

What Languages are Easy to Language-Model? A Perspective from Learning Probabilistic Regular Languages

Nadav Borenstein¹ Anej Svete² Robin Shing Moon Chan² Josef Valvoda¹

Franz Nowak² Isabelle Augenstein¹ Eleanor Chodroff³ Ryan Cotterell²

¹Københavns Universitet ²ETH Zürich ³Universität Zürich

{nb, jval, augenstein}@di.ku.dk eleanor.chodroff@uzh.ch

{asvete, chanr, fnowak, ryan.cotterell}@inf.ethz.ch

Abstract

What can large language models learn? By definition, language models (LM) are distributions over strings. Therefore, an intuitive way of addressing the above question is to formalize it as a matter of learnability of *classes* of distributions over strings. While prior work in this direction focused on assessing the theoretical limits, in contrast, we seek to understand the empirical learnability. Unlike prior empirical work, we evaluate neural LMs on their home turf—learning probabilistic languages—rather than as classifiers of formal languages. In particular, we investigate the learnability of regular LMs (RLMs) by RNN and Transformer LMs. We empirically test the learnability of RLMs as a function of various complexity parameters of the RLM and the hidden state size of the neural LM. We find that the RLM rank, which corresponds to the size of linear space spanned by the logits of its conditional distributions, and the expected length of sampled strings are strong and significant predictors of learnability for both RNNs and Transformers. Several other predictors also reach significance, but with differing patterns between RNNs and Transformers.

1 Introduction

Language models are, definitionally, distributions over strings. However, not all neural LMs are capable of learning—or even representing—all possible distributions. This raises two natural questions: What classes of distributions *can* neural LMs represent and what can they learn from training examples? In terms of the first question, the relationship between recurrent neural networks and more symbolic computational models has been subject to study for over three decades (McCulloch and Pitts, 1943; Kleene, 1956; Siegelmann and Sontag, 1992; Hao et al., 2018; Korsky and Berwick, 2019; Merrill, 2019; Merrill et al., 2020; Hewitt et al., 2020; Chung and Siegelmann, 2021; Merrill et al., 2022; Merrill and Tsilivis, 2022; Svete and Cotterell, 2023; Nowak et al., 2023). Moreover, the

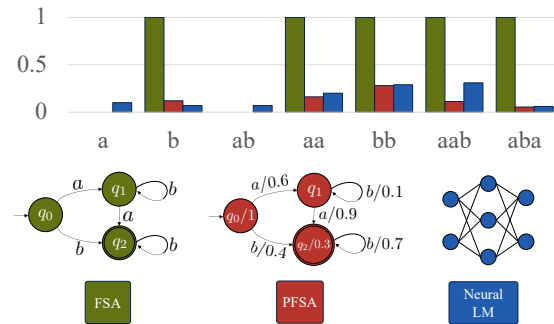


Figure 1: A **finite-state automaton** defines a *set* of strings by assigning string binary weights. A **probabilistic finite-state automaton** and a **neural LM** such as an RNN or a Transformer LM, however, define a *probability distribution* over strings.

prevalence of Transformer-based LMs has led to a recent body of work investigating their representational capacity (e.g., Hahn, 2020; Ebrahimi et al., 2020; Bhattamishra et al., 2020; Merrill and Sabharwal, 2023b). However, almost all of this work is theoretical, i.e., researchers seek theorems that give exact limitations on the capacity of specific neural LMs. While such work provides a good characterization of what neural LMs could, in principle, learn, it does not speak to what LMs can learn in practice.

In contrast to the above, a more empirically minded researcher might prefer to run a series of controlled experiments to determine what neural LMs can and cannot learn. Their goal is to empirically characterize what classes of formal LMs, e.g., probabilistic finite-state automata, *are* learnable with neural LMs in practice, using current best practices. Such work can inform our understanding of what types of languages larger LMs trained on human-written text might represent—specifically, what grammatical structures they can recognize, and how efficiently they can do so. All of the above is crucial for quantifying the practical capabilities, and limits, of neural LMs. While plenty of empirical work has provided insights into the linguistic capabilities of modern LMs using linguistically annotated datasets (e.g., Linzen et al., 2016; Hewitt and Manning, 2019; Jawahar et al., 2019; Liu et al., 2019; Icard, 2020; Manning et al., 2020; Rogers et al., 2021; Belinkov, 2022),

human-annotated datasets give us limited insight into the types of distributions a neural LM can learn because the true distribution that the neural LMs is modeling is often unclear. For instance, fitting an LM to Wikipedia leaves open to interpretation exactly *which* probability distribution over strings the neural LM is modeling. In contrast, learning a probabilistic formal language in a controlled experiment offers an unparalleled level of control.

A closer look at the existing work on the empirical learnability of formal languages (see App. A for an overview) reveals a categorical mismatch between what LMs are, i.e., *probability distributions* over strings, and what learning a formal language means, i.e., classifying strings as members of a specific language, i.e., a *set* of strings (Ebrahimi et al., 2020; Delétang et al., 2023; Wang and Steinert-Threlkeld, 2023). In response, we propose to investigate the practical representation capacity of neural LMs by testing their ability to learn *distributions* over strings. By sampling training corpora from probabilistic finite-state automata (PFSA), we can ask precise questions about the learnability.

We offer an empirical study, sampling datasets of 20k strings from 2100 randomly generated PFSAs, and training 15k RNN and Transformer language models with a varying hidden state size on datasets. Our study is informed by various theoretical results on the representational capacity of RNNs concerning probabilistic finite-state automata. We assess the learnability by approximating the KL divergence between PFSAs and neural LMs. In a regression analysis, we find that a large number of properties of the automaton, e.g., the number of states, the number of transitions, the rank of its emission matrix, and its entropy all contribute to learnability. In particular, the rank of the emission matrix and expected length of the sampled strings are strong predictors of learnability for both RNN and Transformer LMs. Several other predictors also demonstrate relevance in accounting for learnability, but have differing magnitudes and patterns of significance for the RNN and Transformer LMs. Similar to Delétang et al. (2023), we find that RNNs are better suited to modeling formal languages than Transformers.

2 Preliminaries

We begin with an introduction of the relevant mathematical preliminaries, based on Cotterell et al. (2024). An **alphabet** Σ is a finite, non-empty set of

symbols. The **Kleene closure** Σ^* of the alphabet Σ is the set of all strings of the symbols in Σ . The **length** of the string $\mathbf{y} = y_1 \dots y_T \in \Sigma^*$, denoted by $|\mathbf{y}| = T$, is the number of symbols the string contains. A **language model** p is a probability distribution over Σ^* . Two LMs p and q are **equivalent** if $p(\mathbf{y}) = q(\mathbf{y})$ for all $\mathbf{y} \in \Sigma^*$. Most modern LMs define $p(\mathbf{y})$ as a product of conditional probability distributions:

$$p(\mathbf{y}) \stackrel{\text{def}}{=} p(\text{EOS} | \mathbf{y}) \prod_{t=1}^{|\mathbf{y}|} p(y_t | \mathbf{y}_{<t}), \quad (1)$$

where $\text{EOS} \notin \Sigma$ is a special end of sequence symbol. We denote $\bar{\Sigma} \stackrel{\text{def}}{=} \Sigma \cup \{\text{EOS}\}$ and \bar{y} an element of $\bar{\Sigma}$.

2.1 Neural Language Models

Representation-based neural LMs define the conditional distributions $p(\bar{y}_t | \mathbf{y}_{<t})$ through a linearly transformed **hidden state**. In this paper, we focus on **softmax-normalized, representation-based** neural LMs, where a **language encoder** $\mathbf{h}: \Sigma^* \rightarrow \mathbb{R}^D$ (Chan et al., 2024) computes a representation, and the conditional distributions of the neural LMs are defined as follows

$$p(\bar{y}_t | \mathbf{y}_{<t}) \stackrel{\text{def}}{=} \text{softmax}(\mathbf{E}\mathbf{h}(\mathbf{y}_{<t}))_{\bar{y}_t} \quad (2a)$$

$$\stackrel{\text{def}}{=} \frac{\exp(\mathbf{E}\mathbf{h}(\mathbf{y}_{<t}))_{\bar{y}_t}}{\sum_{\bar{y} \in \bar{\Sigma}} \exp(\mathbf{E}\mathbf{h}(\mathbf{y}_{<t}))_{\bar{y}}}. \quad (2b)$$

We will call $\mathbf{E} \in \mathbb{R}^{|\bar{\Sigma}| \times D}$ the **output matrix**.

