

# Learning Transductions and Alignments with RNN Seq2seq Models

**Zhengxiang Wang**

ZHENGXIANG.WANG@STONYBROOK.EDU

*Department of Linguistics*

*Institute for Advanced Computational Science*

*Stony Brook University*

**Editors:** François Coste, Faissal Ouardi and Guillaume Rabusseau

## Abstract

The paper studies the capabilities of Recurrent-Neural-Network sequence to sequence (RNN seq2seq) models in learning four transduction tasks: identity, reversal, total reduplication, and quadratic copying. These transductions are traditionally well studied under finite state transducers and attributed with increasing complexity. We find that RNN seq2seq models are only able to approximate a mapping that fits the training or in-distribution data, instead of learning the underlying functions. Although attention makes learning more efficient and robust, it does not overcome the out-of-distribution generalization limitation. We establish a novel complexity hierarchy for learning the four tasks for attention-less RNN seq2seq models, which may be understood in terms of the complexity hierarchy of formal languages, instead of string transductions. RNN variants also play a role in the results. In particular, we show that Simple RNN seq2seq models cannot count the input length.

**Keywords:** RNNs, sequence to sequence, alignments, string transductions, attention, generalization abilities, complexity hierarchy, formal language theory

## 1. Introduction

Given the black-box nature of neural networks, learning formal languages has emerged as ideal proxy tasks for evaluating the expressive power and generalization capacity of neural networks (Grefenstette et al., 2015; Avcu et al., 2017; Bhattamishra et al., 2020; Delétang et al., 2022; van der Poel et al., 2023). Unlike real-world learning tasks, formal language tasks as well as their underlying functions are typically well-understood and known in advance. This makes possible a more flexible and complete control over data and, as a result, more fine-grained analyses of results obtained in experiments. Moreover, the rich tradition of studying formal languages offers critical insights into interpreting the learning results of neural networks, such as, from automata-theoretic perspectives (Peng et al., 2018; Merrill, 2019b; Ayache et al., 2019; Delétang et al., 2022).

The current paper examines the learning capabilities of Recurrent-Neural-Network sequence to sequence (RNN seq2seq) models in the context of string transductions. It focuses on the abilities of various configurations of RNN seq2seq models to learn a series of four transduction tasks of increasing complexity and generalize to unseen in-distribution and out-of-distribution examples. The four tasks include identity, reversal, total reduplication, and quadratic copying (see §3 for a review). Previous works on RNN seq2seq models have only investigated identity (Grefenstette et al., 2015), reversal (Grefenstette et al., 2015;

Merrill, 2019b), and total reduplication (Nelson et al., 2020), but not quadratic copying. To the best of our knowledge, this paper is the very first study that analyzes the learning capabilities of three major RNN seq2seq model variants, with and without attention, on these four transduction tasks altogether. Both in-distribution and out-of-distribution generalization abilities are studied to better understand models’ generalization capacity, which is central to a more rigorous and interpretable science of machine learning and its reliable application in real world (Doshi-Velez and Kim, 2017). The concept of distribution is closely related to the input sequence lengths (Delétang et al., 2022), as detailed in §3.2.

Through controlled and reproducible experiments, we present the very first comparative and comprehensive results of RNN seq2seq models with various configurations learning the four transduction tasks. These results allow us to draw a wide range of conclusions about the learning capabilities of these models and some contributing factors, such as attention, task complexity, and RNN variants. This study showcases the fruitfulness of probing neural networks with formal string transductions, a nascent approach (Delétang et al., 2022).

The paper proceeds as follows. We review essential technical preliminaries about RNN seq2seq models in §2 and introduce the four learning tasks in §3. Experimental setups, including data, model training details, and evaluation methods, are described in §4. We present the results in §5, which are summarized and further discussed in §6.

The source code, data, model training logs, trained models, and experimental results are open-sourced at <https://github.com/jaaack-wang/rnn-seq2seq-learning>.

## 2. Preliminaries

### 2.1. RNNs

RNNs represent a neural network architecture that utilizes repeated application of a recurrent unit to process a variable-length sequence  $\mathbf{x} = (x_1, \dots, x_T)$ . At each time step  $t$ , the recurrent unit computes a vector  $\mathbf{h}_t \in \mathbb{R}^{D \times 1}$  by taking as inputs the embedding  $\mathbf{e}_t$  of the current input symbol  $x_t$  (via an embedding layer  $\mathbf{E}$ ) and the previous hidden state  $\mathbf{h}_{t-1}$

$$\mathbf{h}_t = f(\mathbf{h}_{t-1}, \mathbf{e}_t) \quad (1)$$

where  $f(\cdot)$  is a non-linear state transition function and varies among different variants of RNNs. The hidden state is commonly initialized as a zero vector. A non-initial hidden state  $\mathbf{h}_t$  may be passed to an output layer to compute the probability distribution of the next symbol  $\hat{x}_{t+1}$  over an output alphabet  $\Gamma$  of size  $N$ , using *softmax*

$$p(\hat{x}_{t+1,i} = 1 \mid x_t, \dots, x_1) = \frac{\exp(\mathbf{w}_i \mathbf{h}_t)}{\sum_{i'=1}^N \exp(\mathbf{w}_{i'} \mathbf{h}_t)} \quad (2)$$

where  $\hat{x}_{t+1,i} = 1$  denotes  $\hat{x}_{t+1}$  being the  $i$ th symbol  $\in \Gamma$  using one-hot encoding and  $\mathbf{w}_i \in \mathbb{R}^{1 \times D}$  is a weight vector associated with that symbol. For the purpose of sequence generation, the embedding  $\hat{\mathbf{e}}_{t+1}$  for  $\hat{x}_{t+1}$  along with  $\mathbf{h}_t$  can be passed as inputs to the recurrent unit to compute the subsequent hidden states and output symbols via the iterative or auto-regressive application of Eq.(1) and Eq.(2).

This study uses three common variants of RNNs (Goldberg, 2016): Simple RNN (SRNN, Elman, 1990), Long Short-term Memory (LSTM, Hochreiter and Schmidhuber, 1997), and

Gated Recurrent Units (GRU, [Cho et al., 2014](#)). The main difference among these three types of RNNs lie in the construction of the recurrent unit, where LSTM and GRU come with additional gating mechanisms to control information flow across time steps, and LSTM has a cell state besides the hidden state. The mathematical details of the state transition functions for each type of RNN are provided in [Appendix A](#). For simplicity and interpretability, all RNNs in this study are single-layered and unidirectional.

## 2.2. RNN seq2seq models

A RNN seq2seq model is an encoder-decoder structure where both the encoder and decoder are RNNs ([Sutskever et al., 2014](#); [Cho et al., 2014](#)). Given a pair of variable-length sequences  $\mathbf{x} = (x_1, \dots, x_T)$  and  $\mathbf{y} = (y_1, \dots, y_{T'})$ , the encoder consumes the input sequence  $\mathbf{x}$  sequentially until the final hidden state  $\mathbf{h}_T^{enc}$  is produced. The decoder takes as initial inputs  $\mathbf{h}_T^{enc}$  and a preset start symbol  $\langle s \rangle$  and is trained to auto-regressively generate an output sequence  $\hat{\mathbf{y}} = (\hat{y}_1, \dots, \hat{y}_{T'})^1$  to approximate  $\mathbf{y}$  as much as possible. A preset end symbol  $\langle /s \rangle$  is also used to signal the termination of generation. Both the start and end symbols are appended to  $\mathbf{x}$  and  $\mathbf{y}$  in our experiments.

In this study, we train three types of RNN seq2seq models, i.e., SRNN seq2seq, GRU seq2seq, and LSTM seq2seq, where the encoders and decoders are RNNs of same variant and with same hidden size. All the models are trained end-to-end by minimizing the cross-entropy loss between  $\hat{\mathbf{y}}$  and  $\mathbf{y}$  through mini-batch gradient descent.

## 2.3. Attention

Attention<sup>2</sup> is a mechanism that allows the decoder in a seq2seq model to access information from all hidden states  $\mathbf{H}^{enc} \in \mathbb{R}^{D \times T}$  of the encoder. It is first proposed to improve the performance of neural machine translation ([Bahdanau et al., 2015](#); [Luong et al., 2015](#)) and has later on been found to be a critical component of the Transformer architecture ([Vaswani et al., 2017](#)). Attention has been hypothesized as external memory resources ([Graves et al., 2014](#)) or a “weighted skip connection” ([Britz et al., 2017](#)) to account for the success of seq2seq models augmented with attention.

Formally, attention typically works as follows. At each decoding time step  $t$ , an attentional weight vector  $\mathbf{a}_t \in \mathbb{R}^{T \times 1}$  can be computed by

$$\mathbf{a}_{t,i} = \frac{\exp(\text{score}(\mathbf{h}_t^{dec}, \mathbf{h}_i^{enc}))}{\sum_{i'=1}^T \exp(\text{score}(\mathbf{h}_t^{dec}, \mathbf{h}_{i'}^{enc}))} \quad (3)$$

where  $\mathbf{a}_{t,i}$  is a scalar weight that corresponds to the  $i_{th}$  hidden state  $\mathbf{h}_i^{enc}$  of the encoder and  $\text{score}$  a function that measures how well  $\mathbf{h}_t^{dec}$  aligns with  $\mathbf{h}_i^{enc}$  for  $i \in \{1, \dots, T\}$ . A context vector  $\mathbf{c}_t \in \mathbb{R}^{D \times 1}$  can be computed by weighing  $\mathbf{H}^{enc}$  with  $\mathbf{a}_t$  through matrix multiplication, then concatenated with the embedding for  $\hat{y}_t$ , and together consumed by the decoder to generate an output. There are many variants of the  $\text{score}$  functions as in [Eq.\(3\)](#) ([Luong et al., 2015](#); [Vaswani et al., 2017](#)). This study uses a simple one as follows

- 
1. This is only for training where  $|\hat{\mathbf{y}}|$  is typically set equal to  $|\mathbf{y}|$  in practice to achieve parallel computation. At inference,  $|\hat{\mathbf{y}}|$  can differ from  $|\mathbf{y}|$ , depending on when the preset end symbol is generated.
  2. Here we only consider the so-called “global attention” where the encoder’s hidden states are all accessible.

$$\text{score}(\mathbf{h}_t^{\text{dec}}, \mathbf{h}_i^{\text{enc}}) = \mathbf{v}_a \tanh(\mathbf{W}_a[\mathbf{h}_t^{\text{dec}}; \mathbf{h}_i^{\text{enc}}]) \quad (4)$$

where  $\mathbf{W}_a \in \mathbb{R}^{D \times 2D}$  and  $\mathbf{v}_a \in \mathbb{R}^{1 \times D}$  are learnt weights to reduce the concatenated hidden states  $[\mathbf{h}_t^{\text{dec}}; \mathbf{h}_i^{\text{enc}}] \in \mathbb{R}^{2D \times 1}$  to an alignment score.  $\tanh$  is a hyperbolic tangent function.

