

An Evolutionary Algorithm that Constructs Recurrent Neural Networks

Peter J. Angeline, Gregory M. Saunders, and Jordan B. Pollack

Abstract—Standard methods for simultaneously inducing the structure and weights of recurrent neural networks limit every task to an assumed class of architectures. Such a simplification is necessary since the interactions between network structure and function are not well understood. *Evolutionary computations*, which include *genetic algorithms* and *evolutionary programming*, are population-based search methods that have shown promise in many similarly complex tasks. This paper argues that genetic algorithms are inappropriate for network acquisition and describes an evolutionary program, called GNARL, that simultaneously acquires both the structure and weights for recurrent networks. GNARL's empirical acquisition method allows for the emergence of complex behaviors and topologies that are potentially excluded by the artificial architectural constraints imposed in standard network induction methods.

I. INTRODUCTION

IN ITS complete form, network induction entails both *parametric* and *structural* learning [1]; i.e., learning both weight values and an appropriate topology of nodes and links. Current connectionist methods to solve this task fall into two broad categories. *Constructive* algorithms initially assume a simple network and add nodes and links as warranted [2]–[8], while *destructive* methods start with a large network and prune off superfluous components [9]–[12]. Though these algorithms address the complete problem of network acquisition, they do so in a highly constrained manner. Generally, constructive and destructive methods limit the available architectures in some way. In some of these methods, once an architecture has been explored and determined to be insufficient, a new architecture is adopted and the old becomes topologically unreachable. Others use only a single predefined structural modification, such as “add a fully connected hidden unit,” to generate successive topologies. Such *structural hill climbing* methods are susceptible to becoming trapped at structural local optima, which places the burden of task induction mostly on the identification of suitable parametric values rather than distributing the burden evenly. In addition, constructive and destructive algorithms investigate only restricted topological subsets rather than the complete class of network architectures. For example, Ash [2] allows only feedforward networks; Fahlman [6] assumes a restricted form of recurrence, and Chen *et al.* [7] explore only fully connected topologies. As a consequence, these algorithms tend to force a task into an

assumed architectural class rather than fitting an appropriate architecture to the task.

The deficiencies of constructive and destructive methods stem from inadequate methods for assigning credit to structural components of a network. The constrained topologies are assumed to limit the complexity of structural and parametric interactions and increase the likelihood of finding a sufficient network to solve the problem. Ideally, such limiting constraints should arise from solving the task rather than be implicit in the algorithm.

This paper presents GNARL, a network induction algorithm that simultaneously acquires both network topology and weight values while making minimal architectural restrictions and avoiding structural hill climbing. The algorithm, described in Section III, is an instance of evolutionary programming [13], [14], a class of evolutionary computations that has been shown to perform well on complex tasks. Section II argues that this class of evolutionary computation is better suited for evolving neural networks than genetic algorithms [15], [16], a more popular variety of evolutionary computations. Finally, Section IV demonstrates GNARL's ability to create recurrent networks for a variety of problems of interest.

II. EVOLVING CONNECTIONIST NETWORKS

Evolutionary computations are a promising collection of algorithms that hold promise for structural and parametric learning of recurrent networks [17]. These algorithms are distinguished by their reliance on a *population* of search space positions, rather than a single position, to locate extrema of a function defined over the search space. During one search cycle, or *generation*, the members of the population are ranked according to a *fitness function*, and those with higher fitness are probabilistically selected to become *parents* in the next generation. New population members, called *offspring*, are created using specialized *reproduction heuristics*. Using the population, reproduction heuristics, and fitness function, evolutionary computations implement nonmonotonic search methods that perform well in complex multimodal environments. Subclasses of evolutionary computations can be distinguished by examining the specific reproduction heuristics employed.

Genetic algorithms (GA's) [15], [16] are a popular form of evolutionary computation that rely chiefly on the reproduction heuristic of *crossover*.¹ This operator forms offspring by

Manuscript received December 23, 1992; revised June 28, 1993.

The authors are with the Laboratory for Artificial Intelligence Research, Computer and Information Science Department, The Ohio State University, Columbus, OH 43210.

IEEE Log Number 9213546.

¹ Genetic algorithms also employ other operators to manipulate the population, including a form of mutation, but their distinguishing feature is a heavy reliance on crossover.




Fig. 1. The dual representation scheme used in genetic algorithms. The interpretation function maps between the elements in recombination space on which the search is performed and the subset of structures that can be evaluated as potential task solutions.

recombining representational components from two members of the population without regard to content. Such a purely structural approach to creating novel population members assumes that components of all parent representations may be freely exchanged without inhibiting the search process.

Various combinations of GA's and connectionist networks have been investigated. Much research concentrates on the acquisition of parameters for fixed network architectures [18]–[21]. Other work allows a variable topology, but disassociates structure acquisition from acquisition of weight values by interweaving a GA search for network topology with a traditional parametric training algorithm [22], [23]. Some studies attempt to coevolve both the topology and weight values within the GA framework, but as in the connectionist systems described above, the network architectures are restricted [24]–[26]. In spite of this collection of studies, current theory from both genetic algorithms and connectionism suggests that GA's are not well-suited for evolving networks. In the following section, the reasons for this mismatch are discussed.

A. Evolving Networks With Genetic Algorithms

Genetic algorithms create new individuals by recombining the representational components of two members of the population. Because of this commitment to structural recombination, GA's typically rely on two distinct representational spaces (Fig. 1). *Recombination space*, usually defined over a set of fixed-length binary strings, is the set of structures to which the genetic operators are applied. It is in this space that the search actually occurs. *Evaluation space*, typically involving a problem-dependent representation, is the set of task-dependent structures that are evaluated on their ability to perform the desired task. When using GA's to evolve networks, the evaluation space is a set of networks. An *interpretation function* maps between these two representational spaces. Any set of finite-length bit strings cannot represent all possible networks, thus the evaluation space is restricted to a predetermined set. By design, the dual representation scheme allows the GA to crossover bit strings without any knowledge of their interpretation as networks. The implicit assumption is that the interpretation function will be defined so that bit strings created by the dynamics of the GA in recombination space will map to successively better networks in evaluation space.

The dual representation of GA's is an important feature for searching in certain environments. For instance, when it is unclear how to search the evaluation space directly, and




Fig. 2. The competing conventions problem [29]. Bit strings A and B map to structurally and computationally equivalent networks that assign the hidden units in different orders. Because the bit strings are distinct, crossover is likely to produce an offspring that contains multiple copies of the same hidden node, yielding a network with less computational ability than either parent.

when there exists an interpretation function such that searching the space of bit strings by crossover leads to good points in evaluation space, then the dual representation is ideal. It is unclear, however, that there exists an interpretation function that makes GA's beneficial for evolving neural networks. Clearly, the choice of interpretation function introduces a strong bias into the search, typically by excluding many potentially interesting and useful networks (another example of forcing the task into an architecture). Moreover, the benefits of having a dual representation hinge on crossover being an appropriate evolutionary operator for the task for some particular interpretation function; otherwise, the need to translate between the dual representations is an unnecessary complication.