Representation-based neural LMs differ in how \mathbf{h} is computed as a function of $\mathbf{y}_{<t}$. In this paper, we consider the two most popular modern language modeling architectures: recurrent neural networks (Elman, 1990), specifically the LSTM variant (Hochreiter and Schmidhuber, 1997), where \mathbf{h} is computed recurrently, and Transformers (Vaswani et al., 2017), where \mathbf{h} is computed with self-attention.

2.2 Regular Language Models

A classic formalism for defining LMs is **probabilistic finite-state automata** (PFSAs), a probabilistic version of finite-state automata that defines string probabilities. Intuitively, a PFSA defines a *finite* number of conditional next-symbol distributions $p(\bar{y} | q)$ based on a finite number of states $q \in Q$ that summarize string prefixes analogous to how the hidden state \mathbf{h} of an RNN summarizes the prefix $y_1 \dots y_t$.

A PFSA moves between its states based on the input symbols according to the transitions defined by a transition relation. It **accepts** a string with the probability equal to the product of the transition weights along the string’s path in the automaton and the last state’s final weight (or the sum over all paths if there are multiple paths accepting the string).¹ The PFSA is **deterministic** (a DPFSA) if the transition relation is a *function* of the current state and symbol, i.e., if, for all $q \in Q$ and $y \in \Sigma$, there exists at most one $q' \in Q$ such that $p(q' | q, y) > 0$. A DPFSA is **minimal** if it has the smallest number of states among all its equivalent DPFSA.² The minimal DPFSA is unique up to the naming of the states.

Definition 2.1. *The LM p is **regular** (an RLM) if there exists an equivalent PFSA.*

Fig. 1 shows an example of an RLM defining a distribution over $\{a, b\}^*$ with $p(ab^n ab^m) = 1 \cdot 0.6 \cdot 0.1^n \cdot 0.9 \cdot 0.7^m \cdot 0.3$ and $p(bb^m) = 1 \cdot 0.4 \cdot 0.7^m \cdot 0.3$.

3 Representing RLMs with Neural LMs

Neural LMs have demonstrated an ability to model human language well. This raises the question of what enables their success in capturing human language and how we can improve them further. However, neural LMs are notoriously challenging to analyze, making it difficult to state any formal claims on what they are (in)capable of modeling. To amend this, a large body of work has linked neural LMs to formal models of computation. DPFSA feature prominently in this line of research (Merrill, 2019; Merrill et al., 2020; Svete and Cotterell, 2023). To facilitate a detailed inspection of how neural LMs can represent DPFSA, we now formalize DPFSA in a way that is particularly easy to connect to softmax-normalized representation-based neural LMs, see App. B for the definition of DPFSA.

3.1 Softmax-Normalized DPFSA

The conditional distributions $p(\bar{y} | q)$ defined by a DPFSA can, in general, be arbitrary distributions over $\bar{\Sigma}$ —a DPFSA therefore defines $|Q|$ distributions, each with $|\bar{\Sigma}| - 1$ degrees of freedom. As we will see, such a parameterization makes the connection to neural LMs, which define conditional distributions in terms of shared parameters of the neural

¹Final weights of states are analogous to the EOS symbol which signals the end of string generation in neural LMs.

²See App. B for more formal details.

network and the output matrix \mathbf{E} , straightforward. To make this connection clear, consider a DPFSA \mathcal{A} with support over Σ^* . Because \mathcal{A} is deterministic, there exists a function $\mathbf{h}_{\mathcal{A}}: \Sigma^* \rightarrow \{0, 1\}^{|Q|}$ that maps every string that the unique state \mathcal{A} enters to a one-hot encoding after reading in the string. Encoding the final weight as EOS, we end up with the following autoregressive formulation of \mathcal{A} ,

$$p(\bar{y}_t | \mathbf{y}_{<t}) \stackrel{\text{def}}{=} \text{softmax}(\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}_{<t}))_{\bar{y}_t}. \quad (3)$$

where $\mathbf{T} \in \mathbb{R}^{|\bar{\Sigma}| \times |Q|}$. Note that it is because we sought to use a softmax that we require support over all of Σ^* . Also, despite its deceiving notation, Eq. (3), is not, in this form, obviously a neural LM as we did not give a neural network that computes $\mathbf{h}_{\mathcal{A}}$ —its computation is performed by a DPFSA. Written as in Eq. (3), we say that a DPFSA is of rank R if the output matrix \mathbf{T} is of rank R .

Minimal representations. If \mathcal{A} is a minimal DPFSA, then the linear subspace spanned by $\{\mathbf{h}_{\mathcal{A}}(\mathbf{y}) | \mathbf{y} \in \Sigma^*\}$ will be a subspace of $\mathbb{R}^{|Q|}$ of dimensionality $|Q|$ due to the one-hot nature of the state encodings. However, if \mathbf{T} is of rank R , then the linear subspace $\{\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) | \mathbf{y} \in \Sigma^*\}$ cannot span a subspace of dimensionality larger than R . Therefore, for a neural LM to be equivalent to the DPFSA, the hidden state’s dimension must be large enough. If that is not the case, the neural LM will naturally *not* be able to match all the conditional distributions defined by the states of the DPFSA, which leads us to the following result.

Theorem 3.1. *Let p be a language model induced by a minimal DPFSA \mathcal{A} with full support. Furthermore, let R be the rank of p when expressed as a representation-based LM (Eq. (3)). Then, any equivalent representation-based LM q must have a hidden state of size at least R .*

Proof. See App. C. ■

Given a rank- R DPFSA \mathcal{A} , Thm. 3.1 says that an equivalent neural LM needs a hidden state of size at least R , establishing a general lower bound on an LM’s hidden state size for equivalence with a DPFSA. A hidden state of size R , however, does not *guarantee* that the neural LM can implement the transitions between DPFSA’s states. The neural LM’s ability to do so depends on the particular architecture implementing the neural LM. In other words, Thm. 3.1 furnishes us with a necessary condition. In the rest of the section,

we briefly recapitulate some known results on the relationship between DPFSAs and the two neural architectures in question—recurrent neural networks and Transformers.

3.2 DPFSAs and Recurrent Neural LMs

RNNs’ connection to DPFSAs is particularly well understood for recurrent neural LMs (e.g., Peng et al., 2018; Merrill, 2019; Merrill et al., 2020; Svete and Cotterell, 2023; Svete et al., 2024). The following theorem from Svete and Cotterell (2023) very concretely summarizes the relationship between DPFSAs and finite-precision RNNs, i.e., ones where the representations \mathbf{h} fall into some finite subset of \mathbb{R}^D .

Theorem 3.2 (Svete and Cotterell (2023), Thm. 4.1). *The classes of finite-precision RNN LMs and DPFSA are equivalent.*

While Thm. 3.2 ensures the existence of a finite-precision RNN LM weakly equivalent to a given DPFSA, it does not describe the size of the hidden state required for doing so. The following theorem makes the relationship more precise: It states that, in general, the size of the hidden state of RNN LMs must scale linearly with the number of states if the DPFSA is full-rank, i.e., of rank $|Q|$.

Theorem 3.3 (Svete and Cotterell (2023), Thms. 5.1 and 5.2). *A family of DPFSAs exists such that the hidden states sizes of equivalent RNN LMs scale linearly with $|Q|$ and $|\Sigma|$.*

Our Thm. 3.1 can be seen as a more precise version of Thm. 3.3, which does not consider the rank of the DPFSA.³

3.3 DPFSAs and Transformer LMs

Statements similar to Thm. 3.2 are more difficult to make for Transformer LMs due to their parallelizable nature (Merrill and Sabharwal, 2023b; Svete and Cotterell, 2024); unlike DPFSAs, Transformers do not compute a sequential inner state. Existing work has connected Transformers to sequential automata like DPFSAs with the empirically successful framework of chain-of-thought reasoning (Pérez et al., 2021; Feng et al., 2023; Merrill and Sabharwal, 2023a) or by assuming a growing number of layers with increasing string length (Liu

³The linear scaling with respect to $|Q|$ is unique to weighted FSAs; more efficient constructions simulating *unweighted* FSAs with RNNs exist (Indyk, 1995; Svete and Cotterell, 2023), highlighting the importance of considering probabilistic models separately.

et al., 2023). However, these results concern binary language recognition rather than the task of language models. To the best of our knowledge, the only work that directly involves language modeling setting connects Transformers to the simple n -gram LMs, a special, well-structured class of DPFSAs.

