Extraction of Rules from Discrete-Time Recurrent Neural Networks*

Christian W. Omlin, C. Lee Giles
NEC Research Institute
4 Independence Way
Princeton, N.J. 08540 USA
Phone: (609) 951-{2691,2642}
E-mail: {omlinc,giles}@research.nj.nec.com

Technical Report CS-TR-3465 and UMIACS-TR-95-54
University of Maryland, College Park, MD 20742

Abstract

The extraction of symbolic knowledge from trained neural networks and the direct encoding of (partial) knowledge into networks prior to training are important issues. They allow the exchange of information between symbolic and connectionist knowledge representations. The focus of this paper is on the quality of the rules that are extracted from recurrent neural networks. Discrete-time recurrent neural networks can be trained to correctly classify strings of a regular language. Rules defining the learned grammar can be extracted from networks in the form of deterministic finite-state automata (DFA's) by applying clustering algorithms in the output space of recurrent state neurons. Our algorithm can extract different finite-state automata that are consistent with a training set from the same network. We compare the generalization performances of these different models and the trained network and we introduce a heuristic that permits us to choose among the consistent DFA's the model which best approximates the learned regular grammar.

Keywords: Recurrent Neural Networks, Grammatical Inference, Regular Languages, Deterministic Finite-State Automata, Rule Extraction, Generalization Performance, Model Selection, Occam's Razor.

*To appear in *Neural Networks*.

1 INTRODUCTION

There has been much interest in the integration and extraction of knowledge from neural networks (Towell et al. 1990; Frasconi et. al. 1991; Hanson & Burr 1991; Giles et al. 1992; Watrous & Kuhn 1992; Fu 1994). One reason is that the lack of understanding of the rules that underly neural network performance has limited their use in some application domains. Another is that some intelligent processing naturally requires the use of symbolic and rule-based knowledge. Recurrent neural networks with discrete-time inputs readily lend themselves to certain types of knowledge encoding and extraction - in particular the ordered triple of a discrete Markov {state; input $\rightarrow$ next-state} process (Giles et al., 1992; Omlin & Giles, 1992). What this paper addresses is the validity and usefulness of extracting this information based on a representation of the neurons' activation.

1.1 Background

Recently, the training of recurrent neural networks that recognize finite state languages has been discussed by several authors (Williams & Zipser 1989; Cleeremans, Servan-Schreiber & McClelland 1989; Elman 1990; Giles et al. 1992, Watrous & Kuhn 1992). The motivation for this work has been multifold; from understanding how these neural networks process natural language to quantifying and understanding their computational characteristics. Cleeremans (1989) and Elman (1990) trained recurrent networks by predicting the next symbol using a truncation of the backward recurrence. Cleeremans (1989) concluded that the hidden unit activations represented past histories and that clusters of these activations can represent the states of the generating automaton. The work reported by Giles et al. (1992) is an extension of Cleeremans (1989) in that complete deterministic finite-state automata are extracted from recurrent networks. Watrous & Kuhn (1992) implemented an alternative approach to state machine extraction. In related work Crutchfield & Young (1991) extract finite-state automata from general dynamical systems though the systems do not have dynamical inputs.

1.2 Motivation

Grammars and automata are intimately related. Grammars generate the languages that automata recognize. The computational power of recurrent networks has been proven by relating these networks to automata. Recurrent networks have recently been shown to be computationally quite powerful, i.e. Turing equivalent (Siegelmann & Sontag 1992). The representational limitations of recurrent network models have been explored by showing their relationships to automata (Goudreau et. al. 1994, Giles et. al. 1995). Thus, it is natural to see how well recurrent networks learn what they can be proven to represent.

This being said, it is not evident that recurrent networks are the best tools for learning regular grammars. However, regular language inference requires that state information to be stored over indefinite periods of time and no feature extraction is necessary for learning. Thus, grammatical inference is a good test bed in which to investigate issues related to the representation of symbolic knowledge in recurrent networks and to explore their computational capabilities. For example, Giles and Omlin (1994) show how grammar learning can be useful in quantifying the computational power of a pruning algorithm of recurrent neural networks. Jim et. al. (1995) show that training noisy recurrent networks with grammatical strings is a good measure of the effect of different type of noise insertion methodologies on networks generalization and convergence performance. The performance of recurrent networks on long temporal sequences is naturally explored by using grammatical strings (Manolios & Fanelli 1994). Frasconi et. al. (1995) used grammar learning and representation to explore a new recurrent network architecture. Furthermore, recent results by Giles et. al. (1994) and Clouse et.al. (1994) imply that neural networks under certain constraints can learn extremely large grammars and may be competitive with other methods.

The extraction of (symbolic) rules from both feed-forward as well as recurrent neural networks has become an active area of research (Cleeremans et al. 1989; Towell et. al. 1990; Giles et al. 1992; Fu 1994). The following question is the focus of our work: How good are the rules extracted from recurrent networks? This is an important one and, to our knowledge, has not been discussed anywhere else in the literature.

1.3 Overview

We train discrete second-order recurrent neural networks to recognize strings of a regular language given a finite set of positive and negative example strings. The DFAs extracted from trained networks depend on a clustering parameter and different, minimized DFAs may be obtained for different values of the clustering parameter. Several among the different DFAs may correctly classify all strings of the training set. Given a set of extracted DFAs which may accept different regular languages, the question arises which DFA best approximates the unknown source grammar, i.e.

which extracted DFA assigns the same label as the unknown source grammar to as many strings as possible. A *model selection heuristic* is introduced which allows us to choose among the different DFAs the one that best models the regular language.

The remainder of this paper is organized as follows: We give a brief introduction ot regular languages and grammatical inference in section 2, followed by a presentation of the neural network architecture and training in section 3; we also include some generalization results of trained networks on long test strings. The rule extraction algorithm and the problem of selecting among the extracted DFAs the best model are discussed in detail in section 4. A summary of the empirical results, followed by a short discussion of related work and possible future research directions conclude this paper.

2 REGULAR LANGUAGES

2.1 Definitions

Regular languages represent the smallest class of formal languages in the Chomsky hierarchy (Hopcroft & Ullman 1979). Regular languages are generated by regular grammars. A regular grammar G is a 4-tuple $G =< S, N, T, P >$ where S is the start symbol, N and T are non-terminal and terminal symbols, respectively, and P are productions of the form $A \rightarrow a$ or $A \rightarrow aB$ where $A, B \,\epsilon\, N$ and $a \,\epsilon\, T$. The regular language generated by G is denoted $L(G)$.