Characterizing tasks for which crossover is a beneficial operator is an open question. Current theory suggests that crossover will tend to recombine short, connected substrings of the bit string representation that correspond to above-average task solutions when evaluated [15], [16]. These substrings are called *building blocks*, making explicit the intuition that larger structures with high fitness are built out of smaller structures with moderate fitness. Crossover tends to be most effective in environments where the fitness of a member of the population is reasonably correlated with the expected ability of its representational components [27]. Environments where this is not true are called *deceptive* [28].

There are three forms of deception when using crossover to evolve connectionist networks. The first involves networks that share both a common topology and common weights. Because the interpretation function may be many-to-one, identical networks need not have the same bit string representation (see Fig. 2). Given two such networks, crossover will tend to create offspring that contain repeated components rather than the full complement of either parent's hidden units. The resulting networks will necessarily perform worse than their parents because they do not possess key computational components for the task. Schaffer *et al.* [29] term this the *competing conventions* problem, and point out that the number of competing conventions grows exponentially with the number of hidden units.

The second form of deception involves two networks with identical topologies but different weights. It is well known that for a given task, a single connectionist topology affords

multiple solutions for a task, each implemented by a unique *distributed representation* spread across the hidden units [30], [31]. While the removal of a small number of nodes has been shown to effect only minor alterations in the performance of a trained network [30], [31], the computational role each node plays in the overall representation of the task solution is determined purely by the presence and strengths of its interconnections. Furthermore, there need be no correlation between distinct distributed representations over a particular network architecture for a given task. This seriously reduces the chance that an arbitrary crossover operation between distinct distributed representations will construct viable offspring *regardless of the interpretation function used*.

Finally, deception can occur when the parents differ topologically. The types of distributed representations that can develop in a network vary widely with the number of hidden units and the network's connectivity. Thus, the distributed representations of topologically distinct networks have a greater chance of being incompatible for recombination into viable offspring.

In short, for crossover to be an appropriate operator for evolving networks, the interpretation function must somehow compensate for all the types of deceptiveness described above. This suggests that the complexity of an appropriate interpretation function may more than rival the complexity of the original learning problem. Thus, the prospect of evolving connectionist networks with crossover appears limited in general, and better results should be expected with representations and reproduction heuristics that respect the uniqueness of a network's distributed representation. Part of this point has been tacitly validated in the genetic algorithm literature by the appearance of non-binary representations for evolving networks (e.g., [32], [33]). Crossover, however, is still commonplace.

B. Networks and Evolutionary Programming

Evolutionary programming (EP) [14], [34] is another form of evolutionary computation more suited to complete network induction. EP systems assume representations that are more natural for the task rather than relying on a singular, general representation as in GA's. Once an appropriate representation is chosen, representation-dependent mutation operators are defined that create offspring within a specific behavioral locus of the parent (see Fig. 3). EP commits to mutation as the sole reproductive operator for searching over a solution space. Recombination operators are not used. Evolutionary programming is preferable to genetic algorithms when there is no sufficient calculus to guide recombination or when separating the search and evaluation spaces does not afford an advantage.

Relatively few previous EP systems have addressed the problem of evolving connectionist networks. Fogel *et al.* [35] investigate training feedforward networks on some classic connectionist problems. McDonnell and Waagen [36] use EP to evolve the connectivity of feedforward networks with a constant number of hidden units by evolving both a weight matrix and a connectivity matrix. Fogel [14], [37] uses EP to induce three-layer fully-connected feedforward networks with




Fig. 3. The evolutionary programming approach to modeling evolution. Unlike genetic algorithms, evolutionary programs perform search in the space of networks. Offspring created by mutation remain within a locus of similarity to their parents.

a variable number of hidden units that employ good strategies for playing tic-tac-toe.

In each of the above studies, the mutation operator alters the parameters of network η by the function:

$$w = w + N(0, \alpha \varepsilon(\eta)) \quad \forall w \in \eta \quad (1)$$

where w is a weight, $\varepsilon(\eta)$ is the error of the network on the task (typically the mean squared error), α is a user-defined proportionality constant, and $N(\mu, \sigma^2)$ is a gaussian variable with mean μ and variance σ^2 . The implementations of structural mutations in these studies differ somewhat. McDonnell and Waagen [36] randomly select a set of positions in the connectivity matrix for feedforward networks and toggle the associated link's presence in the network with a probability based on the variance of the incident nodes' activation over the training set. The structural mutation used by Fogel [14], [37] adds or deletes a single fully connected hidden unit with equal probability. While these methods address complete network induction they too suffer from limiting architectural assumptions.

Evolutionary programming offers distinct advantages over genetic algorithms when evolving networks. First, EP manipulates networks directly, thus obviating the need for a dual representation and a problematic interpretation function. Second, by avoiding recombination between networks when creating offspring, the individuality of each network's distributed representation is respected.

III. THE GNARL ALGORITHM

GNARL, which stands for *GeNeralized Acquisition of Recurrent Links*, is an evolutionary algorithm that nonmonotonically constructs recurrent networks to solve a given task. The name GNARL reflects the types of networks that arise from a generalized network induction algorithm performing both structural and parametric learning. Instead of having uniform or symmetric topologies, the resulting networks have "gnarled" interconnections of hidden units that more accurately reflect constraints inherent in the task.

The general architecture of a GNARL network is straightforward. The input and output nodes are considered to be provided by the task and are immutable by the algorithm; thus each network for a given task always has m_{in} input nodes and m_{out} output nodes. The number of hidden nodes varies from 0 to a user-supplied maximum h_{max} . Bias is optional; if provided in an experiment, it is implemented as an additional input node with constant value one. All non-input nodes employ the




Fig. 4. Sample initial network. The number of input nodes (m_{in}) and number of output nodes (m_{out}) is fixed for a given task. The presence of a bias node ($b = 0$ or 1) as well as the maximum number of hidden units (h_{max}) is set by the user. The initial connectivity is chosen randomly (as indicated in the text). The disconnected hidden node does not affect this particular network's computation, but is available as a resource for structural mutations.

standard sigmoid activation function. Links use real-valued weights, and must obey three restrictions:

- R_1 : There can be no links to an input node.
- R_2 : There can be no links from an output node.
- R_3 : Given two nodes x and y , there is at most one link from x to y .

Thus GNARL networks may have no connections, sparse connections, or full connectivity. Consequently, GNARL's search space is:

$$S = \{\eta: \begin{array}{l} \eta \text{ is a network with real-valued weights,} \\ \eta \text{ satisfies } R_1 - R_3, \\ \eta \text{ has } m_{in} + b \text{ input nodes, where } b = 1 \text{ if a bias} \\ \quad \text{node is provided, and } 0 \text{ otherwise,} \\ \eta \text{ has } m_{out} \text{ output nodes,} \\ \eta \text{ has } i \text{ hidden nodes, } 0 \leq i \leq h_{max} \end{array}\}$$

$R_1 - R_3$ are strictly implementational constraints. Nothing in the algorithm described below hinges on S being pruned by these restrictions.