Theorem 3.4 (Svete and Cotterell (2024), Thm 3.1). *For any n -gram LM p , there exists an equivalent Transformer LM $p_{\mathcal{T}}$ with hard attention.*

Unlike Thm. 3.2, Thm. 3.4 does not provide a complete characterization of Transformer LMs but rather a (loose) lower bound on their capabilities, in line with their weaker connection to sequential automata. We are not aware of any claims analogous to Thm. 3.3 for Transformer LMs; this makes the generally applicable Thm. 3.1 that much more interesting, as it lower-bounds the size of the vectorial representations regardless of the neural architecture used to compute them. This distinction is also mirrored in the empirical part of the paper.

3.4 Beyond Representational Capacity

The theoretical results presented and summarized here capture the fact that the size of the hidden representations in neural LMs must inevitably scale with the size of the DPFSA being simulated or learned (either with its rank, its number of states, or the size of the alphabet). In that sense, the representational capacity of neural LMs can be theoretically described relatively comprehensively in terms of formal models of computation. However, existing theoretical work only considers the question of which distributions can be *represented* by a neural LM. This leaves us with a large gap in understanding which distributions are *learnable* by neural LMs. Compared to pure representational capacity results, formal claims about learning are much more difficult to make due to the dependence on factors such as the learning algorithm and aspects of the training data. Therefore, we test the learnability of DPFSAs *empirically*, allowing us to overcome the above-mentioned difficulties and gain valuable insights into the problem.

4 Practical Learnability of RLMs

Our main goal is to provide a principled study of the ability of neural LMs to learn RLMs. We now describe and justify our experimental design.

Predictor	Interpretation
$ Q $	The number of states
$ \Sigma $	The number of symbols
$ Q \Sigma $	The number of transitions
R	The rank of the output matrix \mathbf{E}
Exp. length	Expected length of strings sampled from PFSA
$\min(Q , \Sigma)$	(Tight) upper bound on R
$H(\mathcal{A})$	Entropy of PFSA
D	The hidden state size of the Elman RNN or Transformer

Table 1: The predictor variables used to estimate KL divergence with their interpretation: the PFSA-related properties are listed first, followed by those of the neural LM (D).

4.1 A Critique of Learning Formal Languages

As discussed in §1, much empirical work has investigated the ability of neural language models to learn formal languages such as those described by finite-state automata, i.e., how well a neural language model can be used to assess membership of individual strings in a set. There are multiple workarounds for this discrepancy. Most solutions involve measuring some sort of *accuracy* of next-symbol prediction. For example, [Suzgun et al. \(2019a,b\)](#) and [Bhattacharya et al. \(2020\)](#) evaluate neural LMs on the next-symbol prediction task, which, intuitively, measures whether all allowed continuations of the string under the formal model achieve a large enough probability under the neural LM. [Delétang et al. \(2023\)](#) evaluate the models by the proportion of correctly predicted tokens (where the approximate argmax of the neural LM prediction has to match the ground-truth label). Unfortunately, all these approaches inevitably shoehorn a neural LM into a sort of classifier, mismatching the type of an LM—a probability distribution—and a language—a set, as illustrated in Fig. 1. Ideally, we would like to measure precisely how the neural LM has learned the *distribution* induced by a formal LM. In this section, we outline and motivate a possible way to approach this challenge.

At a high level, we test the learnability of *random* representation-based RLMs by training neural LMs on strings sampled from randomly generated DPFSAs and measuring the distance between the neural LM and the DPFSA. Crucially, unlike most existing work, we do not have to rely on classification or next-symbol prediction accuracy-based metrics, but rather we directly measure the similarity of distributions which presents a much cleaner way of evaluating model similarity. Concretely, given an DPFSA p and a neural LM q , we measure the KL divergence between the DPFSA and the

Predictor	$\hat{\beta}$	SE	p -value
Intercept	8.72	0.08	< 0.001
$ Q $	0.68	0.15	< 0.001
$ \Sigma $	0.22	0.20	0.26
$ Q \Sigma $	0.23	0.13	0.07
R	4.10	0.10	< 0.001
Exp. len.	3.21	0.19	< 0.05
$\min(Q , \Sigma)$	-0.15	0.26	0.58
$H(\mathcal{A})$	-0.88	0.22	< 0.001
D	-0.63	0.08	< 0.001

Table 2: Estimated coefficients ($\hat{\beta}$), standard errors (SE), and p -values for D_{KL} generated with a linear regression model for RNNs.

neural LM:

$$D_{\text{KL}}(p \parallel q) \stackrel{\text{def}}{=} \sum_{\mathbf{y} \in \Sigma^*} p(\mathbf{y}) \log \frac{p(\mathbf{y})}{q(\mathbf{y})} \quad (4a)$$

$$= H(p, q) - H(p). \quad (4b)$$

The KL divergence is an established and well-understood measure of the distance⁴ between two *distributions*. As such, it lends itself naturally to evaluating the difference between LMs; in our case, measuring how well the neural LM has captured the distribution of the DPFSA. Such a holistic treatment of the difference between two LMs gives us a tangible and interpretable way of understanding how they differ. To compute $D_{\text{KL}}(p \parallel q)$, we use Eq. (4b). We estimate the first term $H(p, q)$ by computing $\hat{H}(p, q)$, the empirical cross-entropy between p and q . The second term can be computed exactly by dynamic programming ([Eisner, 2002](#); [Zmigrod et al., 2021](#)). See App. D.4 for further details on the computation of these evaluation metrics.

4.2 Generating Random DPFSAs

We evaluate neural LMs’ ability to learn *random* representation-based DPFSAs, which we construct as follows. First, we sample the number of states $|Q| \in \{2, 4, 6, 8, 10, 12, 16\}$ and the number of symbols $|\Sigma| \in \{2, 4, 6, 8, 10, 12, 16\}$, both uniformly at random. Then, we select the destination for the outgoing arcs of each of the states—one for each $y \in \Sigma$, i.e., for each $q \in Q$ and $y \in \Sigma$ we randomly choose $q' \in Q$ and add the transition $q \xrightarrow{y} q'$ to \mathcal{A} . Again, the sampling is done uniformly at random. We add weights to the transition function of \mathcal{A} as follows. Finally, we generate a random matrix $\mathbf{T} \in \mathbb{R}^{|\Sigma| \times |Q|}$.

⁴Note, that KL divergence is not a *true* distance, as it is not symmetric and does not fulfill the triangle inequality.

D_{KL}	$ Q $						
	2	4	6	8	10	12	16
RNNs	2.6	6.9	8.3	7.7	8.6	10.7	11.1
Trans.	8.8	13.7	14.4	11.9	13.5	14.4	15.3

Table 3: KL divergence for RNNs and Transformers as a function of the number of states of the DPFSA.

For each $R \in \{1, 2, 4, 6, 8, 10, 12, 16\}$ satisfying $R \leq \min(|Q|, |\Sigma|)$, we define \mathbf{T}^R to be the best rank- R approximation of \mathbf{T} , i.e., $\mathbf{T}^R = \operatorname{argmin}_{\mathbf{M}} \|\mathbf{M} - \mathbf{T}\|_F^2$ subject to the constraint $\operatorname{rank}(\mathbf{M}) = R$. This is an example of the Procrustes problem and can be efficiently computed using SVD. We then set the transition probability of $q \xrightarrow{y} q'$ to $w_{q,y} = \operatorname{softmax}(\mathbf{T}^R_{:,q})_y$. Similarly, we set the final weight $\rho(q) = \operatorname{softmax}(\mathbf{T}^R_{:,q})_{\text{EOS}}$.

As the Transformer LMs have a limited context length, it is undesirable to sample DPFSA that generate strings that are longer than the context length with high probability. This would introduce an artificial discrepancy between the entropy of the DPFSA and the empirical entropy of the truncated strings used to train and evaluate the Transformer models. Therefore, we filter out DPFSA with high expected sequence length. In practice, we filter out DPFSA with an expected length larger than the median value of expected lengths, i.e., half of the DPFSA. See App. D.4 for details about calculating the expected string length of a DPFSA.

This process results in the generation of up to eight random DPFSA, ⁵ all sharing the same Q, Σ , and underlying transition function. They differ, however, in the rank of the matrix \mathbf{T}^R that defines the weights of the transitions. Furthermore, the construction of exactly one transition for each q and y ensures that the DPFSA mirrors the nature of a neural LM, which also defines full-support next-symbol probabilities for any prefix of the string. Altogether, this allows us to precisely control the quantities from Tab. 1 and thus the complexity of the DPFSA. See App. D.1 for additional details.

