FINITE AUTOMATA CAN BE LINEARLY DECODED FROM LANGUAGE-RECOGNIZING RNNS

Anonymous authors
Paper under double-blind review

ABSTRACT

We study the internal representations that a recurrent neural network (RNN) uses while learning to recognize a regular formal language. Specifically, we train an RNN on positive and negative examples from a regular language, and ask if there is a simple decoding function that maps states of this RNN to states of the minimal deterministic finite automaton (MDFA) for the language. Our experiments show that such a decoding function exists, that it is in fact linear, but that it maps states of the RNN not to MDFA states, but to states of an *abstraction* obtained by clustering small sets of MDFA states into "superstates". A qualitative analysis reveals that the abstraction often has a simple interpretation. Overall, the results suggest a strong structural relationship between internal representations used by RNNs and finite automata, and explain the well-known ability of RNNs to recognize formal grammatical structure.

1 INTRODUCTION

Recurrent neural networks (RNNs) seem "unreasonably" effective at modeling patterns in noisy real-world sequences. In particular, they seem effective at recognizing grammatical structure in sequences, as evidenced by their ability to generate structured data, such as source code (C++, LaTeX, etc.), with few syntactic grammatical errors (Karpathy et al., 2015). The ability of RNNs to recognize *formal languages* – sets of strings that possess rigorously defined grammatical structure – is less well-studied. Furthermore, there remains little systematic understanding of *how* RNNs recognize rigorous structure. What is their internal algorithm? We aim to explain this internal algorithm of RNNs through comparison to fundamental concepts in formal languages, namely finite automata and regular languages.

In this paper, we propose a new way of understanding how trained RNNs represent grammatical structure, by comparing them to finite automata that solve the same language recognition task. We ask: Can the internal knowledge representations of RNNs trained to recognize formal languages be easily mapped to the states of automata-theoretic models that are traditionally used to define grammatical structure? Specifically, we investigate this question for the class of *regular languages*, or formal languages accepted by *finite automata* (FAs).

In our experiments, RNNs are trained on a training dataset of positive and negative examples of strings randomly generated from a given formal language. Next, we ask if there exists a *decoding function*: an isomorphism that maps the hidden states of the trained RNN to the states of a canonical FA. Since there exist infinitely many FA that accept the same language, we focus on the *minimal deterministic finite automaton* (MDFA) — the DFA with the smallest possible number of states – that perfectly recognizes the language.

Our experiments, spanning 500 regular languages, suggest that such a decoding function exists and can be understood in terms of a notion of *abstraction* that is fundamental in classical system theory. An *abstraction* $\mathcal{A}$ of a machine $\mathcal{M}$ (either finite-state, like an FA, or infinite-state, like an RNN) is a machine obtained by clustering some of the states of $\mathcal{M}$ into "superstates". Intuitively, an abstraction $\mathcal{A}$ loses some of the discerning power of the original automaton $\mathcal{M}$, and as such recognizes a superset of the language that $\mathcal{M}$ recognizes. We find that the states of an RNN $\mathcal{R}$, trained to recognize a regular language $\mathcal{L}$, can be decoded into states of an abstraction $\mathcal{A}$ of the MDFA for the language, such that with high probability, $\mathcal{A}$ accepts any input string that is accepted by $\mathcal{R}$. Remarkably, a *linear* decoding function suffices to achieve maximal deocding accuracy: allowing nonlinearity in the decoder does not lead to significant gain. Also, we find the abstraction has low "coarseness", in

the sense that only a few of the MDFA states are clustered, and a qualitative analysis reveals that the abstractions often have simple interpretations.

2 Related Work

RNNs have long been known to be excellent at recognizing patterns in text (Kombrink et al., 2011; Karpathy et al., 2015). Extensive work has been done on exploring the expressive power of RNNs. For example, finite RNNs have been shown as capable of simulating a universal Turing machine (Neto et al., 1997). (Funahashi & Nakamura, 1993) showed that the hidden state of an RNN can approximately represent dynamical systems of the same or less dimensional complexity.

Recent work has also explored the relationship between RNN internals and DFAs through a variety of methods. Although there have been multiple attempts at having RNNs learn a DFA structure based on input languages generated from DFAs and push down automata (Firoiu et al., 1998; Gers & Schmidhuber, 2001; Giles et al., 1992; Miclet & de la Higuera, 1996; Omlin & Giles, 1996a), most work has focused on *extracting* a DFA from the hidden states of a learned RNN. Early work in this field (Giles et al., 1991) demonstrated that grammar rules for regular grammars could indeed be extracted from an learned RNN. Other studies (Omlin & Giles, 1996b) tried to directly extract a DFA structure from the internal space of the RNN, often by clustering the hidden state activations from input stimuli, noting the transitions from one state to another given a particular new input stimuli. Clustering was done by a series of methods, such as K-Nearest Neighbor, K-means, Gaussian Mixture Models, and Density Based Spatial Clustering of Applications with Noise (DBSCAN) (Das & Mozer, 1993; Lawrence et al., 2000). Most recently, Weiss et al. (2018) have achieved state-of-the-art accuracy in DFA extraction by utilizing the L* query learning algorithm. Our work is different from these efforts in that we try to relate the RNN to a ground-truth minimal DFA, rather than extracting a machine directly from the RNN's state space.

