

# An Empirical Evaluation of Rule Extraction from Recurrent Neural Networks

Qinglong Wang<sup>1, 2</sup>, Kaixuan Zhang<sup>2</sup>, Alexander G. Ororbia II<sup>2</sup>,  
Xinyu Xing<sup>2</sup>, Xue Liu<sup>1</sup>, C. Lee Giles<sup>2</sup>

<sup>1</sup>McGill University.

<sup>2</sup>Pennsylvania State University.

**Keywords:** Deterministic finite automata, second order recurrent neural networks, rule extraction, deep learning, regular grammars.

## Abstract

Rule extraction from black-box models is critical in domains that require model validation before implementation, as can be the case in credit scoring and medical diagnosis. Though already a challenging problem in statistical learning in general, the difficulty is even greater when highly non-linear, recursive models, such as recurrent neural networks (RNNs), are fit to data. Here, we study the extraction of rules from second order recurrent neural networks trained to recognize the Tomita grammars. We show that production rules can be stably extracted from trained RNNs and that in certain cases the rules outperform the trained RNNs.

## 1 Introduction

Recurrent neural networks (RNNs) have been increasingly adopted for a variety of tasks involving time-varying data, e.g. sentiment analysis, machine translation, image captioning, etc. Despite the impressive performance on these tasks, RNNs are also well-known to be “black-box” models, which makes explaining/interpreting the knowledge acquired by these models difficult or near-impossible. This black-box nature is largely due to the fact that RNNs, much as any neural architecture, store learned knowledge in their weights, which is difficult to inspect, analyze, and verify (Omlin & Giles, 2000).

Given the RNN’s rising popularity in processing time-varying data, we investigate whether and how we might extract knowledge in symbolic form from RNN models that have been trained on symbolic data, in this case a collection of regular grammars. If the information processing procedure of the RNN can be treated as representing knowledge in symbolic form, where a set of rules that govern transitions between symbolic representations are learned, then we can begin to view the RNN as an automated reasoning

process that can be easier to understand. Indeed, for sentiment analysis tasks, recent work (Murdoch & Szlam, 2017) has demonstrated that an RNN is capable of identifying consistently important patterns of words. These words can be viewed as symbolic knowledge and the patterns of these words represents the rules for determining the sentiment. In other work (Dhingra et al., 2017), information about long-term dependencies are also represented in the form of symbolic knowledge to improve the ability of RNNs to handle long-term text data. Also, prior work (Giles et al., 1992; Watrous & Kuhn, 1992; Omlin & Giles, 1996c; Casey, 1996; Jacobsson, 2005) has shown that it is possible to extract deterministic finite automata (DFA) from RNN models trained to perform grammatical inference and that grammatical rules can be stably encoded in second order RNNs (Omlin & Giles, 1996a,b). In these studies, the vector space of an RNN’s hidden layer is first partitioned into finite parts, each treated as the states of a certain DFA. Then, transitions rules between these states are extracted. This paper follows the paradigm of DFA extraction laid out in previous research efforts.

While it has been shown that it is possible to extract DFA from RNNs, it has been argued (Kolen, 1994) that DFA extraction is sensitive to the initial conditions of the hidden layer of RNN. In other words, by viewing an RNN as a nonlinear dynamical system, the value of its hidden layer may exhibit exponential divergence for nearby initial state vectors. As a result, any attempts at partitioning the hidden space may result in forcing the extracted state to split into multiple trajectories independent of the future input sequence. This results in an extracted rule that appears as a nondeterministic state transition, even though underlying dynamical system is completely deterministic (Kolen, 1994).

In this paper, we greatly expand upon previous work in rule extraction from second-order RNNs (Giles et al., 1992) by studying DFA extraction through comprehensive experiments. The main questions that we hope to ask are:

1. What conditions will affect DFA extraction and how sensitive is DFA extraction with respect to these conditions?
2. How well will the extracted DFA perform in comparison with the RNN trained models from which they are extracted?

With respect to the first question, we address the concerns of (Kolen, 1994) by showing that DFA extraction is quite insensitive to the initial conditions of the hidden layer. Moreover, in answering the second question, we find that in most cases the extracted DFA have equal expressive power to the trained RNN models from which the DFA were extracted. Interestingly enough, in certain cases, we observe that extracted DFA even outperform their source RNNs in term of recognition accuracy when processing long sequences. This result is surprising given the difficulty in training RNNs on long sequences, largely due to the vanishing gradient problem (Pascanu et al., 2013), of which a great deal of research has been dedicated to solving (Hochreiter & Schmidhuber, 1997; Cho et al., 2014; Weston et al., 2014; Sukhbaatar et al., 2015; Dhingra et al., 2017). Extracting rules from RNNs also sheds light on an alternative to improve the processing of long pattern sequences.

This work greatly expands upon previous work (Giles et al., 1992). Here, our emphasis is on examining the consistency of DFA extraction. More specifically, we

first train and test RNN models on data sets generated by the seven Tomita grammars (Tomita, 1982a). The RNN models we use have a second-order architecture (Giles et al., 1992). Then we collect the values of hidden layer units of RNN models obtained during the testing phase, and cluster these values. Here we use k-means due to its simplicity and efficiency. We believe other clustering methods could provide similar results. These clustered states and the symbolic inputs are used to form the initial DFA, which may contain equivalent states. Finally, we use a minimization algorithm to minimize the number of states and finalize the minimal DFA.