## 2.4. RNNs versus RNN seq2seq models

RNNs are traditionally associated with finite state automata (Giles et al., 1992; Siegelmann, 1996; Visser, 2000; Chen et al., 2018; Merrill, 2019b). To model transductions, RNNs write an output symbol right after reading an input symbol, just like finite state transducers (FSTs). RNN seq2seq models, however, consume all the input symbols before producing any output symbols. The two distinct ways of processing the inputs and generating the outputs form two different classes of neural networks. Consequently, previous results based on RNNs should not be confused with the current results based on RNN seq2seq models.

That said, both RNNs and RNN seq2seq models are sequential models. When trained with backpropagation and gradient descent, the long-term dependency learning issue identified in RNNs (Bengio et al., 1994) is also an issue for attention-less RNN seq2seq models without the skip connections (Britz et al., 2017) introduced by attention. This is because the backpropagated gradients may either decay or grow exponentially, resulting in the notorious problems of vanishing or exploding gradients (Pascanu et al., 2013; Chandar et al., 2019). When either happens, weight optimization becomes impossible, which prevents learning.

## 3. Learning tasks

### 3.1. Task description and FST characterizations

We are interested in the following four learning tasks, representable by four deterministic string-to-string functions with an input alphabet  $\Sigma$  and an output alphabet  $\Gamma$ : (A) identity; (B) reversal; (C) total reduplication; (D) quadratic copying. For a given string  $w \in \Sigma^*$ ,  $f_A(w) = w$ ,  $f_B(w) = w^R$ ,  $f_C(w) = ww$ , and  $f_D(w) = w^{|w|}$ , where  $w^R$  denotes the reverse of  $w$ . For example, if  $w = abc$ , then  $f_A(abc) = abc$ ,  $f_B(abc) = cba$ ,  $f_C(abc) = abcabc$ , and  $f_D(abc) = (abc)^3 = abcabcabc$ . Let  $|w| = n$ .  $|f_D(w)| = n^2$ . Hence the name of quadratic copying or squaring function for  $f_D$ . For all these functions,  $\Sigma = \Gamma$ .




Figure 1: An example 2-way FST to model total reduplication  $f_C : w \rightarrow ww$ , with  $w$  padded into  $\times w \times$  as the input.  $\lambda$ : empty string;  $+1$ : move right;  $-1$ : move left.

Traditionally, the four tasks are modelled with FSTs (Filiot and Reynier, 2016; Bojanczyk et al., 2019; Dolatian and Heinz, 2020; Rawski et al., 2023). More concretely,  $f_A$  can be easily modelled by a 1-way FST where each input symbol is simply mapped to itself,

whereas a 2-way FST that can read input sequences back and forth is required for modelling  $f_B$  and  $f_C$ , such as the one in Fig 1. To model  $f_D$ , a 2-way FST enhanced with the capability of counting the length of the input string is needed. As these four tasks require FSTs of increasing expressive capacity, they are characterized accordingly (Rawski et al., 2023), with  $f_A$  being a rational function,  $f_B$  and  $f_C$  regular functions, and  $f_D$  a polyregular function (Bojanczyk et al., 2019). Under the FST-theoretic characterizations,  $f_D > f_C > f_B > f_A$ , where  $>$  is a “more complex than” relation. Although  $f_B$  and  $f_C$  are both regular functions, a 2-way FST has to scan any input sequence three times (forward-backward-forward) to model  $f_C$ , instead of two times (forward-backward) to model  $f_B$ .

### 3.2. In-distribution and out-of-distribution

Given the deterministic nature of the four functions above, we define the concept of in-distribution and out-of-distribution in terms of the input sequences. For a model trained on input sequences of lengths  $\mathcal{L}$ , in-distribution input sequences are those whose lengths  $\mathcal{L}' \subseteq \mathcal{L}$ . In the context of this study, input sequences are out-of-distribution if  $\mathcal{L}' \cap \mathcal{L} = \emptyset$ . For example, for a set of strings  $\mathcal{S}$  uniformly sampled from  $\Sigma^l$ , a string  $w$  is out-of-distribution with regard to  $\mathcal{S}$  if  $|w| \neq l$  for the reason that  $w \notin \mathcal{S}$  for every  $\mathcal{S} \subseteq \Sigma^l$ .

Distinguishing in-distribution and out-of-distribution input sequences allows us to examine a trained model’s ability to generalize to examples that are independent and identically distributed and those that are beyond, in relation to the distribution of the training examples. Furthermore, a trained model’s out-of-distribution generalization ability reveals whether the model learns the underlying function or approximates the in-distribution data.

### 3.3. Complexity hypothesis

As discussed in §2.4, RNN seq2seq models take an encoder-decoder structure, where the decoder only writes after the encoder read all the input symbols, unlike the read-and-write operation seen in FSTs or RNNs. Therefore, for a RNN seq2seq model to learn these tasks, the decoder must recall all the input symbols from the information passed from the encoder and select the output symbols in correct alignments with the input symbols. In other words, the four tasks require RNN seq2seq models to learn varying input-target alignments or dependencies. Fig 2 illustrates the conjectured mechanism for learning identity and reversal. Total reduplication and quadratic copying can be learnt in a similar process, as the outputs of these two functions can be seen as the concatenation of multiple identity functions applied in a sequence (Rawski et al., 2023). To learn quadratic copying, the model should additionally be able to count the length of the input sequence.

Accordingly, we propose the following task complexity hierarchy for attention-less RNN seq2seq models: quadratic copying ( $f_D$ )  $>$  total reduplication ( $f_C$ )  $>$  identity ( $f_A$ )  $>$  reversal ( $f_B$ ). Unlike the hierarchy established in §3.1 under FSTs where  $f_B > f_A$ , this hierarchy claims that  $f_A > f_B$ . As can be seen in Fig 2, given the same input sequence of length  $n$ , the input-target dependencies for  $f_A$  are of constant distance, i.e.,  $n$ . However, there are initially  $\frac{n}{2}$  or so shorter dependencies for  $f_B$ , which leads to easier weight optimization with the iterative application of backpropagation and gradient descent (Sutskever et al., 2014). When the long-term dependency learning issue discussed in §2.4 becomes an issue, this results in less complexity in learning  $f_B$  than  $f_A$ . Similarly, since  $f_C$  and  $f_D$  introduces




Figure 2: The conjectured mechanism for RNN seq2seq models learning identity and reversal. The multiple crossings on the top relate to identity; the multiple nested crossings at the bottom relate to reversal.

more and longer dependencies than  $f_A$ , they are necessarily harder. Also, with longer target sequences for given input sequences requires more expressive power to distinguish the right output sequences among exponentially growing possible candidates. For this reason,  $f_D$  is the most complex to learn since its target sequence length grows quadratically as a function of the input length, instead of linearly as for  $f_A$  and  $f_B$ .

The complexity hypothesis proposed above may be understood in terms of the complexity hierarchy of formal languages, if we combine the input and target sequences together and take the transduction tasks as language recognition tasks. For example, if we replace the arrow in Fig 2 with  $\#$  and concatenate the input and target sequences, we have a copy language  $w\#w$  for  $f_A$  and a palindrome language  $w\#w^R$  for  $f_B$ . According to the Chomsky’s hierarchy (Chomsky, 1956), these two languages belong to context sensitive language (CSL) and context free language (CFL), respectively, where CSL is more complex than CFL. Similarly, the two respective languages corresponding to  $f_C$  and  $f_D$  are  $w\#ww$  and  $w\#w^{|w|}$ , where  $w\#ww$  as well as  $w\#w$  is considered a mildly CSL (Joshi, 1985) whereas  $w\#w^{|w|}$  a strict CSL. Obviously, under the complexity hierarchy of formal languages, the languages associated with the four functions also imply that  $f_D > f_C > f_A > f_B$ .

For attentional models, however, we hypothesize that  $f_D > f_C > f_A \geq f_B$ , where  $\geq$  means “no less complex than”. This is because attention allows the decoder to access all the encoder’s hidden states at any decoding time step, potentially alleviating the long-term dependency learning issue identified in the RNNs without attention.

#### 4. Experimental setups

To ensure a fair evaluation of the learnability of the four tasks by RNN seq2seq models in a finite setting, we equipped all models with a decently large parameter size such that the lack of sufficient capacity to fit the training data is less of a bottleneck. For the same reason, we also utilized various training techniques to improve the success rate of convergence for all the models. To make the results comparable across models of varying configurations and across different tasks, the input sequences and the training and evaluation conditions were deliberately set identical for every model trained and evaluated.