The closest piece of related work is by Tiňo et al. (1998). Like our work, this work seeks to relate an RNN state with the state of a DFA. However, the RNN in Tiňo et al. (1998) exactly mimics the DFA; also, the study is carried out in the context of a few specific regular languages that are recognized by automata with 2-3 states. In contrast, our work does not require exact behavioral correspondence between RNNs and DFAs: DFA states are allowed to be *abstracted*, leading to loss of information. Also, the mapping from RNN states to FA states can be approximate, and its accuracy is evaluated quantitatively. We show that this allows us to establish connections between RNNs and DFAs in the setting of a broad class of regular languages that often demand significantly larger automata (with up to 14 states) than those studied by Tiňo et al. (1998).

3 Definitions

We start by introducing some definitions and notation. A *formal language* is a set of strings over a finite alphabet Σ of input symbols. A *Deterministic Finite Automata (DFA)* is a tuple $\mathcal{A} = (Q, \Sigma, \delta, q_0, F)$ where Q is a finite set of states, Σ is a finite alphabet, $\delta : Q \times \Sigma \to Q$ is a deterministic transition function, $q_0 \in Q$ is a starting state and $F \subset Q$ is a set of accepting states. $\mathcal{A}$ reads strings over Σ symbol by symbol, starting from the state q_0 and making state transitions, defined by δ, at each step. It accepts the string if it reaches a final accepting state in F after reading it to the end. The set of strings accepted by a DFA is a special kind of formal language, known as a *regular language*. A regular language $\mathcal{L}$ can be accepted by multiple DFAs; such a DFA $\mathcal{A}_\mathcal{L}$ is *minimal* if there exists no other DFA $\mathcal{A}' \neq \mathcal{A}_\mathcal{L}$ such that $\mathcal{A}'$ exactly recognizes $\mathcal{L}$ and has fewer states than $\mathcal{A}_\mathcal{L}$. It can be shown that this minimal DFA (MDFA), which we denote by $\mathcal{A}^0_\mathcal{L}$, is *unique* (Hopcroft & Ullman, 1979).

Abstractions. A *Nondeterministic Finite Automaton (NFA)* is similar to a DFA, except that the deterministic transition function δ is now a non-deterministic transition relation δ^{NFA}. This means that for a state q in the NFA and $a \in \Sigma$, we have that $\delta^{NFA}(q, a)$ is now a subset of NFA states.

For a given regular language $\mathcal{L}$ we denote by $\mathcal{A}^n_\mathcal{L}$ a *Nondeterministic Finite Automaton (NFA)* with n states that recognizes a superset of the language $\mathcal{L}$. An *abstraction map* is a map $\alpha : Q \to 2^Q$ that combines two states in NFA $\mathcal{A}^n_\mathcal{L}$, resulting in an NFA $\mathcal{A}^{n+1}_\mathcal{L}$; that is, $\mathcal{A}^n_\mathcal{L} \overset{\alpha}{\mapsto} \mathcal{A}^{n+1}_\mathcal{L}$. Since every DFA is also an NFA, we can apply α to the MDFA $\mathcal{A}^0_\mathcal{L}$ to obtain a NFA $\mathcal{A}^1_\mathcal{L}$. Intuitively, α creates a new NFA by combining two states of an existing NFA into a new 'superstate'. An NFA $\mathcal{A}^n_\mathcal{L}$ is an *abstraction* of $\mathcal{A}^0_\mathcal{L}$, if $\mathcal{A}^n_\mathcal{L}$ can be obtained from $\mathcal{A}^0_\mathcal{L}$ by repeated application of the

abstraction map α. Every state of $\mathcal{A}^n_{\mathcal{L}}$ can be viewed as a set of the states of the MDFA $\mathcal{A}^0_{\mathcal{L}}$, i.e $q^n \in Q^n \implies q^n = \{q^0_i\}_{i \in \mathcal{I}}$ with $q^0_i \in Q^0$ for all i.

We define the *coarseness* of an abstraction $\mathcal{A}^n_{\mathcal{L}}$, as the number of applications of α on the MDFA required to arrive at $\mathcal{A}^n_{\mathcal{L}}$. Intuitively, repeated applications of α create NFAs that accept *supersets* of the language $\mathcal{L}$ recognized by the MDFA, and can hence be seen as coarse-grained versions of the original MDFA. The coarsest NFA, given by $\mathcal{A}^{(|Q^0|-1)}_{\mathcal{L}}$, is a NFA with only one accepting node and it accepts all strings on the alphabet Σ.

Given a regular language $\mathcal{L}$, we define $\mathcal{R}_{\mathcal{L}}$ to be a Recurrent Neural Network (RNN) that is trained to recognize the language $\mathcal{L}$, with a certain threshold accuracy. Each RNN $\mathcal{R}_{\mathcal{L}}$ will have a corresponding set of hidden states denoted by H. More details about the RNN are provided in Section 4.1. Note that an RNN can also be viewed as a transition system with 5-tuple $\mathcal{R} = (H, \Sigma, \delta^{\mathcal{R}}, h_0, F^{\mathcal{R}})$, where H is a set of possible 'hidden' states (typically $H \subset \mathbb{R}^K$), $\delta^{\mathcal{R}}$ is the transition function of the trained RNN, h_0 is the initial state of the RNN, and $F^{\mathcal{R}}$ is the set of accepting RNN states. The key distinction between a DFA and a RNN is that the latter has a continuous state space, endowed with a topology and a metric.

Decoding DFA States from RNNs. Inspired by methods in computational neuroscience (Astrand et al., 2014), we can define an *decoding function* or *decoder* $f : H \to Q^0$ as a map from the hidden states of a RNN $\mathcal{R}_{\mathcal{L}}$ to the states of the corresponding (for $\mathcal{L}$) MDFA $\mathcal{A}^0_{\mathcal{L}} = (Q^0, \Sigma^0, \delta^0, q^0_0, F^0)$, such that f preserves transitions: that is, $\delta^0(f(h), a) = f(\delta^{\mathcal{R}}(h, a))$ for all $a \in \Sigma$. We are interested in whether such decoders for MDFA states from trained RNN hidden states exist.

