurs if the network currently consists of more than the initial number of nodes. Thus DGP is temporally dynamic both in the search process and the representation scheme. Evolving variable-length solutions via mutation only has previously been explored a number of times (e.g., [11]). Traditional GP can be seen to primarily rely upon recombination to search the space of possible tree sizes, although the standard mutation operator effectively increases or decreases tree size also.

Whenever an offspring classifier is created and no changes occur to its RBN when undergoing mutation, the parent's numerosity is increased and mutation rate set to the offspring's.

5. EXPERIMENTATION

5.1 Multiplexer

We now apply this discrete version of DGP-XCS (dDGP-XCS) to the well-known multiplexer task. These Boolean functions are defined for binary strings of length $l = x + 2^x$ under which the x bits index into the remaining 2^x bits, returning the value of the indexed bit. The correct classification to a randomly generated input results in a payoff of 1000, otherwise 0.

Figures 1-4 show the performance of the constructed system on the 6-bit multiplexer problem updated asynchronously with $P=800$, $v=10$, $\theta_{GA}=25$, $\beta=0.2$, $T=25$ and $N_{init} = 8$ (6 inputs, one output, one match node). After [34], performance from exploit trials only is recorded (fraction of correct responses are shown), using a 50-point running average, averaged over ten runs.

From Figure 1 it can be seen that a near optimal solution is learnt around 35,000 trials and optimality is observed around trial 58,000. The parameter governing RBN mutation (see Figure 4) declines rapidly until reaching a bottom around 40,000 trials, which is shortly after discovering an optimal solution. For comparative purposes, Figures 5-8 show the performance with the same parameters on the 6-bit multiplexer when updated synchronously. It is shown that the performance is very similar regardless of the updating scheme and that there is thus apparently very little overhead when updating asynchronously, with the possible benefits mentioned above. In both cases, the number of (non-unique) rules initially grows rapidly, before declining to around 650. Furthermore, in both cases, the average degree of connectivity K decreases fractionally, whilst, on average, each network grows approximately one extra node. This behaviour indicates that the evolutionary process is able to identify an appropriate typical topology with which to generate complex behaviour, i.e., in this case a computation. For other tasks, other values of K may prove beneficial; high K may be expected in random number generation, for example. It can be noted that a growth event under which a new node is added into an RBN is essentially neutral here since the new node receives inputs from the existing nodes (or itself) on addition but only provides inputs to other nodes after subsequent connectivity mutations.

Figure 9 provides an illustration of a rule generated whilst solving the 6-bit multiplexer problem when updated asynchronously. There is one new node in addition to the initial eight. The node logic shows to which state each node will transition, given each of the possible inputs. For example, the output node (node 1) has node logic of '10' which is synonymous with a NOT gate where if node 3 is in state '0' then the output node will be set to '1', and if node 3 is in state '1' then the output node will be set to '0'. The logic of node 3 is synonymous with an AND gate, etc.

The rule has a Prediction of 1000.00 and an Error of 0.0, whilst having an Experience of 822, showing that this is a highly accurate rule. Analysis of this RBN rule was undertaken by executing it for each of the sixty-four 6-bit inputs. Each input was run twenty times with $T=25$. The results show that for the majority of environment states the network will return a false match node, preventing it from being added to [M]. However, the network is general as the match node will always return true when the environment states are 110000, 110010, 110100, 110110, 111000, 111010, 111100, and 111110. In all of those cases the output node always advocates action '0'. In addition, there are several environment states for which the match node will only sometimes return true. However, in all cases when the match node does permit the rule to be added to [M], the action advocated will always be consistent. There are four such additional environment states (010000, 010010, 011000, and 011010) for which the rule will match, albeit with a probability less than 50%.

The rule in Figure 9 was then re-run as before, however using a traditional synchronous updating scheme. The results of the match node and output nodes are extremely similar regardless of the updating mechanism. That is, XCS appears to evolve RBN which are very robust to the random nature of the asynchronous updating, meaning they are accurate even for the relatively rare case of all nodes updating concurrently, i.e., the synchronous case.




Figure 1. 6-bit MUX - Performance.