In summary, this work makes the following contributions.

- We conduct a careful experimental study of the factors that influence DFA extraction. Our results show that, despite these factors, DFA can be stably extracted from second-order RNNs. In particular, we find strong evidence that, by adopting a simple clustering method, DFA can be reliably extracted even when the target RNN is only trained using short sequences.
- We explore the impact of network capacity and training time on the RNN’s ability to handle long sequences and find that these factors play key roles. With respect to DFA extraction, however, these factors exhibit only a limited impact. This shows that extracting DFA requires less effort compared to the training of a powerful RNN.
- We investigate a realistic case where “incorrect” DFA are extracted from low capacity second-order RNNs, and demonstrate that, in some cases, these DFA can still outperform the source RNNs when processing long sequences. This sheds light on a possible path to improving an RNN’s ability in handling long sequences—exploiting the DFA’s natural ability to handle infinitely long sequences (which is a challenge for any RNN).

## 2 Background

Recurrent neural networks (RNN) process sequential data by encoding information into the continuous hidden space in an implicit, holistic manner (Elman, 1990). In order to extract rules from this continuous hidden space, it is commonly assumed that the continuous space is approximated by a finite set of states (Jacobsson, 2005). The rule is then referred to as the transitions among the discrete states. A common choice for representation of the extracted rules is a DFA. In the following, we first provide a brief introduction of DFA, followed by an introduction to the target grammars studied. Finally, we present a particular type of RNN – a second-order RNN, which is mainly used in this work.

### 2.1 Deterministic Finite Automata

A finite state machine  $M$  recognizes and generates certain grammar  $G$ , which can be described by a five-tuple  $\{A, S, s_0, F, P\}$ . Here,  $A$  is the input alphabet (a finite, non-empty set of symbols),  $S$  is a finite, non-empty set of states.  $s_0 \in S$  and  $F \in S$

Table 1: Description of seven Tomita grammars

G	Description
1	$1^*$
2	$(10)^*$
3	an odd number of consecutive 1s is always followed by an even number of consecutive 0s
4	any string not containing “000” as a substring
5	even number of 0s and even number of 1s (Giles et al., 1990)
6	the difference between the number of 0s and the number of 1s is a multiple of 3
7	$0^*1^*0^*1^*$

represents the initial state (an element of  $S$ ) and the set of final states (a subset of  $S$ ,  $F$  can be empty).  $P$  denotes a set of production rules (transition function  $P : S \times A \rightarrow S$ ). Every grammar  $G$  also recognizes and generates a corresponding language  $L(G)$ , a set of strings of the symbols from alphabet  $A$ . The simplest automata and its associated grammar are DFA and regular grammars, according to the Chomsky hierarchy of phrase structured grammars (Chomsky, 1956). It is important to realize DFA actually covers a wide range of languages, that is, all languages whose string length and alphabet size are bounded can be recognized and generated by finite state automata (Giles et al., 1992). We refer the reader to a more detailed introduction of regular language and finite state machines (Hopcroft et al., 2006).

## 2.2 Tomita Grammars

We select a set of seven relatively simple grammars, which are originally suggested by Tomita (Tomita, 1982a) and widely studied (Pollack, 1991; Omlin & Giles, 1996c; Watrous & Kuhn, 1992) and use them for an empirical study for extracting rules from RNN. We hypothesize (and note from the work of others) that these simple regular grammars should be learnable. More specifically, the DFA associated with these grammars have between three and six states. These grammars all have  $A = \{0, 1\}$ , and generate an infinite language over  $\{0, 1\}^*$ . Here we denote a finite set of strings  $I$  from regular language  $L(G)$ . Positive examples of the input strings are denoted as  $I_+$  and negative examples as  $I_-$ . We provide a description of positive examples accepted by all seven grammars in Table 1.

The associated DFA for these grammars is shown in the first column in Figure 5. Some of these DFA contain a so-called “garbage state”, that is, a non-final state in which all transition paths lead back to itself. In order to correctly learn this state, RNN must not only learn with positive strings  $I_+$  generated by the grammar, but also negative strings  $I_-$  that are rejected by this grammar.

## 2.3 Second-order Recurrent Neural Networks

Here, we use an RNN constructed with second-order interactions between hidden states and input symbols. More specifically, this second-order RNN has a hidden layer  $H$  containing  $N$  recurrent hidden neurons  $h_i$ , and  $L$  input neurons  $i_l$ . The second-order

interaction is represented as  $w_{ijk}h_j^t i_k^t$ , where  $w_{ijk}$  is a  $N \times N \times L$  real-valued matrix, which modifies a product of the hidden  $h_j$  and input  $i_k$  neurons.  $t$  denotes the  $t$ th discrete time slot. This quadratic form directly represents the state transition diagrams of a state process –  $\{\text{input, state}\} \Rightarrow \{\text{next state}\}$ . More formally, the state transition is defined by following equation:

$$H_i^{t+1} = g\left(\sum_{j,k} W_{ijk} H_j^t I_k^t\right) \quad (1)$$

where  $g$  is a sigmoid discriminant function. Each input string is encoded by one-hot-encoding, and the neural network is constructed with one input neuron for each character in the alphabet of the relevant language. By using one-hot-encoding, we ensure that only one input neuron is activated per discrete time step  $t$ . Note that, when building a second-order RNN, as long as  $L$  is small compared to  $N$ , the complexity of the network only grows as  $O(N^2)$ . Such RNNs have been proved to stably encode finite state machines (Omlin & Giles, 1996a,b) and thus can represent in theory all regular grammars.