Decoding Abstraction States. Now, let $\mathcal{A}^n_{\mathcal{L}}$ be an abstraction of $\mathcal{A}^0_{\mathcal{L}}$, obtained by applying α to $\mathcal{A}^0_{\mathcal{L}}$ repeatedly n times, and let Q^n be the set of states of $\mathcal{A}^n_{\mathcal{L}}$. We can define an *abstraction decoding function* $\hat{f} : H \to Q^n$, by $\hat{f}(h) := (\alpha^{|n|} \circ f)(h)$, that is the composition of f with $\alpha^{|n|}$ [1] applications of α. Given a dataset of input strings $\mathcal{D} \subset \Sigma^*$, we can define the *decoding accuracy* of a map $\hat{f}$ for an abstraction $\mathcal{A}^n_{\mathcal{L}}$ from RNN $\mathcal{R}_{\mathcal{L}}$ by:

$$\rho_{\hat{f}}(\mathcal{R}_{\mathcal{L}}, \mathcal{A}^n_{\mathcal{L}}) = \frac{1}{|\mathcal{D}|} \sum_{w \in \mathcal{D}} \sum_{t=0}^{|w|-1} \frac{\mathbb{1}(\hat{f}(h_{t+1}) = \alpha^{|n|}(\delta^0(f(h_t), a_t)))}{|d|}$$

where $\mathbb{1}(C)$ is the boolean indicator function that evaluates to 1 if condition C is true and to 0 otherwise, and $h_{t+1} = \delta^{\mathcal{R}}(h_t, a_t)$. Note in particular, that for decoding abstraction states we only consider the deterministic transitions δ^0 and $\delta^{\mathcal{R}}$, of the corresponding MDFA and RNN, and not the induced transition relation of the abstraction $\mathcal{A}^n_{\mathcal{L}}$.

Decoding Abstract State Transitions. We now define a measure that takes into account the induced transition relation δ^n of an abstraction $\mathcal{A}^n_{\mathcal{L}}$. Given a decoder $\hat{f}$ with sufficiently high decoding accuracy for some abstraction $\mathcal{A}^n_{\mathcal{L}}$, we are interested in quantifying how well transitions are preserved with respect to the NFA $\mathcal{A}^n_{\mathcal{L}}$, motivating a definition of transitional accuracy of the map $\hat{f}$ from a given RNN for a given abstraction.

More precisely, for a given decoding function $\hat{f}$ and NFA $\mathcal{A}^n_{\mathcal{L}} = (Q^n, \Sigma, \delta^n, q^n_0, F^n)$, we want to check whether $h_{t+1} = \delta^{\mathcal{R}}(h_t, a)$ implies $\hat{f}(h_{t+1}) = \delta^n(\hat{f}(h_t), a)$. We note that in the definition of the decoding function, we only take into account the transitions in the original MDFA $\mathcal{A}^0_{\mathcal{L}}$, and *not* the new transitions in the abstraction $\mathcal{A}^n_{\mathcal{L}}$, unlike we do here. Finally, the *transitional accuracy* of a map $\hat{f}$ for a given RNN and abstraction, with respect to a data-set $\mathcal{D}$, is defined as:

$$\phi_{\hat{f}}(\mathcal{R}_{\mathcal{L}}, \mathcal{A}^n_{\mathcal{L}}) = \frac{1}{|\mathcal{D}|} \sum_{w \in \mathcal{D}} \sum_{t=0}^{|w|-1} \frac{\mathbb{1}(\hat{f}(h_{t+1}) = \delta^n(\hat{f}(h_t), a_t))}{|d|}$$

Our experiments in the next section demonstrate that decoding functions with high decoding and transitional accuracies exist for abstractions with relatively low coarseness.

[1] The function $\alpha^{|n|}$ is the function obtained by taking n compositions of α with itself.

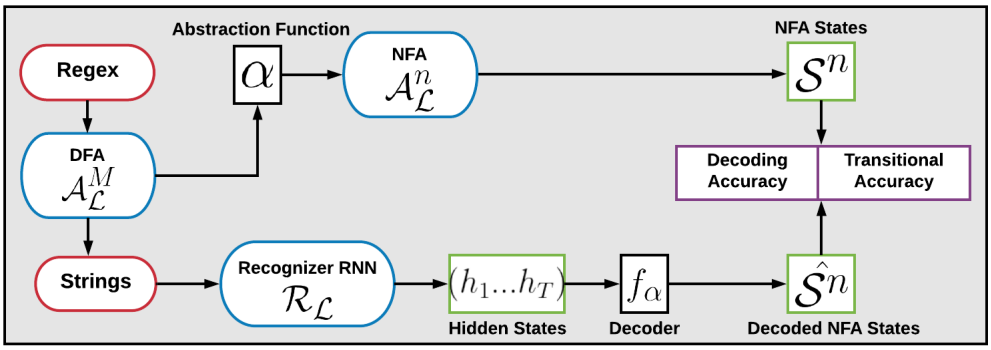

Figure 1: An overview of our experimental setup

4 Experimental Results

Our overall goal is to experimentally test the hypothesis that a high accuracy, low coarseness decoder exists from an RNN to an abstraction of the MDFA. We aim to answer 4 fundamental questions related to the transitional accuracy of $\mathcal{A}_{\mathcal{L}}^{0}$ and $\mathcal{R}$: (1) *How do we choose an appropriate abstraction decoding function $\hat{f}$?* (2) *What necessitates the abstraction function α?* (3) *Can we verify that a low coarseness α and high accuracy $\hat{f}$ exists?* and lastly, (4) *How can we better understand $\mathcal{R}$ in the context of α and $\hat{f}$?*

4.1 Experimental Design

In order to answer the above questions and test our claims with statistical significance, we have designed a flexible framework that facilitates comparisons between MDFA and RNN states, as summarized in Figure 1.