5 Results

5.1 Statistical Evaluation

Following the experimental setup outlined in App. D, we obtained D_{KL} results for 15k RNN and 15k Transformer LMs of differing hidden state dimensions, trained on strings sampled from 2100 random DPFSA with specific sets of complexity

⁵We have that $\{r \mid r \leq \min(|Q|, |\Sigma|), r \in \{1, 2, 4, 6, 8, 10, 12, 16\}\} \leq 8$

Predictor	$\hat{\beta}$	SE	p -value
Intercept	13.5	0.10	< 0.001
$ Q $	0.60	0.19	< 0.01
$ \Sigma $	2.10	0.24	< 0.001
$ Q \Sigma $	0.64	0.16	< 0.001
R	2.99	0.13	< 0.001
Exp. len.	11.70	0.24	< 0.001
$\min(Q , \Sigma)$	-1.36	0.33	< 0.001
$H(\mathcal{A})$	-7.89	0.28	< 0.001
D	-2.76	0.10	< 0.001

Table 4: Estimated coefficients ($\hat{\beta}$), standard errors (SE), and p -values for D_{KL} generated with a linear regression model for Transformers.

parameters. The complexity parameters of DPFSA correspond to the number of states, number of symbols, and the DPFSA rank; additional derivative metrics can also capture overall complexity, including the number of transitions, the expected string length, the upper bound of the DPFSA rank, and the DPFSA entropy (see Tab. 1). We expect the difficulty of learning, quantified as KL divergence, to increase with each complexity parameter and to decrease with the hidden state size.

As shown in Fig. 2, as the number of states and DPFSA rank increase, overall learning difficulty also empirically increases for both RNNs and Transformers. To demonstrate that the observed effect is not wholly reducible to differences in entropy, the relationship between the combined number of states and rank with DPFSA entropy is also presented. (Note that the KL divergence is the LM loss minus the DPFSA entropy.) A more nuanced picture, however, arises in the empirical relationship of alphabet size and DPFSA rank with overall learning difficulty (Fig. 3). Finally, Fig. 4 demonstrates that learning difficulty numerically increases with rank, though a large hidden state size of the neural LM reduce this difficulty to some degree.

To assess the influence of such effects on KL divergence, we implemented a linear regression model for the output of each neural LM. The independent variables were the DPFSA complexity parameters, as well as the neural LM’s hidden state size D . The full list of predictors is shown in Tab. 1. Prior to model fitting, each predictor was standardized using a z -score transformation for an interpretable comparison of the estimated coefficients.

5.2 RNN Findings

As shown in Tab. 2, the KL divergence was significantly influenced by the number of states, DPFSA rank, expected string length, DPFSA entropy, and RNN hidden state size ($|Q|, R, \text{Exp.}$

len., $H(\mathcal{A})$, D). The number of symbols, number of transitions, and the upper bound of R did not reach significance ($|\Sigma|$, $|Q||\Sigma|$, $\min(|Q|, |\Sigma|)$). Of the significant effects, the number of states, expected length, and rank were positive in direction, indicating an increase in KL divergence with an increase in the predictor of interest. Perhaps surprisingly, the DPFSA entropy was negative in influence, indicating a decrease in KL with an increase in DPFSA entropy. Unsurprisingly, the RNN hidden state size was also negative, indicating a decrease in KL with an increase in hidden state size. Overall, the DPFSA rank had the strongest influence on KL divergence, followed by the expected length, the DPFSA entropy, number of states, and RNN hidden state size. The remaining predictors were smaller in influence.

5.3 Transformer Findings

The linear regression model of the Transformers data revealed significant effects of all included predictors on KL divergence (see Tab. 4). Of these, the number of states, number of symbols, number of transitions, DPFSA rank, and expected string length were positive in influence, indicating an increase in KL divergence as the predictor of interest increased ($|Q|$, $|\Sigma|$, $|Q||\Sigma|$, R , Exp. len.). The upper bound of R , DPFSA entropy, and the Transformer hidden state size were negative in influence, indicating a decrease in KL divergence with an increase in the predictor of interest ($H(\mathcal{A})$, D). Of the positive relationships, the expected string length had the largest influence, followed in order by DPFSA rank (R), then number of symbols ($|\Sigma|$). Of the negative relationships, DPFSA entropy had the largest influence, followed by the Transformer hidden state size, then by the upper bound of R .

5.4 Comparing RNN and Transformer LMs

The linear models revealed an overall similar pattern of effects between RNNs and Transformers, albeit with few notable disparities. First, note the overall performance of each neural LM, as indicated by the model intercept for KL divergence—see Tab. 2, Tab. 4, and Fig. 6 in App. E. The model intercept reflects the predicted KL divergence when all other (z -transformed) predictors are equal to 0. RNNs tend to outperform Transformers in this task (RNNs: $\hat{\beta}_0 = 8.72$, Transformers: $\hat{\beta}_0 = 13.5$), demonstrating lower average loss. This difference in performance could be attributed to two main factors: 1) As previous research has shown, RNNs are

better suited to modeling formal languages (Delétang et al., 2023), and 2) Transformers necessitate careful training involving language-specific hyperparameter tuning, which poses a severe computational challenge. Nevertheless, we expect that the trend observed here would persist even if the Transformers were trained under optimal conditions.

While different in magnitude, the number of states, rank, expected string length, DPFSA entropy, and hidden state size all shared the same direction in their influence on KL divergence for each of the RNN and Transformer outputs. Of these, the biggest disparity in magnitude is in the influence of the expected string length and DPFSA entropy, which were significantly larger for Transformers compared to RNNs. Finally, not all predictors reached significance for both RNNs and Transformers. For instance, the number of symbols, number of transitions, and the upper bound of R were significant predictors of KL divergence for Transformers, but not for RNNs.

As shown in Figs. 2, 3, and 4, the observed trends in the DPFSA’s entropy, visualized alongside the neural LMs D_{KL} , suggest that the entropies alone cannot solely account for the observed effect in D_{KL} . In other words, larger D_{KL} values are not fully explained by greater entropy values, highlighting the significance of the rank R on learnability even more clearly. Particularly notable is Fig. 4, which presents similar patterns to those observed in Tab. 2 and Tab. 4. The KL divergence of both model types increases with an increase in the PFSA rank and decreases with an increase in the hidden state size. While somewhat noisier, Fig. 2 and Fig. 3 further demonstrate the weaker influence of the number of states and number of symbols on the KL divergence.

6 Discussion

Implications of Thm. 3.1. Thm. 3.1 concretely quantifies the minimum size of the representation space of *any* neural LM required for the correct representation of regular LMs. While this is special case of the so-called *softmax bottleneck* principle (Chang and McCallum, 2022), it is, to the best of our knowledge, the first result connecting the principle to formal models of computation. Practical implementations of RLMS might use state spaces and alphabets of sizes ranging from tens of thousands to hundreds of thousands (Mohri and Riley, 1999), which is much larger than the representa-

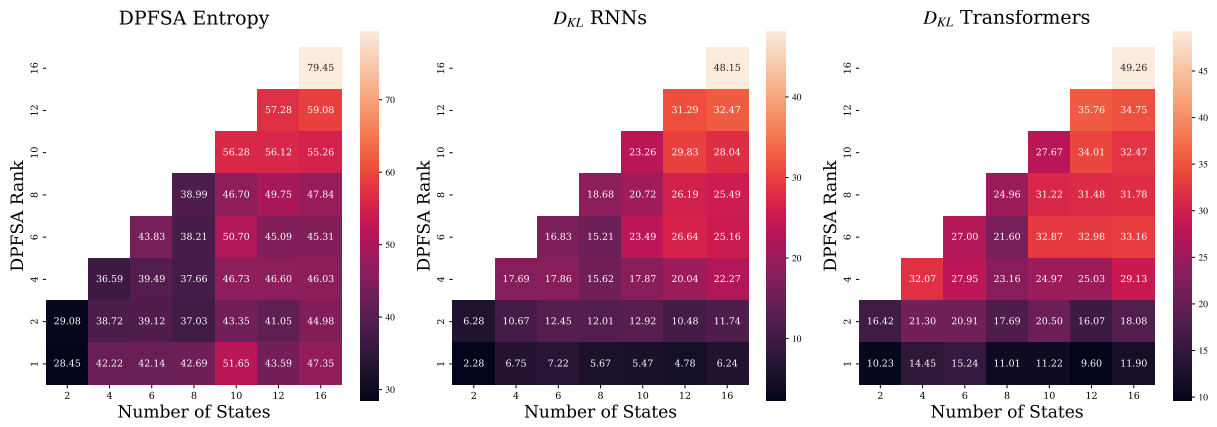


Figure 2: DPFSAs’s entropy H and the D_{KL} between the neural LMs and the DPFSAs (in bits) as a function of $|Q|$ and R .




Figure 3: DPFSAs’s entropy H and the D_{KL} between the neural LMs and the DPFSAs (in bits) as a function of $|\Sigma|$ and R .