To train above second-order RNN, we use the following loss function  $C$  following (Giles et al., 1992):

$$C = \frac{1}{2}(y - H_0^T)^2 \quad (2)$$

where  $C$  is defined by selecting a special “response” neuron  $h_0$ , which is compared to the target label  $y$ . For positive strings,  $y = 1.0$  and  $y = 0.0$  for negative strings.  $h_0^T$  indicates the value of  $h_0$  at time  $T$  after seeing the final input symbol. We adopt *RMSprop* (Tieleman & Hinton, 2012) as the training algorithm.

### 3 DFA Extraction

We introduce our approach to DFA extraction, which largely builds on the research conducted in the 1990’s (please see many of the citations in the bibliography) but note that there has been recent work (Li, 2016). We start by briefly introducing the main ideas behind DFA extraction as well as existing research. We will then examine and identify key factors that affect the quality of each step of the extraction process.

#### 3.1 The DFA Extraction Paradigm

Many methods have been developed to extract knowledge in the form of rule from trained RNNs (Giles et al., 1991, 1992; Omlin & Giles, 1996c; Zeng et al., 1993; Frasconi et al., 1996; Gori et al., 1998). Most of this work can be viewed as roughly following one general DFA extraction process:

1. Collect the hidden activations of RNN when processing every string at every time step. Cluster these hidden activations into different states.
2. Use the clustered states and the alphabet-labeled arcs that connect these states to construct a transition diagram.

### 3. Reduce the diagram to a minimal representation of state transitions.

Previous research effort has largely focused on improving the first two steps. This is largely due to the fact that, for the third step, there already exists a well established minimization algorithm (Hopcroft et al., 2006) for obtaining the minimal representation of DFA.

For the first step, an equipartition-based approach (Giles et al., 1992) was proposed to cluster the hidden space by quantizing the value of a hidden unit to a specified number of bins. For example, if we apply a binary quantization<sup>1</sup> to the vector  $\{0.6, 0.4, 0.2\}$ , we would obtain the encoding  $\{1, 0, 0\}$ . One drawback to this form of quantization is that as the number of hidden units increases, the number of clusters grows exponentially. This computational complexity issue is alleviated if one uses clustering methods that are less sensitive to the dimensionality of data samples, e.g., k-means (Zeng et al., 1993; Frasconi et al., 1996; Gori et al., 1998), hierarchical clustering (Sanfeliu & Alquezar, 1994), and self-organizing maps (Tiño & Šajda, 1995).

In order to construct state transitions for the second step, either breadth-first search (BFS) approaches (Giles et al., 1992) or sampling-based approaches (Tiño & Šajda, 1995) are utilized. BFS approach can construct a transition table relatively consistent but incur high computation cost especially when the size of alphabet increases exponentially. Compared with BFS approach, sampling approach is computationally efficient. However, it introduces inconsistency to the construction of a transition table. For a more detailed exposition of these two classes of methods, we refer the readers to this survey (Jacobsson, 2005).

## 3.2 Factors That Affect DFA Extraction

The efficacy of the different methods used for the first two steps of the process described above rely on the following hypothesis: The state space of a well-trained RNN should already be fairly well-separated, with distinct regions or clusters that represent corresponding states in some DFA. This hypothesis, if true, would greatly ease the process of DFA-extraction. In particular, less effort would be required in the first two steps of DFA extraction if the underlying RNN was constructed to have a well-separated state-space.

With this in mind, we specify the following key factors that affect DFA extraction that also affect representational ability of an RNN.

- Model capacity. An RNN with greater capacity (larger size of hidden layer) can be more expressive and thus more likely to better learn a DFA.
- Training time. A sufficient number of iterations are required in order to ensure convergence (to some local optima).
- Initial conditions of the hidden state. As argued previously (Kolen, 1994), the initial conditions may have significant impact on DFA extraction. In this work, we explore this impact by training several RNN models with random initial hidden

---

<sup>1</sup>Using a threshold value of 0.5, any value greater than 0.5 is assigned to the bin “1”, whereas other values less than or equal to this threshold are assigned to “0”.

activations on all grammars, then examining the extracted DFA from all trained RNN models.

- Choice of state-clustering method. The choice of clustering algorithm is very important, including its hyper-parameter configuration. For example, if k-means or a Gaussian mixture model is adopted, a critical hyper-parameter is the predefined number of clusters.

One could argue that other factors, such as choice of a parameter update rule (e.g., ADAM, RMSProp, etc.) and learning rate, may also influence how well an RNN learns about certain grammar. However, in our experiments, we observe that these latter conditions actually have little and nearly no influence on the final results. Thus, we focus on the factors described in the list above.