A. Selection, Reproduction, and Mutation of Networks

GNARL initializes the population with randomly generated networks (see Fig. 4). The number of hidden nodes for each network is chosen from a uniform distribution over a user-supplied range. The number of initial links is chosen similarly from a second user-supplied range. The incident nodes for each link are chosen in accordance with the structural mutations described below. Once a topology has been chosen, all links are assigned random weights, selected uniformly from the range $[-1, 1]$. There is nothing in this initialization procedure that forces a node to have any incident links, let alone for a path to exist between the input and output nodes. In the experiments below, the number of hidden units for a network in the initial population was selected uniformly between one and five, and the number of initial links varied uniformly between one and 10.

In each generation of search, the networks are first evaluated by a user-supplied fitness function $f: S \rightarrow \mathbf{R}$, where \mathbf{R} represents the reals. Networks scoring in the top 50% are designated as the *parents* of the next generation; all other networks are discarded. This selection method is used in many EP algorithms, although competitive methods of selection have also been investigated [14].

Generating an offspring involves three steps: copying the parent, determining the severity of the mutations to be performed, and finally mutating the copy. Network mutations are separated into two classes, corresponding with the types of

learning discussed in [1]. *Parametric mutations* alter the value of parameters (link weights) currently in the network, whereas *structural mutations* alter the number of hidden nodes and the presence of links in the network, thus altering the space of parameters.

1) *Severity of Mutations*: The severity of a mutation to a given parent, η , is dictated by that network's temperature, $T(\eta)$:

$$T(\eta) = 1 - \frac{f(\eta)}{f_{max}} \quad (2)$$

where f_{max} is the maximum fitness for a given task. Thus, the temperature of a network is determined by how close the network is to being a solution for the task. This measure of the network's performance is used to anneal the structural and parametric similarity between parent and offspring, so that networks with a high temperature are mutated severely, and those with a low temperature are mutated only slightly. This allows a coarse-grained search initially, and a progressively finer-grained search as a network approaches a solution to the task. $T(\eta)$ is related to the concept of temperature in simulated annealing [38] where a higher temperature indirectly increases the variety of states that can be visited by the system.

2) *Parametric Mutation of Networks*: Parametric mutations are accomplished by perturbing each weight w of a network η with gaussian noise, a method motivated by [14], [37]. In that body of work, weights are modified as follows:

$$w = w + N(0, \alpha T(\eta)) \quad \forall w \in \eta \quad (3)$$

where α is a user-defined proportionality constant, and $N(\mu, \sigma^2)$ is a gaussian random variable, as before. While large parametric mutations are occasionally necessary to escape parametric local minima during search, it is more likely they will adversely affect the offspring's ability to perform better than its parent. To compensate, GNARL updates weights using a variant of (3). First, the *instantaneous temperature* \hat{T} of the network is computed:

$$\hat{T}(\eta) = U(0, 1)T(\eta) \quad (4)$$

where $U(0, 1)$ is a uniform random variable over the interval $[0, 1]$. This new temperature, varying from 0 to $T(\eta)$, is then substituted into (3):

$$w = w + N(0, \alpha \hat{T}(\eta)) \quad \forall w \in \eta \quad (5)$$

In essence, this modification lessens the frequency of large parametric mutations without disallowing them completely. In the experiments described below, α is one.

3) *Structural Mutation of Networks*: The structural mutations used by GNARL alter the number of hidden nodes and the connectivity between all nodes, subject to restrictions $R_1 - R_3$ discussed earlier. To avoid radical jumps in fitness from parent to offspring, structural mutations attempt to preserve the behavior of a network. For instance, new links are initialized with zero weight, leaving the behavior of the modified network unchanged. Similarly, hidden units are added to the network without any incident connections. Links must be added by future structural mutations to

determine how to incorporate the new computational unit. Unfortunately, achieving this behavioral continuity between parent and child is not so simple when removing a hidden node or link. Consequently, the deletion of a node involves the complete removal of the node and all incident links with no further modification to compensate for the behavioral change. Similarly, deleting a link removes that parameter from the network.

The selection of which node to remove is uniform over the collection of hidden nodes. Addition or deletion of a link is slightly more complicated. A parameter identifies the likelihood that a link to be added or deleted will be incident with an input or output node. Biasing the link selection process in this way is necessary when there is a large differential between the number of hidden nodes and the number of input or output nodes. This parameter was set to 0.2 in the experiments described in the next section.

Research in [14] and [37] uses the heuristic of adding or deleting at most a single fully connected node per structural mutation. Therefore, it is possible for this method to become trapped at a structural local minima for an indefinite time, although this is less probable than in nonevolutionary algorithms given that several topologies may be present in the population. In order to more effectively search the range of network architectures, GNARL uses a severity of mutation for each separate structural mutation. A unique user-defined interval specifying a range of modification is associated with each of the four structural mutations. Given an interval of $[\Delta_{\min}, \Delta_{\max}]$ for a particular structural mutation, the number of modifications of this type made to an offspring is given by:

$$\Delta_{\min} + [U[0, 1]T(\eta)(\Delta_{\max} - \Delta_{\min})] \quad (6)$$

Thus the number of modifications varies uniformly over a shrinking interval based on the parent network's fitness. In the experiments below, the maximum number of nodes added or deleted was three, while the maximum number of links added or deleted was five. The minimum number for each interval was always one.

B. Fitness of a Network

In evolving networks to perform a task, GNARL does not require an explicit target vector—all that is needed is the feedback given by the fitness function f . But if such a vector is present, as in supervised learning, there are many ways of transforming it into a measure of fitness. For example, given a training set $\{(x_1, y_1), (x_2, y_2), \dots\}$, three possible measures of fitness for a network η are sum of square errors (7), sum of absolute errors (8), and sum of exponential absolute errors (9):

$$\sum_i (y_i - \text{Out}(\eta, x_i))^2 \quad (7)$$

$$\sum_i |y_i - \text{Out}(\eta, x_i)| \quad (8)$$

$$\sum_i e^{|(y_i - \text{Out}(\eta, x_i))|} \quad (9)$$

Furthermore, because GNARL explores the space of networks by mutation and selection, the choice of fitness function does




Fig. 5. An FSA that defines the *enable-trigger* task [39]. The system is given a data stream of bit pairs $\{(a_1, b_1), (a_2, b_2), \dots\}$, and produces an output of 0's and 1's. To capture this system's input/output behavior, a connectionist network must learn to store state indefinitely.

not alter the mechanics of the algorithm. To show GNARL's flexibility, each of these fitness functions will be demonstrated in the experiments below.

