

8, rue du Capitaine Scott

75 015 Paris France

el.Landau,Olivier.Sigaud,Pierre.Gerard}@lifl.fr

<http://miriad.lip6.fr/~landau>

<http://animatlab.lip6.fr/~{sigaud,pgerard}>

Laboratoire d'Informatique Fondamentale de L

Cité Scientifique

59 655 Villeneuve d'Ascq Cedex, France

Sebastien.Picault@lifl.fr

<http://www.lifl.fr/~picault>

In this paper we present ATNoSFERES, a new
an indirect encoding Genetic Algorithm which
ata controllers able to deal with perceptual alia
our ongoing line of research, we compare it wi
sed extension of the most studied Learning Class
gh two benchmark experiments. We focus in p
te generalization, and add special purpose feat
to fulfill that comparison. We then discuss the
state generalization in the experiments studied

Evolutionary Algorithms, Learning Classifier Sys
state generalization, ATN¹

ion

Classifier Systems (LCS) [5] are used to tackl
ptive agents are involved in a sensori-moto
h agents perceive situations through their se
, each representing a perceived feature of the
s is to *learn* the optimal policy – *i.e.* which
, in order to fulfill their goals the best way t
ment Learning (RL) framework [20], the go
imize the scalar rewards it receives from its
by a set of rules – or classifiers – specifying th
e *conditions* concerning the perceived situati

“Augmented Transition Networks”

framework, explicit internal states were added to a pair of the classifiers, e.g. in XCSM [14, 2]. The additional information required to choose actions is non-Markov. The problem of properly setting this is addressed to *Genetic Algorithms* (GA).

Here we extend our first comparison presented in [11] to "ATNoSFERES". The latter is a new system that also models the behavior of agents facing problems in which they have to select actions based on a set of values or vectors of attributes, and have to select actions based on these values. As shown in [12] that such an evolutionary approach is more effective in complex environments; in ATNoSFERES, the goals are defined by a set of conditions (instead of classical LCS learning techniques). In this section, we present the features and properties of the ATNoSFERES system. It relies upon oriented, labeled graphs (§ 2.1) for the representation of the action selection procedure. The specificity of the ATNoSFERES system is a graph from a bitstring (§ 2.2) that can be handled by a Genetic Algorithm of a GA, with additional operators. Then, the ATNoSFERES representation is formally very close to LCS representation in XCSM (§ 3.1). We remind the results of our experiments presented in [12] (§4), and the comparison we made (§ 4.1) between the lack of internal state generalization in ATNoSFERES and the generality of the solutions that were found. A comparison between ATNoSFERES with and without a way to represent internal states (§ 6) allows us to discuss in detail the validity of the ATNoSFERES. In conclusion, we present further additions that could be added to ATNoSFERES to reach an even higher performance (§8).

2. Representation of the ATNoSFERES model

2.1. Representation of the ATNoSFERES model

The ATNoSFERES model is provided by the ATNoSFERES model [11, 18]. The ATNoSFERES model is basically an oriented, labeled graph with a set of nodes and edges. A node (or final) node (see figure 5). Nodes represent states of an automaton.


The ATNoSFERES binds conditions expressed as a set of nodes. The ATNoSFERES is endowed with the ability to generalize conditions. But in ATNoSFERES, the conditions and actions are defined by a set of nodes and edges.

step, the agent crosses an edge:
- updates the set of eligible edges among those starting from the current node.
- An edge is eligible when either it has no conditions on its label or the conditions on its label are simultaneously true.
- An action is randomly chosen in this set. If the set is empty, an action is chosen randomly over all possible actions, the current node is updated, and we do not perform the next two steps.
- The actions on the label of the current edge are sequentially performed on the system. Assuming that only one action can be performed at a time, the last action is actually performed. When the set of eligible edges is empty, an action is chosen randomly.
- The current node becomes the destination of the next edge.
- The process stops when it is at the *End* node (E). This process is a Markov chain model and may never be reached. This appears to be the case in the following experiments (since agents reaching the *End* node thus have a very low fitness, see § 4).

As described how the graphs are used, we now present the experimental results.