### 4.1. Data

To simulate a more realistic learning setting, we set  $\Sigma$  and  $\Gamma$  both to be the 26 lowercase English letters. For every  $l \in \{1, \dots, 30\}$ , we randomly sampled from  $\Sigma^l$  the same number of strings as the input sequences, where duplicates are only allowed for sequences of lengths 1 and 2. The target sequences were obtained by applying the four deterministic functions that represent the tasks. In-distribution strings are those of input lengths 6-15, available in the train, dev (development), and test sets. Out-of-distribution strings are those of input lengths 1-5 and 16-30, available only in the gen (generalization) set.

For the train/dev sets, there are 1,000 input-target pairs per input length. For test/gen sets, the number is 5,000 per input length. The test and gen sets were made five times larger than the train/dev sets for the sake of more reliable evaluations. The four datasets are mutually disjoint. More details about the data can be found in Appendix B.1.

### 4.2. Training details

Models were constructed in PyTorch (Paszke et al., 2019) and run on Google Colaboratory. The single-layered RNN encoder and decoder in each model are both of hidden size 512 and contain an embedding layer of embedding size 128. We initialized all trainable weights using Xavier initialization (Glorot and Bengio, 2010) to reduce the vanishing gradient problem. We applied Adam optimizer (Kingma and Ba, 2015) with  $5e-4$  learning rate and  $1e-5$  L2 weight decay rate. We normalized gradients with maximum norm clipped at 1 (Pascanu et al., 2013) to alleviate or circumvent the exploding gradient problem (Zhang et al., 2020). Details about the model parameter configuration and size are provided in Appendix B.2.

To speed up convergence at the training time, we employed a technique called teacher forcing (Williams and Zipser, 1989) to permit the decoders to access the real next symbols from the target sequences, instead of using the predicted ones as inputs. All models were trained up to 500 epochs with the train/dev sets performances evaluated every 10 epochs. The batch size is 1,000 and every batch only contained input sequences of same lengths to avoid padding, which changes the input-target mapping. Training only stopped if one of the following conditions was met: (1) models trained through all the epochs; (2) the full-sequence accuracy (see §4.3) in the dev set reaches exactly 100% with double precision; (3) the full-sequence accuracy in the train set and dev set exceeds 99.99% and 99.50% simultaneously. Every model was trained and evaluated for 3 runs. In each run, the model with the highest full-sequence accuracy on the dev set was saved and deployed to the test and gen sets. Appendix B.3 provides some important training notes.

### 4.3. Evaluation methods

We used the following three metrics to evaluate how well RNN seq2seq models learn the input-target alignments for the four tasks: full-sequence accuracy, first  $n$ -symbol accuracy, and overlap rate. All these metrics are measured from the initial symbol to the end symbol  $\langle /s \rangle$  of the target sequences against the corresponding output sequences generated to the same lengths (for efficient training within batches). Full-sequence accuracy measures the percentage of the target sequences being correctly generated as a whole, whereas first  $n$ -symbol accuracy measures the average proportion of the first  $n$  symbols being correctly

generated for the target sequences. Overlap rate is the average pairwise overlapping ratio of the output sequences to the target sequences.

These three metrics provide well-rounded measurements of alignments between two sequences. When a more restrictive metric shows a low score and thus becomes less discriminative, there can be a more fine-grained alternative, because neural networks do not learn the input-target alignments strictly from the very left incrementally to the very right, due to the random initialization of model parameters. In this study, we used the full-sequence accuracy as the main metric and reported the last two metrics when necessary.

Table 1: Aggregate full-sequence accuracy (%) across the four learning tasks for models with various configurations. Best results are in **bold** for the test and gen sets.

Task	Dataset	Attentional			Attention-less		
		SRNN	GRU	LSTM	SRNN	GRU	LSTM
Identity	Train	100.00	100.00	100.00	69.74	98.26	100.00
	Test	99.97	<b>100.00</b>	<b>100.00</b>	42.82	70.46	<b>77.57</b>
	Gen	25.52	<b>37.41</b>	36.37	0.00	<b>10.41</b>	10.01
Rev	Train	100.00	100.00	100.00	100.00	100.00	100.00
	Test	<b>99.98</b>	99.87	99.88	<b>99.55</b>	88.46	92.85
	Gen	<b>40.14</b>	23.54	25.79	<b>23.89</b>	19.72	12.42
Total Red	Train	100.00	100.00	99.99	15.22	90.57	93.51
	Test	99.71	<b>99.77</b>	99.64	5.60	50.76	<b>55.17</b>
	Gen	<b>42.34</b>	23.23	20.31	0.00	4.39	<b>6.18</b>
Quad Copy	Train	2.43	79.84	82.73	1.62	49.29	67.29
	Test	1.99	67.75	<b>73.89</b>	0.61	27.76	<b>38.03</b>
	Gen	1.36	<b>8.20</b>	6.07	0.00	<b>0.85</b>	0.18
Average	Train	75.61	94.96	95.68	46.65	84.53	90.19
	Test	75.41	91.85	<b>93.35</b>	37.15	59.36	<b>65.91</b>
	Gen	<b>27.34</b>	23.10	22.13	5.97	<b>8.85</b>	7.20




Figure 3: Test/gen set full-sequence accuracy per input length across the four tasks for the three types of RNN seq2seq models, with and without attention. Test set length range: 6-15; gen set length range: 1-5 & 16-30.



## 5. Results

The main results can be found in Table 1 and Fig 3, which present full-sequence accuracy on aggregate and per-input-length levels, respectively. Related results measured in first  $n$ -symbol accuracy and overlap rate, are in Appendix C for references. Since the results across these three metrics share similar patterns, unless particularly mentioned, findings based on full-sequence accuracy remain valid in terms of the other two metrics. This section also reports results from follow-up experiments conducted under contrastive settings.

Please note that, since we are only interested in the learning potentials of RNN seq2seq models, the results presented in the body text are the overall best results, instead of the average ones. The best results were selected out of the three runs based on the test and gen set performance in full-sequence accuracy, empirically weighted 40% and 60%, respectively. This is consistent with our training procedure, where only the models with best dev set performance are deployed and evaluated, and in line with Grefenstette et al. (2015) and Delétang et al. (2022) who also choose to report best results. Nevertheless, the observations based on the best results are still true for the average results (see Appendix C.1).

### 5.1. Generalization abilities

As can be observed in Table 1, RNN seq2seq models, with or without attention, consistently achieve better aggregate full-sequence accuracy in the test sets than in the gen sets with a large average margin (i.e., at least 31%) for all the tasks. Furthermore, even on the per-input-length level, depicted in Fig 3, the gen set full-sequence accuracy decreases as the difference between the unseen lengths and the nearest seen lengths increases most of the time. These strongly demonstrate that RNN seq2seq models, regardless of attention, are prone to learning a function that fits the training or in-distribution data, instead of the underlying data generation functions, if these functions are learnable at all. Their out-of-distribution generalization abilities are highly limited if longer unseen and untested input sequence lengths are also taken into account. Empirically, out-of-distribution generalization appears to be restricted to data that is more similar to the training data in lengths.

As a result of approximating in-distribution data, the gen set performance is also considerably less correlated with the train set performance, compared to the test set. On the one hand, the train-gen variance is consistently significantly larger than the train-test variance for every model in each run. On the other hand, by measuring the correlation in the standard deviation of full-sequence accuracy of each model across the three runs, we find that the Kendall/Spearman correlation scores for train-test and train-gen are 0.79/0.89 and -0.35/-0.52, respectively, confirming that the train and test set performance is more positively correlated. For instance, Table 7 shows that when attentional models fit the train set nearly perfectly ( $\approx 100.00\%$ ) with a negligible performance variation in the test set ( $\leq 0.12\%$ ), the variation in the gen set, however, ranges from 1.56% to 10.03%.

Moreover, with the follow-up experiments in §5.2 and §5.3, we find that the out-of-distribution generalization is not correlated with the model complexity either. That is to say, the gen set performance is not a monotonic function of the number of parameters of the trained models, even when they all can fit the train/test sets nearly perfectly. A possible explanation for this is that when the models have sufficient capacity to fit the train set, there are potentially many or infinite possible configurations of parameters, which, however, may

lead to very different out-of-distribution performance. In this sense, the out-of-distribution performance is less interpretable than the in-distribution one.

The lack of correlation in performance between the train and gen set may be caused by the selection of the best models during training, which is informed by dev sets that consist of pure in-distribution data. The exact reasons, however, remain to be explored.

## 5.2. Attentional versus attention-less models

The main results show straightforwardly that attention helps models fit the train sets and generalize to the test/gen sets. Attentional models can always achieve much better aggregate full-sequence accuracy in both the train and test sets and have a much smaller train-test variance than the attention-less counterparts. Moreover, attentional models outperform the attention-less models in generalizing to the out-of-distribution examples. In other words, attentional RNN seq2seq models are stronger in-distribution learners with relatively better out-of-distribution generalization abilities, compared to the attention-less ones.

Besides, attention significantly improves learning efficiency, as attentional models tended to converge faster and used less than 50% of the epochs on average than the attention-less models (see Appendix B.3). Furthermore, Fig 3 shows that the test set performance of the attention-less models goes down nearly as a function of the input length, which indicates the need of greater sample complexity for training. To further contrast the learning efficiency between attentional and attention-less models, we conducted a follow-up experiment in total reduplication, which is harder than identity and reversal, but is more feasible than quadratic copying to test for computational reasons. In the experiment, others being same, the attentional models used 1/4 training examples and 1/4 hidden size (i.e., 128), and the attention-less models used 3 times more training data and 3 times more training epochs, compared to the respective original setups. The results in Table 2 show that by using only 1/12 training examples, 1/9 parameter size (see Appendix B.2), and 1/3 training epochs, the attentional models still outperform the attention-less ones in total reduplication.

Table 2: Aggregate full-sequence accuracy (%) for the follow-up experiment in learning total reduplication. Training details are described in text in §5.2.