### 3.3 The DFA Extraction Process

Here, we use an approach similar to (Zeng et al., 1993) to extract DFA from second-order RNNs. To be more specific, we first train second-order RNNs to classify strings generated by each of the seven Tomita grammars (Tomita, 1982b). A desirable outcome of the hypothesis described in the previous section is that, when the hidden space is well-separated, many well-established clustering methods should generate similar results. This allows us to choose our clustering approach based on computational complexity. As a result, we adopt the k-means clustering approach due to its simplicity and efficiency. However, as mentioned before, we must now concern ourselves with choosing an appropriate value of  $K$ .

After clustering the hidden space, we follow the approach taken in Schellhammer et al. (1998) to construct the transition diagram. Specifically, we construct the diagram by counting the number of transitions that have occurred between a state and its subsequent states (given a certain input). For example, given a state  $S_k$  and input symbol  $i$ , we calculate the number of transitions to all states  $\{S\}$  from  $S_k$ , including any self-loops. After obtaining the transition counts, we keep only the most frequent transitions between  $\{S\}$  and  $\{S + 1\}$  given input  $i$  and discard the other less frequent ones in the transition diagram.

It is important to note that when  $K$  is not set properly, certain states may be split differently and inconsistently. As a result, we could observe the situation where  $S_k$  visits multiple different states a similar number of times. This phenomenon is more frequently observed when  $K$  is small. In an extreme case, when the value of  $K$  is set to be even smaller than the minimal number of states of the ground truth DFA, the extraction will never provide the correct DFA. This falsely indicates that the transition learned is more likely to be non-deterministic, while the real case is that the RNN generates  $S_{t+1}$  based on  $S_t$  and  $i$  deterministically. This negative effect can be mitigated when  $K$  is increased beyond a certain value. More specifically, when  $K$  is sufficiently large, the occurrence of “outliers” in the clustering results is less likely, which will reduce the possibility of generating inconsistent transitions. We demonstrate this by evaluating the cluster results using a silhouette coefficient score<sup>2</sup>, calculated under different  $K$  on grammar

---

<sup>2</sup>A higher silhouette coefficient score (Rousseeuw, 1987) reflects better clustering.




Figure 1: Influence of  $K$  on clustering results on grammar 3

3, as shown in Figure 1.

With the constructed transition diagram, we have extracted a DFA that might contain many redundant states. Using the previously described minimization algorithm (Hopcroft et al., 2006), we can then reduce the derived DFA to its minimal representation. Note that this minimization algorithm does not change the performance of the DFA; the unminimized DFA has the same time complexity as the minimized DFA. Note that the DFA extraction method introduced above may be applied to any RNN, regardless of order or manner in which its hidden layers are calculated.

## 4 Experiments

In this section, we empirically study the process of DFA extraction through comprehensive experiments.

### 4.1 Description of Data

To train and test the RNN models, we followed the approach introduced in Giles et al. (1992) and generated string sets. To be specific, we drew strings from an oracle generating random 0 and 1 strings and the grammar specified in Table 1. The end of each string is set to symbol 2, which represents the “stop” symbol (or end-token as in language modeling). For the strings drawn from a grammar, we took them as positive samples while those from that random oracle as negative samples. Note that we verified each string from the random oracle and ensured they are not in the string set represented by that corresponding grammar before treating them as negative samples. It should be noticed that, since each grammar we experimented represents one set of strings with unbounded size, we restricted the length of the strings drawn within a certain range (specified by a lower and upper bound). In our experiment, we set the lower bound equal to 3 and the upper bound equal to 15 for all the grammars. We split the strings generated within the specified range of length for each grammar to build the training set  $D_{train}$  and testing set  $D_{test}$ , then trained and tested the RNNs accordingly.

In order to further the trained RNNs and extracted DFA on longer strings, we build another testing set  $D_{test(200)}$  comprised of strings of length 200 for all grammars. Note



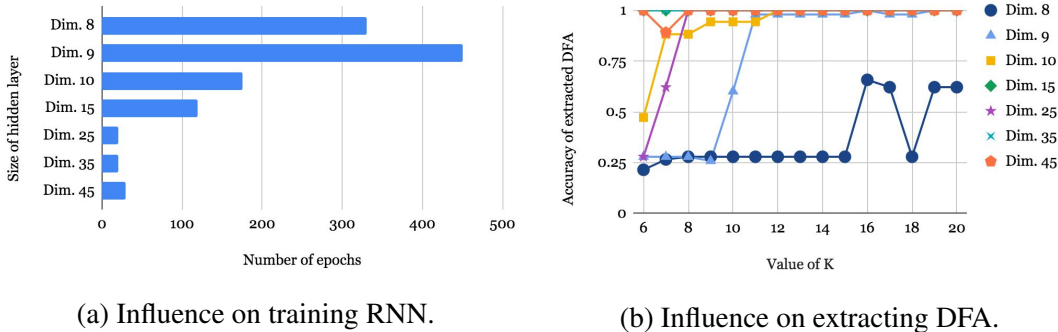


Figure 2: The influence of model capacity on DFA extraction for Grammar 3.

that the complete set of strings with length 200 numbers around  $10^{60}$ . A test set of this size is too expensive and not even necessary for evaluating RNNs or DFA. Therefore, we construct the testing set by randomly sampling 100,000 strings for all grammars. In addition, to preserve the actual balance of positive to negative samples, we sample such that we preserve their original proportions as measured from the original, complete set of length 200 strings. For example, for grammar 5, we sample positive and negative strings with the same ratio of 0.5.