2.2. The graph-building process

The process of building the behaviors is built from a genotype-phenotype mapping. The basic structure containing only the *Start* and *End* nodes is used. Many different evolutionary techniques to automatically generate circuits [8], finite-state machines [2], neural networks [1], etc. Very roughly, we can sketch an opposition between two approaches that use the genotype as an encoding of a solution: Genetic Algorithms [5, 1, 3] or Evolutionary Strategies [11] approaches that use a single structure both as the phenotype and as the genotype, as Genetic Programming [7, 17], Evolutionary Programming [12], developmental program trees, e.g. [6, 4, 16], etc. In the FERES model [10], we try to conciliate advantages of both approaches: on the one hand, since the behavioral phenotype is a path in a graph, we want it to be of any complexity; on the other hand, a fine-grain genotype (a bitstring) to produce it. The search space of the solution space through “blind” genetic search is thus very large.



of the genetic expression we use to produce the graph (Figure 1). The string is first decoded into tokens (Figure 2). The second step is to interpret the tokens as instructions (b) to create nodes, edges, and connections.

Translation is a simple process that reads the genetic code and maps it to a sequence of *tokens* (symbols). It uses the mapping function $\mathcal{G} : \{0, 1\}^n \rightarrow \mathcal{T}$ ($|\mathcal{T}| \leq 2^n$) where \mathcal{T} is the set of tokens. The different roles of which will be described in the next section. The number of available tokens, the genetic code length, and the number of binary substrings of size n (decoded into tokens) are the main parameters of the translator.

Tokens are instructions of the ATNoSFER (Figure 2). They are interpreted one by one, while the graph is built in a stream produced by the translator. The interpreter operates on a stack in which parts of the future graph are stored. The construction of the graph takes place during this interpretation: nodes are created, edges and connections, and connecting them to the existing graph. This is similar to other stack-based languages (*e.g.* Forth, PostScript). The stack can also be directly accessed by some instructions. The interpreter works by other means that only push/pop operations. The interpreter works with a “blind” evolutionary process (i.e. no selection). In a fine-grain genotype, the graph built by the translator is the result of the evolution of the genetic code.

sequence of tokens is meaningful, the graph-building process can accommodate many variations affecting the genotype, thus the grammar is a syntactical constraint on the genetic operators. The grammar is to some extent order-independent and can be applied to very different genotypes.

resulting actions

actions that manipulate the stack

no action, the token is just discarded

swap the two first tokens

push a copy of the first action or condition token

delete the first action or condition token

push a copy of the first node token

delete the first node token^a

pop the token, and push it on the bottom of the stack

take the token from the bottom of the stack, and push it on top

actions that create nodes and connect them with edges

create a new node and push it

create an edge from the first to the second node

label the edge with the set of conditions token

of actions token until the second node,

delete the action and condition tokens that were used

create an edge from the Start node to the first node

label the edge with the set of conditions token

of actions token until the node,

delete the action and condition tokens that were used

create an edge from the first node token to the second node

label the edge with the set of conditions token

of actions token until the node,

delete the action and condition tokens that were used

actions and conditions tokens, specific to the action

push the condition on the stack

push the action on the stack

building language. Here “first” (node, action or condition) token encountered with the first (node, action or condition) token encountered with the first

possible other copies of the node still remain in the stack

be copies of the same node, so it would be a self-loop

nt abilities. Some of these tokens use token

(actions, conditions), that are specific to an a

These token are just pushed onto the stack.


n into an evolutionary framework

ception



ituation

[W ~SW S ~SE]




NE?
NE!

Matching



Selection



ample, the agent, located in a cell of the maze,

ocks in each of the eight surrounding cells. It ha

adjacent cells it should move. From its current

[W ~NW ~W ~SW S ~SE] (token E is true when the

state (node) of its graph, two edges (in bold) ar

their label match the perceptions. One is random

(east) is performed and the current state is upda

the ATNoSFERES model has been applied

to produce controllers for agents.

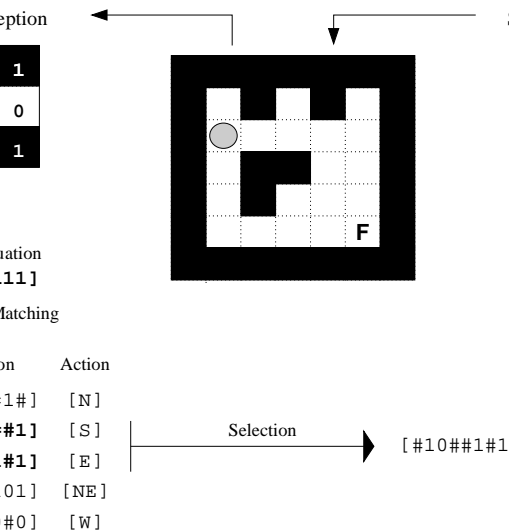
th agent has a bitstring genotype from which

c code depends on the perception abilities of

perform). The fitness of each agent is compu

of may change. Nodes or edges can in fact be any process, as can condition/action labels.