IV. EXPERIMENTS

In this section, GNARL is applied to several problems of interest. The goal in this section is to demonstrate the abilities of the algorithm on problems from language induction to search and collection. The various parameter values for the program are set as described above, unless otherwise noted.

A. Williams' Trigger Problem

As an initial test, GNARL induced a solution for the *enable-trigger* task proposed in [39]. Consider the finite state generator shown in Fig. 5. At each time step the system receives two input bits, (a, b) , representing "enable" and "trigger" signals, respectively. This system begins in state S_1 , and switches to state S_2 only when enabled by $a = 1$. The system remains in S_2 until it is triggered by $b = 1$, at which point it outputs 1 and resets the state to S_1 . So, for instance, on an input stream $\{(0, 0), (0, 1), (1, 1), (0, 1)\}$, the system will output $\{0, 0, 0, 1\}$ and end in S_1 . This simple problem allows an indefinite amount of time to pass between the enable and the trigger inputs; thus no finite length sample of the output stream will indicate the current state of the system. This forces GNARL to develop networks that can preserve state information indefinitely.

The fitness function used in this experiment was the sum of exponential absolute errors (9). Population size was 50 networks with the maximum number of hidden units restricted to six. A bias node was provided in each network in this initial experiment, ensuring that an activation value of 1 was always available. Note that this does not imply that each node had a nonzero bias; links to the bias node had to be acquired by structural mutation.

Training began with all two input strings of length two, shown in Table I. After 118 generations (3000 network evaluations²), GNARL evolved a network that solved this task for the strings in Table I within a tolerance of 0.3 on the output units. The training set was then increased to include all 64 input strings of length three and evolution of the networks was allowed to continue. After an additional 422 generations, GNARL once again found a suitable network. At this point, the difficulty of the task was increased a final time by training on all 256 strings of length four. After another 225 generations

²Number of networks evaluated in a run is $\text{popsize} + (0.5 * \text{generations} * \text{popsize})$, giving 3000 network evaluations for this trial.

TABLE I
INITIAL TRAINING DATA FOR ENABLE-TRIGGER TASK

Input	Target output	Input	Target output
{(0, 0), (0, 1)}	{0,0}	{(1, 0), (0, 0)}	{0, 0}
{(0, 0), (1, 0)}	{0,0}	{(1, 0), (0, 1)}	{0, 1}
{(0, 0), (1, 1)}	{0,0}	{(1, 0), (1, 0)}	{0, 0}
{(0, 1), (0, 0)}	{0, 0}	{(1, 0), (1, 1)}	{0, 1}
{(0, 1), (0, 0)}	{0,0}	{(1, 1), (0, 0)}	{0, 0}
{(0, 1), (0, 1)}	{0,0}	{(1, 1), (0, 1)}	{0, 1}
{(0, 1), (1, 0)}	{0,0}	{(1, 1), (1, 0)}	{0, 0}




Fig. 6. Connectivity of two recurrent networks found in the enable-trigger experiment. (a) The best network of generation 1. (b) The best network of generation 765. This network solves the task for all strings of length eight.

(~ 20000 network evaluations in total) GNARL once again found a network to solve this task, shown in Fig. 6b. Note that there are two completely isolated nodes. Given the fitness function used in this experiment, the two isolated nodes do not effect the network's viability. To investigate the generalization of this network, it was tested over all 4096 unique strings of length six. The outputs were rounded off to the nearest integer, testing only the network's separation of the strings. The network performed correctly on 99.5% of this novel set, generating incorrect responses for only 20 strings.

Figure 7 shows the connectivity of the population member with the best fitness for each generation over the course of the run. Initially, the best network is sparsely-connected and remains sparsely-connected throughout most of the run. At about generation 400, the size and connectivity increases dramatically only to be overtaken by the relatively sparse architecture shown in Fig. 6(b) on the final generation. Apparently, this sparsely connected network evolved more quickly than the full architectures that were best in earlier generations. The oscillations between different network architectures throughout the run reflects the development of such competing architectures in the population.

B. Inducing Regular Languages

A current topic of research in the connectionist community is the induction of finite state automata (FSA's) by networks with second-order recurrent connections. For instance, Pollack [40] trains sequential cascaded networks (SCN's) over a test set of languages, provided in [41] and shown in Table II, using a variation of backpropagation. An interesting result of this work is that the number of states used by the network to implement finite state behavior is potentially infinite. Other studies using the training sets in [41] have investigated various network architectures and training methods, as well as algorithms for extracting FSA's from the trained architectures [42]–[45].

An explicit collection of positive and negative examples, shown in Table III, that pose specific difficulties for inducing




Fig. 7. Different network topologies explored by GNARL during the first 540 generations on the enable-trigger problem. The presence of a link between node i and j at generation g is indicated by a dot at position $(g, 10*i + j)$ in the graph. Note that because node 3 is the output node, there are no connections from it throughout the run. The arrow designates the point of transition between the first two training sets.

TABLE II
REGULAR LANGUAGES TO BE INTRODUCED

Language	Description
1	1*
2	(1 0)
3	No odd length 0 strings anytime after an odd length 1 string
4	No more than two 0's in a row
5	An even number of 10's and 01's, pairwise
6	(Number of 1's-number of 0's) mod 3 = 0
7	0*1*0*1*

the intended languages is offered in [41]. Notice that the training sets are unbalanced, incomplete, and vary widely in their ability to strictly define the intended regular language. GNARL's ability to learn and generalize from these training sets was compared against the training results reported for the second-order architecture used in [42]. Notice that all the languages in Table II require recurrent connections in order to induce the language completely. The type of recurrence needed for each language varies. For instance, languages 1 through 4 require an incorrect input be remembered indefinitely, forcing the network to develop an analog version of a trap state. Networks for language 6, however, must parse and count individual inputs, potentially changing state from accept to reject or vice versa on each successive input.

The results obtained in [42] are summarized in Table IV. The table shows the number of networks evaluated to learn the training set and the accuracy of generalization for the learned network to the intended regular language. Accuracy is measured as the percentage of strings of length 10 or less that are correctly classified by the network. For comparison,

TABLE III
TRAINING SETS FOR THE LANGUAGES OF TABLE II FROM [41]