## 4.2 The Influence of Model Capacity

In following experiment, we first measure the influence of model capacity, i.e. the size  $N$  of hidden layer of RNN models, on learning the target DFA. Specifically, we measure the training time needed for RNNs with different hidden layer sizes to reach perfect testing accuracy on the testing set  $D_{test}$  for all grammars. Note that due to space limit, we only show this measure for grammar 3 in Figure 2a. Similar effects are observed with other grammars. It is clear from Figure 2a that it takes less training for an RNN with a larger capacity  $N$  to converge. This is what we would expect; an RNN with larger capacity, in general, has more expressive power and thus can better fit the data. In particular, when  $N$  is greater than 25, an RNN rapidly converges within 30 epochs.

Next, we evaluate how stably we can extract correct DFA from these trained RNN models. Specifically, we compare the classification accuracy on  $D_{test}$  of the extracted DFA when increasing  $K$  from 6 to 20. As shown in Figure 2b, when  $N > 10$ , the correct DFA can be reliably extracted in almost all cases of  $K$  from 6 to 20. On the contrary, when  $N = 8$ , the extraction fails for all  $K$  within the same range. In addition, when the  $N$  is 9 and 10, successful extraction is only observed when  $K$  is larger than 12 for both cases. These results also indicate that DFA extraction is more likely to succeed when  $K$  is set to larger values. This observation is consistent with the results reported in Zeng et al. (1993).

The above experimental results indicate that RNNs with larger capacity are more likely to automatically form a reasonably well-separated state space. As a result, the extraction of DFA is less sensitive to the hidden state clustering step of the process.

Table 2: Influence of training time on DFA extraction and RNN performance.

Grammar	Classification errors reached under different training epochs.							
G3	Epoch	10	20	30	40	50	60	70
	RNN	8451	53	45	1037	67	0	0
	DFA	99771	0	0	0	0	0	0
G4	Epcoh	5	10	15	20	25	30	35
	RNN	34390	28179	4	279	5	2	6
	DFA	49762	0	0	0	0	0	0
G5	Epoch	600	610	620	630	640	650	700
	RNN	7981	7944	7941	7953	0	0	0
	DFA	8035	8035	8035	8035	0	0	0

### 4.3 The Influence of Training Time

In this part, we evaluate the classification performance of both trained RNNs and extracted DFA when processing longer strings. More specifically, we measure the classification errors made by both RNNs and DFA on the test set  $D_{test(200)}$ , as shown in Table 2. For example, with respect to grammar 3, we train seven RNNs with different training epochs (increasing from 10 to 70). The training performance of each as a function of epoch is displayed in Table 2. Seven DFA are then extracted, the classification errors also shown in Table 2. Due to the space restriction, here we only show the results obtained for grammars 3,4 and 5. Similar results were observed with grammars 1,2,6 and 7.

As expected, as the training time increases, RNNs tend to make fewer classification errors. In particular, for grammars 3 and 5, the trained RNNs reach zero error. However, for grammar 4, the RNN still makes a few mistakes even when it has already obtained perfect training accuracy. We also observe that the correct DFA can sometimes be extracted even when the RNN has not yet fully reached zero training error (20th epoch and 15th epoch for grammar 3 and 4, respectively). This indicates that the hidden state space learned in the early stages of training (before the RNN is fully trained) can still be sufficient for a clustering method to recognize each individual state. This observation implies that less effort is needed to extract the correct DFA from an “imperfect” RNN than in training a “perfect” RNN.

Two other interesting observations can be made with respect to grammar 5. First, it takes much longer to both train an RNN to achieve perfect classification performance and extract the correct DFA from it. Second, the correct DFA can only be successfully extracted when the source RNN quits making any mistakes on the test sets. The difficulty behind training on grammar 5 might be explained through examination of the “differences” between the positive and negative strings generated by the grammar. More specifically, by flipping 1 bit of a string from 0 to 1, or vice versa, any positive (negative) string can be converted to a negative (positive) string. In order to learn the small differences, an RNN needs significantly more training time. The second observation may be explained by noting that, before reaching 640 epochs, RNNs make a nearly constant number of errors. This clearly indicates that the RNN is stuck at a certain local minima (also verified in Figure 3a). While the training of RNN is trapped in this

minima, the state space does not start to form the correct partition. However, after 640 epochs, the model escapes this minima and finally converges to a better one, resulting in a state space that is separated correctly.

