

Verification of Recurrent Neural Networks Through Rule Extraction*

Qinglong Wang,¹ Kaixuan Zhang,² Xue Liu,¹ C. Lee Giles,²

¹School of Computer Science, McGill University, Canada

²Information Sciences and Technology, Pennsylvania State University, USA

Abstract

The verification problem for neural networks is verifying whether a neural network will suffer from adversarial samples, or approximating the maximal allowed scale of adversarial perturbation that can be endured. While most prior work contributes to verifying feed-forward networks, little has been explored for verifying recurrent networks. This is due to the existence of a more rigorous constraint on the perturbation space for sequential data, and the lack of a proper metric for measuring the perturbation. In this work, we address these challenges by proposing a metric which measures the distance between strings, and use deterministic finite automata (DFA) to represent a rigorous oracle which examines if the generated adversarial samples violate certain constraints on a perturbation. More specifically, we empirically show that certain recurrent networks allow relatively stable DFA extraction. As such, DFAs extracted from these recurrent networks can serve as a surrogate oracle for when the ground truth DFA is unknown. We apply our verification mechanism to several widely used recurrent networks on a set of the Tomita grammars. The results demonstrate that only a few models remain robust against adversarial samples. In addition, we show that for grammars with different levels of complexity, there is also a difference in the difficulty of robust learning of these grammars.

1 Introduction

Verification for neural networks is crucial for validating deep learning techniques in security critical applications. However, the black-box nature of neural networks makes inspection, analysis, and verification of their captured knowledge difficult or near-impossible [23]. Moreover, the complicated architecture of neural networks also make these models vulnerable to adversarial attacks [27] – a synthetic sample generated by slightly modifying a source sample in order to trick a neural network into “believing” this modified sample belongs to an incorrect class with high confidence.

Most prior work on neural network verification has been on verifying feed-forward neural networks using mixed-integer linear programming (MILP) [3, 9, 21, 28] and Satisfiability Modulo Theories (SMT) [2, 6, 18]. Specifically, these approaches can either verify if a neural network can

remain robust to a constrained perturbation applied to an input, or approximate the maximal allowed scale of the perturbation that can be tolerated. To apply these verification approaches, two critical requirements need to be satisfied. One is that an adversarial sample should be recognized by a hypothetical oracle that is very similar or even identical to its source sample. Another requirement is that the adversarial perturbation must be of small enough scale to avoid being detected by the oracle.

Depending on applications, there are different ways to set up the oracle and various distance metrics that measure the scale of an adversarial perturbation. For image recognition, fortunately, it is not challenging to satisfy the two requirements mentioned above. More specifically, in this scenario, a human is usually assumed to be the oracle and adversarial images must avoid a straightforward visual inspection. However, it is neither realistic nor efficient to assign a human oracle. As such, the oracle in this case can be simply replaced by the ground truth labels of source images. As for the distance metrics, various L_p norms ($p = 0, 2, \infty$) have been widely adopted in prior work [2, 27, 30] on adversarial sample problem. The convenience brought by image recognition has made this application as the benchmark for much verification work [7, 10, 22, 28].

When dealing with sequential data, e.g. natural language, programming code, DNA sequence, etc., however these requirements are challenging to satisfy. This is mainly due to the lack of appropriate oracles and proper distance metrics. For instance in sentiment analysis, it has been shown that even the change of a single word is sufficient to fool a recurrent neural network (RNN) [24]. However, the adversarial sentence presented in this work [24] contains grammatical errors. This indicates that for sequential data, the adversarial samples need not only be negligible, but also satisfy certain grammatical or semantic constraints. Unfortunately, it is very challenging to formulate these constraints and construct an oracle with these constraints. Since RNNs are often used for processing sequential data, the difficulty of verifying sequential data has consequents which limits research work on verifying RNNs.

Here, we propose to use deterministic finite automata (DFA) as the oracle. There exists much prior work on relating RNNs to DFA. Our line of research aims at extracting rules from RNNs, where extracted rules are usually ex-

*Some of the work discussed here can be found at [32] and [31]
Copyright © 2018, AAAI Spring Symposium on Verification of Neural Networks (VNN19). All rights reserved.

pressed by a DFA. Furthermore, we design a distance metric – average edit distance – for measuring and constraining the perturbations applied to strings generated by regular grammars. Since it is very difficult, if not impossible, to design comprehensive distance metrics for real-world sequential data, we propose this work as a steppingstone for verifying RNNs that can be built for more sophisticated applications and can have extracted rules. In summary, this work makes the following contributions:

- We propose a distance metric for measuring the scale of adversarial perturbations applied to strings generated by regular grammars. We show that the average edit distance can also describe the complexity of regular grammars.
- We empirically study the factors that influence DFA extraction, and conduct a careful experimental study of evaluating and comparing different recurrent networks for DFA extraction on the Tomita grammars [29]. Our results show that, despite these factors, DFA can be stably extracted from second-order RNNs [13]. In addition, among all RNNs investigated, RNNs with strong quadratic (or approximate quadratic) forms of hidden layer interaction provide the most accurate and stable DFA extraction for all of the Tomita grammars.
- We demonstrate that using DFA can evaluate the adversarial accuracy of different RNNs on Tomita grammars. The experiments show the difference between the robustness of RNNs and the difference in the difficulty of robust learning of grammars with different complexity.

2 Verification Framework for RNNs

The verification problem for neural networks is typically formulated as a MILP or SMT problem. Our work is closely related to prior work on verifying feed-forward neural networks [9, 28] and propose the following formulation for verifying recurrent networks.