Classifier Systems



perceives the presence/absence (resp. 1/0) of black cells (considered clockwise, starting with the north cell), the agent perceives $[01010111]$. Within the list of classifiers, the LCS first selects those matching the current situation. In this case, the selected classifier is $[#10##1#1]$ and the corresponding action is performed.

In the introduction, the problems tackled by LCS are those where the situations are defined by several attributes of the environment. A LCS has to build a model of the system as shown in figure 3. Within this model, the “#” symbols in the condition parts of the classifiers *don't care* symbols make it possible to use the same classifier for several situations. Indeed, a *don't care* symbol in the condition part of a classifier means that the classifier is not sensitive to the value of the considered attribute.

FERES model, a Pittsburgh style LCS evolves a list of classifiers. In the *Michigan* style, the GA evolves a list of classifiers of a single agent. Here, this list is updated and modified. A fitness is associated to each classifier kept. Thus Michigan style LCS use GA to evolve classifiers that are improved during the life time of the agent. Utility functions that depend on scalar rewards are defined in the RL framework [20]. In early LCS [5], the fitness was defined directly on the classifier. After having defined a very simple fitness function, we found much more efficient to define the fitness on the utility prediction. The resulting system, called *Michigan* LCS, was used to solve Markov problems.


Simple Condition-Action classifiers does not ensure optimal behavior in perceptually aliased problems. It can happen that the current perception does not allow to always choose the optimal action: as soon as the perception changes in different states, it will choose the same action. This is not appropriate in some of these states (see figure 1). To solve these problems, it is necessary to introduce internal states. A simple way to probabilistically link classifiers in order to solve these problems. In contrast, Lanzi [14] proposed XCSM, where M is the number of XCS with explicit internal states. XCSM memory is composed of several bits that explicitly represent the internal state. The memory register provides XCSM with a list of internal perceptions. Thus, dealing with perceptually aliased information from the past experience of the agent, each classifier contains four parts (see table 6): an external condition, an internal condition about the internal state of the agent in the environment and an internal action. The external condition and the internal action contain a list of bits in the memory register. In order to be selected, the classifier must match with both the external and internal conditions. The LCS performs the corresponding action in the environment.

al states in XCSM and make ATNoSFERES
 ng. Thus it is natural to compare ATNoSFE
 ATN are characterized by several informatio
 classifiers: the source and destination nodes
 valent to the internal condition and the in
 ated to the edges correspond to the external
 ions associated to the edges correspond to th

Experiments

Conceptual aliasing problem

ose was to compare the evolutionary use
 respect to their ability to deal with non-Ma
 that comparison, we experimented our mo
 which [14] provides empirical results obtaine



10 environment. **F** represents the goal to reach
 of the maze; a few cells are unambiguous (S_i) b
 al situations may require either similar actions or
 $\{e,4\}$ but go south in S_{2_3})

Experimental setup

duce an experimental setup close to that use
 ent, with regards to the specificities of our m

ent involves the following steps:

population with $N = 300$ agents with random generation, build the graph of each agent and


individuals with higher fitness (namely, 20 % of new ones by crossing over the parents. The mutations (with a 1% rate) and insertions or deletions (with a 1% rate) on the bitstring of the offspring.

process with the new generation.

evaluate the individuals, they are put into the environment in the grid, and they have to find the food. The agent cannot perceive the food, and it can move in one step; when this action is incompatible with the environment (e.g., if the west cell contains an obstacle), it is simply not done (the agent stays on the same cell). Its fitness is calculated as $F + 2 * R$ (F : fitness for the run; D : number of cells visited during the run; K : time steps spent on exploring the environment before the food is found; R : remaining time if the food is not found within the time limit). Thus, the selection pressure encourages agents to find the food quickly. At each time step, the current cell is added to the explored cells (in order to compute D and K). The term R is added to the fitness of an already-known cell is still preferred to explore new cells. Each agent is evaluated 4 times starting from a random position. The total fitness is the sum of the fitnesses computed for all agents. With $B = 30$ and a 20 time steps limit, the experiments reported here were carried out on various environments of size 540 bits. The original population genotype space of the experiment has been bounded by 10,000 generations to reach high enough fitness values.