We randomly generate a regular expression that specifies a language $\mathcal{L}$. Let the minimal DFA for recognizing $\mathcal{L}$ be $\mathcal{A}_{\mathcal{L}}^{0}$. We then use $\mathcal{A}_{\mathcal{L}}^{0}$ to randomly generate a training dataset $\mathcal{D} \equiv \mathcal{D}_{+} \bigcup \mathcal{D}_{-}$ of positive ($x_{+} \in \mathcal{L}$) and negative ($x_{-} \notin \mathcal{L}$) example strings from $\mathcal{A}_{\mathcal{L}}^{0}$ (see Appendix for details). We use this dataset to train an RNN $\mathcal{R}_{\mathcal{L}}$ on the language recognition task: given an input string $x \in \Sigma^{*}$, is $x \in \mathcal{L}$? Thus, we have two language recognition models, each of which is a state transition system. We next pass a dataset of test strings through both models, and extract the state trajectories for each. Given a length T input string $\mathbf{x} = (x_1, x_2, x_t, ..., x_T) \in \mathcal{D}$, let the categorical states generated by MDFA $\mathcal{A}_{\mathcal{L}}^{0}$ be denoted by $\mathbf{s}(\mathbf{x}) = (s_0, s_1, s_t, ..., s_T)$ and let the continuous states generated by the trained RNN $\mathcal{R}_{\mathcal{L}}$ be $\mathbf{h} = (h_0, h_1, h_t, ..., h_T)$. The recorded state trajectories for all input strings $\mathbf{x} \in \mathcal{D}$ are used as inputs into our analysis of how MDFA states $\mathbf{s}$ relate to RNN states $\mathbf{h}$. For our experiments, we sample a total of ~ 500 different regular expressions, and thus perform an analysis of ~ 500 recognizer minimal DFAs and ~ 500 trained recognizer RNNs.

4.2 Learning an Accurate Decoder

As mentioned in section 4.1, we must first determine what is a reasonable form for the decoders f and $\hat{f}$ to insure high accuracy on the decoding task. Figure 2 shows decoding accuracy $\mathbb{E}_{\mathcal{D}}[\rho_f(\mathcal{R}, \mathcal{A}_{\mathcal{L}}^{0})|f]$ for several different decoding functions f. We tested two linear classifiers (Multinomial Logistic Regression and Linear Support Vector Machines (SVM)) and two non-linear classifiers (SVM with an RBF kernel, Multi-Layer Perceptrons). In order to evaluate whether decoding accuracy varies significantly amongst all decoders, we use an F-test. Results are shown in Figure 2. Surprisingly, we find there to be no statistical difference: the nonlinear decoders achieve no additional accuracy beyond the simpler linear decoders. We also observe in our experiments that as the size of the MDFA M increases, the decoding accuracy decreases for all decoders in a similar manner. Figure 3a shows this relationship for the multinomial logistic regression classifier.

Taken together, these results have several implications. First, we find that a highly expressive non-linear decoder, does not yield any increase in decoding accuracy, even as we scale up in MDFA complexity. From this, and the fact that we did extensive hyperparameter search in training decoders,

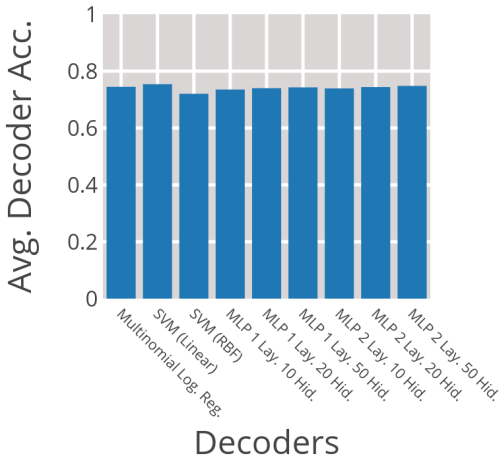

Figure 2: Average decoding accuracy on all MDFAs for Linear and Nonlinear decoders. Linear decoders are the two leftmost while the rest are non-linear. There is no statistical difference between the decoder accuracies.

we can safely conclude that the decoder models we have chosen are indeed expressive enough. Second, we find that decoding accuracy for MDFA states is not very high. These two observations suggest the need for a different interpretation of the internal representation of the trained RNN.

4.3 Why Abstractions are Necessary

Given the observations above, how is the hidden state space of the trained RNN organized? One hypothesis that is consistent with the observations above is that the trained RNN somehow reflect some coarse-grained abstraction of the MDFA states, rather than the MDFA states themselves.[2]

In order to test this hypothesis, we propose a simple greedy algorithm to find an abstraction mapping α: (a) given an NFA $A^n_{\mathcal{L}}$ with n unique states in $\mathcal{S}^n$, consider all $(n-1)$-partitions of $\mathcal{S}^{n-1}$ (i.e. two NFA states s, s' have merged into a single superstate $\{s, s'\}$); (b) select the partition with the highest (linear) decoding accuracy; (c) Repeat this iterative merging process until only a 2-partition remains. We note that this algorithm does not *explicitly* take into consideration the transitions between states which are essential to evaluating $\phi_{\hat{f}}(\mathcal{R}_{\mathcal{L}}, \mathcal{A}^n_{\mathcal{L}})$. Instead, the transitions are taken into account *implicitly* while learning the decoder f at each iteration of the abstraction algorithm.