Dataset	Attentional			Attention-less		
	SRNN	GRU	LSTM	SRNN	GRU	LSTM
Train	100.00	100.00	100.00	94.99	100.00	100.00
Test	99.20	99.53	99.58	84.93	90.21	91.86
Gen	35.20	14.07	19.37	0.00	5.10	4.54

## 5.3. Task complexity

The complexity hypothesis formulated in §3.3 for the four learning tasks is borne out for attention-less models. Table 1 and Fig 3 show clear evidence that for each type of attention-less models, quadratic copying is more complex than total reduplication, which is more

complex than identity, with reversal being the least complex, on both aggregate and per-input-length levels. Given the identical training and evaluation conditions for the four tasks indicated in §4, we argue that the complexity of learning these tasks is the major, if not only, attributable reason for the observed performance difference.

For attentional models, it is only clear that quadratic copying remains the most complex task to learn. For the rest tasks, however, the results do not distinguish their relative complexity informatively enough as full-sequence accuracy in the train and test sets all nears 100.00%, due to apparent overparameterization. To verify the related complexity hierarchy proposed in §3.3, we re-trained all the attentional models with three decreasing hidden sizes (i.e., 16, 32, 64, since §5.2 suggests that 128 is already too big) while keeping the rest original setups unchanged. The results in Fig 4 show that total reduplication is more complex than identity and reversal, since it requires more parameters or greater model complexity than the other two tasks. That said, Fig 4 also shows that when the hidden size is 16, the attentional SRNN and LSTM models appear to learn identity better than reversal, which contradicts with our proposal for unclear reasons.




Figure 4: Test/gen set full-sequence accuracy across the three tasks (except quadratic copying) for the three attentional RNN seq2seq models with two reduced hidden sizes: 16 and 32. The results for the hidden size 64 are not included because of the near 100% and thus less informative train/test performance for all the tasks.

#### 5.4. RNN variants

Overall, when attention is not used, GRU and LSTM are expectedly and significantly more expressive than SRNN for all tasks other than reversal. This probably results from the additional built-in gating mechanisms that improve long-term memory (Hochreiter and Schmidhuber, 1997; Chung et al., 2014) and thus reduce the sample complexity to learn the input-target alignments. For reversal, we find that when the models fit the train set with (nearly) 100% full-sequence accuracy, SRNN consistently shows better generalization abilities for reversal, both in-distribution and out-of-distribution, than GRU and LSTM, regardless of attention. In particular, when attention is used and hidden size  $\geq 32$  (see the follow-up experiments in §5.2 and §5.3), the worst performing SRNN always outperforms the best performing GRU and LSTM in the gen sets with other conditions being equal. For identity and total reduplication, however, such a consistent picture does not occur.

Previous research (Weiss et al., 2018; Merrill, 2019a; Delétang et al., 2022) has shown that LSTM can learn counting more effectively than SRNN and GRU. The results for quadratic copying largely support this claim, but also show that GRU may also learn counting in a similar way to LSTM in the seq2seq construction. To make sure that this is not due to the much smaller parameter size in the SRNN seq2seq models (see Appendix B.2), we re-trained them, both attentional and attention-less, with significantly increased model complexity. However, the best train set full-sequence accuracy among them is only 3.43% (see Appendix D.2), implying that the increased model complexity helps little.

Moreover, we notice that for quadratic copying, attentional SRNN models have unusually high and overall best first  $n$ -symbol accuracy and overlap rate (see Appendix C.2), despite rather low full-sequence accuracy. To confirm that SRNN does not count, we tested all the previously saved attentional models for quadratic copying on the mapping  $w \rightarrow w^{40}$  (i.e., making 40 copies), which is not present in the data for quadratic copying. Not surprisingly, full-sequence accuracy is always 0.00% for all the models. However, the best SRNN model achieves 96.20/80.81% first  $n$ -symbol accuracy in the test/gen set, while the same figure for the GRU and LSTM models never exceeds 26.83/23.66% (see Appendix D.2). This means that the attentional SRNN models simply have learnt somehow periodically repeating the input sequences, which leads to the accidentally high first  $n$ -symbol accuracy and overlap rate for quadratic copying. Their inability to generate the end symbol  $\langle s \rangle$  at the right timing indicates their inability to count the input length. We encourage future studies to explore why they approximated a periodic repetition function instead.

## 6. Discussion and Conclusion

This study investigated how well the three major types of RNN seq2seq models, with and without attention, learn four transduction tasks that require learning varying alignments or dependencies between the input and target sequences. Through unified training/evaluation conditions and comprehensive experiments, we compared the learning results across tasks, different model configurations, and test/gen set performance etc., and highlighted factors that influence the learning capabilities and generalization capacity of RNN seq2seq models. Unlike previous research, the input alphabet  $\Sigma$  for our experiments has 26 unique symbols, instead of binary, making our results more relevant to real-world tasks that concern, say, morpho-phonological transductions. The major findings are further discussed below.

**Generalization abilities.** For our experiments, RNN seq2seq models, regardless of attention, tend to approximate the training or in-distribution data, instead of learning the underlying functions. This makes their out-of-distribution generalization abilities restricted, which, however, may not be surprising. For an auto-regressive model like RNN seq2seq models to learn these transductions with bounded precision or parameter size, the probability of the correct output decreases exponentially as a function of the output length  $n$ , i.e.,  $P(\text{target}) = (1 - \varepsilon)^n$ , where  $\varepsilon$  is the expected error rate. As  $n$  increases indefinitely, the probability of correct generation eventually becomes infinitesimal and indistinguishably small. It is true for both training and testing where strings of (unseen) longer lengths are generally more difficult to fit and generalize to, as evidenced by the results in §5.

**Attention.** Attention greatly improves the learning efficiency, by reducing both model complexity and sample complexity, as well as the learning robustness, in terms of the gen-

eralization performance and overfitting problem. However, attention does not overcome the out-of-distribution generalization limitation, since it does not change the auto-regressive nature of the models. Nonetheless, the impressive learning efficiency accelerated by attention echoes its original motivation, i.e., “learning to align” (Bahdanau et al., 2015).

**Task complexity.** For the four tasks: identity ( $f_A$ ), reversal ( $f_B$ ), total reduplication ( $f_C$ ), quadratic copying ( $f_D$ ), we established the following task complexity hierarchy for attention-less RNN seq2seq models:  $f_D > f_C > f_A > f_B$ . This differs from the traditional FST-theoretic viewpoint, which sees  $f_B$  strictly more complex than  $f_A$ . As discussed in §3.3,  $f_B$  is easier than  $f_A$  for RNN seq2seq models because learning  $f_B$  contains many initially shorter input-target dependencies, instead of constantly-distanced dependencies for  $f_A$ . This makes iteratively optimizing the model parameters easier with backpropagation and results in less complexity for  $f_B$ . Interestingly, we also show in §3.3 that this result may be understood under the complexity hierarchy of formal languages, if the four transduction tasks are re-framed as language recognition tasks. However, for attentional models, we find that  $f_D > f_C > f_B > f_A$ , which appears to align with the FST characterizations but contradicts with our expectation that  $f_A$  and  $f_B$  are comparably complex with attention. Constrained by time and resources, we leave the related investigation to future studies.

**RNN variants.** The effect of RNN variants on the seq2seq models is a complicated one and interacts with other factors, e.g., attention and the task to learn. When attention is not used, GRU and LSTM are generally more expressive than SRNN. The only exception is reversal, which SRNN consistently outperforms GRU and LSTM in the test/gen set, regardless of attention. Please note that, this exception is only true when SRNN seq2seq models fit the related train set with (nearly) 100% full-sequence accuracy. Moreover, for quadratic copying and regardless of attention, both GRU and LSTM can count the input length while SRNN cannot, which arguably is not a matter of model parameter size.

Since all the tasks in the study require the decoder to recall the entire input symbols at any decoding time step, an important step to test the generality of our conclusions is to experiment with tasks that do not. Inspired by one of the reviewers, we re-run the main experiments on ascending sorting ( $f_E$ ), and descending sorting ( $f_F$ ), which sort the input sequence in an ascending and descending lexicographic order, respectively. For example, for any string  $w \in \{abc, acb, bac, bca, cab, cba\}$ ,  $f_E(w) = abc$  and  $f_F(w) = cba$ , which crucially means that these two tasks cannot be learnt through any static input-target alignments. Instead, an easier and viable way of learning is through counting the occurrences of each symbol  $\sigma \in \Sigma$  for the input sequence in the encoder and then retrieving each  $\sigma$ , ascendingly or descendingly, according to the previously encoded counts to generate the output sequence in the decoder. The results are shown in Table 3, which, reaffirm the limited out-of-distribution generalization abilities of RNN seq2seq models, despite their nearly perfect in-distribution performance. Moreover, attention is significantly beneficial for SRNN seq2seq models, but less so for GRU and LSTM models, probably because GRU and LSTM can learn the two sorting tasks through counting even without attention, which SRNN cannot. These explanations are consistent with our previous conclusions.

The results presented above showcase the fruitfulness of utilizing formal language tasks to probe the learning capabilities of neural networks in the context of string transductions. A few key insights or lessons, which are more or less commonsense but worth emphasizing are as follows. (1) Conclusions about learnability should be carefully made, since many

Table 3: Aggregate full-sequence accuracy (%) for the two sorting tasks.