We first denote the domain of all regular strings as $\mathcal{X} = \Sigma^*$, where Σ is the alphabet for regular strings. Then we denote the oracle by λ , which can process any $x \in \mathcal{X}$ and produce a classification decision $\lambda(x) \in \mathcal{Y}$. The set of strings classified by λ having the same label of y is denoted by \mathcal{X}_y , i.e. $\mathcal{X}_y = \{x \mid \lambda(x) = y\}$. We assume there is a distance metric denoted by d (detailed introduction of our defined distance metric is provided in Section 3.) to measure the distance between strings. Let $\mathcal{P}(x)$ denote the set of all possible strings generated by perturbing x with respect to a certain distance constraint, i.e. $\mathcal{P}(x) = \{x' \mid d(x, x') \leq d(x, \mathcal{X}_\mu)\}$, where $\mu \in \mathcal{Y} \setminus \lambda(x)$, represents any label that is different from $\lambda(x)$. The above constraint indicates that the allowed perturbation to x must not lead to a different classification result made by λ .

Similarly, a RNN f can process any $x \in \mathcal{X}$ and produce a vector of classification scores, i.e. $f(x) \in \mathbb{R}^{|\mathcal{Y}|}$ and $\sum_{i=1}^{|\mathcal{Y}|} f_i(x) = 1$. Then we say f is robust or locally invariant [20] with respect to \hat{x} if and only if finding a x' that satisfying (1) is infeasible:

$$(x' \in \mathcal{X} \cap \mathcal{P}(\hat{x}) \cap \mathcal{X}_{\lambda(\hat{x})}) \wedge (f_{\lambda(\hat{x})}(x') < \max_{\mu \in \mathcal{Y} \setminus \lambda(\hat{x})} f_{\mu}(x')) \quad (1)$$

To describe the relation between f and λ from a global perspective, we adapt the local invariance property described above to determine the equivalence [20] between f and λ . More formally, we say there exists an equivalence relation between f and λ if it is infeasible to find a x that satisfies the following:

$$(x \in \mathcal{X}) \wedge (\arg \max_{\mu \in \mathcal{Y}} f_{\mu}(x) \neq \lambda(x)). \quad (2)$$

As discussed in Section 1, λ and d play two crucial roles in our verification framework. As such, it is important to have some λ with high “quality” to represent the oracle. Our prior work [32] demonstrated that for certain RNNs, DFAs with high classification accuracy can be extracted in a relatively stable manner. As such, we use DFAs as oracles for verifying recurrent networks. In addition, it needs to be noted that equation (2) provides a way to evaluate the fidelity (introduced in Section 4.2) of an extracted DFA regarding its source RNN. This is important since analyzing an extracted DFA with high fidelity can be more tractable than analyzing its complicated source RNN. Also, our previous work [31] defined the *average edit distance* that measures the difference between different sets of strings. From this, it appears that the average edit distance can also be applied to our verification framework. In the following sections, we will introduce the average edit distance, followed by our empirical study on extracting DFAs from various RNNs, and verifying these RNNs with DFAs.

3 Average Edit Distance

Our definition of average edit distance is an extension of the common definition of *edit* distance, which measures the minimum number of operations – insertion, deletion and substitution of one symbol from a string – needed to covert a string into another [5]. We use this definition for the problem of measuring the difference between sets of strings. One particular application of this metric is for evaluating the complexity of regular grammars [31], where we measure the difference between the sets of strings that are accepted and rejected by a regular grammar.

3.1 Definition of Average Edit Distance

Without loss of generality, we consider for simplicity only two sets of strings with different labels. Given a string $x \in X_y^N$ and a string $x' \in X_{y'}^N$, where X_y^N and $X_{y'}^N$ denote the sets of strings with length N and labels y and y' , respectively. The edit distance between x and the set of strings in $X_{y'}^N$ can be expressed as:

$$\bar{d}_e(x, X_{y'}^N) = \min_{x' \in X_{y'}^N} d_e(x, x'). \quad (3)$$

We then have the following definition.

Definition 1 (Average Edit Distance). The average edit distance $D(\mathcal{X}_y, \mathcal{X}_{y'})$ between two sets of strings \mathcal{X}_y and $\mathcal{X}_{y'}$ is:

$$D(\mathcal{X}_y, \mathcal{X}_{y'}) = \frac{1}{2} \cdot \lim_{N \rightarrow \infty} \left(\frac{1}{|X_y^N|} D_y^N + \frac{1}{|X_{y'}^N|} D_{y'}^N \right), \quad (4)$$

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 [13]
6	the difference between the number of 0s and the number of 1s is a multiple of 3
7	$0^*1^*0^*1^*$

Table 1: Descriptions of Tomita grammars.

Table 2: Average edit distance for Tomita grammars.

N	Grammar							
	G1	G2	G3	G4	G5	G6	G7	
D^N	8	2.51	2.51	1.13	1.16	1.00	1.00	1.17
	10	3.00	3.00	1.18	1.16	1.00	1.00	1.31
	12	3.50	3.50	1.24	1.18	1.00	1.00	1.51
	14	4.00	4.00	1.30	1.22	1.00	1.00	1.75

where D_y^N and $D_{y'}^N$ denote $\sum_{x \in X_y^N} \bar{d}_e(x, X_{y'}^N)$ and $\sum_{x' \in X_{y'}^N} \bar{d}_e(x', X_y^N)$, respectively.