The abstraction algorithm is greedy in the sense that we may not find the globally optimal partition (i.e. with the highest decoding accuracy and lowest coarseness), but an exhaustive search over all partitions is computationally intractable. The greedy method we have proposed has $\mathcal{O}(M^2)$ complexity instead, and in practice gives satisfactory results. Despite it being greedy, we note that the resulting sequence of clusterings are stable with respect to randomly chosen initial conditions and model parameters. Recognizer RNNs with a different number of hidden units result in clustering sequences that are consistent with each other in the critical first few abstractions.

4.4 Decoding Abstractions and Transitions

Once an abstraction α has been found, we can define this α as an abstraction function $\alpha(\mathcal{A}^n_{\mathcal{L}}) = \mathcal{A}^{n+1}_{\mathcal{L}}$ for $0 \leq n \leq |Q| - 1$. We must then evaluate whether the learned abstraction decoder $\hat{f}$ is of high accuracy, and whether the α found is of low coarseness. Results showing the relationship between high transitional accuracy $\phi_{\hat{f}}(\mathcal{R}_{\mathcal{L}}, \mathcal{A}^n_{\mathcal{L}})$ and low coarseness are presented in Figure 4.

Our fundamental work shows a large scale analysis of how RNNs $\mathcal{R}_{\mathcal{L}}$ relate to abstracted NFAs $\mathcal{A}^n_{\mathcal{L}}$ for hundreds of minimal DFAs, many of which are much larger and more complex than DFAs typically used in the literature. By evaluating the transition accuracy between $\mathcal{R}$ and $\mathcal{A}^n_{\mathcal{L}}$ we empirically validate our claim. We show that there does exist high accuracy decoder from $\mathcal{R}$ to an abstracted NFA $\mathcal{A}^n_{\mathcal{L}}$.

[2]This idea can be motivated by recasting the regular expression for (say) E-Mails into a hierarchical grammar with production rules.

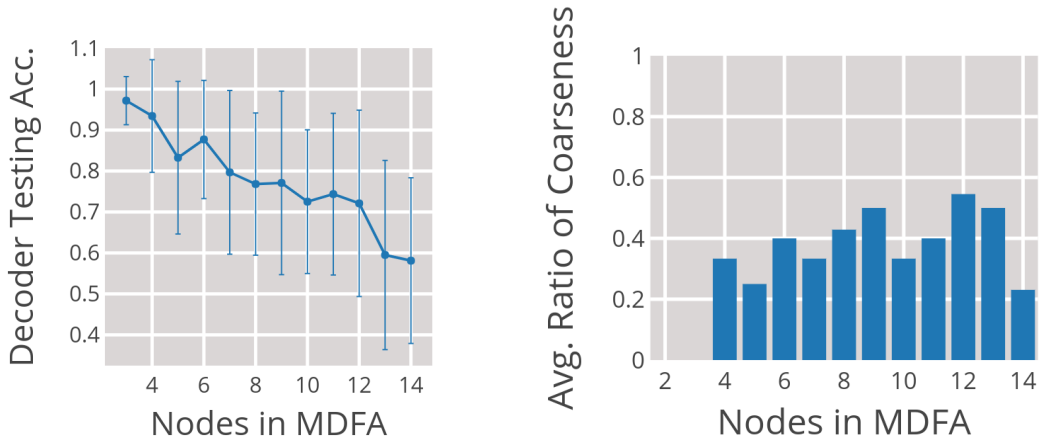

Figure 3: 3a (left): Average linear decoding accuracy as a function of M in the MDFA. 3b (right): Average ratio of coarseness that must be created relative to M in MDFA to achieve 90% testing Accuracy. Suggesting high decoding accuracy with low coarseness.

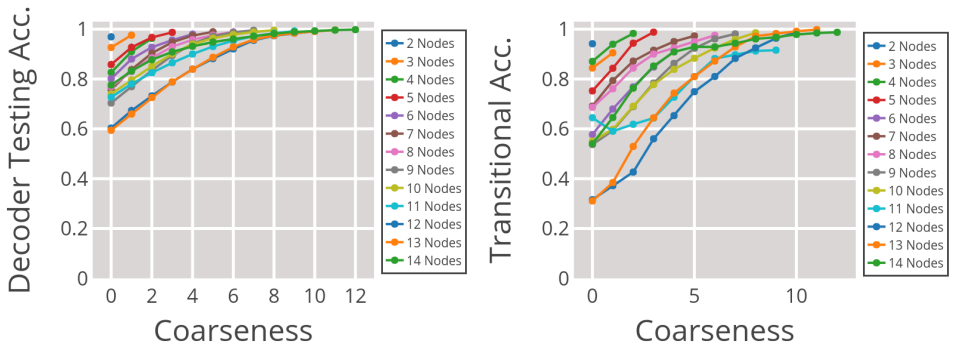

Figure 4: 4a (left): Average linear decoder testing accuracy as a function of coarseness (The number of times α is applied), sorted by the number of nodes in the MDFA. 4b (right): Average transitional accuracy vs. coarseness, sorted by the number of nodes in the MDFA while using a linear decoder

4.5 Interpreting the RNN Hidden State Space with respect to the Minimal DFA

A byproduct of showing that there exists a high accuracy $\phi_{\hat{f}}(\mathcal{R}_{\mathcal{L}}, \mathcal{A}^n_{\mathcal{L}})$ with low coarseness α allows us to gain a unique interpretation of $\mathcal{H}$ with respect to $\mathcal{A}^0_{\mathcal{L}}$. Using α and f to relate the two, we uncover an interpretation of how $\mathcal{R}$ organizes $\mathcal{H}$ with respect to $\mathcal{A}^n_{\mathcal{L}}$ $\forall\ n \in [M]$. We can then determine the appropriate level of abstraction the network uses to accomplish the logical language recognition task in relation to the underlying MDFA. We provide two example 'real-world' DFAs to illustrate this interpretation and show several interesting patterns.