Task	Dataset	Attentional			Attention-less		
		SRNN	GRU	LSTM	SRNN	GRU	LSTM
Ascend	Train	100.00	100.00	100.00	37.28	100.00	100.00
	Test	99.03	99.69	<b>99.73</b>	6.48	99.50	<b>99.74</b>
	Gen	10.89	31.06	<b>31.43</b>	0.02	<b>42.72</b>	35.66
Descend	Train	100.00	100.00	100.00	24.01	100.00	100.00
	Test	99.05	<b>99.78</b>	99.69	0.49	99.19	<b>99.66</b>
	Gen	14.65	31.12	<b>32.35</b>	0.00	34.33	<b>37.08</b>

factors play a role in the final performance of trained neural networks, such as sample complexity, model complexity, training conditions, etc. For instance, in our main experiments, attention-less SRNN models barely learnt total reduplication, in line with Nelson et al. (2020), but the follow-up experiment in §5.2 shows that this is a result of insufficient training resources. (2) Reliably evaluating the empirical learning capabilities of neural networks requires well-rounded evaluation metrics. Had this study only looked at full-sequence accuracy or first  $n$ -symbol accuracy, it would become unlikely to discover that attentional SRNN seq2seq models learnt a somehow periodic repetition function when trained for quadratic copying. (3) Task complexity is strongly tied to the structure of the learner, as demonstrated by the relative complexity between identity and reversal for FSTs and attention-less RNN seq2seq models. To make this point stronger, we show in Appendix D.3 that RNNs of few hundred parameters can seemingly learn identity and easily generalize to input sequences of 10,000 symbols, while RNN seq2seq models of millions of parameters cannot. Therefore, we discourage over-interpretations of our results beyond the context of this study.

Besides some unexplained puzzles brought up here, good continuations of the current research may include experimenting with (1) other types of seq2seq models, such as CNN seq2seq (Gehring et al., 2017) and transformer (Vaswani et al., 2017); (2) Tape-RNN, which show promising generalization results in various transduction tasks (Delétang et al., 2022); (3) and other novel transduction tasks. We hope that our study encourages more works at the intersection of neural networks and formal language theory in the future.

## Acknowledgments

The current research would not be initiated and successfully continued without the guidance and inspirations from Prof. Jeffrey Heinz. I am deeply grateful to the three anonymous reviewers for their constructive comments. I also thank Prof. Jordan Kodner, William Oliver, Sarah Payne, Nicholas Behrje who read through the early draft and provided helpful feedback. Parts of the work have been presented at various occasions at Stony Brook University, Yale University, University of Pennsylvania, and George Mason University as a talk or poster over the past few months, so my thanks also go for the audiences there. Moreover, I thank the inspirations from Imane Machouahi Rifi.

## References

- Enes Avcu, Chihiro Shibata, and Jeffrey Heinz. Subregular complexity and deep learning, 2017. URL <https://arxiv.org/abs/1705.05940>.
- Stéphane Ayache, Rémi Eyraud, and Noé Goudian. Explaining black boxes on sequential data using weighted automata. In Olgierd Unold, Witold Dyrka, and Wojciech Wieczorek, editors, *Proceedings of The 14th International Conference on Grammatical Inference 2018*, volume 93 of *Proceedings of Machine Learning Research*, pages 81–103. PMLR, feb 2019. URL <https://proceedings.mlr.press/v93/ayache19a.html>.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1409.0473>.
- Yoshua Bengio, Patrice Simard, and Paolo Frasconi. Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, 5(2):157–166, 1994. doi: 10.1109/72.279181.
- Satwik Bhattamishra, Kabir Ahuja, and Navin Goyal. On the Ability and Limitations of Transformers to Recognize Formal Languages. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 7096–7116, Online, November 2020. Association for Computational Linguistics. doi: 10.18653/v1/2020.emnlp-main.576. URL <https://aclanthology.org/2020.emnlp-main.576>.
- Mikolaj Bojanczyk, Sandra Kiefer, and Nathan Lhote. String-to-String Interpretations With Polynomial-Size Output. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming (ICALP 2019)*, volume 132 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 106:1–106:14, Dagstuhl, Germany, 2019. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. ISBN 978-3-95977-109-2. doi: 10.4230/LIPIcs.ICALP.2019.106. URL <http://drops.dagstuhl.de/opus/volltexte/2019/10682>.
- Denny Britz, Anna Goldie, Minh-Thang Luong, and Quoc Le. Massive exploration of neural machine translation architectures. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1442–1451, Copenhagen, Denmark, September 2017. Association for Computational Linguistics. doi: 10.18653/v1/D17-1151. URL <https://aclanthology.org/D17-1151>.
- Sarath Chandar, Chinnadhurai Sankar, Eugene Vorontsov, Samira Ebrahimi Kahou, and Yoshua Bengio. Towards non-saturating recurrent units for modelling long-term dependencies. In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence and Thirty-First Innovative Applications of Artificial Intelligence Conference and Ninth AAAI Symposium on Educational Advances in Artificial Intelligence, AAAI’19/IAAI’19/EAAI’19*. AAAI Press, 2019. ISBN 978-1-57735-809-1. doi: 10.1609/aaai.v33i01.33013280. URL <https://doi.org/10.1609/aaai.v33i01.33013280>.

- Yining Chen, Sorcha Gilroy, Andreas Maletti, Jonathan May, and Kevin Knight. Recurrent neural networks as weighted language recognizers. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2261–2271, New Orleans, Louisiana, June 2018. Association for Computational Linguistics. doi: 10.18653/v1/N18-1205. URL <https://aclanthology.org/N18-1205>.
- Kyunghyun Cho, Bart van Merriënboer, Caglar Gulcehre, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. Learning phrase representations using RNN encoder–decoder for statistical machine translation. In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1724–1734, Doha, Qatar, October 2014. Association for Computational Linguistics. doi: 10.3115/v1/D14-1179. URL <https://aclanthology.org/D14-1179>.
- Noam Chomsky. Three models for the description of language. *IRE Transactions on Information Theory*, 2(3):113–124, 1956. doi: 10.1109/TIT.1956.1056813.
- Junyoung Chung, Caglar Gulcehre, KyungHyun Cho, and Yoshua Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling, 2014. URL <https://arxiv.org/abs/1412.3555>.
- Grégoire Delétang, Anian Ruoss, Jordi Grau-Moya, Tim Genewein, Li Kevin Wenliang, Elliot Catt, Chris Cundy, Marcus Hutter, Shane Legg, Joel Veness, and Pedro A. Ortega. Neural networks and the chomsky hierarchy, 2022. URL <https://arxiv.org/abs/2207.02098>.
- Hossep Dolatian and Jeffrey Heinz. Computing and classifying reduplication with 2-way finite-state transducers. *Journal of Language Modelling*, 8(1):179–250, Sep. 2020. doi: 10.15398/jlm.v8i1.245. URL <https://jlm.ipipan.waw.pl/index.php/JLM/article/view/245>.
- Finale Doshi-Velez and Been Kim. Towards a rigorous science of interpretable machine learning, 2017. URL <https://arxiv.org/abs/1702.08608>.
- Jeffrey L. Elman. Finding structure in time. *Cognitive Science*, 14(2):179–211, 1990. doi: [https://doi.org/10.1207/s15516709cog1402\\_1](https://doi.org/10.1207/s15516709cog1402_1). URL [https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog1402\\_1](https://onlinelibrary.wiley.com/doi/abs/10.1207/s15516709cog1402_1).
- Emmanuel Filiot and Pierre-Alain Reynier. Transducers, logic and algebra for functions of finite words. *ACM SIGLOG News*, 3(3):4–19, aug 2016. doi: 10.1145/2984450.2984453. URL <https://doi.org/10.1145/2984450.2984453>.
- Jonas Gehring, Michael Auli, David Grangier, Denis Yarats, and Yann N. Dauphin. Convolutional sequence to sequence learning. In Doina Precup and Yee Whye Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1243–1252. PMLR, 06–11 Aug 2017. URL <https://proceedings.mlr.press/v70/gehring17a.html>.



- C. L. Giles, C. B. Miller, D. Chen, H. H. Chen, G. Z. Sun, and Y. C. Lee. Learning and Extracting Finite State Automata with Second-Order Recurrent Neural Networks. *Neural Computation*, 4(3):393–405, 05 1992. ISSN 0899-7667. doi: 10.1162/neco.1992.4.3.393. URL <https://doi.org/10.1162/neco.1992.4.3.393>.
- Xavier Glorot and Yoshua Bengio. Understanding the difficulty of training deep feedforward neural networks. In Yee Whye Teh and Mike Titterton, editors, *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, volume 9 of *Proceedings of Machine Learning Research*, pages 249–256, Chia Laguna Resort, Sardinia, Italy, 13–15 May 2010. PMLR.
- Yoav Goldberg. A primer on neural network models for natural language processing. *Journal of Artificial Intelligence Research*, 57:345–420, 2016. doi: <https://doi.org/10.1613/jair.4992>. URL <https://jair.org/index.php/jair/article/view/11030>.
- Alex Graves, Greg Wayne, and Ivo Danihelka. Neural Turing machines, 2014. URL <https://arxiv.org/abs/1410.5401>.
- Edward Grefenstette, Karl Moritz Hermann, Mustafa Suleyman, and Phil Blunsom. Learning to transduce with unbounded memory. In C. Cortes, N. Lawrence, D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 28. Curran Associates, Inc., 2015. URL <https://proceedings.neurips.cc/paper/2015/file/b9d487a30398d42ecff55c228ed5652b-Paper.pdf>.
- Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural Comput.*, 9(8):1735–1780, nov 1997. ISSN 0899-7667. doi: 10.1162/neco.1997.9.8.1735. URL <https://doi.org/10.1162/neco.1997.9.8.1735>.
- Aravind K. Joshi. *Tree adjoining grammars: How much context-sensitivity is required to provide reasonable structural descriptions?*, page 206–250. Studies in Natural Language Processing. Cambridge University Press, 1985. doi: 10.1017/CBO9780511597855.007.
- Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In Yoshua Bengio and Yann LeCun, editors, *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015. URL <http://arxiv.org/abs/1412.6980>.
- Thang Luong, Hieu Pham, and Christopher D. Manning. Effective approaches to attention-based neural machine translation. In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1412–1421, Lisbon, Portugal, September 2015. Association for Computational Linguistics. doi: 10.18653/v1/D15-1166. URL <https://aclanthology.org/D15-1166>.
- William Merrill. Sequential neural networks as automata. In *Proceedings of the Workshop on Deep Learning and Formal Languages: Building Bridges*, pages 1–13, Florence, August 2019a. Association for Computational Linguistics. doi: 10.18653/v1/W19-3901. URL <https://aclanthology.org/W19-3901>.