#### 4.4 The Influence of Initial States & Clustering Configuration


In the following experiments, we examine if a DFA can be stably extracted under random initial conditions. Specifically, for each grammar, we randomly assign an initial value to the hidden activations, i.e.  $H_{0:N}^0$  at  $t_0$  time step, within the interval of  $[0.0, 1.0]$ . We repeat this random initialization ten times (training ten different RNNs) for each grammar. Furthermore, we vary the value of  $K$  for the k-means clustering algorithm, measuring the classification performance of each extracted DFA, and counting the number of times the correct DFA is extracted (only DFA achieving 100% accuracy are regarded as correct). Through this procedure, we hope to uncover the relationship between the initial condition of the RNN’s hidden layer as well as the clustering algorithm’s meta-parameter  $K$  and DFA extraction.

As previously discussed, training an RNN properly is critical for successful DFA extraction. In Figure 3a, we show the mean and variance of the training loss obtained when training each RNN with 10 times of random initialization of hidden activation for all grammars. It is clear from Figure 3a that, except for grammar 5 and 6, RNNs trained on other grammars rapidly converge. For grammar 5 and 6, RNNs need much more training time while having much larger variance of training loss. Recall above discussion in Section 4.3, this is a clearer indication that the training of these RNNs is trapped to different local optima with different initial activation. However, when given sufficient training, all RNNs trained on all grammars converge on the training set and reach 100% accuracy on  $D_{test}$ . In addition, once these RNNs converge to small training loss, the variance reduce to almost 0. This indicates that, with sufficient training and reasonable capacity, RNN’s training is relatively insensitive to the hidden layer’s initialization.


Given the RNNs trained as described above, we then vary  $K$  as we extract DFA from these models. Similarly, we report the mean and variance of the classification accuracy obtained on  $D_{test}$  from all extracted DFA in Figure 3b. To be more specific, for each grammar, under each random initialization of the model’s hidden layer, we run the extraction process 13 times, varying  $K$  in the range from 3 to 15. In total, we conduct 130 rounds of DFA extraction from the ten trained RNNs for each grammar.

As shown in Figure 3b, when  $K$  is set to small values (below 8), except for grammar 1 and 2, the extracted DFA on other grammars have not only poor classification accuracy, but also relatively large variance. In this case, it is difficult to determine whether random initialization of hidden activation or  $K$  have stronger impact on the extraction. When  $K$  is set to a sufficiently large value, however, the variance is significantly reduced while the classification accuracy is greatly improved. This indicates that a sufficiently large  $K$  can offset the impact of initial states.

Beside of showing the classification performance obtained by extracted DFA, we further measure the number of times that correct DFA can be extracted under different  $K$  in Figure 4. Among all 130 rounds of extraction on each grammar, we observe that the correct DFA is successfully extracted with highest success rate of 100% (on



(a) Mean and variance of training loss of RNNs on all grammars.



(b) Mean and variance of testing accuracy of extracted DFA with varying K on all grammars.

Figure 3: Influence of random initialized hidden activations and clustering configuration on training RNN and extracting DFA.




Figure 4: Histograms of the classification performance of extracted DFA on all grammars

grammar 1), lowest success rate of 50% (on grammar 3) and averaged success rate of 75% among all grammars. The reason for the worse extraction results obtained on

Grammar	Correct	Actual
1		
2		
3		
4		
5		
6		
7		

Figure 5: Visualization of ground truth DFA and extracted DFA for all grammars

grammar 3 can be explained by visualizing the extracted DFA in Figure 5.

In Figure 5, we see that all extracted DFA can correctly recognize their associated positive strings, except for ones have length smaller than the minimal length we set. Recall that for all grammars, we generate strings while constraining the minimal string length. The visualization indicates that the extracted DFA not only accurately represent the target grammar that generated the string samples, but also obey the constraint on length. In order for the extracted DFA to satisfy the minimal length constraint, extra states are required, as shown in right panel of Figure 5. Especially, for grammar 3, 4 and 7, the correct DFA contain 5, 4 and 5 states, while the corresponding extracted DFA have 6, 5 and 7 states, respectively. Recall that in above experiments, the minimal value