Figure 4: DPFSAs’s entropy H and the D_{KL} between the neural LMs and the DPFSAs (in bits) as a function of D and R .

tions used by most modern neural LMs, which tend to be in the order of a few thousand dimensions (Groeneveld et al., 2024). The high performance of much smaller neural LMs on similar datasets indicates that those LMs are indeed low-rank and can thus be approximated well using smaller hidden representations. Nevertheless, Thm. 3.1 provides an interesting limitation on what distributions neural LMs of finite size can represent and points out the limitations of parameter sharing in representing formal models of computation; while neural LMs are good at approximating such models of computation, their inability to represent them *exactly* implies that, with increasing string lengths, their errors will unavoidably accumulate. This leads to poor length generalization often observed in prior work (Weiss et al., 2018; Suzgun et al., 2019b; Bhattamishra et al., 2020; Delétang et al., 2023).

Takeaways from the empirical results. The empirical results in §5.2 complement the theoretical discussion from §3 and the growing field of literature characterizing the representational capacity of neural LMs. In line with the theoretical setting and in contrast to related work, our approach directly evaluates the KL divergence between neural LMs and DPFSAs, instead of relying on classification or next-token prediction accuracy measures. Comparing distributions over strings offers a more holistic view of a neural LM’s overall ability to emulate DPFSAs allowing us to provide compelling insights into what aspects of distributions affect the learnability of formal LMs by controlling for various properties of the DPFSAs being learned. Neatly, the observed effects of the rank and the neural LM hidden state size on KL divergence align with the theoretical results derived in Thm. 3.1: Namely, as the rank of a DPFSA grows, a larger hidden state is required in the neural LM to model the DPFSA’s language appropriately. The dependence of the performance on the rank of the DPFSA demonstrates the utility of formal language theory in providing interpretable insights into the learning abilities of neural LMs.

Another noteworthy observation is the negative influence of the DPFSA entropy on the KL divergence, indicating that neural LMs tend to better learn RLMS with higher intrinsic randomness. A possible explanation for this counter-intuitive result is that when a DPFSA is truly random, its simulation is trivial. However, it is more challenging to accurately learn RLMS with more

underlying latent structures.

Extensions. We focus on the learnability of deterministic PFSA. This makes the theoretical results from §3 particularly interpretable. Extensions to the non-deterministic automata, however, are an interesting next step. Note that in this case, the PFSA rank analysis is slightly more nuanced. A non-deterministic PFSA can, at any point, be in any of the $|Q|$ states (with a probability distribution over them), meaning that the probability of the next symbol is a convex combination of the individual conditional probability distributions (not their logits). This makes the analysis trickier and less interpretable; we leave it for future work to make the current exposition more concise. A further interesting follow-up is also the study of the learnability of (deterministic) *context-free* LMs represented by probabilistic pushdown automata (PPDAs). PPDAs augment PFSA by implementing a stack that gives the automaton infinitely many configurations. Despite the infinitely many configurations, controlling for their rank analogously to the rank of a PFSA could elucidate how efficiently they are representable by neural LMs.

7 Conclusion

We provide a comprehensive empirical study of the learnability of DPFSAs by neural LMs. More concretely, we investigate how well RNN and Transformer LMs learn to match the distributions over strings generated by DPFSAs of varying complexity. To this end, we operationalized learning difficulty using the KL divergence between such distributions over strings; this measure holistically captures the similarity between the neural LM and the targeted regular LM. In an regression analysis of KL divergence variation, several complexity parameters of the DPFSA reached significance, along with the hidden state size of the neural LM. Of note were the DPFSA rank and the expected length of sampled strings, which corresponded to significant increases in KL divergence across the RNN and Transformer results. We establish that for equivalence, a neural LM’s hidden state size is theoretically lower-bounded by the DPFSA’s rank. This is consistent with the results of our controlled experiments. Our results demonstrate that using probabilistic formal language theory can help us generate insights into what neural LMs can and cannot learn. Nevertheless, our findings call for further theoretical investigations closer to practical applications.

Limitations

We point out some limitations of the presented study. First and foremost, the finite-state automata considered in this paper are limited by their bounded memory. Additionally, to keep our work concise and results self-contained, we focus only on deterministic DPFSAs. Similar and more comprehensive investigations could of course include non-deterministic automata and languages higher up on the Chomsky hierarchy, such as context-free LMs, or even context-sensitive LMs. Our experiments also omit the effect of training dataset size, which might be an interesting quantity to consider when training neural LMs. We leave those considerations to future work.

Moreover, due to computational constraints and the substantial computation load imposed by our experiments, we could not fine-tune our models with language-specific hyperparameters, which are particularly important for transformers. For the same reason, we had to refrain from optimizing larger and more capable models. However, we believe that this should not impair the validity of our results, as the trend we observed would hold even with optimal training.

Acknowledgements

Ryan Cotterell acknowledges support from the Swiss National Science Foundation (SNSF) as part of the “The Forgotten Role of Inductive Bias in Interpretability” project. Anej Svete is supported by the ETH AI Center Doctoral Fellowship. Josef Valvoda is funded by the Nordic Programme for Interdisciplinary Research Grant 105178 and the Danish National Research Foundation Grant no. DNRF169. This research was partially funded by a DFF Sapere Aude research leader grant under grant agreement No 0171-00034B, the Danish-Israeli Study Foundation in Memory of Josef and Regine Nachemsohn, and the Privacy Black & White project, a UCPH Data+ Grant. This work was further supported by the Pioneer Centre for AI, DNRF grant number P1.

References

Ekin Akyürek, Bailin Wang, Yoon Kim, and Jacob Andreas. 2024. [In-context language learning: Architectures and algorithms](#).

Yonatan Belinkov. 2022. [Probing classifiers: Promises,](#)

[shortcomings, and advances](#). *Computational Linguistics*, 48(1):207–219.

Satwik Bhattamishra, Kabir Ahuja, and Navin Goyal. 2020. [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. Association for Computational Linguistics.

Robin SM Chan, Reda Boumasmoud, Anej Svete, Yuxin Ren, Qipeng Guo, Zhijing Jin, Shauli Ravfogel, Mrinmaya Sachan, Bernhard Schölkopf, Mennatallah El-Assady, and Ryan Cotterell. 2024. [On affine homotopy between language encoders](#).

Haw-Shiuan Chang and Andrew McCallum. 2022. [Softmax bottleneck makes language models unable to represent multi-mode word distributions](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 8048–8073, Dublin, Ireland. Association for Computational Linguistics.

Stephen Chung and Hava Siegelmann. 2021. [Turing completeness of bounded-precision recurrent neural networks](#). In *Advances in Neural Information Processing Systems*, volume 34, pages 28431–28441. Curran Associates, Inc.

Ryan Cotterell, Anej Svete, Clara Meister, Tianyu Liu, and Li Du. 2024. [Formal aspects of language modeling](#).

Gregoire 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. 2023. [Neural networks and the Chomsky hierarchy](#). In *The Eleventh International Conference on Learning Representations*.

A. K. Dewdney. 1977. [Threshold matrices and the state assignment problem for neural nets](#). In *Proceedings of the 8th SouthEastern Conference on Combinatorics, Graph Theory and Computing*, pages 227–245, Baton Rouge, La, USA.

Javid Ebrahimi, Dhruv Gelda, and Wei Zhang. 2020. [How can self-attention networks recognize Dyck-n languages?](#) In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 4301–4306, Online. Association for Computational Linguistics.

Jason Eisner. 2002. [Parameter estimation for probabilistic finite-state transducers](#). In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 1–8, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.

Jeffrey L. Elman. 1990. [Finding structure in time](#). *Cognitive Science*, 14(2):179–211.