If we let \mathcal{X}_y and $\mathcal{X}_{y'}$ represent the sets of accepted and rejected strings for a certain regular grammar, then $D(\mathcal{X}_y, \mathcal{X}_{y'})$ essentially reflects the complexity of this grammar [31]. In particular, a grammar with higher complexity (hence smaller average edit distance) will be more challenging for robust learning (we show this result in Section 5). In the following, we show the case of how to use the average edit distance to categorize a certain set of regular grammars.

3.2 Average Edit Distance for Tomita Grammars

Tomita grammars [29] denote a set of seven regular grammars and have been widely adopted in the study of DFA extraction for recurrent networks. These grammars all have alphabet $\Sigma = \{0, 1\}$, and generate an infinite language over $\{0, 1\}^*$. A description of the Tomita grammars is provided in Table 1. For a more detailed introduction of Tomita grammars, please see Tomita’s early work [29].

Using Definition 1, we can calculate the average edit distance for the Tomita grammars. As shown in Table 2, different Tomita grammars have different values and changing trends of average edit distance as we increase the length of strings. More specifically, as N increases, the average edit distance of grammars 1, 2 and 7 monotonically increases, while for other grammars, their average edit distance increases at a slower rate (grammar 3, 4) or remain constant (grammar 5, 6). These observations allow us to categorize Tomita grammars into the following three classes. Detailed discussion and calculation of the average edit distance for each grammar is provided in [31].

- (a) For grammar 1, 2 and 7, $D(G_{1,2,7}) = \infty$;
- (b) For grammar 3 and 4, $D(G_{3,4}) > 1$;
- (c) For grammar 5 and 6, $D(G_{5,6}) = 1$.

4 Rule Extraction for Recurrent Networks

Rule extraction for recurrent networks essentially describes the process of developing or finding a rule that approximates the behaviors of a target RNN [17]. More formally, given a RNN denoted as a function $f : \mathcal{X} \rightarrow \mathcal{Y}$ where \mathcal{X} is the data space, \mathcal{Y} is the target space, and a data set $B = \{X, Y\}$ with n samples $X \in \mathcal{X}^n$ and $Y \in \mathcal{Y}^n$. Let r denote a rule which is also a function with its data and target space identical to that of f . The rule extraction problem is to find a function $\mathcal{L} : (\mathcal{X} \rightarrow \mathcal{Y}) \times (\mathcal{X}^n \times \mathcal{Y}^n) \rightarrow (\mathcal{X} \rightarrow \mathcal{Y})$ such that \mathcal{L} takes as input f and B , then outputs a rule r .

There are three key components in the above formulation - the extraction algorithm \mathcal{L} , a recurrent network f , and the underlying data sets B . In our previous study [31, 32], we investigated each component for their effect on the performance DFA extraction. More specifically, we empirically studied that when applying a quantization-based rule extraction algorithm to a second-order RNN [13], what conditions will affect DFA extraction and how sensitive DFA extraction is with respect to these conditions [32]. With respect to this question, we are interested in uncovering the relationship between different conditions. For instance, what is the influence of the initial condition of the RNN’s hidden layer and the configuration of a particular quantization algorithm on DFA extraction. Specially, through our empirical study, we address the concerns of [19] by showing that DFA extraction is very insensitive to the initial conditions of the hidden layer.

In addition, we also investigate how DFA extraction will be affected when we apply it to different recurrent networks trained on data sets with different levels of complexity. More specifically, when the underlying data sets are generated by Tomita grammars, we denote by B_G a data set generated by a grammar G . Then in our evaluation framework, we fix the extraction method \mathcal{L} as a quantization-based method and evaluate the performance obtained by \mathcal{L} when its input, i.e. B_G and f trained on B_G^{train} ¹, vary across different grammars and different recurrent networks respectively. It is important to note that by comparing the extraction performance obtained by a given model across different grammars, we then examine for DFA extraction, how sensitive each model is with respect to the underlying data.

In the following, we introduce the rule extraction algorithm adopted in our previous work and the metrics proposed to evaluate the performance of DFA extraction.

4.1 Quantization-Based DFA Extraction

Quantization-based DFA extraction methods have been the most frequently used in previous work [14, 17, 23, 26, 35]. In these methods rules are constructed based on the hidden layers – ensembles of hidden neurons – of a RNN, and are also referred to as compositional approaches [17]. Also, it is commonly assumed that the vector space of a RNN’s hidden layer can be approximated by a finite set of discrete states, where each rule refers to the transitions between states. As

¹Data set B_G is split into a training set B_G^{train} and a test set B_G^{test} as typically done for supervised learning.

such, a generic compositional approach can be described by the following basic steps:

1. Given a trained RNN, collect the values of a RNN’s hidden layers when processing every sequence at every time step. Then quantize the collected hidden values into different states. This quantization is usually implemented with clustering methods. One such method that has been widely adopted is k-means clustering [11, 26, 32, 35]. In this study, we also use k-means due to its simplicity and computational efficiency.

2. Then use the quantized states and the alphabet-labeled arcs that connect these states to construct a transition diagram. Here we follow [26, 32] and count the number of transitions observed between states. Then we only preserve the more frequently observed transitions.

3. Next, reduce the diagram to a minimal representation of state transitions with a standard and efficient DFA minimization algorithm [16] which has been broadly adopted in previous work for minimizing DFAs extracted from different recurrent networks and for other DFA minimization.

There are other DFA extraction approaches, e.g., pedagogical approaches which construct rules by regarding the target RNN as a black box and build a DFA by only querying the outputs of this RNN for certain inputs. These approaches can be effectively applied to regular languages with small alphabet sizes. However, for RNNs which perform complicated analysis when processing sophisticated data, the extraction process becomes extremely slow [33]. This survey [17] has a more detailed introduction of various rule extraction methods.