- William Merrill. Sequential neural networks as automata, 2019b. URL <https://arxiv.org/abs/1906.01615>.
- Max Nelson, Hossep Dolatian, Jonathan Rawski, and Brandon Prickett. Probing RNN encoder-decoder generalization of subregular functions using reduplication. In *Proceedings of the Society for Computation in Linguistics 2020*, pages 167–178, New York, New York, January 2020. Association for Computational Linguistics. URL <https://aclanthology.org/2020.scil-1.22>.
- Razvan Pascanu, Tomas Mikolov, and Yoshua Bengio. On the difficulty of training recurrent neural networks. In Sanjoy Dasgupta and David McAllester, editors, *Proceedings of the 30th International Conference on Machine Learning*, volume 28 of *Proceedings of Machine Learning Research*, pages 1310–1318, Atlanta, Georgia, USA, 17–19 Jun 2013. PMLR. URL <https://proceedings.mlr.press/v28/pascanu13.html>.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc., 2019. URL <https://proceedings.neurips.cc/paper/2019/file/bdbca288fee7f92f2bfa9f7012727740-Paper.pdf>.
- Hao Peng, Roy Schwartz, Sam Thomson, and Noah A. Smith. Rational recurrences. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1203–1214, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1152. URL <https://aclanthology.org/D18-1152>.
- Jonathan Rawski, Hossep Dolatian, Jeffrey Heinz, and Eric Raimy. Regular and polyregular theories of reduplication. *Glossa: a journal of general linguistics*, 8(1), 2023. doi: 10.16995/glossa.8885. URL <https://doi.org/10.16995/glossa.8885>.
- Hava T. Siegelmann. Recurrent neural networks and finite automata. *Computational Intelligence*, 12(4):567–574, 1996. doi: <https://doi.org/10.1111/j.1467-8640.1996.tb00277.x>. URL <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8640.1996.tb00277.x>.
- Ilya Sutskever, Oriol Vinyals, and Quoc V Le. Sequence to sequence learning with neural networks. In Z. Ghahramani, M. Welling, C. Cortes, N. Lawrence, and K.Q. Weinberger, editors, *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc., 2014. URL <https://proceedings.neurips.cc/paper/2014/file/a14ac55a4f27472c5d894ec1c3c743d2-Paper.pdf>.
- Sam van der Poel, Dakotah Lambert, Kalina Kostyszyn, Tiantian Gao, Rahul Verma, Derek Andersen, Joanne Chau, Emily Peterson, Cody St. Clair, Paul Fodor, Chihiro Shibata,

and Jeffrey Heinz. Mlregtest: A benchmark for the machine learning of regular languages, 2023.

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett, editors, *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc., 2017. URL <https://proceedings.neurips.cc/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf>.

Ingmar Visser. Hidden markov model interpretations of neural networks. *Behavioral and Brain Sciences*, 23:494 – 495, 08 2000. doi: 10.1017/S0140525X00513351.

Gail Weiss, Yoav Goldberg, and Eran Yahav. On the practical computational power of finite precision RNNs for language recognition. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 740–745, Melbourne, Australia, July 2018. Association for Computational Linguistics. doi: 10.18653/v1/P18-2117. URL <https://aclanthology.org/P18-2117>.

Ronald J. Williams and David Zipser. A learning algorithm for continually running fully recurrent neural networks. *Neural Computation*, 1(2):270–280, 1989. doi: 10.1162/neco.1989.1.2.270.

Jingzhao Zhang, Tianxing He, Suvrit Sra, and Ali Jadbabaie. Why gradient clipping accelerates training: A theoretical justification for adaptivity. In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net, 2020. URL <https://openreview.net/forum?id=BJgnXpVYwS>.

## A. Mathematical details of the three RNN variants

For SRNN,  $f$  is typically a sigmoid or hyperbolic tangent function. We used the latter.

$$\mathbf{h}_t = \tanh(\mathbf{W}[\mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}) \tag{5}$$

where  $[\cdot; \cdot]$  denotes vector concatenation along the last axis. Since the codomain of  $\tanh = (-1, 1)$ , its exact output may be seen as a signal with a relevancy score  $\in [0, 1)$  and a polarity ( $\pm$ ) to compute the next output or the next hidden state.

For LSTM,  $f$  takes what is known as a cell state  $\mathbf{c}_t$  in addition to the hidden state  $\mathbf{h}_t$ . At each time  $t$ , both states are updated as follows.

$$\mathbf{h}_t = o_t * \tanh(\mathbf{c}_t) \tag{6a}$$

$$\mathbf{c}_t = f_t * \mathbf{c}_{t-1} + i_t * \tilde{\mathbf{c}}_t \tag{6b}$$

$$\tilde{\mathbf{c}}_t = \tanh(\mathbf{W}_c[\mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}_c) \tag{6c}$$

$$o_t = \text{sigmoid}(\mathbf{W}_o[\mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}_o) \tag{6d}$$

$$f_t = \text{sigmoid}(\mathbf{W}_f[\mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}_f) \tag{6e}$$

$$i_t = \text{sigmoid}(\mathbf{W}_i[\mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}_i) \tag{6f}$$

where Eq.(6d-6f) are the equations for the output gate  $o_t$ , forget gate  $f_t$ , and input gate  $i_t$ , respectively, all of which are scalar values  $\in (0, 1)$ , thanks to *sigmoid*. Because of this numerical property, they are said to act as a filter to control the information flow inside the recurrent unit. The update of  $\mathbf{h}_t$  is determined by  $o_t$  and  $\mathbf{c}_t$ , whose update is determined by weighing  $\mathbf{c}_{t-1}$  and the temporarily updated cell state  $\tilde{\mathbf{c}}_t$  with  $f_t$  and  $i_t$ , respectively.

Inspired by LSTM, GRU takes a simplified state transition function as follows.

$$\mathbf{h}_t = z_t * \mathbf{h}_{t-1} + (1 - z_t) * \tilde{\mathbf{h}}_t \quad (7a)$$

$$\tilde{\mathbf{h}}_t = \tanh(\mathbf{W}_h[r_t * \mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}_h) \quad (7b)$$

$$z_t = \text{sigmoid}(\mathbf{W}_z[\mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}_z) \quad (7c)$$

$$r_t = \text{sigmoid}(\mathbf{W}_r[\mathbf{h}_{t-1}; \mathbf{e}_t] + \mathbf{b}_r) \quad (7d)$$

where  $z_t$  is the update gate,  $r_t$  the reset gate, and  $\tilde{\mathbf{h}}_t$  the temporarily updated  $\mathbf{h}_t$ . Here,  $\tilde{\mathbf{h}}_t$  takes the place of both  $\mathbf{h}_t$  and  $\mathbf{c}_t$  as in LSTM, and  $z_t$  replaces LSTM’s  $f_t$  and  $i_t$ .  $\mathbf{h}_t$  is updated by weighing  $\mathbf{h}_{t-1}$  and  $\tilde{\mathbf{h}}_t$  with  $z_t$  and  $(1 - z_t)$ , respectively.

Table 4: Summary of data size for the four learning tasks in the main experiments.

Dataset	Input length	# of pairs per length	# of pairs
Train	6-15	1,000	10,000
Dev	6-15	1,000	10,000
Test	6-15	5,000	50,000
Gen	1-5 & 16-30	5,000	100,000

## B. Experimental details

### B.1. Data

Table 4 summarizes the data size for the four learning tasks in the main experiments. As mentioned in §4.1, duplicates are only allowed for sequences of lengths 1 and 2 in the gen sets to make the later evaluations unified and simpler. For strings of length  $l$  (1 or 2), they were randomly sampled from a list of strings  $\Sigma^l \times m$ , where  $m$  is the smallest multiplier to repeat  $\Sigma^l$   $m$  times into a list such that the size of the list  $> 5000$ . For strings of longer length  $l^+$ , we uniformly sampled  $l^+$  symbols from  $\Sigma$  with replacement to create a string and repeated this process until we had enough unique strings of length  $l^+$ .

We reused the original train/dev/test/gen sets for all the follow-up experiments and evaluations with two exceptions. (1) For the follow-up experiment in §5.2, the attentional models were trained on 1/4 of data from the original train/dev sets and evaluated on the original test/gen sets. The attention-less models were trained on the train/dev sets that are three times larger than the original ones, with the four datasets re-created from scratch following the same data generation procedure to ensure each dataset remains disjoint from one another. (2) For the follow-up evaluation on the mapping  $w \rightarrow w^{40}$  in §5.4 as well as the follow-up experiments in sorting in §6, all the input sequences were directly from the corresponding original datasets. The target sequences were re-created accordingly.

## B.2. Model parameter configuration and size

Table 5 summarizes the parameter configuration and size of models for all the experiments.

Table 5: Summary of model parameter configuration and size. H/E: hidden size or embedding size. NA is for models not trained. In the “Where” row, §5 refers to the main experiments with the rest denoting the respective follow-up experiments.

	Attentional						Attention-less	
Where	§5.4	§5 & §6	§5.2	§5.3	§5.3	§5.3	§5.4	§5 & §6
H/E	640/384	512/128	128/128	64/128	32/128	16/128	1024/384	512/128
SRNN	2,582,812	1,466,396	126,236	46,236	21,596	13,116	5,037,084	1,204,252
GRU	NA	3,305,500	291,100	104,092	44,380	22,972	NA	2,519,068
LSTM	NA	4,225,052	373,532	133,020	55,772	27,900	NA	3,176,476

Table 6: Average number of epochs used over the three runs for each model configuration across the four tasks in the main experiments.