- Guhao Feng, Bohang Zhang, Yuntian Gu, Haotian Ye, Di He, and Liwei Wang. 2023. [Towards revealing the mystery behind chain of thought: A theoretical perspective](#). *arXiv preprint arXiv:2305.15408*.
- U. Grenander. 1967. *Syntax-controlled Probabilities*. Division of Applied Mathematics, Brown University.
- Dirk Groeneveld, Iz Beltagy, Pete Walsh, Akshita Bhagia, Rodney Kinney, Oyvind Tafjord, Ananya Harsh Jha, Hamish Ivison, Ian Magnusson, Yizhong Wang, Shane Arora, David Atkinson, Russell Authur, Khyathi Raghavi Chandu, Arman Cohan, Jennifer Dumas, Yanai Elazar, Yuling Gu, Jack Hessel, Tushar Khot, William Merrill, Jacob Morrison, Niklas Muenighoff, Aakanksha Naik, Crystal Nam, Matthew E. Peters, Valentina Pyatkin, Abhilasha Ravichander, Dustin Schwenk, Saurabh Shah, Will Smith, Emma Strubell, Nishant Subramani, Mitchell Wortsman, Pradeep Dasigi, Nathan Lambert, Kyle Richardson, Luke Zettlemoyer, Jesse Dodge, Kyle Lo, Luca Soldaini, Noah A. Smith, and Hannaneh Hajishirzi. 2024. [Olmo: Accelerating the science of language models](#).
- Michael Hahn. 2020. [Theoretical limitations of self-attention in neural sequence models](#). *Transactions of the Association for Computational Linguistics*, 8:156–171.
- Yiding Hao, William Merrill, Dana Angluin, Robert Frank, Noah Amsel, Andrew Benz, and Simon Mendelsohn. 2018. [Context-free transductions with neural stacks](#). In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 306–315, Brussels, Belgium. Association for Computational Linguistics.
- John Hewitt, Michael Hahn, Surya Ganguli, Percy Liang, and Christopher D. Manning. 2020. [RNNs can generate bounded hierarchical languages with optimal memory](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1978–2010, Online. Association for Computational Linguistics.
- John Hewitt and Christopher D. Manning. 2019. [A structural probe for finding syntax in word representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4129–4138, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Thomas F. Icard. 2020. [Calibrating generative models: The probabilistic Chomsky–Schützenberger hierarchy](#). *Journal of Mathematical Psychology*, 95:102308.
- P. Indyk. 1995. [Optimal simulation of automata by neural nets](#). In *STACS 95*, pages 337–348, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Ganesh Jawahar, Benoît Sagot, and Djamé Seddah. 2019. [What does BERT learn about the structure of language?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3651–3657, Florence, Italy. Association for Computational Linguistics.
- Jaap Jumelet and Willem Zuidema. 2023. [Transparency at the source: Evaluating and interpreting language models with access to the true distribution](#). In *Findings of the Association for Computational Linguistics: EMNLP 2023*, pages 4354–4369, Singapore. Association for Computational Linguistics.
- Diederik P. Kingma and Jimmy Ba. 2014. [Adam: A method for stochastic optimization](#). *arXiv preprint arXiv:1412.6980*.
- S. C. Kleene. 1956. [Representation of events in nerve nets and finite automata](#). In C. E. Shannon and J. McCarthy, editors, *Automata Studies. (AM-34), Volume 34*, pages 3–42. Princeton University Press, Princeton.
- Samuel A. Korsky and Robert C. Berwick. 2019. [On the computational power of RNNs](#). *CoRR*, abs/1906.06349.
- Brenden Lake and Marco Baroni. 2018. [Generalization without systematicity: On the compositional skills of sequence-to-sequence recurrent networks](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2873–2882. PMLR.
- Tal Linzen, Emmanuel Dupoux, and Yoav Goldberg. 2016. [Assessing the ability of LSTMs to learn syntax-sensitive dependencies](#). *Transactions of the Association for Computational Linguistics*, 4:521–535.
- Bingbin Liu, Jordan T. Ash, Surbhi Goel, Akshay Krishnamurthy, and Cyril Zhang. 2023. [Transformers learn shortcuts to automata](#). *The International Conference on Learning Representations (ICLR)*.
- Nelson F. Liu, Matt Gardner, Yonatan Belinkov, Matthew E. Peters, and Noah A. Smith. 2019. [Linguistic knowledge and transferability of contextual representations](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 1073–1094, Minneapolis, Minnesota. Association for Computational Linguistics.
- Ilya Loshchilov and Frank Hutter. 2018. [Fixing weight decay regularization in Adam](#). *arXiv.org*.
- Luca Malagutti, Andrius Buinovskij, Anej Svete, Clara Meister, Afra Amini, and Ryan Cotterell. 2024. [The role of \$n\$ -gram smoothing in the age of neural networks](#).
- Christopher D. Manning, Kevin Clark, John Hewitt, Urvashi Khandelwal, and Omer Levy. 2020. [Emergent linguistic structure in artificial neural networks](#)

- trained by self-supervision. *Proceedings of the National Academy of Sciences*, 117(48):30046–30054.
- Warren S. McCulloch and Walter Pitts. 1943. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5(4):115–133.
- William Merrill. 2019. Sequential neural networks as automata. In *Proceedings of the Workshop on Deep Learning and Formal Languages: Building Bridges*, pages 1–13, Florence. Association for Computational Linguistics.
- William Merrill. 2023. Formal languages and the NLP black box. In *Developments in Language Theory*, pages 1–8, Cham. Springer Nature Switzerland.
- William Merrill and Ashish Sabharwal. 2023a. The expressive power of transformers with chain of thought. *arXiv preprint arXiv:2310.07923*.
- William Merrill and Ashish Sabharwal. 2023b. The parallelism tradeoff: Limitations of log-precision transformers. *Transactions of the Association for Computational Linguistics*, 11:531–545.
- William Merrill, Ashish Sabharwal, and Noah A. Smith. 2022. Saturated transformers are constant-depth threshold circuits. *Transactions of the Association for Computational Linguistics*, 10:843–856.
- William Merrill and Nikolaos Tsilivis. 2022. Extracting finite automata from RNNs using state merging. *arXiv preprint arXiv:2201.12451*.
- William Merrill, Gail Weiss, Yoav Goldberg, Roy Schwartz, Noah A. Smith, and Eran Yahav. 2020. A formal hierarchy of RNN architectures. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 443–459, Online. Association for Computational Linguistics.
- Mehryar Mohri and Michael Riley. 1999. Integrated context-dependent networks in very large vocabulary speech recognition. In *6th European Conference on Speech Communication and Technology*, pages 811–814.
- Richard Montague. 1970. Universal grammar. *Theoria*.
- Franz Nowak, Anej Svete, Li Du, and Ryan Cotterell. 2023. On the representational capacity of recurrent neural language models. In *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 7011–7034, Singapore. Association for Computational Linguistics.
- Hao Peng, Roy Schwartz, Sam Thomson, and Noah A. Smith. 2018. Rational recurrences. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 1203–1214, Brussels, Belgium. Association for Computational Linguistics.
- Jorge Pérez, Pablo Barceló, and Javier Marinkovic. 2021. Attention is Turing-complete. *Journal of Machine Learning Research*, 22(75):1–35.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2021. A primer in BERTology: What we know about how BERT works. *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Joan Andreu Sánchez, Martha Alicia Rocha, Verónica Romero, and Mauricio Villegas. 2018. On the derivational entropy of left-to-right probabilistic finite-state automata and hidden Markov models. *Computational Linguistics*, 44(1):17–37.
- Hava T. Siegelmann and Eduardo D. Sontag. 1992. On the computational power of neural nets. In *Proceedings of the Fifth Annual Workshop on Computational Learning Theory*, COLT '92, page 440–449, New York, NY, USA. Association for Computing Machinery.
- Lena Strobl, William Merrill, Gail Weiss, David Chiang, and Dana Angluin. 2023. Transformers as recognizers of formal languages: A survey on expressivity. *arXiv preprint arXiv:2311.00208*.
- Mirac Suzgun, Yonatan Belinkov, Stuart Shieber, and Sebastian Gehrmann. 2019a. LSTM networks can perform dynamic counting. In *Proceedings of the Workshop on Deep Learning and Formal Languages: Building Bridges*, pages 44–54, Florence. Association for Computational Linguistics.
- Mirac Suzgun, Yonatan Belinkov, and Stuart M. Shieber. 2019b. On evaluating the generalization of LSTM models in formal languages. In *Proceedings of the Society for Computation in Linguistics (SCiL) 2019*, pages 277–286.
- Anej Svete and Ryan Cotterell. 2023. Recurrent neural language models as probabilistic finite-state automata. *Proceedings of the 2023 Conference on Empirical Methods in Natural Language Processing*, pages 8069–8086.
- Anej Svete and Ryan Cotterell. 2024. Transformers can represent n -gram language models. *arXiv preprint arXiv:2404.14994*.
- Anej Svete, Franz Nowak, Anisha Mohamed Sahabdeen, and Ryan Cotterell. 2024. Lower bounds on the expressivity of recurrent neural language models. *Proceedings of the 2024 Conference of the North American Chapter of the Association for Computational Linguistics*.
- Josef Valvoda, Naomi Saphra, Jonathan Rawski, Adina Williams, and Ryan Cotterell. 2022. Benchmarking

- compositionality with formal languages. In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 6007–6018, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.
- Shunjie Wang and Shane Steinert-Threlkeld. 2023. [Evaluating transformer’s ability to learn mildly context-sensitive languages](#). In *Proceedings of the 6th BlackboxNLP Workshop: Analyzing and Interpreting Neural Networks for NLP*, pages 271–283, Singapore. Association for Computational Linguistics.
- Gail Weiss, Yoav Goldberg, and Eran Yahav. 2018. [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. Association for Computational Linguistics.
- Jennifer C. White and Ryan Cotterell. 2021. [Examining the inductive bias of neural language models with artificial languages](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 454–463, Online. Association for Computational Linguistics.
- Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W. Cohen. 2018. [Breaking the softmax bottleneck: A high-rank RNN language model](#). In *International Conference on Learning Representations*.
- Ran Zmigrod, Tim Vieira, and Ryan Cotterell. 2021. [Efficient computation of expectations under spanning tree distributions](#). *Transactions of the Association for Computational Linguistics*, 9:675–690.