Associated with each regular language L is a deterministic finite-state automaton (DFA) M which is an acceptor for the language $L(G)$, i.e. $L(G) = L(M)$. DFA M accepts only strings which are a member of the regular language $L(G)$. Formally, a DFA M is a 5-tuple $M =< \Sigma, Q, R, F, \delta >$ where $\Sigma = \{a_1, \ldots, a_k\}$ is the alphabet of the language L, $Q = \{s_1, \ldots, s_M\}$ is a set of states, $R \,\epsilon\, Q$ is the start state, $F \subseteq Q$ is a set of accepting states and $\delta : Q \times \Sigma \rightarrow Q$ define state transitions in M. A string x is accepted by the DFA M and hence is a member of the regular language $L(M)$ if an accepting state is reached after the string x has been read by M. Alternatively, a DFA M can also be considered a generator which generates the regular language $L(M)$.

2.2 Grammatical Inference

Grammatical inference, the problem of inferring grammar(s) from samples of strings of an unknown regular language, is a NP-complete problem (Gold 1978; Angluin & Smith 1983). The problem of grammatical inference can be stated informally as follows: Given a finite set of strings along with a label indicating whether or not a string is a member of the language, find the rules which define the language. For example, we may be given the following two sets of positive and negative example strings:

positive examples	**negative examples**
10010	100010
1011101	001110001
00	000
111	100110111000
100100100111	1001001000111
1010011	100001
001111	00001111
101010	1010001010

In this example the rule might be to reject strings which contain three or more consecutive 0's and to accept all other strings. It is intuitively clear that the inferred grammar will better model the underlying language if more example strings - both positive and negative - are provided as input to the inference algorithm.

Our algorithm extracts rules from (trained) networks in the form of DFAs, which can be readily transformed into a regular grammar where each state of the DFA corresponds to a non-terminal symbol and the DFA transitions correspond to production rules.

3 RECURRENT NEURAL NETWORK

3.1 Architecture

Recurrent neural networks have been shown to have powerful capabilities for modeling many computational structures; an excellent discussion of recurrent neural network models and references can be found in (Hertz, Krogh & Palmer 1991). To learn grammars, we use a second-order recurrent neural network (Lee et al. 1986; Giles et al. 1990; Sun et al. 1990; Pollack 1991). The network architecture is illustrated in figure 1. This net has N recurrent hidden neurons labeled S_j; L special, nonrecurrent input neurons labeled I_k; and $N^2 \times L$ real-valued weights labeled W_{ijk}. As long as the number of input neurons is small compared to the number of hidden neurons, the complexity of the network only grows as $O(N^2)$, the same as a linear network. We refer to the values of the hidden neurons collectively as a state *vector* $\mathbf{S}$ in the finite N-dimensional space $[0,1]^N$. Note that the weights W_{ijk} modify a product of the hidden S_j and input I_k neurons. This quadratic form directly represents the state transition diagrams of a state process — $\{input, state\} \Rightarrow \{nextstate\}$. This recurrent network accepts a time-ordered sequence of inputs and evolves with dynamics defined by the following equations:

$$S_i^{(t+1)} = g(\Xi_i), \qquad \Xi_i \equiv \sum_{j,k} W_{ijk} S_j^{(t)} I_k^{(t)},$$

where g is a sigmoid discriminant function. Each input string is encoded into the input neurons one character per discrete time step t. The above equation is then evaluated for each hidden neuron S_i to compute the next state vector $\mathbf{S}$ of the hidden neurons at the next time step $t+1$. With unary encoding the neural network is constructed with one input neuron for each character in the alphabet of the relevant language. This condition might be too restrictive for grammars with large alphabets.

3.2 Training Procedure

A special recurrent neuron is selected as the network's output neuron. A network learns by comparing its actual output with the desired classification at the end of a string. Weight updates are computed using a second-order form of the RTRL net of Williams (1989) which computes the partial derivative of a quadratic error function with respect to all weights W_{ijk}.

3.3 Network Performance

We consider a regular language over the alphabet $\{0,1\}$ that is recognized by a randomly generated minimal 10-state DFA. ¿From a data set consisting of 1,000 alternating positive and negative example strings starting with strings of length 0, we formed five different training sets. The first training set consisted of all 1,000 example strings; the remaining training sets consisted of the first 700, 300, 100 and 50 strings, respectively, of the 1,000 example strings. This was done in order to achieve a variety of generalization performances for networks of fixed size.

All the weights were initialized to random values in the interval [-1.0, 1.0]. The learning rate α and the momentum η were both set to 0.5.

We used an incremental learning strategy where a network is first trained on a small subset of all training strings ("working set"); this subset only contains the shortest strings. When a network has either learned to correctly classify all the strings of a working set or a maximum number of epochs has expired for the current working set, new strings from the training set are added to the working set and the network is trained on this expanded working set. The working set is expanded until a network correctly classifies all the strings of the training set. This incremental training heuristic aims at overcoming the problems that arise from training on long strings with a gradient descent learning algorithm (Bengio et al. 1994).

We trained 25 networks with eight state neurons on each of the five training sets and performed a statistical analysis on the networks' average generalization performances (table 1) on three test sets consisting of alls string of length up to 15 (65k strings), 1,000 randomly strings of length 100, and 1,000 randomly strings of length 1,000. We observe that the average generalization performance improves with increasing training set size which contain strings of increasing length; the 90% confidence intervals are non-overlapping. This result is not surprising since networks trained on smaller data sets have less information and thus learn a poorer model of the source grammar. Networks generally perform poorer when tested on strings that are longer compared to the training strings; a network's internal DFA representation becomes increasingly unstable with increasing string length which leads to string misclassification.

The next section addresses the issue of how DFA extraction can be made computationally feasible and the quality of the extracted rules, i.e. how well do extracted DFAs generalize compared to the network's generalization performance; we also compare the performance of DFAs that are extracted with different quantization levels from the same network.

4 RULE EXTRACTION

4.1 Algorithm

We describe our heuristic for extracting rules from recurrent networks in the form of DFAs (Giles et al. 1992). Different approaches are described in Cleeremans (1989) and in Watrous & Kuhn (1992). The heuristic is not restricted to second-order networks and can be applied to any discrete-time neural network. The algorithm used to extract a finite-state automaton from a network is based on the observation that the outputs of the recurrent state neurons of a trained network tend to cluster as shown in figure 2. The outputs of a trained 4-neuron recurrent network which is tested on a small test set are shown as two-dimensional projections into the (S_i, S_j)-plane for all possible pairs (S_i, S_j) (6 projections). Our hypothesis is that collections of these clusters correspond to the states of the finite-state automaton the network has learned. Furthermore, these clusters tend to be well separated in networks which learned a given grammar well as measured by their generalization performance on unseen strings. For the DFA extraction, we consider network states where the output of the response neuron S_0 is larger than 0.5 to correspond to accepting DFA states; otherwise, the current network state corresponds to a rejecting DFA state. The problem of DFA extraction is thus reduced to identifying clusters in the output space $[0, 1]^N$ of all state neurons. We use a dynamical state space exploration which identifies the DFA states and at the same time avoids exploration of the entire space which is computationally not feasible.