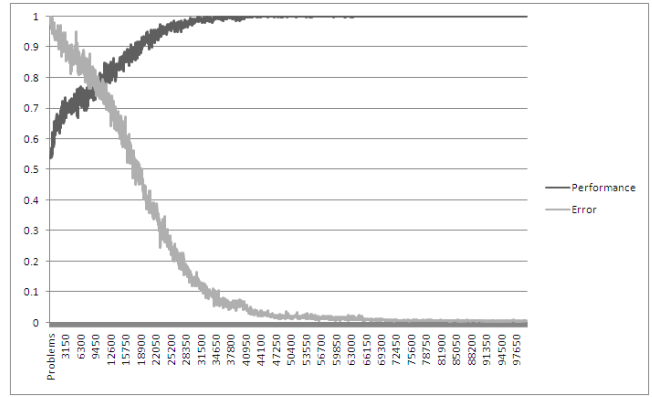


Figure 5. 6-bit MUX Synchronous - Performance.




Figure 2. 6-bit MUX - Size.

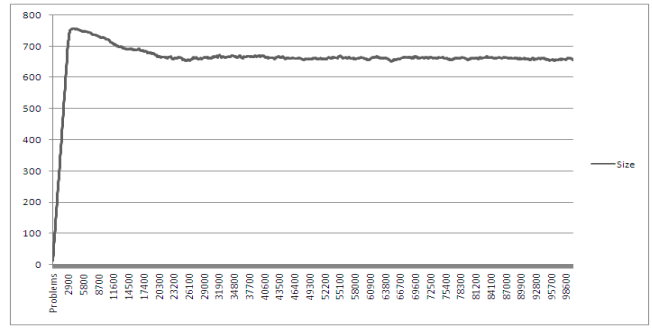


Figure 6. 6-bit MUX Synchronous - Size.




Figure 3. 6-bit MUX - Topology.




Figure 7. 6-bit MUX Synchronous - Topology.




Figure 4. 6-bit MUX - Mutation.





Figure 8. 6-bit MUX Synchronous - Mutation.



Prediction: 1000.0. Prediction Error: 0.0. Numerosity: 1. Experience: 822.
Average Set Size: 70.10529. Time-stamp: 99947.

	Node Logic:	Connections:
Node 0 (M):	10011000100000001110011010101000	7, 4, 0, 3, 1
Node 1 (O):	10	3
Node 2 (I):	00011111	11, 2, 5
Node 3 (I):	0001	12, 2
Node 4 (I):	11101110	13, 6, 3
Node 5 (I):	0110110100001010	14, 2, 7, 6
Node 6 (I):	0001011101010101	15, 5, 2, 3
Node 7 (I):	0100	16, 3
Node 8 (N):	00010111	3, 1, 5

Figure 9. dDGP-XCS 6-bit MP asynchronous rule.

5.2 Maze Environments

In addition to the single-step multiplexer problems, the created dDGP-XCS is applied to versions of three well-known multi-step maze environments: Woods1 (see Figure 10), Maze4 (see Figure 11), and Woods101 (see Figure 12).

*	*	*	*	*
*	*	*	*	*
O	O	F	*	*
O	O	O	*	*
O	O	O	*	*

Figure 10. Woods1 Environment.

O	O	O	O	O	O	O	O
O	*	*	O	*	*	F	O
O	O	*	*	O	*	*	O
O	O	*	O	*	*	O	O
O	*	*	*	*	*	*	O
O	O	*	O	*	*	*	O
O	*	*	*	*	O	*	O
O	O	O	O	O	O	O	O

Figure 11. Maze4 Environment.

O	O	O	O	O	O	O
O	*	*	*	*	*	O
O	*	O	*	O	*	O
O	*	O	F	O	*	O
O	O	O	O	O	O	O

Figure 12. Woods101 Environment.

Each cell in the maze environments is encoded with two binary bits, where white space is represented as a '*', obstacles as 'O', and food as 'F'. Furthermore, actions are encoded as an integer as shown in Figure 13. The task is simply to find the shortest path to the food (F) given a random start point. Obstacles (O) represent cells which cannot be occupied. In Woods1 the optimal number of steps to the food is 1.7, in Maze4 optimal is 3.5 steps, and in Woods101 it is 2.9.

Cell	Binary Encoding	Actions		
*	00	7	0	1
O	01	6		2
F	11	5	4	3

Figure 13. Maze Encoding.