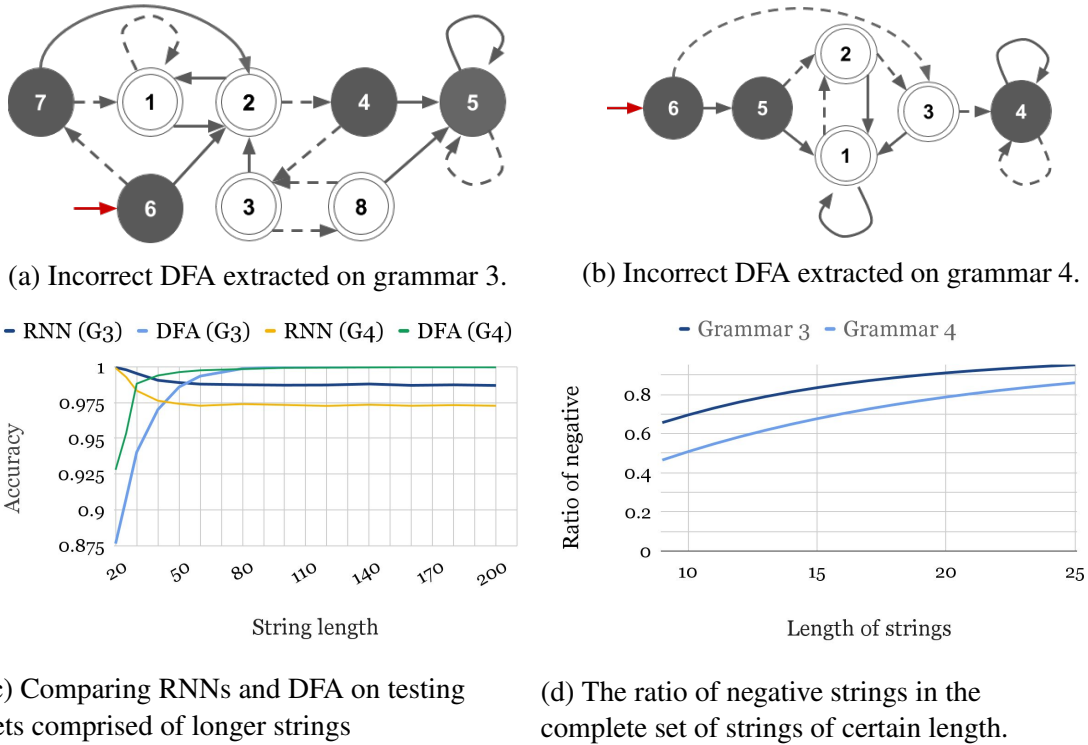


Figure 6: Evaluation of small-capacity RNNs and their associated incorrect DFA for Tomita grammars 3 and 4.

of  $K$  is set to 3 consistently for all grammars. As a result, this setting of  $K$  causes many failures of extraction for these grammars. As shown in Figure 3b, when  $K$  is below 8, the averaged classification accuracy of the extracted DFA are relatively lower in comparison with DFA extracted from other grammars.

### 4.5 Comparison between Low-Capacity RNNs and Extracted DFA

As discussed in Section 4.2, RNNs with larger capacity can learn a DFA better. In practice, it is usually not possible to know the appropriate capacity when constructing an RNN. As a result, it is possible that a smaller RNN can be well-trained on short sequences, but will generalize poorly when confronted with long sequences. Above experiments suggest one solution of extracting a DFA from a trained RNN model, given that DFA extraction is relatively stable and a DFA can maintain its expressive power when processing long strings. In reality, however, it is impractical to assume that the ground truth DFA can be obtained to evaluate the extracted ones, which are possible to be incorrect. In following experiments, we empirically compare some RNNs and their “incorrectly” extracted DFA. Here we demonstrate the results on grammar 3 and 4 due to space constraint. These grammars are selected as we observed in the experiments that the RNNs trained on these grammars are more sensitive to model capacity.

We first construct two RNNs with 9 hidden neurons and have them trained to reach 100% accuracy on data set  $D_{test}$ . Their associated incorrect DFA extracted, as shown in Figure 6a and 6b, achieve 93% and 98% accuracy on  $D_{test}$  respectively. We next eval-

uate these RNNs and their incorrect DFA using multiple testing sets with the number of samples fixed at 100,000 and string length varying from 20 to 200. The sampling of positive and negative strings is similar to what was described in Section.4.1.

RNN test-set performance is shown in Figure. 6c. We observe that on these test sets composed of longer strings, RNNs make more classification errors. This may due to the fact that as the string length increases, the ratio of negative strings to positive ones also increases (shown in Figure.6d). This would mean that an RNN processes more negative strings, which can be interpreted as “noisy” samples, and as a result, would generate more false positive errors. On the other hand, for the DFA associated with these RNNs, fewer and fewer mistakes are made as the number of negative strings increases. This might be the result of the fact that these incorrect DFA generate their own regular language  $L_{3'}$  and  $L_{4'}$  respectively, which are quite similar to the target languages  $L_3$  and  $L_4$ . As a result, many of the negative strings rejected by the extracted DFA are also rejected by the correct DFA. As more and more negative strings are sampled, this overlapping behavior gradually dominates the testing sets. These results demonstrate that, in certain cases, it is possible to extract a DFA which does not fully represent the target RNN and yet still outperforms the RNN when processing longer sequences. Given this result, one possible path to improving an RNN’s ability to handle longer sequences might lie in the exploitation of this useful DFA behavior.

## 5 Conclusion

In this paper, we conducted a careful experimental study of the extraction of deterministic finite automata from second-order recurrent neural networks. We identified the factors that influence the reliability of the extraction process and were able to show that, despite these factors, the automata can still be stably extracted even when the neural model is trained only using short sequences. Our experiments also show that while model capacity does indeed strongly damage the neural network’s ability to handle longer sequences, this hardly affects the extraction process. Furthermore, the automata extracted from low-capacity second order RNNs, in some cases, actually outperform the RNN trained model when processing sequences longer than what were seen during training. Our findings imply that one potential pathway to improving an RNN’s ability to learn longer-term dependencies might be through the exploitation of the DFA’s natural ability to handle infinitely long sequences and that it would be interesting to exploit transfer learning in this area.

## 6 Acknowledgements

We gratefully acknowledge partial support from the College of Information Sciences and Technology.