The extraction algorithm divides the output of each of the N state neurons into q intervals of equal size, yielding q^N partitions in the space of outputs of the state neurons. We also refer to q as the *quantization level.*

Starting in some defined initial network state, the trained weights of the network will cause the current network state to follow a continuous state trajectory as symbols are presented as input. The algorithm considers all strings in alphabetical order starting with length 1. This procedure defines a search tree with the initial state as its root; the number of successors of each node is equal to the number of symbols in the input alphabet. Links between nodes correspond to transitions between DFA states. The search is performed in breadth-first order. When a transition is made into another (not necessarily different) partition, then all paths from that new network state are followed for subsequent input symbol creating new states in the DFA for each new input symbol, subject to the following three conditions: (1) When a previously visited partition is reached, then only the new transition is defined between the previous and the current partition; no new DFA state is created. This corresponds to pruning the search tree at that node. (2) In general, many different state vectors belong to the same partition. We choose the state vector which first reached a partition as the new initial state for all subsequent symbols. This condition only applies when two or more transitions from a partition to another partition are extracted. (3) When a transition leads from a partition immediately to the same partition, then a loop is created and the search tree is also pruned at that node. The algorithm terminates when no new partitions are visited for the first time - hence no new DFA states are created - and all possible transitions for all DFA states have been extracted.

The extraction would seem computationally infeasible if all q^N partitions had to be visited. However, the clusters corresponding to DFA states are often local and can be covered with fewer partitions. Also, we aim to extract DFAs with the smallest possible quantization level which explains the training data. These properties along with the pruning of the search tree makes DFA extraction feasible even for DFAs with larger alphabets.

Clearly, the extracted DFA depends on the quantization level q chosen, i.e., in general, different DFA will be extracted for different values of q. Furthermore, because of condition 2, different DFA may be extracted depending on the order in which the successors of a node in the search tree are visited. In general, however, these distinctions are not significant because of the subsequent standard minimization algorithm (Hopcroft & Ullman 1978). This algorithm yields a unique, minimal representation of the extracted DFA. Many different DFA extracted thus collapse into *equivalence classes*.

4.2 Example

We will illustrate in detail how the extraction algorithm builds the corresponding DFA using a simple example. Assume a recurrent network with 2 state neurons and 2 input neurons is trained on some data set. The network state vector has dimension 2 and the range of possible values of S_0 and S_1 can be represented as a unit square in the (S_1, S_2)-plane. For reasons of simplicity, we choose quantization level $q = 3$, i.e. the output of each of the two state neurons is divided into 3 equal length intervals, defining $3^2 = 9$ discrete partitions. Each of these partitions corresponds to a hypothetical state in the unknown DFA. We will assign labels *1, 2, 3, ...* to the partitions in the order in which they are visited for the first time.

The start state is defined by the initial network state vector used for training the network; it lies in partition *1* (figure 3a). The state vector reached in partition *1* falls within the accepting region (the output of the response neuron is larger than 0.5); thus, this initial states is marked as an accepting state (shaded circle). The start state represents the totally unknown DFA. On input '0' and '1', the network makes a transition into partitions *1* and *2*, respectively. This causes the creation of a transition to a new accepting DFA state 2 and a transition from state 1 to itself. In the next step, transitions from partitions *2* into partitions *3* and *4* on input '0' and '1', respectively, occur and the resulting partial DFA is shown in figure 3b. The DFA in figure 3c shows the current knowledge about the DFA after all state transitions from states 3 and 4 have been extracted from the network. In the last step, only one more new state is created (figure 3d). Since no known state has left any undefined transitions, the extraction algorithm terminates. Notice that not all partitions have been assigned to DFA states. The algorithm usually only visits a subset of all available partitions for the DFA extraction. Many more partitions are reached when large test sets (especially when they contain many long strings) are used (e.g. when measuring the generalization performance on a large test set). The extracted DFA and its unique, minimized representation are shown in figure 4. The DFAs accept all strings which do not contain three consecutive 0's (notice that both DFAs accept exactly the same regular language).

If several DFAs are extracted with different quantization levels q_i, then one or more of the extracted DFA M_{q_i} may be consistent with the given training set, i.e. several DFA M_{q_i} may correctly classify the training set. Clearly, we need to make a choice between different consistent DFA. The heuristic for choosing a DFA will be discussed next.

4.3 Model Selection

Let M denote the unknown DFA and $L(M)$ the language accepted by M. By choosing a particular quantization level q_i, we extract a minimized finite-state automaton, which we consider to be a *hypothesis* M_{q_i} for the grammar to be inferred. A DFA M is called a *consistent DFA* if it correctly classifies all strings of the training set; otherwise, it is called an *inconsistent* model of the unknown source grammar. Given a set of consistent hypotheses M_{q_1}, M_{q_2}, $\ldots$, M_{q_Q} we need criteria that allow us to choose the hypothesis that best approximates the unknown language $L(M)$. We refer to the process of choosing a DFA M_{q_i} as *model selection.* A possible heuristic for model selection would be to split a given data set into two disjoint sets (training and testing set), to train the network on the training set and to test the network's generalization performance on the test set. However, by disregarding a subset of the original data set for training, we may be eliminating valuable data from the training set which would improve the network's generalization performance if the entire data set were used for training. Clearly, we wish to make a model selection based solely on simple properties of the extracted DFAs and not resort to a test set.

The minimized DFAs which were extracted from a trained network for quantization levels $q = 3$, $q = 6$ and $q = 8$ are shown in figure 5. All three DFAs were consistent with the training set, i.e. they correctly classified all strings in the training set. Which DFA best models the unknown source M ?

To answer this question, we extracted DFAs with quantization levels $q = \{2, \ldots, 10\}$ from the 25 trained networks for each of the five training sets and measured the DFAs generalization performances on the same three test sets as before. In some cases, DFAs that were extracted with quantization levels 2 or 3 were inconsistent with a network's training set; these DFAs were discarded.

In all our experiments, the consistent, minimized DFAs that were extracted with the smallest quantization levels also had the fewest number of states. Following Occam's Razor, we measured the generalization performances of these smallest consistent DFAs and compared them to the performances of the remaining DFAs. The results of a statistical analysis are shown in table 2. We observe that the smallest consistent DFAs generally outperform the remaining consistent DFAs. In the case where DFAs were extracted from networks trained on only 50 strings, the smallest consistent DFAs are outperformed by the remaining consistent DFAs. However, this is statistically not significant since the confidence intervals for the DFAs average performance overlap.