Language	Positive instances	Negative instances
1	ϵ , 1, 11, 111, 1111, 11111, 111111, 1111111, 11111111	0, 10, 01, 00, 011, 110, 000, 11111110, 10111111
2	ϵ , 10, 1010, 101010, 10101010, 101010101010	1, 0, 11, 00, 01, 101, 100, 1001010, 10110, 110101010
3	ϵ , 1, 0, 01, 11, 00, 100, 110, 111, 000, 100100, 110000011100001, 111101100010011100	10, 101, 010, 110, 1011, 10001, 111010, 1001000, 11111000, 0111001101, 11011100110 000, 11000 0001 000000000, 00000, 0000, 11111000011, 1101010000010111 1101010000010111, 1010010001
4	ϵ , 1, 0, 10, 01, 00, 100100, 001111110100, 0100100100, 11100, 010	1, 0, 111, 010, 000000000, 1000, 01, 10, 1110010100, 01011111110, 0001, 011
5	ϵ , 11, 00, 001, 0101, 1010, 1000111101, 1001100001111010, 111111, 0000	1, 0, 11, 00, 101, 011, 11001, 1111, 00000000, 010111, 1011101111, 1001001001
6	ϵ , 10, 01, 1100, 111, 000000, 0111101111, 100100100	1010, 00110011000, 0101010101, 1011010, 10101, 010100, 101001, 100100110101
7	ϵ , 1, 0, 10, 01, 11111, 000, 00110011, 0101, 0000100001111, 00100, 01111011111, 00	

TABLE IV
SPEED AND GENERALIZATION RESULTS REPORTED
BY [42] FOR LEARNING THE DATA SETS OF TABLE III

Language	Average evaluations	Average % accuracy	Fewest evaluations	Best % accuracy
1	3033.8	88.98	28	100.0
2	4522.6	91.18	807	100.0
3	12326.8	64.87	442	78.31
4	4393.2	42.50	60	60.92
5	1587.2	44.94	368	66.83
6	2137.6	23.19	306	46.21
7	2969.0	36.97	373	55.74

the table lists both the average and best performance of the five runs reported in [42].

This experiment used a population of 50 networks, each limited to at most eight hidden units. Each run lasted at most 1000 generations, allowing a maximum of 25050 networks to be evaluated for a single data set. Two experiments were run for each data set, one using the sum of absolute errors (SAE) and the other using sum of square errors (SSE). The error for a particular string was computed only for the final output of the network after the entire string plus three trailing "null" symbols had been entered, one input per time step. The concatenation of the trailing null symbols was used to identify the end of the string and allow input of the null string, a method also used in [42]. Each network had a single input and output and no bias node was provided. The three possible logical inputs for this task, 0, 1, and null, were represented by activations of -1, 1, and 0, respectively. The tolerance for the output value was 0.1, as in [42].

TABLE V
SPEED AND GENERALIZATION RESULTS FOR GNARL TO TRAIN
RECURRENT NETWORKS TO RECOGNIZE THE DATA SETS OF TABLE III

Language	Evaluations (SAE)	% accuracy (SAE)	Evaluations (SSE)	% accuracy (SSE)
1	3975	100.00	5300	99.27
2	5400	96.34	13975	73.33
3	25050†	58.87	18650	68.00
4	15775	92.57*	21850	57.15
5	25050†	49.39	22325	51.25
6	21475	55.59*	25050†	44.11
7	12200	71.37*	25050†	31.46




Fig. 8. The number of network evaluations required to learn the seven datasets of Table 3. GNARL (using both SAE and SSE fitness measures) compared to the average number of evaluations for the five runs described in [42].

Table V shows, for both fitness functions, the number of evaluations until convergence and the accuracy of the best evolved network. Only four of the runs, each of those denoted by a '+' in the table, failed to produce a network with the specified tolerance in the allotted 1000 generations. In the runs using SAE, the two runs that did not converge had not separated a few elements of the associated training set and appeared to be far from discovering a network that could correctly classify the complete training set. Both of the uncompleted runs using SSE successfully separated the data sets but had not done so to the 0.1 tolerance within the 1000 generation limit. Fig. 8 compares the number of evaluations by GNARL to the average number of evaluations reported in [42]. As the graph shows, GNARL consistently evaluates more networks, but not a disproportionate number. Considering that the space of networks being searched by GNARL is much larger than the space being searched by [42], these numbers appear to be within a tolerable increase.

The graph of Fig. 9 compares the accuracy of the GNARL networks to the average accuracy found in [42] over five runs. The GNARL networks consistently exceeded the average accuracy found in [42].

These results demonstrate GNARL's ability to simultaneously acquire the topology and weights of recurrent networks, and that this can be done within a comparable number of network evaluations as training a network with static architecture on the same task. GNARL also appears to generalize better consistently, possibly due to its selective inclusion and exclusion of some links.




Fig. 9. Percentage accuracy of evolved networks on languages in Table II. GNARL (using SAE and SSE fitness measures) compared to average accuracy of the five runs in [42].




Fig. 10. The ant problem. The trail is connected initially, but becomes progressively more difficult to follow. The underlying 2-D grid is toroidal, so that position "A" is the first break in the trail—it is simple to reach this point. Positions "B" and "C" indicate the only two positions along the trail where the ant discovered in run 1 behaves differently from the 5-state FSA of [46] (see Fig. 13).

C. The Ant Problem

GNARL was tested on a complex search and collection task—the *Tracker* task described in [46], and further investigated in [47]. In this problem, a simulated ant is placed on a two-dimensional toroidal grid that contains a trail of food. The ant traverses the grid, collecting any food it contacts along the way. The goal of the task is to discover an ant that collects the maximum number of pieces of food in a given time period (see Fig. 10).

Following [46], each ant is controlled by a network with two input nodes and four output nodes (Fig. 11). The first input node denotes the presence of food in the square directly in front of the ant; the second denotes the absence of food in this same square, restricting the possible legal inputs to the network to (1, 0) or (0, 1). Each of the four output units corresponds to a unique action: move forward one step, turn left 90°, turn right 90°, or no-op. At each step, the action whose corresponding output node has maximum activation is performed. As in the original study [46], no-op allows the ant to remain at a fixed position while activation flows along recurrent connections. Fitness is defined as the number of grid positions cleared within 200 time steps. The task is difficult because simple networks can perform surprisingly well; the network shown in Fig. 11 collects 42 pieces of food before spinning endlessly at position A (in Fig. 10), illustrating a very high local maximum in the search space.

The experiment used a population of 100 networks, each limited to at most nine hidden units, and did not provide a bias node. In the first run (2090 generations), GNARL found






Fig. 11. The semantics of the I/O units for the ant network. The first input node denotes the presence of food in the square directly in front of the ant; the second denotes the absence of food in this same square. This particular network finds 42 pieces of food before spinning endlessly in place at position P, illustrating a very high local maximum in the search space.



(a)



(b)

Fig. 12. The Tracker Task, first run. (a) The best network in the initial population. Nodes 0 and 1 are input, nodes 5–8 are output, and nodes 2–4 are hidden nodes. (b) Network induced by GNARL after 2090 generations. Forward links are dashed; bidirectional links and loops are solid. The light gray connection between nodes 8 and 13 is the sole backlink. This network clears the trail in 319 epochs.

a network [see Fig. 12(b)] that clears 81 grid positions within the 200 time steps. When this ant is run for an additional 119 time steps, it successfully clears the entire trail. To understand how the network traverses the path of food, consider the simple FSA shown in Fig. 13, hand-crafted in [46] as an approximate solution to the problem. This simple machine receives a score of 81 in the allotted 200 time steps, and clears the entire trail only five time steps faster than the network in Fig. 12(b). A step by step comparison indicates there is only a slight difference between the two. GNARL's evolved network follows the general strategy embodied by this FSA at all but two places, marked as positions B and C in Fig. 10. Here the evolved network makes a few additional moves, accounting for the slightly longer completion time.