a behavioral graph obtained by the best individual in the experiment. It has also been represented in a L-shaped

environment. The graph described in figure 5 has the food source in the top-left corner, it first reaches the horizontal corridor



#	0	#	#	#	#	#
0	#	#	#	#	#	#
1	#	1	#	#	#	#
#	0	#	#	#	#	#
#	#	1	1	#	#	#

Fig. 6. A LCS-like representation graph on figure 5. E: external conditions, IC: internal conditions, IA: internal actions.

the best individual in a experiment

It goes straight to the food. This is a nearly clear distinction between the bottom of cells $S_{\{1,2\}_n}$, the top of vertical corridors (North corridor ($E \rightarrow S_8, S_{\{4,5\}_n}$) and the crucial ($\neg N \neg NW$).

Results of the first experiments

We presented a discussion resulting from the comparison of ATNoSFERES and XCSM. The main points we made were the following:

Representation: While XCSM produces a complex model in which the size of the external conditions part is chosen in advance, ATNoSFERES builds a graph where the number of nodes and labels on the edges can be minimal (that we focus on the best agent only). Hence, ATNoSFERES can be minimal while XCSM model cannot.

Learning and Classifier Selection: One important feature of ATNoSFERES is that the forces of classifier selection are more powerful than in XCSM. In order to remedy the fact that ATNoSFERES is not able to select the best classifier, it is necessary to include into the fitness function information about the actual behavior of the agent.

will be discussed next: ATNoSFERES cannot
on rules that can be fired whatever the internal
with an internal condition composed of “#

in XCSM, a # in the internal condition allows
the internal state represented by the memor
its action regardless of the internal state. Or
been chosen in those first experiments (see ta
dealing with a default behavior, since connect
between two nodes (i.e. two internal states). V
6.

Its given in [12] showed that the behavior ob
ES was not completely optimal, and that obt
require a major structural change in the graph w
antage.

Experiments

Addressing the need for state generalization

ded from the previous discussion by assuming
to deal with a default behavior, regardless of
age of XCSM over our model. Therefore we
add an internal state generalization property
graph building language, with a new *defaultSelfC*

Resulting actions

create nodes and connect them with edges

creates an edge from all the already present no
labels the edges with the set of conditions token
of actions token until the first node,
deletes the action and condition tokens that w

on to the graph building language (see table 1
node token encountered while going down the sta

that, since only the already created nodes will be used, the generalization depends on the time when the nodes are created. For example, if a *defaultSelfConnect* comes before the nodes are created, then the next nodes will not be created by this instruction.

In figure 7, self-connecting transitions on nodes are defined by int in internal condition and/or internal action. As the figure shows, if only the transition in node n is defined to a completely specified internal condition, and the transition is present in all nodes, it is equivalent to a completely specified internal condition and internal action.

Stick experimental setup

In experiments on Maze10, the best fitness was achieved by the theoretic fitness. Since these results are very consistent, they do not provide a large enough opportunity to evaluate the efficiency of the new encoding. Therefore, we used a maze (see figure 8) where the advantage of the new encoding should be more significant with respect to the old one.

The maze starts from the top cells of any of the vertical candles. The agent can only move to adjacent locations only, because we are focusing on the local search rather than searching for a general behavior to solve the maze. While going south along the “candles” from the top, if the agent finds an empty cell on the side, the agent can determine whether it is in the left or right part of the candle. This is done once it reaches the bottom of the candles. The agent can then move left or right and far-left and far-right candles.

In the state management strategy we have envisioned, the agent uses the following. The agent just needs one bit of information. If the agent sees an empty space on its left hand side, the bit represents whether it is in the left part or the right part of the candle. This information suffices to choose the right part of the candle. At the bottom of the candles. Given this internal state, the agent has the ability to generalize on the internal state value. The agent can move to all the $\{S_{a_i}\}$ cells. In all those cells, which represent the same conceptual conditions, the agent must go south, i.e. *if the agent is in the left part or the right part of the candle, it must go south*.

andlestick environment. \mathbf{F} represents the goal to
 res; 5 cells are unambiguous (S_i). In the other
 may require either similar actions or different o
 o east in $S_{g_ \{6,7,8,9,a\}}$)

NoSFERES, this means that it must follow the
 rnal state is.