## References

Casey, M. (1996). The dynamics of discrete-time computation, with application to

- recurrent neural networks and finite state machine extraction. *Neural computation*, 8(6), 1135–1178.
- Cho, K., Van Merriënboer, B., Bahdanau, D., & Bengio, Y. (2014). On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Chomsky, N. (1956). Three models for the description of language. *IRE Transactions on information theory*, 2(3), 113–124.
- Dhingra, B., Yang, Z., Cohen, W. W., & Salakhutdinov, R. (2017). Linguistic knowledge as memory for recurrent neural networks. *arXiv preprint arXiv:1703.02620*.
- Elman, J. L. (1990). Finding structure in time. *Cognitive science*, 14(2), 179–211.
- Frasconi, P., Gori, M., Maggini, M., & Soda, G. (1996). Representation of finite state automata in recurrent radial basis function networks. *Machine Learning*, 23(1), 5–32.
- Giles, C. L., Chen, D., Miller, C., Chen, H., Sun, G., & Lee, Y. (1991). Second-order recurrent neural networks for grammatical inference. In *Neural networks, 1991., ijcnn-91-seattle international joint conference on* (Vol. 2, pp. 273–281).
- 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 networks. *Neural Computation*, 4(3), 393–405.
- Giles, C. L., Sun, G.-Z., Chen, H.-H., Lee, Y.-C., & Chen, D. (1990). Higher order recurrent networks and grammatical inference. In *Advances in neural information processing systems* (pp. 380–387).
- Gori, M., Maggini, M., Martinelli, E., & Soda, G. (1998). Inductive inference from noisy examples using the hybrid finite state filter. *IEEE Transactions on Neural Networks*, 9(3), 571–575.
- Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- Hopcroft, J. E., Motwani, R., & Ullman, J. D. (2006). Automata theory, languages, and computation. *International Edition*, 24.
- Jacobsson, H. (2005). Rule extraction from recurrent neural networks: a taxonomy and review. *Neural Computation*, 17(6), 1223–1263.
- Kolen, J. F. (1994). Fool’s gold: Extracting finite state machines from recurrent network dynamics. In *Advances in neural information processing systems* (pp. 501–508).
- Li, . P. J. C., K. (2016). The kernel adaptive autoregressive-moving-average algorithm. *IEEE transactions on neural networks and learning systems*, 27(2), 334-346.



- Murdoch, W. J., & Szlam, A. (2017). Automatic rule extraction from long short term memory networks. *arXiv preprint arXiv:1702.02540*.
- Omlin, C. W., & Giles, C. L. (1996a). Constructing deterministic finite-state automata in recurrent neural networks. *Journal of the ACM*, 43(6), 937–972.
- Omlin, C. W., & Giles, C. L. (1996b). Constructing deterministic finite-state automata in recurrent neural networks. *Neural Computation*, 8(4), 675–696.
- Omlin, C. W., & Giles, C. L. (1996c). Extraction of rules from discrete-time recurrent neural networks. *Neural networks*, 9(1), 41–52.
- Omlin, C. W., & Giles, C. L. (2000). Symbolic knowledge representation in recurrent neural networks: Insights from theoretical models of computation. *Knowledge based neurocomputing*, 63–115.
- Pascanu, R., Mikolov, T., & Bengio, Y. (2013). On the difficulty of training recurrent neural networks. In *International conference on machine learning* (pp. 1310–1318).
- Pollack, J. B. (1991). The induction of dynamical recognizers. *Machine learning*, 7(2), 227–252.
- Rousseeuw, P. J. (1987). Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20, 53–65.
- Sanfeliu, A., & Alquezar, R. (1994). Active grammatical inference: a new learning methodology. In *in shape, structure and pattern recognition, d. dori and a. bruckstein (eds.), world scientific pub.*
- Schellhammer, I., Diederich, J., Towsey, M., & Brugman, C. (1998). Knowledge extraction and recurrent neural networks: An analysis of an elman network trained on a natural language learning task. In *Proceedings of the joint conferences on new methods in language processing and computational natural language learning* (pp. 73–78).
- Sukhbaatar, S., Weston, J., Fergus, R., et al. (2015). End-to-end memory networks. In *Advances in neural information processing systems* (pp. 2440–2448).
- Tieleman, T., & Hinton, G. (2012). Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural networks for machine learning*, 4(2), 26–31.
- Tiño, P., & Šajda, J. (1995). Learning and extracting initial mealy automata with a modular neural network model. *Neural Computation*, 7(4), 822–844.
- Tomita, M. (1982a). Dynamic construction of finite automata from example using hill-climbing. *Proceedings of the Fourth Annual Cognitive Science Conference*, 105–108.
- Tomita, M. (1982b). Dynamic construction of finite-state automata from examples using hill-climbing. In (pp. 105–108).

- Watrous, R. L., & Kuhn, G. M. (1992). Induction of finite-state automata using second-order recurrent networks. In *Advances in neural information processing systems* (pp. 309–317).
- Weston, J., Chopra, S., & Bordes, A. (2014). Memory networks. *arXiv preprint arXiv:1410.3916*.
- Zeng, Z., Goodman, R. M., & Smyth, P. (1993). Learning finite state machines with self-clustering recurrent networks. *Neural Computation*, 5(6), 976–990.