4.2 Evaluation Metrics for DFA Extraction

Here, we evaluate the performance of DFA extraction by measuring the quality of extracted DFAs. To be more specific, we introduce three metrics: (1) the accuracy of an extracted DFA when it is tested on the test set for a particular grammar; (2) the success rate from different random trials of extracting DFAs that are identical to the ground truth DFA associated with a particular grammar, which should then perform perfectly on the test set generated by that grammar; (3) the fidelity of an extracted DFA from its source RNN when evaluated on the test set for a particular grammar. These metrics quantitatively measure the abilities of different recurrent networks for learning different grammars. In particular, the first metric reflects the abilities of different recurrent networks for learning “good” DFAs, and has been frequently adopted in much research [12, 25, 33]. The second metric, which is more rigorous, reflects the abilities of these models to learn correct DFAs. The third metric describes how similar an extracted DFA behaves with respect to the RNN from which the DFA is extracted. In the following, we formally introduce these metrics. It is important to note that our evaluation framework is agnostic to the underlying extraction method since we impose no constraints on \mathcal{L} .

Given a model m (which can either be a RNN f or a DFA r) and a data set $B = \{X, Y\}$ consisting of samples x and

their corresponding labels y_x . Let X_i denote the set of samples with the same label i , i.e., $X_i = \{x \in X \mid y_x = i\}$. Then X can be decomposed into disjoint subsets X_i . Similarly, let X_i^m denote the set of samples classified by m as having the label i , i.e., $X_i^m = \{x \in X \mid m(x) = i\}$. Then we have the following metrics for evaluating the performance of DFA extraction.

Accuracy. The accuracy $A(m, B)$ of model m on data set B is defined as :

$$A(m, B) = \frac{1}{|X|} \sum_{i=1}^{|\mathcal{Y}|} |X_i^m \cap X_i|, \quad (5)$$

where $|\cdot|$ denotes the cardinality of a set. To evaluate the accuracy of an extracted DFA f on regular strings, we use $A(r, B) = \sum_{i=1}^2 |X_i^r \cap X_i|/|X|$.

Rate of Success. The rate of success $S(r, B, T)$ of DFA extraction on data set B over T trails is defined as:

$$S(r, B, T) = \frac{1}{T} \sum_{t=1}^T \mathbb{1}_{A(r_t, B)=1}, \quad (6)$$

where r_t is the DFA extracted in the t -th trial. Correspondingly, the average accuracy of extracted DFAs is calculated by averaging $A(r_t, B)$ over T trials.

Fidelity. The fidelity $F(m_1, m_2, X)$ of two models m_1 and m_2 on a data set X is defined as:

$$F(m_1, m_2, X) = \frac{1}{|X|} \sum_{i=1}^{|\mathcal{Y}|} |X_i^{m_1} \cap X_i^{m_2}|. \quad (7)$$

Let $F(f, r, X)$ denote the fidelity of an extracted DFA r regarding its source RNN f on X , it is easy to derive that $F(f, r, X) = 1 - |X_1^f \Delta X_1^r|/|X|$. Here X_1^f and X_1^r denote the sets of strings classified as positive by f and r , respectively, and Δ denotes the symmetric difference of two sets.

In the following section, these three metrics are used to evaluate the DFA extraction performance for various recurrent networks.

5 Experiments

Here we present the experiment results of investigating the effect of various conditions (shown in Table 3) on DFA extraction performance and verifying adversarial accuracy on different recurrent models. We first demonstrate that rule extraction performance is relatively stable for second-order RNN regardless of several varying conditions. Next, we show the evaluation results when we apply DFA extraction to different types of recurrent networks trained on data sets with different levels of complexity. Then, we present the results of verifying recurrent networks with DFAs.

Conditions	Description
Data Complexity	Complexity of Tomita grammars
Model	◦ Elman RNN, Second-order RNN, MI-RNN, GRU, LSTM
	◦ Randomly initialized hidden activation
	◦ Size of the hidden layer
	◦ Training epochs
Quantization	K for k -means clustering

Table 3: Conditions that affect DFA extraction.

5.1 Evaluation of DFA Extraction for Second-order RNN

Due to space constraints, we only present the extraction results of randomly initializing the hidden layer of second-order RNNs, and varying the pre-specified K for k -means clustering. These two factors have been shown to be more influential than other conditions [32], all shown in Table 3. The extraction performance is evaluated by the average accuracy of extracted DFAs and the rate of success in DFA extraction. Discussion of the fidelity tests for second-order RNN and other recurrent networks is provided in Section 5.2.


We followed [14, 32] and generated string sets by drawing strings from an oracle that generates random 0 and 1 strings for a grammar specified in Table 1. 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 each grammar in our experiments represents one set of strings with unbounded size. As such we restricted the length of generated strings as previously specified [32]. We split the strings generated for each grammar to generate the training and test sets. Both data sets were used to train and test the RNNs accordingly, while only the test sets not used were used for evaluating extracted DFAs.

We perform 130 trials of DFA extraction for each RNN on every grammar to comprehensively evaluate the performance of the DFA extraction. In particular, given a RNN and the data set generated by a grammar, we vary two factors – the initial value of the hidden vector of this RNN (randomly initialized for 10 times.²) and the pre-specified value of K for k -means clustering in the range from 3 to 15.