These simulation results lead us to conjecture that the DFA M_q which best models the unknown DFA M is the smallest consistent DFA. M_q can be found by extracting DFAs $M_2, \ldots, M_{q-1}, M_q$ in that order. The best model is the first consistent DFA M_q. We further hypothesize that there always exists a value q such that M_q is the smallest consistent DFA. Assuming that we succeed in extracting a DFA that is consistent with the training set, then such a smallest q exists by the well-ordering principle. Our intuition leads us to believe that large DFAs tend to overfit the given training data set. They do not capture the structure of the unknown source grammar well leading to poorer generalization performance.

5 CONCLUSIONS

A heuristic has been introduced that permits the inference of good grammars from trained networks. The rules of learned grammars are extracted in the form of deterministic finite-state automata (DFAs). The problem of rule extraction is reduced to a dynamical state exploration and cluster analysis in the N-dimensional output space of recurrent state neurons. Clusters of state neuron activations are found through a simple partitioning procedure, parameterized by a *partitioning* or *quantization level* q. Other methods such as hierarchical clustering or unsupervised learning algorithms could be used instead. Although different DFAs may be extracted for different values of the quantization level q, a standard minimization algorithm (Hopcroft & Ullman 1979) can be applied yielding unique, minimal representations of the DFAs. Furthermore, these DFAs may be identical over ranges of the quantization level q.

We have observed that different DFAs of minimal representation can be extracted from a trained network depending on the parameter q. It is quite common that several extracted DFA are consistent with the training set, i.e. several DFAs correctly classify the training set. Generalization tests on the different minimized DFA have shown that there exists among the consistent DFA one which shows the best generalization performance. Thus, it becomes necessary to perform a *model selection* among several DFA candidates. In our experiments, the smallest consistent DFA always outperformed the trained network and all other extracted DFAs. It was also always the DFA that was extracted with the smallest quantization level and that was consistent with the training set.

Our rule extraction algorithm depends on dividing the continuous state space of recurrent networks with sigmoidal discriminant functions into discrete partitions. Different partitionings may cause different, minimal DFAs to be extracted from a network which makes a model selection necessary. Two recently proposed methods avoid that problem.

Zeng et al. (1993) define a discrete state space *prior* to training rather than prior to the rule extraction. The sigmoidal discriminant function is still used for computing the weight updates. For string classification, the sigmoidal discriminant function $g(x)$ becomes a step function $h(x) = g(x_i)$ for $x_i \leq x < x_{i+1}$ where $g(x_i)$ is the original sigmoidal discriminant function. This yields q^N discrete state vectors in $[0,1]^N$, each of which potentially corresponds to a DFA state.

Das & Mozer (1994) have proposed a self-clustering learning method which assumes that the unknown source grammar requires a finite number of discrete internal states that are corrupted due to inaccuracies in the weights. Their networks learn to recover these states with maximum a posteriori probability. These clusters changes during the course of learning.

Both proposed methods report improved network generalization performance based on the stability of the internal DFA representation. However, they fail to address the issue of training performance, i.e. DFAs which are readily learned by networks without state space clustering may not be learned as easily by either one of the two clustering methods.

Our results raise the issue of could improved performance be obtained from some recursive combination of rule extraction and rule insertion. We have developed a simple algorithm which maps a (partial) description of a DFA into a *second-order* recurrent neural network by programming some of the weights accordingly (Omlin & Giles 1992). The architecture of the network is defined by the amount of prior information about the otherwise unknown DFA. Significant learning time improvements were achieved by training networks with prior knowledge. By continuously inserting and extracting knowledge from a network, starting with little or no prior information, the size of the network would change after each rule insertion/training/rule extraction cycle and be determined by the current partial knowledge of the DFA, i.e. the extracted symbolic knowledge would control the growth and decay of the network architecture. Thus, this symbolically-guided training procedure could lead to faster training and better generalization performance.

6 ACKNOWLEGDMENT

We would like to acknowledge useful discussions with E.B. Baum, K. Lang and L. Valiant.

7 REFERENCES

Angluin, D. & Smith, C.H. (1983). Inductive Inference: Theory and Methods. *ACM Computing Surveys.* **15**. 237-269.

Bengio, Y., Simard, P., Frasconi, P. (1994). Learning Long-Term Dependencies with Gradient Descent is Difficult. *IEEE Transactions on Neural Networks.* **1**. 157-166.

Cleeremans, A., Servan-Schreiber, D. & McClelland, J. (1989). Finite State Automata and Simple Recurrent Recurrent Networks. *Neural Computation.* **1**, 372-381.

Clouse, D.S., Giles, C.L., Horne, B.G. & Cottrell, G.W. (1994). Learning Large DeBruijn Automata with Feed-Forward Neural Networks. Technical Report CS94-398, Department of Computer Science and Engineering, University of California at San Diego, La Jolla.

Crutchfield, J.P. & Young, K. (1991). Computation at the Onset of Chaos. In W.H. Zurek (Ed.), *Proceedings of the 1988 Workshop on Complexity, Entropy and the Physics of Information*, (pp.223-269). Redwood City: Addison-Wesley.

Das, S. & Mozer, M.C. (1994). A Unified Gradient-Descent/Clustering Architecture for Finite State Machine Induction. In Cowan, J.W., Tesauro, G. & Alspector J. (Eds.), *Advances in Neural Information Processing Systems 6*, (pp. 19-26). San Mateo: Morgan Kaufmann Publishers.

Elman, J.L. (1990). Finding Structure in Time. *Cognitive Science*. **14**. 179-211.

Frasconi, P., Gori, M., Maggini, M. & Soda, G. (1991). A Unified Approach for Integrating Explicit Knowledge and Learning by Example in Recurrent Networks. In 1991 IEEE INNS International Joint Conference on Neural Networks, Seattle , (pp.811-816). Piscataway: IEEE Press.

Frasconi, P., Gori, M., Maggini, M. & Soda, G. (1995). Representation of Finite State Automata in Recurrent Radial Basis Function Networks. *Machine Learning*. In press.

Fu, L. (1994). Rule Generation from Neural Networks. *IEEE Transactions on Neural Networks*. **24**. 1114-1124.

Giles, C.L., Sun, G.Z., Chen, H.H., Lee, Y.C., & Chen, D. (1990). Higher Order Recurrent Networks & Grammatical Inference. In D.S. Touretzky (Ed), *Advances in Neural Information Systems 2* (pp. 380-387). San Mateo: Morgan Kaufmann Publishers.

Giles, C.L., Miller, C.B., Chen, D., Chen, H.H., Sun, G.Z. & Lee, Y.C. (1992). Learning and Extracting Finite State Automata with Second-Order Recurrent Neural Network. *Neural Computation*, **4**. 393-405.

Giles, C.L., Omlin, C.W. (1992). Inserting Rules into Recurrent Neural Networks. In S. Kung, F. Fallside, J.A. Sorenson & C. Kamm (Eds.), *Neural Networks for Signal Processing II, Proceedings of The 1992 IEEE Workshop*, (pp.13-22). Piscataway: IEEE Press.