f this property of the 12-Candlestick maze
 rnal state management strategy presented ab
 nce. Any other strategy implies that the agen
 choose the correct directions. Indeed, disambig
 right corners like the strategy obtained in Ma
 r with respect to the optimal path. But, as it
 of that paper, even if this optimal behavior
 using the *defaultSelfConnect* token, it is nev
 t.

ntal setup is similar to the one in our previo
 r the genetic encoding and the fact that th
 e empty cells. Like before, we need at least
 me in order to observe the influence of our
 encodings. In the first encoding, we use the n
 replaced by the *defaultSelfConnect* token.

evaluated 6 times starting in each of the 1
 ion is the one of Maze10 experiments. In
 a 30 time steps limit, the fitness is 6732. A
 made the experiments for 5 different initial
 0 and 540 bits). During the experiments, the
 lengths of the genotypes.

estick experiment results


respectively show the best solutions found
connect token. The fitness of both these indivi
 2. The missing points are lost when the age
 foot of the far-right candle (see figure 8, c



0	#	#	1	#	#	1
#	#	#	#	#	#	1
#	#	#	#	#	0	0
0	#	#	#	0	1	#
#	#	1	0	#	#	1
#	#	#	#	#	1	0
#	#	#	#	#	#	#

stick. Graph of the best
defaultSelfConnect to-
 appears 3 times, after the 2
 , so there are 3 identical
 ges for both nodes.

Fig. 10. A LCS-like re
 graph on figure 9. The t
 correspond to self-conn
 graph. EC, IC, EA, IA




EC						
E	NE	N	NW	W	SW	
#	#	#	#	#	#	
#	#	#	#	#	0	
#	#	#	#	#	#	
#	#	0	#	#	1	
#	#	1	0	#	1	
#	#	#	#	#	0	
0	0	#	#	0	#	
#	#	#	#	#	0	
#	#	1	#	#	#	
#	#	#	#	#	#	

stick. Graph of the best
 g the *defaultSelfConnect*

Fig. 12. A LCS-like re
 graph on figure 11. E
 figure 6.

Summary of 12-Candlestick results

ows, the average fitness obtained with the
 Candlestick experiment is significantly better
 that token. The difference between the perf



estick. Average fitness
self token





Fig. 14. Maze10. Average
without self token

to the different initial lengths of the bitstring. We can clearly see from these figures that the selective advantage to our agents.



estick. Fitness without
token





Fig. 16. 12-Candlestick
faultSelfConnect token

of making these experiments on the 12-Candlestick. The *faultSelfConnect* token is working well and property was designed. But, since the 12-Candlestick was designed for the use of the *defaultSelfConnect* token, it is not clear if or not this property can be generalized to other environments. We will present more significant results.

Maze10 experimental setup

The experimental setup is that of our previous Maze10 experiments. The only change is for the genetic encoding. Like in the previous

res 17 and 19. The best fitness for both genet
defaultSelfConnect token) are close to each o
 r a maximum of 4500. This best fitness was
 st, and only one time for the second, over th
 tic encoding. As in our previous Maze10 exp
 t when reaching the NE corner when coming
 her points are also lost when the agent goes
 ng through the cell on top of it, instead of



EC						
E	NE	N	NW	W	SW	S
1	#	0	#	#	#	#
0	#	#	#	#	0	#
#	1	#	#	#	#	#
0	0	#	#	0	0	#
1	#	#	#	#	#	#
0	#	0	#	#	#	#
0	#	#	#	#	#	#
#	#	0	#	#	0	#
1	#	#	#	1	#	#

Graph of the best individ-
defaultSelfConnect token for
 n appears 1 time, and is
 le only.

Fig. 18. A LCS-like re
 graph on figure 17. E
 figure 6.

Discussion

nd 14, it is clear that ATNoSFERES obtains b
 destick and on Maze10 with the *defaultSelfC*

behavior is never reached, however (the bes
 um). The best individual on the 12-Candles
 e *defaultSelfConnect* token as expected, so
 cified internal state. This result seems to s
 ding to which being able to generalize on th
 ve such behavioral problems, and results on
 rality of the assumption.

#	0	#	1	#	#	#
0	#	0	#	#	#	#
1	#	0	#	#	#	#
#	#	#	#	#	#	#
#	#	1	#	#	1	#

Fig. 20. A LCS-like re-graph on figure 19. EC figure 6.