Accuracy of Extracted DFA for Second-order RNN. As shown in Figure 1a, for a sufficiently well trained (100.0% accuracy on the test set) second-order RNN, the initial value of hidden layer has significant influence on the extraction performance when k is set to small values. This impact can be gradually alleviated when K increases. We observe that when k is sufficiently large, the influence of randomly initializing the hidden layer is negligible.

Rate of Success for Second-order RNN. Besides showing the accuracy of the extracted DFAs, we further measure

²For each trial, we select a different seed for generating the initial hidden activations randomly.



(a) Mean and variance of testing accuracy of extracted DFA with varying K for second-order RNN on the Tomita grammars [32].



(b) Histograms of the classification accuracy of extracted DFAs for second-order RNN on the Tomita grammars [32].


Figure 1: Extraction performance for second-order RNN.

the success rate of extraction for second-order RNNs in Figure 1b. More specifically, the success rate of extraction is the percentage of DFAs with 100.0% accuracy among all DFAs extracted for each grammar under different settings of K and random initializations. From all 130 rounds of extraction for each grammar, we observe that the correct DFA successfully extracted with highest success rate of 100.0% is on grammar 1, the lowest success rate of 50.0% on grammar 3, and an averaged success rate of 75.0% among all grammars. These results indicate that DFA extraction is relatively stable for a second-order RN on most grammars.


5.2 Evaluation of DFA Extraction for Different RNNs

Though we empirically investigated DFA extraction for second-order RNNs, it is not clear if DFA extraction can be effectively applied to different recurrent networks, and what is the cause for the inconsistent extraction performance observed across different grammars. To address these questions, we empirically evaluate DFA extraction performance for other recurrent networks for Elman RNN [8], multiplicative integration recurrent neuron networks (MI-RNN) [34], long-short-term-memory networks (LSTM) [15] and gated-recurrent-unit networks (GRU) [4]. We also show in Figure 2a that the complexity of different Tomita grammars is the underlying reason for the inconsistent extraction performance.

Our experiment setup is the same as that for second-order RNNs. In particular, for every pair of a recurrent network and a grammar, we conducted 10 trials with random initialization of the hidden layer of that RNN, and apply DFA extraction for this RNN multiple times by ranging K from



(a) Average accuracy of DFAs extracted from recurrent networks on the Tomita grammars. Left vertical axis: average edit distance. Right vertical axis: average accuracy of extracted DFAs [31].



(b) Histograms of the classification accuracy of extracted DFAs for different RNNs on all grammars [31].

Figure 2: Extraction performance for different RNNs.

3 to 15. We tested and recorded the accuracy of each extracted DFA using the same test set constructed for evaluating all corresponding recurrent networks. The extraction performance is then evaluated based on results obtained from these trials. This we believe alleviates the impact of different recurrent networks being sensitive to certain initial state settings and clustering configurations. Also, we used recurrent networks with approximately the same number of weight and bias parameters regardless of their different architectures.

Accuracy of Extracted DFA for Different RNNs. In Figure 2a, we plot the average accuracy of 130 DFAs extracted from each model trained on each grammar, and the average edit distance of each grammar calculated by setting $N = 20$. As shown in Figure 2a, except for second-order RNN and MI-RNN, the average accuracy obtained by DFAs extracted from each model decreases as the average edit distance of grammars decreases. This indicates that it is generally more difficult for recurrent networks to learn a grammar with a higher level of complexity.

Rate of Success for Different RNNs. The results for evaluating and comparing different RNN models on their rate of success in extracting the correct DFAs associated with the Tomita grammars are shown in Figure 2b. We find that on grammars with lower complexity, all models are capable of producing the correct DFAs. In particular, all models achieve

Grammar	Model	RNN Evaluation		Fidelity
		Clean	Noisy	
3	2nd-RNN	1.00	0.99	1.00
	Elman-RNN	1.00	0.99	1.00
	MI-RNN	1.00	0.99	0.98
	GRU	1.00	0.99	1.00
	LSTM	1.00	0.99	1.00
4	2nd-RNN	1.00	0.99	0.99
	Elman-RNN	1.00	0.99	0.93
	MI-RNN	0.99	0.99	0.69
	GRU	0.99	0.99	0.99
	LSTM	1.00	0.99	0.89
7	2nd-RNN	1.00	0.99	1.00
	Elman-RNN	1.00	0.99	1.00
	MI-RNN	0.99	0.99	0.40
	GRU	1.00	0.99	0.99
	LSTM	1.00	0.99	1.00

Table 4: Fidelity test results for different recurrent networks on grammar 3, 4 and 7. Columns “Noisy” and “Clean” present the results for evaluating RNNs on the data sets with and without noise, respectively.




Figure 3: Fidelity test of MI-RNN regarding K (for k -means) ranging from 6 to 30 on grammar 3, 4 and 7.

much higher success rates on grammar 1. This may due to the fact that the DFA associated with grammar 1 has the fewest number of states (two states) and simplest state transitions among all other DFAs. Thus, the hidden vector space for all RNN models is much easier to separate during training and identify during extraction. As for other grammars with lower complexity, their associated DFAs have both a larger number of states and more complicated state transitions. For grammars with higher levels of complexity, the second-order RNN enables a much more accurate and stable DFA extraction. Also, for the most part, the MI-RNN provides the second best extraction performance, especially, for grammars 5 and 6, which have the highest complexity. In this case only the second-order RNN and MI-RNN are able to extract correct DFAs, while all other models fail.