Giles, C.L., Horne, B.G & Lin, T. (1994). Learning a Class of Large Finite State Machines with a Recurrent Neural Network. Technical Report UMIACS-TR-94-94, Institute for Advanced Computer Studies, University of Maryland, College Park.

Giles, C.L. & Omlin, C.W. (1994). Pruning Recurrent Neural Networks for Improved Generalization Performance. *IEEE Transactions on Neural Networks.* **5**(5). 848-851.

Giles, C.L., Chen, D., Sun, G.Z., Chen, H.H., Lee, Y.C. & Goudreau, M.W. (1995). Constructive Learning of Recurrent Neural Networks: Limitations of Recurrent Casade Correlation and a Simple Solution. *IEEE Transactions on Neural Networks.* In press.

Gold, E.M. (1978). Complexity of Automaton Identification from Given Data. *Information and Control*, **37**. 302-320.

Goudreau, M.W., Giles, C.L., Chakradhar, S.T. & Chen, D. (1994). First-Order Vs. Second-Order Single Layer Recurrent Networks. *IEEE Transactions on Neural Networks.* **5**(3). 511-513.

Hanson, S. & Burr, D. (1991). What Connectionist Models Learn: Learning and Representation in Connectionist Networks. In R. Mammone & Y.Zeevi (Eds.), *Neural Networks: Theory and Applications.* Boston: Academic Press.

Hertz, J., Krogh, A. & Palmer, R.G. (1991). *Introduction to the Theory of Neural Computation.* Redwood City: Addison-Wesley.

Hopcroft, J.E. & Ullman, J.F. (1979). *Introduction to Automata Theory, Languages and Computation*, Reading: Addison-Wesley.

Jim, K., Horne, B.G. & Giles, C.L. (1995). Effects of Noise on Convergence and Generalization in Recurrent Networks. In Tesauro, G., Touretzky, D. & Leen, T. (Eds.), *Advances in Neural Information Processing Systems 7.* In press.

Lee, Y.C., Doolen, G., Chen, H.H., Sun, G.Z., Maxwell, T., Lee, H.Y. & Giles, C.L. (1986). Machine Learning Using a Higher Order Correlational Network. *Physica D*, **22**. 276-306.

Manolios, P. & Fanelli, R. (1994). First Order Recurrent Neural Networks and Deterministic Finite State Automata. *Neural Computation.* **6**(6). 1154-1172.

Omlin, C.W. & Giles, C.L., (1992). Training Second-Order Recurrent Neural Networks using Hints. In Sleeman, D. & Edwards, P. (Eds.), *Proceedings of the Ninth International Conference on Machine Learning*, (pp. 363-368). San Mateo: Morgan Kaufmann Publishers.

Pollack, J.B. (1991). The Induction of Dynamical Recognizers. *Machine Learning.* **7**. 227-252.

Siegelmann, H.T. & Sontag, E.D. (1992). On the Computational Power of Neural Nets. In Proceedings of the Fifth ACM Workshop on Computational Learning Theory (pp.440-449). New York: ACM.

Sun, G.Z., Chen, H.H., Giles, C.L., Lee, Y.C. & Chen, D. (1990). Connectionist Pushdown Automata that Learn Context-Free Grammars. *Proceedings of the International Joint Conference on Neural Networks* (pp. 577-580). Hillsdale: Lawrence Erlbaum.

Towell, G.G., Craven, M.W. & Shavlik, J.W., (1990). Constructive Induction using Knowledge-Based Neural Networks. In L.A. Birnbaum, G.C. Collins (Eds.), *Proceedings of the Eighth International Machine Learning Workshop* (p.213). San Mateo: Morgan Kaufmann Publishers.

Watrous, R.L. & Kuhn, G.M. (1992). Induction of Finite-State Languages Using Second-Order Recurrent Networks. *Neural Computation.* **4**. 406-414.

Williams, R.J. & Zipser, D. (1989). A Learning Algorithm for Continually Running Fully Recurrent Neural Networks. *Neural Computation.* **1**. 270-280.

Zeng, Z., Goodman, R., Smyth, P. (1993). Learning Finite State Machines with Self-Clustering Recurrent Networks. *Neural Computation.* **5**. 976-990.