Figures 14-17 show the performance of dDGP-XCS in the Woods1 environment. The parameters used are identical to those applied in the aforementioned multiplexer experiments, except that $N_{mit} = 20$ (16 inputs, 3 outputs, one match node). As can be seen from Figure 14, optimality is observed around 2,500 trials. This roughly matches the performance of neural-XCS using self-adaptive constructivism [16] and is faster than XCS with Lisp-S expression conditions [20]. Figure 15 shows that there is an average of 745 (non-unique) rules evolved. Figure 15 shows that on average the networks add one extra node (from the original 20) and the average number of connections decreases slightly. Figure 17 shows that the mutation rate declines rapidly by 2,800 trials, shortly after the optimal solution is learnt.

Figures 18-21 present the performance of dDGP-XCS in the Maze4 environment. The parameters used are identical to those in the Woods1 environment, however a bigger population limit of $P=2000$ is used, reflecting the bigger search space. Optimality is observed around trial 23,000, which is again similar to the performance observed in [16] using a neural-XCS. The average number of rules evolved is around 1,800. The average number of nodes in the networks also increases by almost one, and the average number of connections declines slightly from 3. The parameter governing RBN mutation declines rapidly after 4,000 trials, before finally stabilising after 15,000 trials.

The Woods101 maze is a non-Markov environment that includes two positions which are identically sensed, but require different optimal actions. To solve this maze optimally, a form of memory must be utilised. Optimal performance has previously been achieved in Woods 101 through the addition of a memory register mechanism in XCS [21] and by a neural-LCS using recurrent links [6]. Here we explore and extend the hypothesis of inherent memory existing within synchronous RBN due to different possible routes to a basin of attraction [37] for the asynchronous case by simply not resetting the node states on each step.




Figure 14. Woods1 - Performance.




Figure 18. Maze4 - Performance.




Figure 15. Woods1 - Size.




Figure 19. Maze4 - Size.




Figure 16. Woods1 - Topology.




Figure 20. Maze4 - Topology.




Figure 17. Woods1 - Mutation.




Figure 21. Maze4 - Mutation

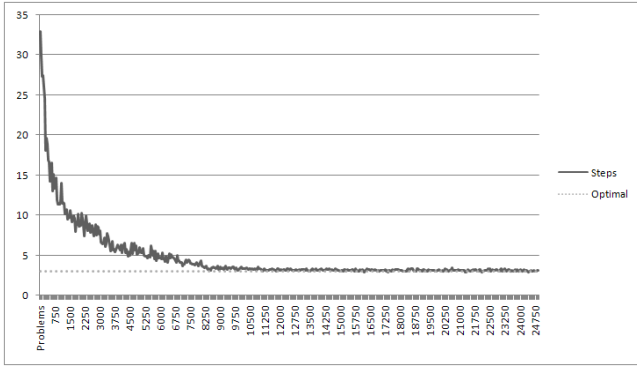


Figure 22. Woods101 - Performance.




Figure 23. Woods101 - Size.




Figure 24. Woods101 - Topology.




Figure 25. Woods101 - Mutation.

That is, nodes are initialised at random for the initial random placing in the maze but thereafter they are not reset for each subsequent matching cycle. The parameters used are identical to those applied in the Maze4 environment. As can be seen from Figure 22, dDGP-XCS, without node resets, is able to achieve optimal performance in Woods101 after approximately 12,000 trials. Optimal performance is unattainable however when the nodes are reset randomly between matching (not shown), proving that the system is exploiting the potential for memory within asynchronous RBN here. The mechanism works within XCS because rules/networks experience each input but need not match on each cycle. Hence for the ambiguous states they remain accurate for the payoff received on providing the action but do so having processed the previous input in an appropriate way without matching.

6. CONCLUSIONS

In this paper a form of XCS has been presented with which to design asynchronous random Boolean networks. It has been shown that XCS is able to design ensembles of RBN that collectively solve a computational task under a reinforcement learning scheme. In particular, it has been shown possible to exploit the inherent dynamics of the representation scheme to solve a non-Markov maze, i.e., without extra mechanisms. Current research is exploring the possibilities of DGP as a general representation scheme by which to solve complex problems with LCS.