Fidelity of Extracted DFAs for Different RNNs. When both a RNN and the DFA extracted from that particular RNN obtain 100% accuracy on the test set, it is trivial to check that the fidelity has a value of 1. In order to evaluate the fidelity

in a more realistic scenario, we inject noise into the training sets by randomly selecting several training samples and flipping their labels. To avoid causing a RNN to be severely biased by this noise, we limit the number of samples selected to 4, i.e., 2 for positive strings and 2 for negative strings. We then train a RNN on the noisy training set and extract DFAs accordingly. The fidelity is calculated as discussed in Section 4.2. Table 4 shows the results for fidelity tests for all recurrent networks on grammars 3, 4 and 7 from a single trial.³ Specifically, during the extraction, we set the value of K to 20 since as shown in Figure 1a a larger K is more likely to provide an accurate DFA. As shown in Table 4, for most recurrent networks the high fidelity values obtained by the extracted DFAs from these models across three grammars indicate that these networks can effectively tolerate training set noise. An exception is MI-RNN, from which the extracted DFAs have consistently the lowest fidelity values across three grammars. This shows that MI-RNN is more sensitive to the training set noise. This results in reducing its overall classification performance and causing worse extraction performance. To better illustrate this effect, we show in Figure 3 how the fidelity varies as we increase the value of K from 6 to 30. As K increases, the DFAs extracted from MI-RNN on these grammars have better accuracy on the clean data sets. As such, the similarity between extracted DFAs and their source RNNs increases. This result indicates that a proper setting of K is important for DFA extraction that is both accurate and faithful.

In general, Elman-RNN obtains the worst extraction performance on most grammars, while DFAs extracted from second-order RNN and MI-RNN have consistently higher accuracy and rate of success across all grammars. The former may be due to the simple recurrent architecture of Elman-RNN, which possibly limits its ability to fully capture complicated symbolic knowledge. The better extraction performance of second-order RNN and MI-RNN raises questions regarding the quadratic interaction between input and hidden layers used by these models and whether such an interaction could improve other models’ DFA extraction, an interesting question for future work.

5.3 Verification of Recurrent Networks with DFAs

Following the verification framework described in Section 2, we present in the following experiment the results of verifying recurrent networks with DFAs. It is important to note that when selecting a ground truth DFA as the oracle, we can comprehensively examine the robustness or adversarial accuracy [28] of a certain RNN with respect to small-scale perturbations. If we select an extracted DFA as the oracle, then our verification framework can be adopted for examining the fidelity of the extracted DFA. The latter case has been demonstrated with a simplified case study shown in the previous section, here we focus on the former case.

Given a well trained RNN and a ground truth DFA associated with the grammar used for training this RNN, our verification task mainly focuses on the local invariance prop-

³Other grammars are not shown due to space constraints.

Algorithm 1 Adversarial Accuracy Verification

Input: RNN f ; Extracted DFA r ; String length N ;
Number of samples T ; Allowed perturbed distance d ;
Output: Adversarial accuracy γ ;

- 1: Randomly generate T samples X with length N , and $X = \{x_i \mid r(x_i) = f(x_i) = p\}$, where p denotes the positive label;
- 2: $count \leftarrow 0$;
- 3: **for** $i = 1$ to T **do**
- 4: Generate samples O_{x_i} from x_i satisfying $O_{x_i} = \{x_j \mid d_e(x_i, x_j) \leq d\}$;
- 5: **for** $j = 1$ to $|O_{x_i}|$ **do**
- 6: **if** $r(x_j) = n$ **then**
- 7: **Continue**;
- 8: **else if** $f(x_{ij}) = n$ (where n denotes the negative label) **then**
- 9: $count \leftarrow count + 1$;
- 10: **Break**;
- 11: **end if**
- 12: **end for**
- 13: **end for**
- 14: $\gamma \leftarrow 1 - count/N$;
- 15: **return** γ ;

erty [20] of the RNN. More specifically, we verify the case when a small-scale perturbation is applied to a positive string x , whether a RNN will produce a negative label while the DFA still classifies x as positive⁴. Here we only use grammar 3, 4 and 7 for the verification task. This is because for other grammars, it is easy to check that given a positive string x , almost all strings with the edit distance to x equals 1 belong to the negative class. This means that for grammar 1, 2, 5 and 6, all their positive samples lie on the decision boundary hence the perturbation space is rather limited. While for grammar 3, 4 and 7, it is easier to find adversarial samples with small perturbed edit distance. As such, in the following experiments, we set the maximal allowed perturbed edit distance as 1 to satisfy the constraint mentioned in Section 2.

Since all recurrent networks have been sufficiently well trained on short strings that make up the training and test sets, we verified these models’ adversarial accuracy on long strings. It is well known that recurrent networks have difficulty capturing long-term dependencies. As such, we randomly sampled strings with length 200 to construct the verification data sets. All sampled strings were examined to be correctly classified by both a target recurrent network and the ground truth DFA for grammar 3, 4 and 7. Since the number of strings increases exponentially as their length increases, we randomly sampled 100 positive and 100 negative strings for 30 trials for verification. This allows us to better approximate the ideal results by exploring the entire data space. Based on the verification framework introduced in Section 2, we design the verification algorithm for a single

⁴We also report the results for verifying the local invariance property of a RNN with negative strings.

G	y	γ	RNN				
			2nd	Elman	MI-RNN	GRU	LSTM
3	1	$\bar{\gamma}_+$	1.00	3.96e-2	1.00	1.00	1.00
	0	$\bar{\gamma}_-$	1.00	1.00	1.00	1.00	0.96
4	1	$\bar{\gamma}_+$	1.00	1.00	1.00	1.00	1.00
	0	$\bar{\gamma}_-$	1.00	1.00	1.00	1.00	1.00
7	1	$\bar{\gamma}_+$	1.00	0.99	1.00	0.99	0.98
	0	$\bar{\gamma}_-$	1.00	1.00	1.00	1.00	1.00

Table 5: Verification results for positive and negative strings with the length of 200.