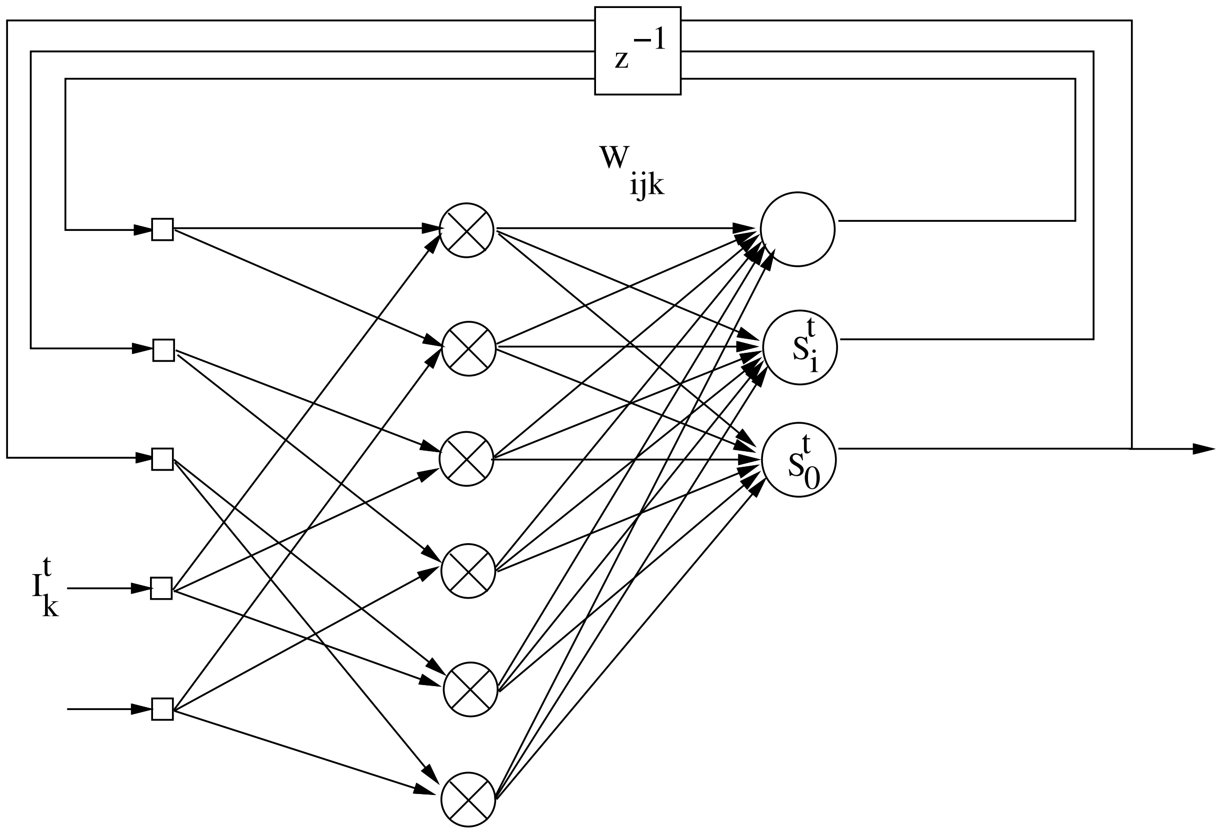

Figure 1: **Second-Order, Single Layer Recurrent Neural Network**: $S_i^{(t)}$ and $I_k^{(t)}$ represent the values of the i^{th} and k^{th} state and input neuron, respectively, at time t. The blocks marked '*' represent the operation $W_{ijk} \times S_j^{(t)} \times I_k^{(t)}$.

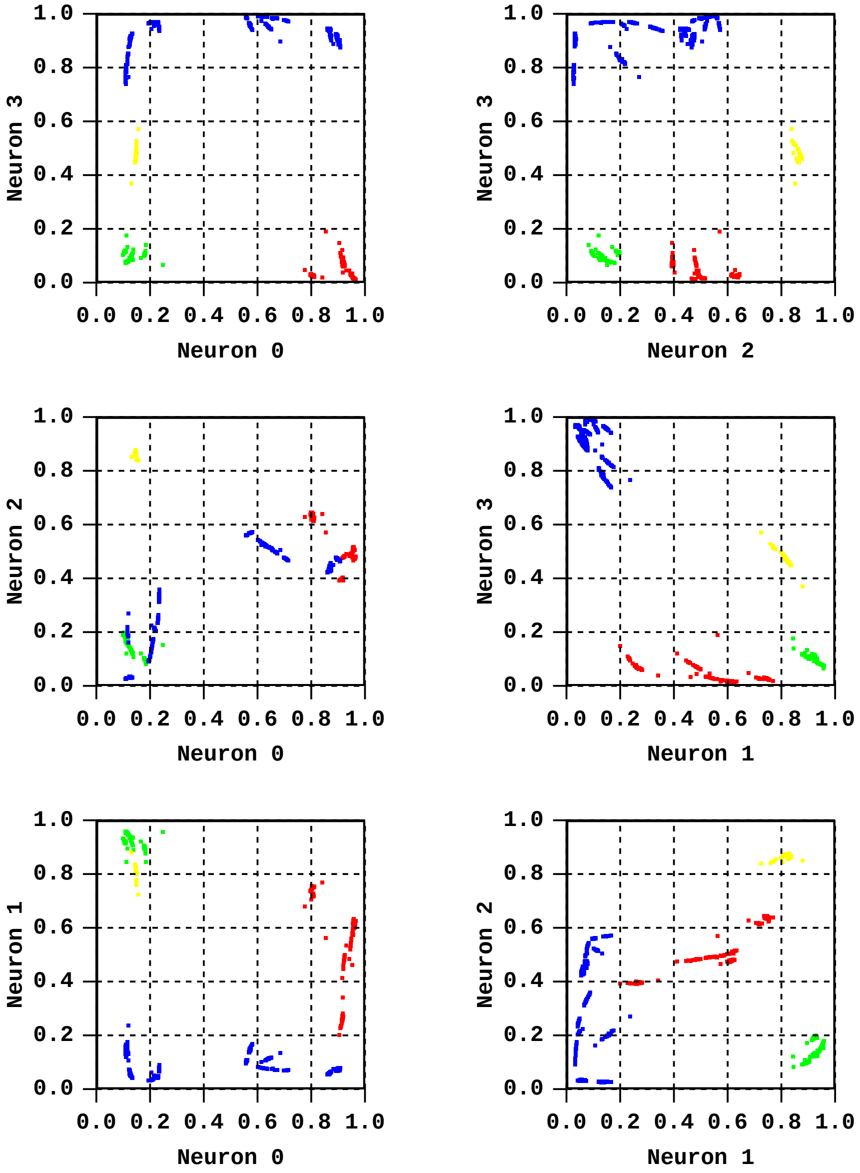

Figure 2: **State Clustering**: A recurrent network with 4 state neurons was trained to accept only strings which do not contain three consecutive 0's. The network builds an internal representation of the learned language in the space of its output neurons. The two-dimensional projection of the state space $[0,1]^4$ into the (S_i, S_j)-plane for all possible pairs (S_i, S_j) are shown (6 projections). The figures show that the state vectors reached for the strings of the test set tend to form clusters. We hypothesize that these cluster correspond to states of the learned DFA and transitions between clusters correspond to state transitions in the DFA.