	Attentional			Attention-less		
Task	SRNN	GRU	LSTM	SRNN	GRU	LSTM
Identity	60	93	103	500	500	500
Rev	30	93	90	483	500	500
Total Red	83	180	176	500	500	500
Quad Copy	500	500	500	500	500	500
Average	168	216	217	496	500	500

## B.3. Training notes

Constrained by computational resources, we were unable to experiment with different sets of hyperparameters, various optimization methods, and so on. Therefore, we can in no way guarantee that those negative results (e.g., out-of-distribution generalization, SRNN seqseq models can not learn quadratic copying) shown in the paper are not due to our choices of training protocol and model configurations. We only **assume** that our training and model decisions are reasonably effective and that the obtained results are reasonably representative of the learning capabilities of generic RNN seq2seq models, at least in an empirical sense.

Nevertheless, our training is mostly effective. Table 6 provides details about the average number of epochs used over the three runs for each type of model from the main experiments. It shows that when some tasks are easy to learn, particularly for attentional models learning tasks other than quadratic copying, the models converge within 180 epochs on average or less, without unnecessarily running through all epochs. When the tasks are difficult to learn, even when the train set can be overfitted with 100% full-sequence accuracy, the selection of the best model based on the dev set performance serves as a grid search for the best model out of 50 evaluated models across 500 epochs. With three runs, we try to ensure that

the final obtained models are indicative of their respective learning potentials for each task. Finally, the extensive follow-up experiments also help reduce the risks of us being misguided by insufficient training conditions, besides making more fine-grained observations.

Table 7: Average full-sequence accuracy (%) with standard deviation on the aggregate level across the four learning tasks for models with various configurations. Best average results are in **bold** for the test and gen sets.

Task	Dataset	Attentional			Attention-less		
		SRNN	GRU	LSTM	SRNN	GRU	LSTM
Identity	Train	100.00 $\pm$ 0.01	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	33.05 $\pm$ 31.81	98.22 $\pm$ 0.21	99.99 $\pm$ 0.02
	Test	99.88 $\pm$ 0.12	<b>100.00</b> $\pm$ 0.00	99.99 $\pm$ 0.01	15.09 $\pm$ 24.05	70.62 $\pm$ 1.49	<b>76.97</b> $\pm$ 0.86
	Gen	18.76 $\pm$ 5.99	33.75 $\pm$ 5.69	<b>33.92</b> $\pm$ 2.14	0.00 $\pm$ 0.00	9.67 $\pm$ 1.20	<b>10.02</b> $\pm$ 0.06
Rev	Train	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00	100.00 $\pm$ 0.00
	Test	<b>99.98</b> $\pm$ 0.00	99.93 $\pm$ 0.05	99.90 $\pm$ 0.04	<b>99.57</b> $\pm$ 0.02	88.32 $\pm$ 0.45	92.66 $\pm$ 0.22
	Gen	<b>37.51</b> $\pm$ 2.97	21.73 $\pm$ 1.56	24.10 $\pm$ 1.57	<b>19.92</b> $\pm$ 4.61	19.04 $\pm$ 0.66	12.44 $\pm$ 0.21
Total Red	Train	100.00 $\pm$ 0.01	100.00 $\pm$ 0.01	100.00 $\pm$ 0.01	12.26 $\pm$ 2.96	86.99 $\pm$ 3.32	92.14 $\pm$ 1.22
	Test	99.79 $\pm$ 0.07	<b>99.84</b> $\pm$ 0.07	99.77 $\pm$ 0.12	4.25 $\pm$ 1.20	48.86 $\pm$ 1.65	<b>54.50</b> $\pm$ 0.59
	Gen	<b>34.98</b> $\pm$ 10.03	19.21 $\pm$ 4.18	19.17 $\pm$ 1.64	0.00 $\pm$ 0.00	3.77 $\pm$ 0.76	<b>4.81</b> $\pm$ 1.27
Quad Copy	Train	1.73 $\pm$ 0.62	71.12 $\pm$ 8.15	66.11 $\pm$ 14.99	1.01 $\pm$ 0.53	44.91 $\pm$ 5.45	56.86 $\pm$ 9.34
	Test	1.52 $\pm$ 0.43	60.25 $\pm$ 6.87	<b>61.04</b> $\pm$ 11.18	0.37 $\pm$ 0.21	24.12 $\pm$ 3.19	<b>34.21</b> $\pm$ 3.82
	Gen	1.62 $\pm$ 0.23	5.20 $\pm$ 2.61	<b>6.41</b> $\pm$ 0.59	0.00 $\pm$ 0.00	<b>0.74</b> $\pm$ 0.39	0.13 $\pm$ 0.11
Average	Train	75.43 $\pm$ 0.16	92.78 $\pm$ 2.04	91.53 $\pm$ 3.75	36.58 $\pm$ 8.82	82.53 $\pm$ 2.24	87.25 $\pm$ 2.64
	Test	75.29 $\pm$ 0.16	90.00 $\pm$ 1.75	<b>90.17</b> $\pm$ 2.84	29.82 $\pm$ 6.37	57.98 $\pm$ 1.69	<b>64.59</b> $\pm$ 1.37
	Gen	<b>23.22</b> $\pm$ 4.80	19.97 $\pm$ 3.51	20.90 $\pm$ 1.49	4.98 $\pm$ 1.15	<b>8.30</b> $\pm$ 0.75	6.85 $\pm$ 0.41

## C. Results: main experiments

### C.1. Average results: full-sequence accuracy

Table 7 provides the average results in full-sequence accuracy with standard deviation based on the three runs. Clearly, except few trivial differences, such as which type of model has the best average gen set performance for identity and quadratic copying, Table 7 is in line with Table 1. The observations we made from Table 1 remain valid in terms of Table 7. While Table 7 additionally shows how model performance varies across different runs, the interpretation of the variation may not be meaningful, since such variation also easily “varies” with changes in hyperparameters. As mentioned in the beginning of §5, this study is only interested in the learning potentials of RNN seq2seq models.

The average results in the other two metrics, i.e., first  $n$ -symbol accuracy and overlap rate, show similar data patterns as the corresponding best results, which are reported in the following two sections. To view the complete sets of results obtained in the paper, raw or summarized, please check the GitHub repository for the study at <https://github.com/jaaack-wang/rnn-seq2seq-learning>.

### C.2. Best aggregate results measured in other metrics

Table 8 and Table 9 show the aggregate best results measured in first  $n$ -symbol accuracy and overlap rate, respectively. The results were selected from the same runs consistent with Table 1. Best results are in **bold** for the test and gen sets.

Table 8: Aggregate first  $n$ -symbol accuracy (%) across the four learning tasks for models with various configurations.

Task	Dataset	Attentional			Attention-less		
		SRNN	GRU	LSTM	SRNN	GRU	LSTM
Identity	Train	100.00	100.00	100.00	79.49	99.05	99.98
	Test	99.99	<b>100.00</b>	<b>100.00</b>	59.55	85.64	<b>90.58</b>
	Gen	65.51	<b>74.83</b>	71.89	6.90	43.44	<b>50.19</b>
Rev	Train	100.00	100.00	100.00	100.00	100.00	100.00
	Test	<b>99.99</b>	99.93	99.93	<b>99.94</b>	96.26	97.66
	Gen	<b>64.01</b>	35.52	36.88	<b>69.69</b>	45.58	53.39
Total Red	Train	100.00	100.00	100.00	36.66	93.99	96.09
	Test	99.89	<b>99.90</b>	99.83	24.27	66.85	<b>70.68</b>
	Gen	<b>73.43</b>	50.15	50.40	2.52	20.60	<b>21.44</b>
Quad Copy	Train	87.67	99.70	98.73	17.86	83.48	92.65
	Test	85.39	<b>98.49</b>	97.01	10.94	49.85	<b>60.52</b>
	Gen	<b>76.42</b>	59.33	29.68	0.85	<b>12.43</b>	9.86
Average	Train	96.92	99.93	99.68	58.50	94.13	97.18
	Test	96.32	<b>99.58</b>	99.19	48.68	74.65	<b>79.86</b>
	Gen	<b>69.84</b>	54.96	47.21	19.99	30.51	<b>33.72</b>

Table 9: Aggregate overlap rate (%) across the four learning tasks for models with various configurations.

Task	Dataset	Attentional			Attention-less		
		SRNN	GRU	LSTM	SRNN	GRU	LSTM
Identity	Train	100.00	100.00	100.00	89.31	99.63	99.99
	Test	<b>100.00</b>	<b>100.00</b>	<b>100.00</b>	76.51	92.71	<b>95.56</b>
	Gen	68.61	<b>80.31</b>	78.09	12.34	50.99	<b>58.55</b>
Rev	Train	100.00	100.00	100.00	100.00	100.00	100.00
	Test	<b>100.00</b>	99.97	99.97	<b>99.97</b>	98.02	98.92
	Gen	<b>69.36</b>	39.65	41.95	<b>72.77</b>	53.57	61.99
Total Red	Train	100.00	100.00	100.00	59.39	97.80	98.74
	Test	<b>99.96</b>	<b>99.96</b>	99.95	47.89	83.30	<b>85.98</b>
	Gen	<b>83.09</b>	58.00	54.39	8.49	30.96	<b>32.28</b>
Quad Copy	Train	98.45	99.71	98.85	29.82	89.65	95.76
	Test	98.22	<b>98.60</b>	97.91	23.24	64.59	<b>76.24</b>
	Gen	<b>88.02</b>	64.30	40.64	6.16	<b>19.53</b>	16.54
Average	Train	99.61	99.93	99.71	69.63	96.77	98.62
	Test	99.54	<b>99.63</b>	99.46	61.90	84.65	<b>89.17</b>
	Gen	<b>77.27</b>	60.56	53.77	24.94	38.76	<b>42.34</b>

### C.3. Best per-input-length results measured in other metrics

Fig 5 and Fig 6 show the main results measured in first  $n$ -symbol accuracy and overlap rate on the per-input-length level, respectively. Please note that, the reason why the best

attentional SRNN model outperforms the related GRU and LSTM counterparts in terms of the overall test/gen set performance is not because it learns quadratic copying. Instead, attentional SRNN models all ended up learning a function that more or less repeats the input sequences without knowing when exactly to produce the end symbol  $\langle s \rangle$ . Please check §5.4 for the reasoning.