Figure 4: Adversarial accuracy of an Elman-RNN for strings with varying length (from 100 to 200) on grammar 3.

trial as shown in Algorithm 1.

The γ obtained from 30 trials of verifying positive (negative) strings are averaged and denoted as $\bar{\gamma}_+$ ($\bar{\gamma}_-$). The results presented in Table 5 indicate the different levels of robustness obtained by different recurrent networks. In particular, second-order RNN and MI-RNN are most robust with no adversarial samples identified, while other recurrent networks suffer from adversarial samples to a different extent. These results are consistent with the results reported previously in Section 5.2 and [31]. Of note, the Elman-RNN obtains the lowest adversarial accuracy when verified for positive strings from grammar 3. To understand the reason for this worst result, we show in Figure 4 how the adversarial accuracy of an Elman-RNN changes when for verification the length of strings sampled changes. This indicates that an Elman-RNN cannot generalize to long strings and may cause it to more likely suffer from adversarial attacks.

It can also be observed from Table 5 that the difference between recurrent networks’ robustness against adversarial samples may result from a difference between the underlying grammars. Specifically, grammar 4 enables better learning than grammar 3, even though these two grammars have similar levels of complexity. As for grammar 7, although it has the lowest complexity in comparison with grammar 3 and 4, there are effective adversarial samples identified for most recurrent networks. This indicates that this grammar is prone to overfitting the recurrent networks since the data sets for this grammar is very imbalanced for positive and negative samples.

6 Conclusions and Discussion

Here we propose to verify recurrent networks with DFA extraction. We extend the verification framework proposed in prior work for feed-forward neural networks to accommodate what the rigorous requirements for verification of recurrent networks. In particular, we empirically study DFA extraction on various recurrent networks. We show that for certain recurrent networks, their extracted DFAs have such an accuracy that they can be regarded as surrogates for their ground truth counterparts to be used in the verification task. We also show through a case study that our verification framework can also be adopted for examining the equivalence between an extracted DFA and its source RNN using a fidelity metric. In addition, we define an average edit distance metric that is suitable for measuring the adversarial perturbation applied to strings generated by regular grammars. These results are then used in an experimental study of verification for several different recurrent networks. The results demonstrate that while all recurrent networks can sufficiently learn short strings generated by the different Tomita grammars, only certain RNN models can generalize to long strings without suffering from adversarial samples.

Future work would include employing a DFA-based verification for model refinement and conducting more efficient fidelity tests between an extracted DFA and the source recurrent network. Specifically, since a DFA is usually much easier for formal analysis, we could efficiently identify certain implicit weaknesses of a RNN by using a DFA extracted from that RNN to generate specific adversarial samples. The generated adversarial samples could then be used for refining the source RNN. In addition, as discussed in Section 2, for DFA-based verification for a RNN, it is crucial to extract a DFA that is faithful to the source RNN. Indeed, this fidelity requirement is critical for not only verification but also for explainability. A comprehensive fidelity test can be very challenging for recurrent networks since the dimension of sequential data expands exponentially. This also raises a limitation on the computational efficiency of conducting verification for this work. This is largely due to the difficulty of computing edit distance, which results in solving an old NP-hard problem [1, 5]. As such, future work could be in exploring more efficient approximation algorithms.

References

- [1] Backurs, A., and Indyk, P. 2015. Edit distance cannot be computed in strongly subquadratic time (unless seth is false). In *Proceedings of the forty-seventh annual ACM symposium on Theory of computing*, 51–58. ACM.
- [2] Carlini, N.; Katz, G.; Barrett, C.; and Dill, D. L. 2017. Ground-truth adversarial examples. *CoRR* abs/1709.10207.
- [3] Cheng, C.; Nührenberg, G.; and Ruess, H. 2017. Maximum resilience of artificial neural networks. In *Automated Technology for Verification and Analysis - 15th International Symposium, ATVA 2017, Pune, India, October 3-6, 2017, Proceedings*, 251–268.
- [4] Cho, K.; Van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014. On the properties of neural machine trans-