We present the clustering sequences of two regular expressions that have real-world interpretations, namely the Dates and Simple Emails languages that recognize simple dates and simple emails respectively. Figure 6 displays the Dates (top) with its clustering sequence superimposed on the MDFA in the form of a dendrogram. The dendrogram can be read in a top-down fashion, which displays the membership of the MDFA states and the sequence of abstractions up to n. A question then arises: *How should one pick a correct level of abstraction n?*. The answer can be seen in the corresponding accuracies $\rho_{\hat{f}}(\mathcal{R}_{\mathcal{L}}, \mathcal{A}^n_{\mathcal{L}})$ in Figure 7.

As n increases and the number of total NFA states decreases, the LDC prediction task obviously gets easier (100% accuracy when the number of NFA states $Q^{|Q|-1}$ is 1), and hence it is important to

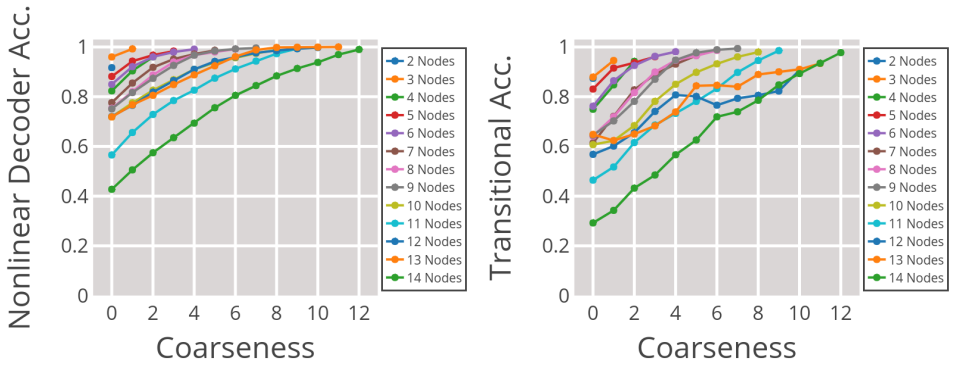

Figure 5: 5a (left): Average nonlinear testing accuracy as a function of coarseness, sorted by the number of nodes in the MDFA. 5b (right): Average transitional accuracy vs coarseness, sorted by the number of nodes in the MDFA while using a non-linear decoder

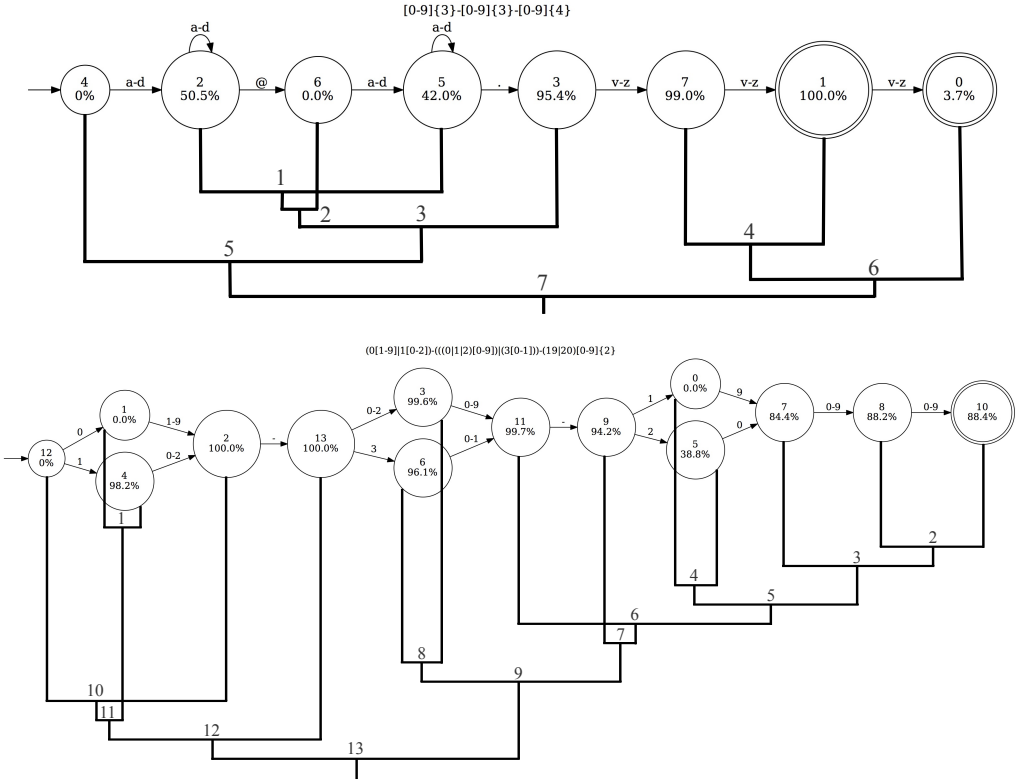

Figure 6: 6a (Top): The MDFA of the SIMPLE EMAILS language with a dendrogram representing the the sequence of abstractions created while using a linear decoder. Showing the initial abstractions are those of the same pattern [a-d]*. 6b (Bottom) The MDFA of the DATES language with a dendrogram representing the the sequence of abstractions created while using a linear decoder. Showing the initial abstractions are those representing states that represent the same moment in time.

consider how to choose the number of abstractions in the final partition. We typically set a threshold for $\rho_{\hat{f}}(\mathcal{R}_{\mathcal{L}}, \mathcal{A}^n_{\mathcal{L}})$ and select the minimum n required to achieve the threshold accuracy.