7. REFERENCES

- [1] Ahluwalia, M. & Bull, L. (1999) A Genetic Programming Classifier System. In W. Banzhaf et al. (Eds) Proceedings of the Genetic and Evolutionary Computation Conference – GECCO-99. Morgan Kaufmann, pp11-18.
- [2] Andre, D., Koza, J.R., Bennett, F.H. & Keane, M. (1999) Genetic Programming III. MIT Press.
- [3] Banzhaf, W. (1993) Genetic Programming for Pedestrians. In S. Forrest (Ed) Proceedings of the Fifth International Conference on Genetic Algorithms. Morgan Kaufmann, pp628.
- [4] Bull, L. (2002) On Using Constructivism in Neural Classifier Systems. In J. Merelo et al. (Eds) Parallel Problem Solving from Nature - PPSN VII. Springer Verlag, pp558-567.
- [5] Bull, L., Hurst, J. & Tomlinson, A. (2000) Self-Adaptive Mutation in Classifier System Controllers. In J-A. Meyer et al. (Eds) From Animals to Animats 6. MIT Press, pp460-468.
- [6] Bull, L. & Hurst, J. (2003) A Neural Learning Classifier System with Self-Adaptive Constructivism. In IEEE Congress on Evolutionary Computation. IEEE Press. Vol.2, pp991- 997.
- [7] Butz, M. & Wilson, S.W. (2001) An Algorithmic Description of XCS. In Lanzi, P. L., Stolzmann, W., and S. W. Wilson (Eds), Advances in Learning Classifier Systems. Third International Workshop (IWLCS-2000), Lecture Notes in Artificial Intelligence (LNAI-1996). Berlin: Springer-Verlag.