- lation: Encoder-decoder approaches. In *Proceedings of SSST@EMNLP 2014, Eighth Workshop on Syntax, Semantics and Structure in Statistical Translation, Doha, Qatar, 25 October 2014*, 103–111.
- [5] De la Higuera, C. 2010. *Grammatical inference: learning automata and grammars*. Cambridge University Press.
- [6] Ehlers, R. 2017a. Formal verification of piece-wise linear feed-forward neural networks. In *Automated Technology for Verification and Analysis - 15th International Symposium, ATVA 2017, Pune, India, October 3-6, 2017, Proceedings*, 269–286.
- [7] Ehlers, R. 2017b. Formal verification of piece-wise linear feed-forward neural networks. In *International Symposium on Automated Technology for Verification and Analysis*, 269–286. Springer.
- [8] Elman, J. L. 1990. Finding structure in time. *Cognitive science* 14(2):179–211.
- [9] Fischetti, M., and Jo, J. 2017a. Deep neural networks as 0-1 mixed integer linear programs: A feasibility study. *CoRR* abs/1712.06174.
- [10] Fischetti, M., and Jo, J. 2017b. Deep neural networks as 0-1 mixed integer linear programs: A feasibility study. *arXiv preprint arXiv:1712.06174*.
- [11] Frasconi, P.; Gori, M.; Maggini, M.; and Soda, G. 1996. Representation of finite state automata in recurrent radial basis function networks. *Machine Learning* 23(1):5–32.
- [12] Frosst, N., and Hinton, G. E. 2017. Distilling a neural network into a soft decision tree. In *Proceedings of the First International Workshop on Comprehensibility and Explanation in AI and ML 2017 co-located with 16th International Conference of the Italian Association for Artificial Intelligence (AI*IA 2017), Bari, Italy, November 16th and 17th, 2017*.
- [13] Giles, C. L.; Sun, G.-Z.; Chen, H.-H.; Lee, Y.-C.; and Chen, D. 1990. Higher order recurrent networks and grammatical inference. In *Advances in neural information processing systems*, 380–387.
- [14] Giles, C. L.; Miller, C. B.; Chen, D.; Chen, H.-H.; Sun, G.-Z.; and Lee, Y.-C. 1992. Learning and extracting finite state automata with second-order recurrent neural networks. *Neural Computation* 4(3):393–405.
- [15] Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural computation* 9(8):1735–1780.
- [16] Hopcroft, J. E.; Motwani, R.; and Ullman, J. D. 2003. *Introduction to automata theory, languages, and computation - international edition (2. ed)*.
- [17] Jacobsson, H. 2005. Rule extraction from recurrent neural networks: Ataxonomy and review. *Neural Computation* 17(6):1223–1263.
- [18] Katz, G.; Barrett, C. W.; Dill, D. L.; Julian, K.; and Kochenderfer, M. J. 2017. Reluplex: An efficient SMT solver for verifying deep neural networks. In *Computer Aided Verification - 29th International Conference, CAV 2017, Heidelberg, Germany, July 24-28, 2017, Proceedings, Part I*, 97–117.
- [19] Kolen, J. F. 1994. Fool’s gold: Extracting finite state machines from recurrent network dynamics. In *Advances in neural information processing systems*, 501–508.
- [20] Leofante, F.; Narodytska, N.; Pulina, L.; and Tacchella, A. 2018. Automated verification of neural networks: Advances, challenges and perspectives. *arXiv preprint arXiv:1805.09938*.
- [21] Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2017. Towards deep learning models resistant to adversarial attacks. *CoRR* abs/1706.06083.
- [22] Narodytska, N.; Kasiviswanathan, S. P.; Ryzhyk, L.; Sagiv, M.; and Walsh, T. 2017. Verifying properties of binarized deep neural networks. *arXiv preprint arXiv:1709.06662*.
- [23] Omlin, C. W., and Giles, C. L. 2000. Symbolic knowledge representation in recurrent neural networks: Insights from theoretical models of computation. *Knowledge based neurocomputing* 63–115.
- [24] Papernot, N.; McDaniel, P. D.; Swami, A.; and Harang, R. E. 2016. Crafting adversarial input sequences for recurrent neural networks. In *2016 IEEE Military Communications Conference, MILCOM 2016, Baltimore, MD, USA, November 1-3, 2016*, 49–54.
- [25] Ribeiro, M. T.; Singh, S.; and Guestrin, C. 2016. ”why should I trust you?”: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, San Francisco, CA, USA, August 13-17, 2016*, 1135–1144.
- [26] Schellhammer, I.; Diederich, J.; Towsey, M.; and 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*, 73–78. Association for Computational Linguistics.
- [27] Szegedy, C.; Zaremba, W.; Sutskever, I.; Bruna, J.; Erhan, D.; Goodfellow, I. J.; and Fergus, R. 2013. Intriguing properties of neural networks. *CoRR* abs/1312.6199.
- [28] Tjeng, V.; Xiao, K.; and Tedrake, R. 2017. Evaluating robustness of neural networks with mixed integer programming. *arXiv preprint arXiv:1711.07356*.
- [29] Tomita, M. 1982. Dynamic construction of finite automata from example using hill-climbing. *Proceedings of the Fourth Annual Cognitive Science Conference* 105–108.
- [30] Wang, Q.; Guo, W.; Zhang, K.; II, A. G. O.; Xing, X.; Liu, X.; and Giles, C. L. 2017. Adversary resistant deep neural networks with an application to malware detection. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Halifax, NS, Canada, August 13 - 17, 2017*, 1145–1153.

- [31] Wang, Q.; Zhang, K.; Ororbia II, A. G.; Xing, X.; Liu, X.; and Giles, C. L. 2018a. A comparison of rule extraction for different recurrent neural network models and grammatical complexity. *In submission*.
- [32] Wang, Q.; Zhang, K.; Ororbia II, A. G.; Xing, X.; Liu, X.; and Giles, C. L. 2018b. An empirical evaluation of rule extraction from recurrent neural networks. *Neural Computation* 30(9):2568–2591.
- [33] Weiss, G.; Goldberg, Y.; and Yahav, E. 2018. Extracting automata from recurrent neural networks using queries and counterexamples. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, 5247–5256. PMLR.
- [34] Wu, Y.; Zhang, S.; Zhang, Y.; Bengio, Y.; and Salakhutdinov, R. 2016. On multiplicative integration with recurrent neural networks. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems 2016, December 5-10, 2016, Barcelona, Spain*, 2856–2864.
- [35] Zeng, Z.; Goodman, R. M.; and Smyth, P. 1993. Learning finite state machines with self-clustering recurrent networks. *Neural Computation* 5(6):976–990.