A Additional Related Work

A.1 Representational Capacity of Neural LMs

Much of the theoretical work has investigated the representational capacity of various neural LM architectures (Merrill, 2023; Strobl et al., 2023). Regular languages (and, to a lesser extent, regular LMs) have been linked to neural LMs particularly often, especially to recurrent neural LMs, but similar connections have also been made for Transformers (Merrill, 2019; Merrill et al., 2020; Liu et al., 2023). Distinctively interesting are the bounds on the space requirements for emulating FSAs (Dewdney, 1977; Indyk, 1995; Hewitt et al., 2020; Svete and Cotterell, 2023). This work bridges the theoretical work with practice, tests its applicability, and uses its insights for an informed study of the practical representational capacity of neural LMs.

A.2 Learning Formal Languages

Work similar to ours in spirit is that of Jumelet and Zuidema (2023), where the authors train and evaluate neural LMs with probabilistic context-free grammars. They use the underlying data-generating distribution (the probabilistic grammar) to evaluate how well the model has learned the distribution. Moreover, the knowledge of grammar allows them to probe the model for the encodings of individual constituents, similar to how we probe for the states of the automaton. In contrast to our work, however, Jumelet and Zuidema (2023) focus on learning human-language-based grammars, which do not provide a holistic picture of the representability of general formal LMs by neural LMs.

Delétang et al. (2023) provide a comprehensive survey of the learnability of diverse formal languages. Unlike us, they focus on learning discrete languages, particularly from the perspective of learning algorithms and investigating LMs’ inductive biases. They formulate this as a *transduction*—a string-to-string mapping. They arrive at interesting results showing that popular neural LMs are hard to place on the standard Chomsky hierarchy of languages. This can partly be explained by the mismatch of the training task—transduction—and the probabilistic nature of a neural LM, since the probabilistic Chomsky hierarchy is known to differ from the discrete one (Icard, 2020). In contrast to our work, Delétang et al. (2023) also only consider a limited set of hand-picked languages which, while providing algorithmic insights into how LMs work, do not extensively probe the learnability of the language classes.

Testing the compositional generalization of NNs, Valvoda et al. (2022) sample an infinite number of finite languages. Thereby they can draw conclusions about the learnability of an entire class of languages—sub-regular ones encoded by subsequential finite state transducers. Their work connects Montague’s theory of compositional generalization (Montague, 1970) with the popular SCAN benchmark of compositional behavior (Lake and Baroni, 2018). Unlike our work, they investigate deterministic transducers and seq2seq models.

Another similar work is that of White and Cotterell (2021), who use artificial languages to identify the biases of neural LMs. By modifying a base grammar, they experiment with the learnability of 64 languages. Unlike us, their work focuses solely on topological aspects of the language, which limits their findings to observations over the word order.

In a different line of work, Akyürek et al. (2024) evaluate neural LMs’ abilities to learn regular languages *in context*. Rather than learning one particular distribution from the training dataset, they train neural LMs to model the language of any finite-state automaton given a number of samples from it—that is, to infer the generating mechanism from the context. They consider only discrete languages (even though their generative setup is probabilistic) and due to the in-context learning setting, they do not analyze the dynamics of the neural LM implementing individual languages.

B Probabilistic Finite-state Automata

We begin by more formally defining the notion of probabilistic finite-state automata (PFSAs), which were only informally introduced in §2.

Definition B.1. A *probabilistic finite-state automaton* (PFSA) is a 5-tuple $(\Sigma, Q, \delta, \lambda, \rho)$ where Σ is an alphabet, Q a finite set of states, $\delta \subseteq Q \times \Sigma \times [0, 1] \times Q$ a finite set of weighted transitions and $\lambda, \rho: Q \rightarrow [0, 1]$ the initial and final weighting functions. Moreover, δ, λ and ρ are required to satisfy

that $\sum_{q \in Q} \lambda(q) = 1$, and, for all $q \in Q$, $\sum_{(q,y,w,q') \in \delta} w + \rho(q) = 1$. We denote $(q, y, w, q') \in \delta$ with $q \xrightarrow{y/w} q'$.

Definition B.2. A **path** π in a PFSA \mathcal{A} is a string of consecutive transitions $q_1 \xrightarrow{y_1/w_1} q_2, \dots, q_N \xrightarrow{y_N/w_N} q_{N+1}$. A path π 's **length** $|\pi|$ is the number of transitions in it and its **scan** $s(\pi)$ the concatenation of the symbols on them. We denote with $\Pi(\mathcal{A})$ the set of all paths in \mathcal{A} and with $\Pi(\mathcal{A}, \mathbf{y})$ the set of all paths that scan $\mathbf{y} \in \Sigma^*$.

The weights of the transitions along a path are multiplicatively combined to form the weight of the path. The weights of all the paths scanning the same string are combined additively to form the weight of that string.

Definition B.3. The **path weight** of $\pi \in \Pi(\mathcal{A})$ is $w(\pi) = \lambda(q_1) \left[\prod_{n=1}^N w_n \right] \rho(q_{N+1})$. The **stringsum** of $\mathbf{y} \in \Sigma^*$ is $\mathcal{A}(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\pi \in \Pi(\mathcal{A}, \mathbf{y})} w(\pi)$.

It is easy to see that the final weights $\rho(q)$ play an analogous role to the EOS symbol in the context of autoregressive LMs—they both correspond to the probabilities of ending the generation of the string.

Definition B.4. A PFSA $\mathcal{A} = (\Sigma, Q, \delta, \lambda, \rho)$ is **deterministic** if $|\{q \mid \lambda(q) > 0\}| = 1$ and, for every $q \in Q, y \in \Sigma$, there is at most one $q' \in Q$ such that $q \xrightarrow{y/w} q' \in \delta$ with $w > 0$.

In general, there can be infinitely many PFSA that define a given FSLM. However, in the deterministic case, there is a unique minimal DPFSA.

Definition B.5. A DPFSA $\mathcal{A} = (\Sigma, Q, \delta, \lambda, \rho)$ is **minimal** for the FSLM p if there is no equivalent DPFSA $\mathcal{A}' = (\Sigma, Q', \lambda', \rho', \delta')$ with $|Q'| < |Q|$.

C Proofs of Theoretical Results

Theorem 3.1. Let p be a language model induced by a minimal DPFSA \mathcal{A} with full support. Furthermore, let R be the rank of p when expressed as a representation-based LM (Eq. (3)). Then, any equivalent representation-based LM q must have a hidden state of size at least R .

Proof. This proof is effectively a restatement of the softmax bottleneck (Yang et al., 2018). Because the language model p has full support, we may transform it into a softmax-normalized, representation-based LM. By supposition, after this transformation, p 's output matrix \mathbf{T} , as in §3.1, is of rank R . For q to be equivalent to p , it has to hold that

$$\text{softmax}(\mathbf{E}\mathbf{h}(\mathbf{y})) = \text{softmax}(\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y})), \quad (5)$$

for all $\mathbf{y} \in \Sigma^*$ where \mathbf{E} 's q 's output matrix. Next, define

$$Z_1(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\bar{y} \in \bar{\Sigma}} \exp((\mathbf{E}\mathbf{h}(\mathbf{y}))_{\bar{y}}), \quad (6a)$$

$$Z_2(\mathbf{y}) \stackrel{\text{def}}{=} \sum_{\bar{y} \in \bar{\Sigma}} \exp((\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}))_{\bar{y}}). \quad (6b)$$