Fig. 14 illustrates the strategy the network uses to implement the FSA by showing the state of the output units of the network over three different sets. Each point is a triple of the form (*move*, *right*, *left*).³ Figure 14(a) shows the result of supplying to the network 200 "food" inputs—a fixed point that executes

³No-op is not shown because it was never used in the final network.




Fig. 13. FSA hand-crafted for the Tracker task in [46]. The large arrow indicates the initial state. This simple system implements the strategy "move forward if there is food in front of you, otherwise turn right four times, looking for food. If food is found while turning, pursue it, otherwise, move forward one step and repeat." This FSA traverses the entire trail in 314 steps, and gets a score of 81 in the allotted 200 time steps.




Fig. 14. Limit behavior of the network that clears the trail in 319 steps. Graphs show the state of the output units Move, Right, Left. (a) A fixed-point attractor that results for sequence of 500 "food" signals; (b) A limit cycle attractor that results when a sequence of 500 "no food" signals is given to network; (c) All states visited while traversing the trail; (d) The path of the ant on an empty grid. The z axis represents time. Note that x is fixed, and y increases monotonically at a fixed rate. The large jumps in y position are artifacts of the toroidal grid.

"Move." Figure 14(b) shows the sequence of states reached when 200 "no food" signals are supplied to the network—a collection of points describing a limit cycle of length five that repeatedly executes the sequence "Right, Right, Right, Right, Move." These two attractors determine the response of the network to the task [Fig. 14(c), (d)]; the additional points in Fig. 14(c) are transients encountered as the network alternates between these attractors. The differences in the number of steps required to clear the trail between the FSA of Fig. 13 and GNARL's network arise due to the state of the hidden units when transferring from the "food" attractor to the "no food" attractor.

However, not all evolved network behaviors are so simple as to approximate an FSA [40]. In a second run (1595 generations) GNARL induced a network that cleared 82 grid points within the 200 time steps. Fig. 15 demonstrates the behavior of this network. Once again, the "food" attractor, shown in Fig. 15(a), is a single point in the space that always executes "Move." The "no food" behavior, however, is not an




Fig. 15. Limit behavior of the network of the second run. Graphs show the state of the output units Move, Right, Left. (a) Fixed point attractor that results for sequence of 3500 "food" signals; (b) A limit cycle attractor that results when a sequence of 3500 "no food" signals is given to the network; (c) All states visited while traversing the trail; (d) The path of the ant on an empty grid. The z axis represents time. The ant's path is comprised of a set of "railroad tracks." Along each track, tick marks represent back and forth movement. At the junctures between tracks, a more complicated movement occurs. There are no artifacts of the toroidal grid in this plot, all are actual movements [cf. Fig. 14(d)].

FSA; instead, it is a quasiperiodic trajectory of points shaped like a "D" in output space [see Fig. 15(b)]. The placement of the "D" is in the "Move / Right" corner of the space and encodes a complex alternation between these two operations [see Fig. 15(d)].

In contrast, research in [46] uses a genetic algorithm on a population of 65 536 bit strings with a direct encoding to evolve only the weights of a neural network with five hidden units to solve this task. The particular network architecture in [46] uses Boolean threshold logic for the hidden units and an identity activation function for the output units. The first GNARL network was discovered after evaluating a total of 104 600 networks, while the second was found after evaluating 79 850. The experiment reported in [46] discovered a comparable network after about 17 generations. Given [46] used a population size of 65 536 and replaced 95% of the population each generation, the total number of network evaluations to acquire the equivalent network was 1 123 942. This is 10.74 and 14.07 times the number of networks evaluated by GNARL in the two runs. In spite of the differences between the two studies, this significant reduction in the number of evaluations provides empirical evidence that crossover may not be best suited to the evolution of networks.

V. CONCLUSION

Allowing the task to specify an appropriate architecture for its solution should, in principle, be the defining aspect of the complete network induction problem. By restricting the space of networks explored, constructive, destructive, and genetic algorithms only partially address the problem of topology acquisition. GNARL's architectural constraints

$R_1 - R_3$ similarly reduce the search space, but to a less extensive degree. Furthermore, none of these constraints is necessary, and their removal would affect only ease of implementation. In fact, no assumed features of GNARL's networks are essential for the algorithm's operation. GNARL could even use nondifferentiable activation functions, a constraint necessary for backpropagation.

GNARL's minimal representational constraints would be meaningless if not complemented by appropriate search dynamics to traverse the space of networks. First, unlike constructive and destructive algorithms, GNARL permits a non-monotonic search over the space of network topologies. Consider that in monotonic search algorithms, the questions of *when* and *how* to modify structure take on great significance because a premature topological change cannot be undone. In contrast, GNARL can revisit a particular architecture at any point, but for the architecture to be propagated it must confer an advantage over other competing topologies. Such a non-linear traversal of the space is imperative for acquiring appropriate solutions because the efficacy of the various architectures changes as the parametric values are modified.

GNARL allows multiple structural manipulations to a network within a single mutation. As discussed earlier, constructive and destructive algorithms define a unit of modification; e.g., "add a fully connected hidden node." Since such singular structural modifications create a one-unit structural horizon beyond which no information is available, such algorithms may easily fixate on an architecture that is better than networks one modification step away, but worse than those two or more steps distant. In GNARL, several nodes and links can be added or deleted with each mutation, the range being determined by user-specified limits and the current ability of the network. This simultaneous modification of the structural and parametric modifications based on fitness allows the algorithm to discover appropriate networks quickly, especially in comparison with evolutionary techniques that do not respect the uniqueness of distributed representations.

Finally, as in all evolutionary computations, GNARL maintains a population of structures during the search. This allows the algorithm to investigate several differing architectures in parallel while avoiding over-commitment to a particular network topology.

These search dynamics, combined with GNARL's minimal representational constraints, make the algorithm extremely versatile. Of course, if topological constraints are known *a priori*, they should be incorporated into the search. But such constraints should be introduced as part of the task specification rather than being built into the search algorithm. Since the only requirement on a fitness function f is that $f : \mathbf{S} \rightarrow \mathbf{R}$, diverse criteria can be used to rate a network's performance. For instance, the first two experiments described above evaluated networks based on a desired input/output mapping; the Tracker task experiment, however, considered overall network performance, not specific mappings. Other criteria could also be introduced, including specific structural constraints (e.g., minimal number of hidden units or links) as well as constraints on generalization. In some cases, strong task restrictions can even be implicit in simple fitness functions [48].