Graph of the best individual-*defaultSelfConnect* token. 010 and 100 is actually

examination of the best individuals obtained that our assumption must be refined. In *size10*, presented in figure 17, uses the *defaultSelfConnect* token has been used many times to a transition in the graph, or none at all (if no node is interpreted).

in figure 7, since the interpretation of that connecting transitions only to the nodes already always result in the equivalent of full generation. In particular, when there is only one self-connect interpretation of the *defaultSelfConnect* token is interpreted as the internal condition part followed by an action.

however ATNoSFERES can obtain one self-connect token. But doing so requires that the conditions on the bitstring are fulfilled. Thus a self-connect can happen without the *defaultSelfConnect* token.

that ATNoSFERES gets a better performance than without it just because having self-connect is a special property and having that token significantly helps to have that property.

This more detailed study seems to reveal that it is necessary to solve a non-Markov problem, it is the possibility to specify (condition, action) transitions to change its internal state.

f a comparison between XCSM and ATNoS
aper the importance of the ability to represe
order to do so, we have introduced a new
a self-connecting transition to all the nodes
ve also presented a new maze experiment spe
is able to generalize on the internal state.
nts have shown that the performance of our
th this addition. But, while this result seem
being able to generalize on the internal sta
tive algorithms, a closer examination of wha
ments reveals that our *defaultSelfConnect* to
nt purpose than just generalizing on the in
ave another interesting property than the on

able to conclude more accurately on the rela
generalization property and the stable inter
s will be necessary. We believe that going i
son between XCSM and ATNoSFERES on th
will help identifying further what is really r
rior. In particular, we should try to assess th
on the internal condition and on the interna
o specialized tokens representing each prope

ONG, *An Analysis of the Behavior of a Class of*
thesis, Dept. of Computer and Communication
1975.

, A. J. OWENS, AND M. J. WALSH, *Artificial*
Evolution, John Wiley & Sons, 1966.

BERG, *Genetic Algorithms in Search, Optimization*
dison-Wesley, 1989.

Neural Network Synthesis Using Cellular Encodi
a.D. thesis, ENS Lyon – Université Lyon I, 1994

ND, *Adaptation in Natural and Artificial Systems*
Applications to Biology, Control, and Artificial I
an Press, Ann Arbor, MI, 1975.

. 151–170.

D. E. GOLDBERG, D. B. FOGEL, AND R. L. ...
Learning 1996: Proceedings of the First Annual C...
A, 1996, MIT Press.

AND S. PICAULT, *Stack-Based Gene Expression*
1, LIP6, Paris, 2002.

S. PICAULT, AND A. DROGOUL, *ATNoSFERES*
Behaviors, in Proceedings of the AISB'01 Symp
Multi-Agent Systems, 2001.

S. PICAULT, O. SIGAUD, AND P. GÉRARD, *A C*
S and XCSM, in Langdon et al. [13], pp. 926–93

N, E. CANTU-PAZ, K. MATHIAS, R. ROY, D
HNAN, V. HONAVAR, G. RUDOLPH, J. WEGENE
SCHULTZ, J. F. MILLER, E. BURKE, AND
of the Genetic and Evolutionary Computation C
ork, july 9-13 2002, Morgan Kaufmann.

An Analysis of the Memory Mechanism of XC
Genetic Programming Conference, 1998.

AYER, *Mathematical Models for Cellular Interacti*
], *Journal of theoretical biology*, 18 (1968).

L. SPECTOR, *Evolving Graphs and Networks u*
Report, in Koza [9], pp. 117–124.

NA, *Strongly Typed Genetic Programming*, in EV
1995.

ND S. LANDAU, *Ethogenetics and the Evolution*
Proceedings of the 5th World Multi-Conference
Informatics (SCI'01), N. Callaos, S. Esquivel,
pp. 528–533.

EFEL, *Evolution and Optimum Seeking*, John W

N AND A. G. BARTO, *Reinforcement Learning*, a
idge, MA, 1998.

ON AND L. BULL, *CXCS*, in *Learning Classifier S*
Applications, P. Lanzi, W. Stolzmann, and S. Wi
lberg, 2000, pp. 194–208.

N, *ZCS, a Zeroth level Classifier System*, *Evolutio*
1–18.

N, *Classifier Fitness Based on Accuracy*, *Evolutio*
149–175.

DS, *Transition Networks Grammars for Natural*
ions of the Association for the Computational M

ving Artificial Neural Networks, Proceedings of t