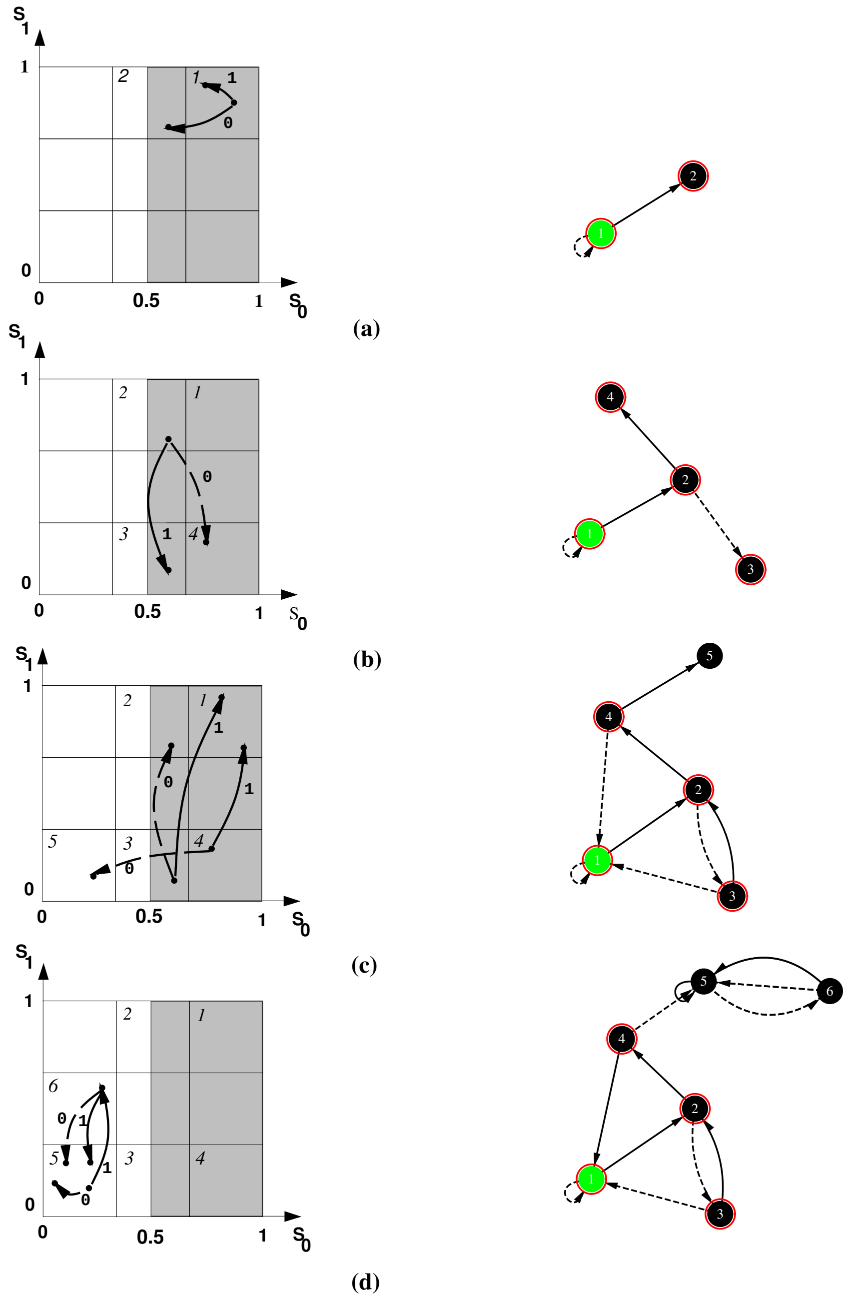

Figure 3: **DFA Extraction Algorithm**:Example of extraction of a DFA from a recurrent network with 2 state neurons. The state space is represented as a unit square in the (S_0, S_1)-plane. The output range of each state neuron has been divided into 3 intervals of equal length resulting in 9 partitions in the networks state space. The figures show the transitions performed between partitions and the (partial) extracted DFA at different stages of the extraction algorithm. (a) The initial state 1 and all possible transitions (b) all transitions from state 2 (c) all transitions from states 3 and 4 (d) all possible transitions from states 5 and 6.

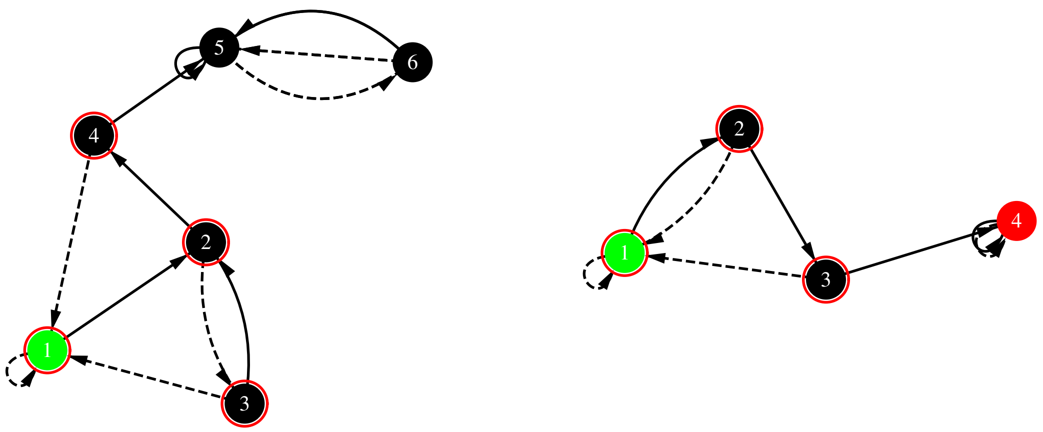

Figure 4: **Minimization of DFAs**: (a) extracted DFA and (b) its unique representation obtained through minimization. Both DFA's accept the same language, consisting of all strings which do not contain the substring '000'.