The dynamics of the algorithms guided by the task constraints represented in the fitness function allow GNARL to empirically determine an appropriate architecture. Over time, the continual cycle of *test-prune-reproduce* will constrain the population to only those architectures that have acquired the task most rapidly. Inappropriate networks will not be indefinitely competitive and will be removed from the population eventually.

Complete network induction must be approached with respect to the complex interaction between network topology, parametric values, and task performance. By fixing topology, gradient descent methods can be used to discover appropriate solutions. But the relationship between network structure and task performance is not well understood, and there is no "backpropagation" through the space of network architectures. Instead, the network induction problem is approached with heuristics that, as described above, often restrict the available architectures, the dynamics of the search mechanism, or both. Artificial architectural constraints (such as "feedforwardness") or overly constrained search mechanisms can impede the induction of entire classes of behaviors, while forced structural liberties (such as assumed full recurrence) may unnecessarily increase structural complexity or learning time. By relying on a simple stochastic process, GNARL strikes a middle ground between these two extremes, allowing the network's complexity and behavior to emerge in response to the requirements of the task.

ACKNOWLEDGMENT

This research has been partially supported by ONR grants N00014-92-J-1195 and N00014-93-1-0059. We are indebted to Ed Large, Dave Stucki, and especially John Kolen for proofreading help and discussions during the development of this research. We would also like to thank our anonymous reviewers, David Fogel, and the attendees of Connectfest '92 for feedback on preliminary versions of this work.

REFERENCES

- [1] A. G. Barto, "Connectionist learning for control," in *Neural Networks for Control*, W. T. Miller III, R. S. Sutton, and P. J. Werbos, Eds. Cambridge, MA: MIT Press, 1990, pp. 5-58.
- [2] T. Ash, "Dynamic node creation in backpropagation networks," *Connection Science*, vol. 1, no. 4, pp. 365-375, 1989.
- [3] M. Frean, "The upstart algorithm: A method for constructing and training feed-forward neural networks," Technical Report Preprint 89/469, Edinburgh Physics Dept, 1990.
- [4] S. J. Hanson, "Meiosis networks," in *Advances in Neural Information Processing Systems 2*, D. Touretzky, Ed. San Mateo, CA: Morgan Kaufmann, 1990, pp. 533-541.
- [5] S. E. Fahlman and C. Lebiere, "The cascade-correlation architecture," in *Advances in Neural Information Processing Systems 2*, D. Touretzky, Ed. San Mateo, CA: Morgan Kaufmann, 1990, pp. 524-532.
- [6] S. Fahlman, "The recurrent cascade-correlation architecture," In *Advances in Neural Information Processing System 3*, R. Lippmann, J. Moody, and D. Touretzky, Eds. San Mateo, CA: Morgan Kaufmann, 1991, pp. 190-196.
- [7] D. Chen, C. Giles, G. Sun, H. Chen, Y. Less, and M. Goudreau, "Constructive learning of recurrent neural networks," *IEEE International Conference on Neural Networks*, vol. 3, 1993, pp. 1196-1201.
- [8] M. R. Azimi-Sadjadi, S. Sheedvash and F. O. Trujillo, "Recursive dynamic node creation in multilayer neural networks," *IEEE Transactions on Neural Networks*, vol. 4, no. 2, pp. 242-256, 1993.