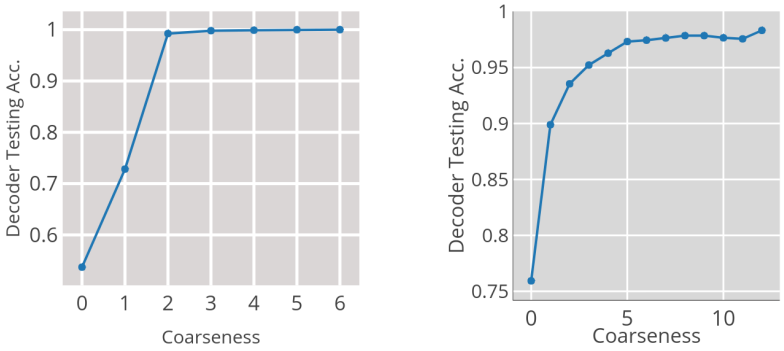

Figure 7: 7a (Left): Linear decoder accuracies as a function of coarseness for the SIMPLE EMAILS language in Figure 6a. 7b (Right): Linear decoder accuracies as a function of coarseness for the DATES language corresponding to Figure 6a.

Consider the first two abstractions of the SIMPLE EMAILS DFA. We notice that both states 2 and 5 represent the pattern matching task `[a-d]*`, because they are agglomerated by the algorithm. Once two abstractions have been made, the decoder accuracy is at a sufficient point, as seen in Figure 7. This suggests that the collection of hidden states for the two states are not linearly separable. One possible and very likely reason for this is the network has learned an abstraction of the pattern `[a-d]*` and uses the same hidden state space regardless of location in string to recognize this pattern, which has been indicated in past work(Karpathy et al., 2015). This intuitive example demonstrates the RNN's capability to learn and abstract patterns from the DFA. This makes intuitive sense because $\mathcal{R}_{\mathcal{L}}$ does not have any direct access to $\mathcal{A}^0_{\mathcal{L}}$, only to samples generated from $\mathcal{A}^0_{\mathcal{L}}$. The flexibility of RNNs allows such abstractions to be created easily.

The second major pattern that arises can be seen in the dendrogram in the bottom row of Figure 6. We notice that, generally, multiple states that represent the same location in the input string get merged (1 and 4, 3 and 6, 0 and 5). The SIMPLE EMAILS dendrogram shows patterns that are location-independent, while the fixed length pattern in the DATES regex shows location-dependent patterns. We also notice that the algorithm tends to agglomerate states that are within close sequential proximity to each other in the DFA, again indicating location-dependent hierarchical priors. Overall, our new interpretation of $\mathcal{H}$ reveals some new intuitions, empirically backed by our decoding and transitional accuracy scores, regarding how the RNN $\mathcal{R}_{\mathcal{L}}$ structures the hidden state space $\mathcal{H}$ in the task of language recognition.

5 CONCLUSIONS

We have studied how RNNs trained to recognize regular formal languages represent knowledge in their hidden state. Specifically, we have asked if this internal representation can be decoded into canonical, minimal DFA that exactly recognizes the language, and can therefore be seen to be the "ground truth". We have shown that a linear function does a remarkably good job at performing such a decoding. Critically, however, this decoder maps states of the RNN not to MDFA states, but to states of an *abstraction* obtained by clustering small sets of MDFA states into "abstractions". Overall, the results suggest a strong structural relationship between internal representations used by RNNs and finite automata, and explain the well-known ability of RNNs to recognize formal grammatical structure.

We see our work as a fundamental step in the larger effort to study how neural networks learn formal logical concepts. We intend to explore more complex and richer classes of formal languages, such as context-free languages and recursively enumerable languages, and their neural analogs.

REFERENCES

Elaine Astrand, Pierre Enel, Guilhem Ibos, Peter Ford Dominey, Pierre Baraduc, and Suliann Ben Hamed. Comparison of classifiers for decoding sensory and cognitive information from prefrontal neuronal populations. *PLOS ONE*, 9(1):pages 1–14, 01 2014. .

Sreerupa Das and Michael Mozer. A unified gradient-descent/clustering architecture for finite state machine induction. In *Advances in Neural Information Processing Systems 6, NIPS 1993, Denver, Colorado, USA*, pp. 19–26, .

Laura Firoiu, Tim Oates, and Paul R. Cohen. Learning deterministic finite automaton with a recurrent neural network. In *Grammatical Inference, 4th International Colloquium, ICGI-98, Ames, Iowa, USA, July 12-14, 1998, Proceedings*, pp. 90–101, 1998. .

Ken-ichi Funahashi and Yuichi Nakamura. Approximation of dynamical systems by continuous time recurrent neural networks. *Neural Networks*, 6(6):801–806, 1993. .

Felix A. Gers and Jürgen Schmidhuber. LSTM recurrent networks learn simple context-free and context-sensitive languages. *IEEE Trans. Neural Networks*, 12(6):pages 1333–1340, 2001. .