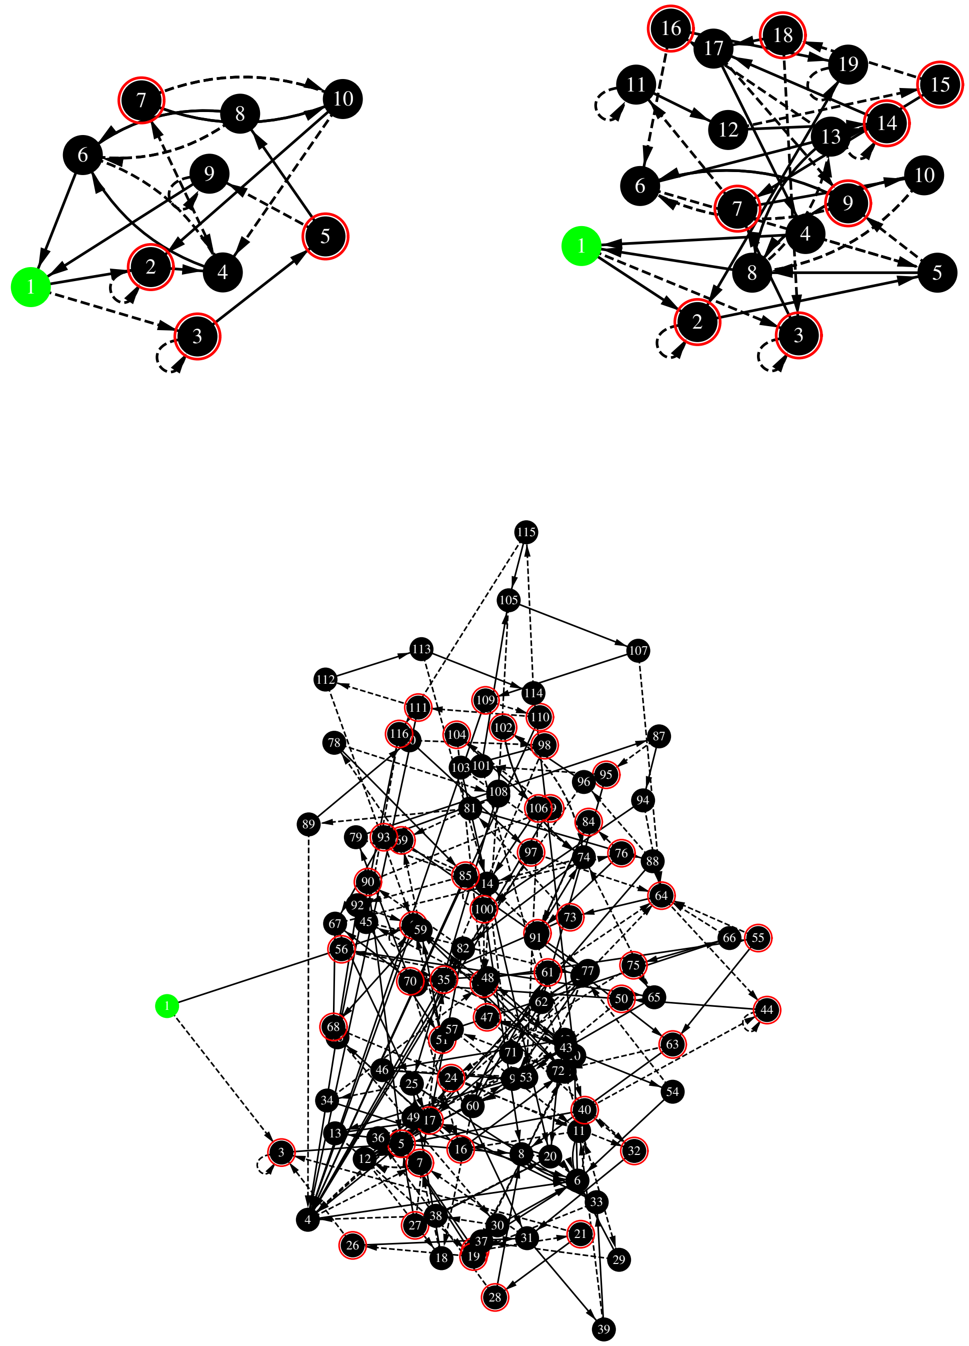

Figure 5: **Model Selection**: Minimized DFA's extracted from a trained network with quantization levels (a) q=3 (b) q=6 and (c) q=8. The smallest DFA was used to generate the training set consisting of the first 1,000 positive and negative example strings in alphabetical order. Thus, it correctly classifies all strings. DFA's M_6 and M_8 make generalization errors with M_6 showing a better generalization performance than M_8.

test set	training set size	average μ	90% confidence interval
strings of length 1-15	1000	0.84%	$0.67\% \leq \mu \leq 1.00\%$
	700	1.63%	$1.26\% \leq \mu \leq 1.99\%$
	300	9.74%	$8.34\% \leq \mu \leq 11.14\%$
	100	38.09%	$36.23\% \leq \mu \leq 39.96\%$
	50	57.96%	$56.29\% \leq \mu \leq 59.64\%$
1000 random strings of length 100	1000	4.61%	$3.57\% \leq \mu \leq 5.66\%$
	700	7.07%	$5.67\% \leq \mu \leq 8.47\%$
	300	26.23%	$23.38\% \leq \mu \leq 29.07\%$
	100	48.50%	$47.47\% \leq \mu \leq 49.52\%$
	50	61.17%	$59.36\% \leq \mu \leq 62.98\%$
1000 random strings of length 1000	1000	4.59%	$3.58\% \leq \mu \leq 5.59\%$
	700	7.68%	$6.08\% \leq \mu \leq 9.28\%$
	300	25.89%	$23.02\% \leq \mu \leq 28.76\%$
	100	47.47%	$46.43\% \leq \mu \leq 48.50\%$
	50	60.24%	$58.05\% \leq \mu \leq 62.43\%$

Table 1: **Network Generalization Performance**: The table shows some statistics (t-student) on the generalization performance of networks with 8 recurrent state neurons on two different test sets. The training sets contained the first 1000, 700, 300, 100, and 50 strings in alphabetical order; starting with strings of length 1 the strings alternated between positive and negative examples. The average performance and the 90% confidence intervals clearly show how the network performances improve with increasing training set size - and thus string lengths - particularly for the test set containing strings of lengths 100 and 1000. The networks generally perform poorer on longer strings - maximum string length 15 compared to string length 100 and 1000 - due to the increasinginstability of the internal DFA representation with increasing string length, leading to misclassification of strings.

test set	training set size	smallest consistent DFA		other consistent DFAs	
		average μ	90% confidence interval	average μ	90% confidence interval
strings of length 1-15	1000	0.01%	$0.00\% \leq \mu \leq 0.03\%$	0.18%	$0.13\% \leq \mu \leq 0.23\%$
	700	0.10%	$0.03\% \leq \mu \leq 0.16\%$	0.52%	$0.45\% \leq \mu \leq 0.60\%$
	300	3.21%	$2.00\% \leq \mu \leq 4.42\%$	4.93%	$4.44\% \leq \mu \leq 5.43\%$
	100	38.19%	$36.77\% \leq \mu \leq 39.62\%$	38.63%	$38.19\% \leq \mu \leq 39.07\%$
	50	39.53%	$38.68\% \leq \mu \leq 40.39\%$	38.54%	$38.10\% \leq \mu \leq 38.99\%$
1000 random strings of length 100	1000	0.02%	$0.00\% \leq \mu \leq 0.04\%$	2.00%	$1.50\% \leq \mu \leq 2.49\%$
	700	1.35%	$0.54\% \leq \mu \leq 2.16\%$	3.24%	$2.81\% \leq \mu \leq 3.50\%$
	300	12.00%	$8.47\% \leq \mu \leq 15.48\%$	19.84%	$18.51\% \leq \mu \leq 21.18\%$
	100	38.52%	$37.30\% \leq \mu \leq 39.74\%$	38.63%	$38.17\ \% \leq \mu \leq 39.09\%$
	50	42.91%	$41.73\% \leq \mu \leq 44.09\%$	42.23%	$41.72\% \leq \mu \leq 42.75\%$
1000 random strings of length 1000	1000	0.02%	$0.00\% \leq \mu \leq 0.06\%$	4.41%	$3.25\% \leq \mu \leq 5.57\%$
	700	0.92%	$0.22\% \leq \mu \leq 1.62\%$	6.79%	$5.52\% \leq \mu \leq 8.06\%$
	300	12.39%	$8.84\% \leq \mu \leq 15.94\%$	24.93%	$22.42\% \leq \mu \leq 27.43\%$
	100	37.38%	$36.47\% \leq \mu \leq 38.30\%$	40.82%	$39.11\% \leq \mu \leq 42.53\%$
	50	41.70%	$40.48\% \leq \mu \leq 42.93\%$	40.97%	$40.42\% \leq \mu \leq 41.53\%$

Table 2: **DFA Generalization Performance**: The table shows some statistics (t-student) on the generalization performance of consistent DFAs on test sets containing strings of varying lengths (DFAs which were inconsistent with the training set were disregarded); the DFAs were extracted from trained networks with quantization levels $q = \{2, 3, \ldots, 10\}$. The average performance and 90% confidence show that the DFAs extracted with the smallest quantization levels have increasingly superior performance over the remaining consistent DFAs as the length of the test strings increases. This becomes the case even for networks that were trained on small data sets (100 strings) and thus necessarily learned a poorer model of the random 10-state DFA than the networks that were trained on larger data sets.