- [8] Di, J. & Lala, P.K. (2007) Cellular Array-based Delay-insensitive Asynchronous Circuits Design and Test for Nanocomputing Systems. *Journal of Electronic Testing: Theory and Applications*, Volume 23, Issue 2-3, pp175-192.
- [9] Di Paolo, E. A. (2001) Rhythmic and Non-rhythmic Attractors in Asynchronous Random Boolean Networks. *Biosystems*, 59(3):185–195
- [10] Ferreira, C. (2006) *Gene Expression Programming*. Springer.
- [11] Fogel, L. J., Owens, A.J. & Walsh, M.J. (1965) Artificial Intelligence Through A Simulation of Evolution. In M. Maxfield, A. Callahan & L.J. Fogel (Eds) *Biophysics and Cybernetic Systems: Proceedings of the 2nd Cybernetic Sciences Symposium*. Spartan Books, pp131-155.
- [12] Gershenson, C. (2002) Classification of Random Boolean Networks. In R.K. Standish, M. Bedau & H. Abbass (Eds) *Artificial Life VIII*. MIT Press, pp1-8.
- [13] Harvey, I. & Bossomaier, T. (1997) Time out of Joint: Attractors in Asynchronous Random Boolean Networks. In P. Husbands & I. Harvey (Eds) *Proceedings of the Fourth European Artificial Life Conference*. MIT Press, pp67-75.
- [14] Holland, J.H. (1975) *Adaptation in Natural and Artificial Systems*. University of Michigan Press.
- [15] Holland, J.H. (1976) Adaptation. In R. Rosen & F.M. Snell (Eds) *Progress in Theoretical Biology 4*. Plenum, pp263-293.
- [16] Howard, D., Bull, L. & Lanzi, P-L. (2008) Self-Adaptive Constructivism in Neural XCS and XCSF. In M. Keijzer et al. (eds) *GECCO-2008: Proceedings of the Genetic and Evolutionary Computation Conference*. ACM Press, pp1389-1396.
- [17] Ingerson, T & Buvel, R. (1984) Structure in Asynchronous Cellular Automata. *Physica D* 10 (1-2): 59-68.
- [18] Kauffman, S. A. (1993) *The Origins of Order: Self-Organization and Selection in Evolution*. Oxford.
- [19] Koza, J.R. (1992) *Genetic Programming*. MIT Press.
- [20] Lanzi, P-L & Perrucci, A. (1999) Extending the Representation of Classifier Conditions Part II: From Messy Coding to S-Expressions. In W. Banzhaf, J. Daida, A.E. Eiben, M.H. Garzon, V. Honavar, M. Jakiela & R.E. Smith (Eds) *Proceedings of the Genetic and Evolutionary Computation Conference – GECCO-99*. Morgan Kaufmann, pp345-352.
- [21] Lanzi, P-L & Wilson, S.W. (2000) Optimal Classifier System Performance in Non-Markov Environments. *Evolutionary Computation*, 8(4), pp393-418.
- [22] Mesot, B. & Teuscher, C. (2005) Deducing Local Rules for Solving Global Tasks with Random Boolean Networks. *Physica D* 211(1-2):88-106.
- [23] Miller, J. (1999) An Empirical Study of the Efficiency of Learning Boolean Functions using a Cartesian Genetic Programming Approach. In W. Banzhaf, J. Daida, A.E. Eiben, M.H. Garzon, V. Honavar, M. Jakiela & R.E. Smith (Eds) *Proceedings of the Genetic and Evolutionary Computation Conference – GECCO-99*. Morgan Kaufmann, pp1135-1142.
- [24] Mitchell, M., Hraber P., & Crutchfield J. (1993) Revisiting the Edge of Chaos: Evolving Cellular Automata to Perform Computations. *Complex Systems* 7: 83-130.
- [25] Packard, N. (1988) Adaptation Toward the Edge of Chaos. In J. Kelso, A. Mandell & M. Shlesinger (Eds) *Dynamic Patterns in Complex Systems*. World Scientific, pp293-301.
- [26] Pujol, J. & Poli, R. (1998) Efficient Evolution of Asymmetric Recurrent Neural Networks using a PDGP-inspired two-dimensional Representation. In W. Banzhaf, R. Poli, M. Schoenauer & T.C. Fogarty (Eds) *Proceedings of the First European Workshop on Genetic Programming*. Springer, pp130-141.
- [27] Schmidt, M. & Lipson, H. (2007) Comparison of Tree and Graph Encodings as Function of Problem Complexity. In *Proceedings of the Genetic and Evolutionary Computation Conference – GECCO-07*. ACM Press, pp1674-1679.
- [28] Sipper, M. (1997) *Evolution of Parallel Cellular Machines*. Springer.
- [29] Sipper, M. & Ruppim, E. (1997) Co-evolving architectures for cellular machines. *Physica D*, vol. 99, pp428-441.
- [30] Teller, A. & Veloso, M. (1996) Neural Programming and an Internal Reinforcement Policy. In J.R. Koza (Ed) *Late Breaking Papers at the Genetic Programming 1996 Conference*, Stanford University, pp186-192.
- [31] Valenzuela-Rendon, M. (1991) The Fuzzy Classifier System: a Classifier System for Continuously Varying Variables. In L. Booker & R. Belew (Eds) *Proceedings of the Fourth International Conference on Genetic Algorithms*. Morgan Kaufmann, pp346-353.
- [32] Von Neumann, J. (1966) *The Theory of Self-Reproducing Automata*. University of Illinois.
- [33] Werner, T. & Akella, V. (1997) Asynchronous processor survey. *Computer (USA)*, 30(11):67-76, November 1997.
- [34] Wilson, S.W. (1995) Classifier Fitness Based on Accuracy. *Evolutionary Comp.* 3: 149-175.
- [35] Wilson, S.W. (2000) Get Real! XCS with Continuous-Valued Inputs. In P-L. Lanzi, W. Stolzmann & S.W. Wilson (Eds) *Learning Classifier Systems: From Foundations to Applications*. Springer, pp209-222.
- [36] Wilson, S.W. (2008) Classifier Conditions using Gene Expression Programming. *Learning Classifier Systems. International Workshops, IWLCS 2006-2007, Revised Selected Papers. Lecture Notes in Artificial Intelligence (LNAI-4998)*, J. Bacardit, E. Bernado-Mansilla, and M. V. Butz (Eds). Berlin, Springer-Verlag, pp206-217.
- [37] Wuensche, A. (2004) Basins of attraction in network dynamics: A conceptual framework for biomolecular networks. In *Modularity in Development and Evolution (Eds.)*. G. Schlosser and G.P. Wagner. Chicago, University Press, pp288-311.