C. Lee Giles, Clifford B. Miller, Dong Chen, Guo-Zheng Sun, Hsing-Hen Chen, and Yee-Chun Lee. Extracting and learning an unknown grammar with recurrent neural networks. In *Advances in Neural Information Processing Systems 4, [NIPS Conference, Denver, Colorado, USA, December 2-5, 1991]*, pp. 317–324, .

C. Lee Giles, Clifford B. Miller, Dong Chen, Hsing-Hen Chen, Guo-Zheng Sun, and Yee-Chun Lee. Learning and extracting finite state automata with second-order recurrent neural networks. *Neural Computation*, 4(3):pages 393–405, 1992. .

John E. Hopcroft and Jeffrey D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, 1979. .

Andrej Karpathy, Justin Johnson, and Fei-Fei Li. Visualizing and understanding recurrent networks. *Computing Research Repository, CoRR*, abs/1506.02078, .

Stefan Kombrink, Tomas Mikolov, Martin Karafiát, and Lukás Burget. Recurrent neural network based language modeling in meeting recognition. pp. pages 2877–2880, .

Steve Lawrence, C. Lee Giles, and Sandiway Fong. Natural language grammatical inference with recurrent neural networks. *IEEE Transactions on Knowledge and Data Engineering*, 12(1):pages 126–140, 2000. .

Laurent Miclet and Colin de la Higuera. *Grammatical Inference: Learning Syntax from Sentences*. Springer, .

Anders Møller. dk.brics.automaton – finite-state automata and regular expressions for Java, . `http://www.brics.dk/automaton/`.

João Pedro Guerreiro Neto, Hava T. Siegelmann, José Félix Costa, and Carmen Paz Suárez Araujo. Turing universality of neural nets (revisited). pp. pages 361–366, 1997. .

Christian W. Omlin and C. Lee Giles. Constructing deterministic finite-state automata in recurrent neural networks. *Journal of the Association of Computing Machinery, JACM*, 43(6):pages 937–972, 1996a. .

Christian W. Omlin and C. Lee Giles. Extraction of rules from discrete-time recurrent neural networks. *Neural Networks*, 9(1):41–52, 1996b. .

Peter Tiňo, Bill G. Horne, C. Lee Giles, and Pete C. Collingwood. Chapter 6 - finite state machines and recurrent neural networks — automata and dynamical systems approaches. In Omid Omidvar and Judith Dayhoff (eds.), *Neural Networks and Pattern Recognition*, pp. 171 – 219. Academic Press, San Diego, 1998. ISBN 978-0-12-526420-4. .

Gail Weiss, Yoav Goldberg, and Eran Yahav. Extracting automata from recurrent neural networks using queries and counterexamples. In *Proceedings of the 35th International Conference on Machine Learning, ICML 2018, Stockholmsmässan, Stockholm, Sweden, July 10-15, 2018*, pp. 5244–5253, .

A DATASET GENERATION

In order to generate a wide variety of strings that are both accepted and rejected by the DFA corresponding to a given regex R, we use the Xeger Java library, built atop the `dk.brics.automaton` library Møller (2017). The Xeger library, given a regular expression, generates strings that are accepted by the regular expression's corresponding DFA. However, there is no standard method to generate examples that would be rejected by the DFA. These rejected examples need to be diverse to properly train an acceptor/rejector model: if the rejected examples are completely different from the accepted examples, the model will not be able to discern between similar input strings, even if one is an accepted string and the other is a rejected string. However, if the rejected examples were too similar to the accepted examples, the model would not be able to make a judgment on a completely new string that does not resemble any input string seen during training. In other words we want the rejected strings to be drawn from two distinct distributions, one similar and one independent compared to the distribution of the accepted strings. In order to achieve this, we generate negative examples in two ways: First, we randomly swap two characters in an accepted example enough times until we no longer have an accepted string. And secondly, we take an accepted string and randomly shuffle the characters, adding it to our dataset if the resulting string is indeed rejected.

In our experiments we typically generate 1000 training examples with a 50:50 accept/reject ratio. When applicable we generate strings of varying length capped at some constant, for example with the SIMPLE EMAILS language we generate strings of at most 20 characters.

A.1 EXAMPLE REGULAR EXPRESSIONS AND CORRESPONDING DFAS USED

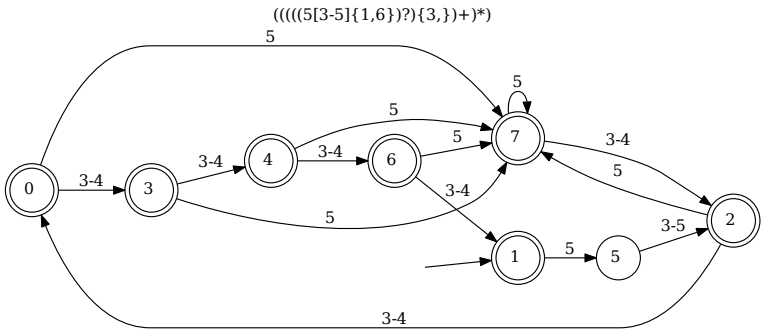

Figure 8: An example of a typical regex and DFA randomly generated from our framework.

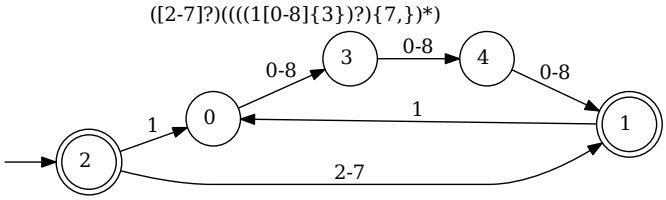

Figure 9: An example of a typical regex and DFA randomly generated from our framework.