- [9] M. Mozer and P. Smolensky, "Skeletonization: A technique for trimming the fat from a network via relevance assessment," in *Advances in Neural Information Processing Systems 1*, D. Touretzky, Ed. San Mateo, CA: Morgan Kaufmann, 1989, pp. 107-115.
- [10] Y. L. Cun, J. Denker, and S. Solla, "Optimal brain damage," in *Advances in Neural Information Processing Systems 2*, D. Touretzky, Ed. San Mateo, CA: Morgan Kaufmann, 1990.
- [11] B. Hassibi and D. G. Stork, "Second order derivatives for network pruning: Optimal brain surgeon," in *Advances in Neural Information Processing Systems 5*, S. J. Hanson, J. D. Cowan, and C. L. Giles, Eds. San Mateo, CA: Morgan Kaufmann, 1993, pp. 164-171.
- [12] C. W. Omlin and C. L. Giles, "Pruning recurrent neural networks for improved generalization performance," Tech. Report No. 93-6, Computer Science Department, Rensselaer Polytechnic Institute, April 1993.
- [13] L. J. Fogel, A. J. Owens, and M. J. Walsh, *Artificial Intelligence through Simulated Evolution*, New York: John Wiley & Sons, 1966.
- [14] D. B. Fogel, "Evolving artificial intelligence," Ph.D. thesis, University of California, San Diego, 1992.
- [15] J. H. Holland, *Adaptation in Natural and Artificial Systems*, Ann Arbor, MI: The University of Michigan Press, 1975.
- [16] D. E. Goldberg, *Genetic Algorithms in Search, Optimization, and Machine Learning*, Reading, MA: Addison-Wesley Publishing Company, Inc., 1989.
- [17] D. B. Fogel, "An introduction to simulated evolutionary optimization," *IEEE Transactions on Neural Networks: Special Issue on Evolutionary Programming*, D. B. Fogel and L. J. Fogel, Eds. In press.
- [18] A. P. Wieland, "Evolving neural network controllers for unstable systems," in *IEEE International Joint Conference on Neural Networks*, Seattle, WA: IEEE Press, 1990, pp. II-667-II-673.
- [19] D. Montana and L. Davis, "Training feedforward neural networks using genetic algorithms," in *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, San Mateo, CA: Morgan Kaufmann, 1989, pp. 762-767.
- [20] D. Whitley, T. Starkweather, and C. Bogart, "Genetic algorithms and neural networks: Optimizing connections and connectivity," *Parallel Computing*, vol. 14, pp. 347-361, 1990.
- [21] R. D. Beer and J. C. Gallagher, "Evolving dynamical neural networks for adaptive behavior," *Adaptive Behavior*, vol. 1, no. 1, pp. 91-122, 1992.
- [22] G. F. Miller, P. M. Todd, and S. U. Hegde, "Designing neural networks using genetic algorithms," in *Proceedings of the Third International Conference on Genetic Algorithms*, J. D. Schaffer, Ed. San Mateo, CA: Morgan Kaufmann, 1989, pp. 379-384.
- [23] R. K. Belew, J. McInerney, and N. N. Schraudolph, "Evolving networks: Using the genetic algorithm with connectionist learning," in *Artificial Life II: Proceedings of the Workshop on Artificial Life*, C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, Eds. Reading, MA: Addison-Wesley, 1992, pp. 511-547.
- [24] J. Torreale, "Temporal processing with recurrent networks: An evolutionary approach," in *Fourth International Conference on Genetic Algorithms*, R. K. Belew and L. B. Booker, Eds. San Mateo, CA: Morgan Kaufmann, 1991, pp. 555-561.
- [25] M. A. Potter, "A genetic cascade-correlation learning algorithm," in *Proceedings of COGANN-92 International Workshop on Combinations of Genetic Algorithms and Neural Networks*, L. D. Whitley and J. D. Schaffer, Eds. Los Alamitos, CA: IEEE Computer Society Press, 1992.
- [26] N. Karunanithi, R. Das, and D. Whitley, "Genetic cascade learning for neural networks," in *Proceedings of COGANN-92 International Workshop on Combinations of Genetic Algorithms and Neural Networks*, L. D. Whitley and J. D. Schaffer, Eds. Los Alamitos, CA: IEEE Computer Society Press, 1992.
- [27] D. E. Goldberg, "Genetic algorithms and Walsh functions: Part 2, Deception and its analysis," *Complex Systems* vol. 3, pp. 153-171, 1989.
- [28] D. E. Goldberg, "Genetic algorithms and Walsh functions: Part 1, A gentle introduction," *Complex Systems*, vol. 3, pp. 129-152, 1989.
- [29] J. D. Schaffer, L. D. Whitley, and L. J. Eshelman, "Combinations of genetic algorithms and neural networks: A survey of the state of the art," in *Proceedings of COGANN-92 International Workshop on Combinations of Genetic Algorithms and Neural Networks*, L. D. Whitley and J. D. Schaffer, Eds. Los Alamitos, CA: IEEE Computer Society Press, 1992.
- [30] G. E. Hinton, J. L. McClelland, and D. E. Rumelhart, "Distributed representations," in *Parallel Distributed Processing: Explorations in the Microstructure of Cognition, Volume 1: Foundations*, D. E. Rumelhart and J. L. McClelland, Eds. Cambridge, MA: MIT Press, 1986, pp. 77-109.
- [31] T. J. Sejnowski and C. R. Rosenberg, "Parallel networks that learn to pronounce english text," *Complex Systems*, vol. 1, pp. 145-168, 1987.
- [32] J. Koza and J. Rice, "Genetic generation of both the weights and architecture for a neural network," in *IEEE International Joint Conference on Neural Networks*, Seattle, WA: IEEE Press, 1991, pp. II-397-II-404.
- [33] R. Collins and D. Jefferson, "An artificial neural network representation for artificial organisms," in *Parallel Problem Solving from Nature*, H. P. Schwefel and R. Manner, Eds. Heidelberg: Springer-Verlag, 1991.
- [34] D. B. Fogel, "A brief history of simulated evolution," in *Proceedings of the First Annual Conference on Evolutionary Programming*, D. B. Fogel and W. Atmar, Eds. La Jolla, CA: Evolutionary Programming Society, 1992.
- [35] D. B. Fogel, L. J. Fogel, and V. W. Porto, "Evolving neural networks," *Biological Cybernetics*, vol. 63, pp. 487-493, 1990.
- [36] J. R. McDonnell and D. Waagen, "Determining neural network connectivity using evolutionary programming," in *Twenty-fifth Asilomar Conferences on Signals, Systems, and Computers*, Monterey, CA, 1992.
- [37] D. B. Fogel, "Using evolutionary programming to create neural networks that are capable of playing Tic-Tac-Toe," in *International Conference on Neural Networks*, San Francisco, CA: IEEE Press, 1993, pp. 875-880.
- [38] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi, "Optimization by simulated annealing," *Science*, vol. 220, pp. 671-680, 1983.
- [39] R. J. Williams, *Adaptive State Representation and Estimation Using Recurrent Connectionist Networks*, Cambridge, MA: MIT Press, 1990, chap. 4, pp. 97-114.
- [40] J. B. Pollack, "The induction of dynamical recognizers," *Machine Learning*, vol. 7, pp. 227-252, 1991.
- [41] M. Tomita, "Dynamic construction of finite automata from examples using hill-climbing," in *Proceedings of the Fourth Annual Conference of the Cognitive Science Society*, Ann Arbor, MI, 1982, pp. 105-108.
- [42] R. L. Watrous and G. M. Kuhn, "Induction of finite-state automata using second-order recurrent networks," in *Advances in Neural Information Processing 4*, J. E. Moody, S. J. Hanson, and R. P. Lippmann, Eds. San Mateo, CA: Morgan Kaufmann, 1992, pp. 309-316.
- [43] C. L. Giles, G. Z. Sun, H. H. Chen, Y. C. Lee, and D. Chen, "Higher order recurrent networks & grammatical inference," in *Advances in Neural Information Processing Systems 2*, D. S. Touretzky, Ed. San Mateo, CA: Morgan Kaufmann, 1990, pp. 380-387.
- [44] C. L. Giles, C. B. Miller, D. Chen, G. Z. Sun, H. H. Chen, and Y. C. Lee, "Extracting and learning an unknown grammar with recurrent neural networks," in *Advances in Neural Information Processing 4*, J. E. Moody, S. J. Hanson, and R. P. Lippmann, Eds. San Mateo, CA: Morgan Kaufmann, 1992, pp. 317-324.
- [45] Z. Zeng, R. M. Goodman, and P. Smyth, "Learning finite state machines with self-clustering recurrent networks," *Neural Computation*, in press.
- [46] D. Jefferson, R. Collins, C. Cooper, M. Dyer, M. Flowers, R. Korf, C. Taylor, and A. Wang, "Evolution as a theme in artificial life: The genesys/tracker system," in *Artificial Life II: Proceedings of the Workshop on Artificial Life*, C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, Eds. Reading, MA: Addison-Wesley, 1992, pp. 549-577.
- [47] J. Koza, "Genetic evolution and co-evolution of computer programs," in *Artificial Life II: Proceedings of the Workshop on Artificial Life*, C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, Eds. Reading, MA: Addison-Wesley, 1992, pp. 603-629.
- [48] P. J. Angeline and J. B. Pollack, "Competitive environments evolve better solutions for complex tasks," in *Genetic Algorithms: Proceedings of the Fifth International Conference (GA93)*, S. Forrest, Ed. San Mateo, CA: Morgan Kaufmann, 1993, pp. 264-270.



Peter J. Angeline received the B.S. degree in mathematics in 1984 from Carnegie-Mellon university, and the M.S. degree in computer science from The Ohio State University in 1989. He is currently a Ph.D. candidate in computer science at The Ohio State University. His research interests include evolutionary algorithms, emergent computation, machine learning, artificial intelligence, and artificial life.



Gregory M. Saunders received the B.S. degree in mathematics in 1988 from Ohio State University, Columbus, Ohio. From 1988 to 1991 he held a National Science Foundation fellowship, and in 1989 received the M.S. degree in computer science from Ohio State University. He is currently a Ph.D. candidate in computer science at Ohio State University, Columbus, Ohio. His research interests include evolutionary computation, neural networks, and connectionist methods of behavior-based control.



Jordan B. Pollack received the Ph.D. degree in 1987 from the University of Illinois and is now an assistant professor in the Computer and Information Sciences Department of The Ohio State University, where his interests span cognitive science, artificial intelligence, connectionist and neural networks, dynamical systems, and artificial life.