Figure 5: Test/gen set first  $n$ -symbol accuracy per input length across the four tasks for the three types of RNN seq2seq models, with and without attention.




Figure 6: Test/gen set overlap rate per input length across the four tasks for the three types of RNN seq2seq models, with and without attention.

## D. Results: follow-up experiments

### D.1. Overview

Follow-up experiments with RNN seq2seq models were conducted in the following sections: (1) §5.2, where we contrasted attentional models with attention-less ones for their efficiency in learning total reduplication; (2) §5.3, where we further examined the relative complexity among identity, reversal, and total reduplication for attentional models; (3) §5.4, where we studied the possibility of SRNN seq2seq models, with or without attention, learning quadratic copying; (4) and §6, where we experimented with ascending and descending sorting, to test the generality of the findings based on the four tasks in the main experiments.



These follow-up experiments all follow the same training and evaluation procedures as the main experiments. That is to say, for each experiment, we trained and evaluated each model for three runs and selected the best results for interpretation. Details about data can be found in Appendix B.1. Details about model parameter configuration and size are in Appendix B.2. For more details, please check the GitHub repository for the paper.

In what follows, we provide some further details about the follow-up experiments in §5.4, mentioned above. We also introduce another follow-up experiment with RNNs trained for identity mentioned in §6, but not above. Both are for the self-containedness of the paper.

## D.2. Attentional SRNN seq2seq models trained for quadratic copying

Table 10 shows the performance of the best enlarged SRNN models trained and evaluated on the original datasets created for quadratic copying. Clearly, neither fit the train set and the one without attention is not even learning anything because its overlap rate is just about random guess, which equals  $\frac{1}{|\Sigma|} = \frac{1}{26} = 3.85\%$ , or  $\frac{1}{28} = 3.57\%$  plus the two edge symbols. This is despite the attention-less SRNN model having more than 5 million parameters, the largest model trained in the study. We intended to make the attentional SRNN model larger than the related GRU or LSTM model, but it was impossible without modifying the model structure. Increasing hidden size too much exceeded the maximum 40G GPU RAM on Google Colaboratory due to too much computation for softmax in the attention layer.

Table 10: The performance across the three metrics (%) for the best enlarged SRNN models trained for quadratic copying.

Dataset	Attentional			Attention-less		
	Full-seq	First $n$ -symbol	Overlap	Full-seq	First $n$ -symbol	Overlap
Train	3.43	92.43	98.65	0.00	0.05	3.80
Test	3.00	90.92	98.53	0.00	0.05	3.81
Gen	2.79	84.23	92.82	0.00	0.19	3.68

In line with Tables 8 and 9, Table 10 also shows unusually high first  $n$ -symbol accuracy and overlap rate for attentional SRNN models trained for quadratic copying. Moreover, we found that during training, all the six attentional SRNN models, including the three enlarged ones, could easily reach nearly 100.00% first  $n$ -symbol accuracy (and thus overlap rate) within 100 epochs and that first  $n$ -symbol accuracy is not correlated with full-sequence accuracy, unlike all the rest models we trained that all show a rough correlation. Fig 7 is an example training log of a non-enlarged attentional SRNN model for quadratic copying.

The poor correlation is probably because attentional SRNN models are approximating a somehow periodic function that repeats the input sequence for either a fixed or unbounded number of times, which conflicts with learning quadratic copying, which requires a bounded and yet varied number of repetitions of the input sequence according to the input length. The fact that the trained attentional SRNN models tend to have a unusually high first  $n$ -symbol accuracy means that the approximated periodic function, if perfectly learnt, should

ideally repeat any given input sequence more than 30 times, i.e., the maximum input length available in the gen set.

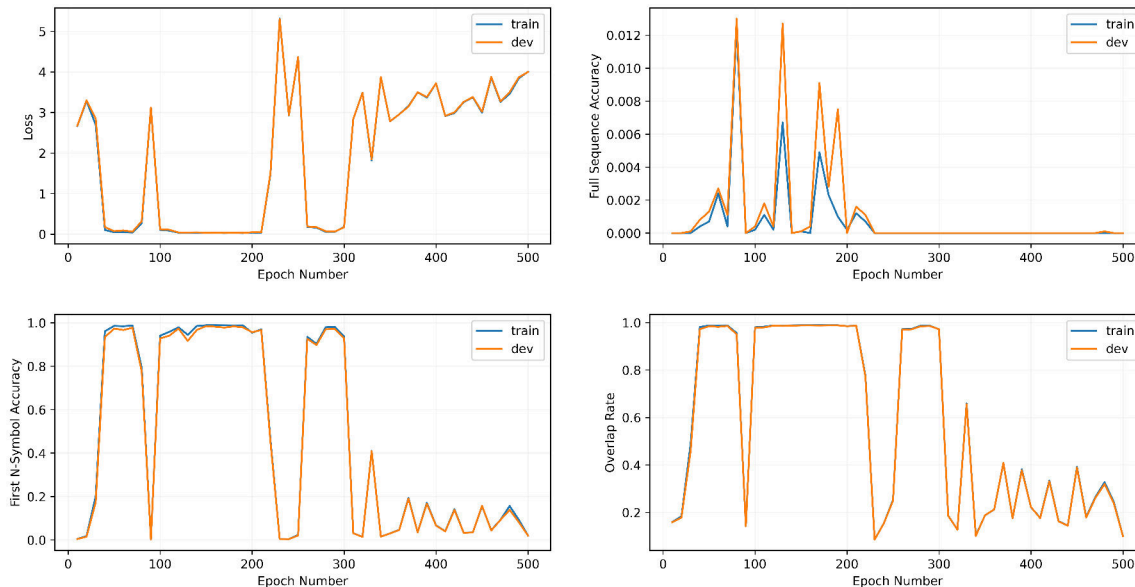


Figure 7: Training log of the third attentional SRNN model for quadratic copying in the main experiment, i.e., the “worst” SRNN model in Table 11.

Therefore, we tested all the saved attentional models trained for quadratic copying on the mapping  $w \rightarrow w^{40}$ . The results in Table 11 shows that even the “worst-performing” SRNN model outperforms the “best-performing” GRU/LSTM model in terms of first  $n$ -symbol accuracy, given that full-sequence accuracy is always 0.00%. Since the saved attentional SRNN models were selected based on the best dev set full-sequence accuracy, we expect that all these SRNN models would have shown much better first  $n$ -symbol accuracy on  $w \rightarrow w^{40}$  if the best first  $n$ -symbol accuracy was used instead. As can be seen in Fig 7, the train-dev variance in first  $n$ -symbol accuracy is negligible and thus the test set performance should come close. That said, the exact function learnt by SRNN models remains unknown.

Table 11: The test/gen set first  $n$ -symbol accuracy (%) for all the attentional models trained for quadratic copying across three runs on the mapping  $w \rightarrow w^{40}$ .

Model	Test			Gen		
	Run#1	Run#2	Run#3	Run#1	Run#2	Run#3
SRNN	67.95	84.16	68.33	67.07	68.42	30.89
SRNN <sub>Large</sub>	84.86	82.14	96.20	62.89	71.70	80.81
GRU	26.42	25.49	26.82	23.66	10.67	14.15
LSTM	26.83	25.51	25.52	6.07	8.72	7.56

### D.3. RNNs of few hundred parameters learning identity

The RNNs we trained here are essentially the decoder part of RNN seq2seq models without attention. The major difference between them is that RNNs work just like FSTs, instead of auto-regressively, so they produce an output symbol immediately after consuming an input symbol, until the entire input sequence is read.

We trained SRNN, GRU, and LSTM with embedding size 3 and hidden size 4 and the parameter size for these three models are only 244, 316, and 352, respectively. The training conditions were mostly identical to the main experiments, except for the much larger learning rate, i.e., 1e-2. Apart from 150,000 strings of lengths 1-30 from the original test and gen sets, we also evaluated the trained models on 10,000 random unique strings of length 10,000. All the models we trained achieved 100% full-sequence accuracy on these unseen 160,000 strings. Since RNNs produce as many symbols as they consume, the result also means that these models can transduce all the prefixes of those 10,000 strings, ranging from length 1 to 9,999, successfully. See: <https://github.com/jaaack-wang/RNNs-learn-identity>.

The experiment presented above simply serves as a proof of concept. The result is actually not surprising because for RNNs to learn identity, they only need to produce whatever they consume. Although RNNs can learn identity much better than RNN seq2seq models, it is easy to prove that unidirectional RNNs cannot learn at all functions that require 2-way FSTs, because they cannot look ahead.<sup>3</sup> The point here is to show that there is a strong connection between task complexity and the structure of the learner and that RNNs and RNN seq2seq models represent two classes of neural networks (also see §2.4). Other than the aforementioned point, we do not make any more claims here.

---

3. Please note that, here we only refer to canonical RNNs that work like FSTs. RNNs can be trained to consume all input symbols first and then use the final hidden state to auto-regressively produce output symbols. However, RNNs trained in this way can be seen as a special case of attention-less RNN seq2seq models where encoder and decoder share the same weights. RNNs trained in [Delétang et al. \(2022\)](#) are similar to this, but instead of doing auto-regression, they let the models consume as many dummy tokens as the target sequence after they consumed the entire real input sequence, during which the models produce an output symbol once they consume a dummy token, just like a FST. Nevertheless, neither of these two special types of “RNNs” are considered as RNNs in the sense of this study.