For all $\bar{y} \in \bar{\Sigma}$, due to the additive invariance property of the softmax, we have the following

$$\text{softmax}(\mathbf{E}\mathbf{h}(\mathbf{y}) - \log Z_1(\mathbf{y}) \cdot \mathbf{1})_{\bar{y}} = \exp(\mathbf{E}\mathbf{h}(\mathbf{y}) - \log Z_1(\mathbf{y}) \cdot \mathbf{1})_{\bar{y}} \quad (7a)$$

$$= \exp(\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) - \log Z_2(\mathbf{y}) \cdot \mathbf{1})_{\bar{y}} \quad (7b)$$

$$= \text{softmax}(\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) - \log Z_2(\mathbf{y}) \cdot \mathbf{1})_{\bar{y}}, \quad (7c)$$

where $\mathbf{1}$ is the vector of all ones $\in \mathbb{R}^{|\bar{\Sigma}|}$. It follows that, by taking logs, we have

$$\mathbf{E}\mathbf{h}(\mathbf{y}) - \log Z_1(\mathbf{y}) \cdot \mathbf{1} = \mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) - \log Z_2(\mathbf{y}) \cdot \mathbf{1}, \quad (8)$$

and, furthermore, that

$$\mathbf{E}\mathbf{h}(\mathbf{y}) = \mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) + \log \frac{Z_1(\mathbf{y})}{Z_2(\mathbf{y})} \cdot \mathbf{1}. \quad (9)$$

From Eq. (9), it follows that

$$\text{span} \left(\left\{ \mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) + \log \frac{Z_1(\mathbf{y})}{Z_2(\mathbf{y})} \cdot \mathbf{1} \mid \mathbf{y} \in \Sigma^* \right\} \right) = \text{span} (\{\mathbf{E}\mathbf{h}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\}). \quad (10)$$

Noting that $\text{span} \left(\left\{ \log \frac{Z_2(\mathbf{y})}{Z_1(\mathbf{y})} \cdot \mathbf{1} \mid \mathbf{y} \in \Sigma^* \right\} \right) = \text{span}(\{\mathbf{1}\})$, we can write

$$\text{span} (\{\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\}) \oplus \text{span} (\{\mathbf{1}\}) = \text{span} (\{\mathbf{E}\mathbf{h}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\}), \quad (11)$$

where \oplus is the direct sum, and we note that $\dim(\text{span} (\{\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\})) = R$ exactly because \mathcal{A} 's minimality implies $\dim(\text{span} (\{\mathbf{h}_{\mathcal{A}}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\})) = |Q|$. Next Eq. (11), in turn, implies

$$\underbrace{\text{span} (\{\mathbf{T}\mathbf{h}_{\mathcal{A}}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\})}_{\text{dimensionality } R} \subseteq \text{span} (\{\mathbf{E}\mathbf{h}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\}). \quad (12)$$

Thus, we arrive at a lower bound on the rank of \mathbf{E} . Why is it a lower bound? Because for an arbitrary representation-based LM, we do not know $\dim(\text{span} (\{\mathbf{h}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\}))$. However, we do know it is lower-bounded by $\dim(\text{span} (\{\mathbf{E}\mathbf{h}(\mathbf{y}) \mid \mathbf{y} \in \Sigma^*\}))$. Thus, $\text{rank}(\mathbf{E}) \geq R$. Because \mathbf{E} 's rank is bounded above by D , we have $D \geq R$, as desired. ■

D Experimental Details

D.1 Sampling DPFSAs of varying complexity

The DPFSAs we used in our experiments were sampled with $|Q| \in \{2, 4, 6, 8, 10, 12, 16\}$ over alphabets alphabets of sizes $|\Sigma| \in \{2, 4, 6, 8, 10, 12, 16\}$. Given a sampled DPFSA \mathcal{A} with $|Q|$ states over an alphabet Σ , we randomly set its unweighted transition function. That is, for each $q \in Q$ and $y \in \Sigma$ we randomly choose $q' \in Q$ and add the transition $q \xrightarrow{y} q'$ to \mathcal{A} .

We add weights to the transition function of \mathcal{A} as follows. We generate a random matrix $\mathbf{T} \in \mathbb{R}^{|\Sigma| \times |Q|}$, $\mathbf{T}_{i,j} \sim \mathcal{N}(\mu = 0, \sigma^2 = 4)$, and define $R_{\max} = \text{rank}(\mathbf{T})$ (note that $R_{\max} \leq \min(|Q|, |\Sigma|)$). For each $R \in \{r \mid r \leq \min(|Q|, |\Sigma|), r \in \{1, 2, 4, 6, 8, 10, 12, 16\}\}$, we define $\mathbf{T}^R \stackrel{\text{def}}{=} \text{SVD}(\mathbf{T}, R)$, where $\text{SVD}(\mathbf{T}, R)$ is the operation of reducing the rank of \mathbf{T} to R using SVD. Next, we add weights to the transition function of \mathcal{A} by replacing each unweighted transition $q \xrightarrow{y} q'$ with $q \xrightarrow{y/w_{q,y}} q'$, where $w_{q,y} = \text{softmax}(\mathbf{T}_{:,q}^R)_y$. Finally, we set $\rho(q) = \text{softmax}(\mathbf{T}_{:,q}^R)_{\text{EOS}}$. This process results with the generation of up to eight⁶ random DPFSAs, all sharing the same Q, Σ and underlying transition function. They differ, however, in the rank of the matrix \mathbf{T}^R that defines the weights of the transitions.

As the Transformer LMs have a limited context length, it is undesirable to sample DPFSAs that are likely to generate strings longer than the context length. This may unfairly hinder the Transformer LMs ability to learn such languages. We, therefore, filter out DPFSAs with an expected string length larger than the median value of expected string lengths (in practice, 46 symbols), i.e., half of the DPFSAs. See App. D.4 for details about how the expected generation length of DPFSAs is calculated.

D.2 Generating the Data

For a given DPFSA \mathcal{A} , we sample 20k random strings, terminating the generation process of each string when EOS is sampled. We divide the dataset into train and test splits, such that no string is shared between the sets, and the test set has at least 2k strings. We truncate the strings to 256 symbols to accommodate the limited context length of the Transformer model we used. Fig. 5 shows a histogram of the average length of strings generated for each DPFSA.

⁶Note that we have $|\{r \mid r \leq \min(|Q|, |\Sigma|), r \in \{1, 2, 4, 6, 8, 10, 12, 16\}\}| \leq 8$.




Figure 5: The statistics of the training dataset.

D.3 Training the Neural LMs

We train neural LMs on our dataset using the following procedure, repeated 2100 times:

1. Sample a random DPFSA \mathcal{A} with $|Q|$, $|\Sigma|$ and rank R using the process described in App. D.1.
2. Sample 20k strings from \mathcal{A} and split them to train set and test set using the process described in App. D.2.
3. For each $D \in \{2, 4, 6, 8, 10, 12, 16\}$ train an RNN on the train set strings.
4. For each $D \in \{2, 4, 6, 8, 10, 12, 16\}$ train a Transformer model on the train set strings.
5. Compute the D_{KL} between \mathcal{A} and the trained neural LMs on the test set strings.

We train the RNN and Transformer models using the following hyperparameters:

- **RNNs:** We use a unidirectional LSTM with four hidden layers, each with 64-dimensional hidden states and an embedding size of 64. We trained each model for two epochs using a batch size of 32 and a learning rate of 0.001, an Adam optimizer with default settings, and a standard cross-entropy loss (Kingma and Ba, 2014). We did not tie the weights of the word embeddings.
- **Transformers:** We use the GPT-2 model architecture (Radford et al., 2019) with six attention layers, each with four attention heads and 128-dimensional representations. We use an embedding size of 64 and an input context length of 256. We trained each model for two epochs using a batch size of 32 and a learning rate of 0.001, an AdamW (Loshchilov and Hutter, 2018) optimizer with default settings, and a standard cross-entropy loss. We did not tie the weights of the word embeddings.

D.4 Derivations

D.4.1 Entropy of a DPFSAs

Let $\mathcal{A} = (\Sigma, Q, \delta, \lambda, \rho)$ be a softmax-normalized DPFSAs with the parameters matrix $\mathbf{T} \in \mathbb{R}^{|\Sigma| \times |Q|}$. The entropy of \mathcal{A} is defined as

$$H(\mathcal{A}) \stackrel{\text{def}}{=} - \sum_{\mathbf{y} \in \Sigma^*} \mathcal{A}(\mathbf{y}) \log \mathcal{A}(\mathbf{y}). \quad (13)$$

Grenander (1967, Thm. 4.2) provides an efficient method of computing $H(\mathcal{A})$, which we summarize in the following, based on Sánchez et al.’s (2018) exposition.

In deterministic PFSAs, Eq. (13) can equivalently be computed by summing over the *paths* in \mathcal{A} , $\Pi(\mathcal{A})$, rather than the strings $\mathbf{y} \in \Sigma^*$; as \mathcal{A} is deterministic, there exists a single path π in $\Pi(\mathcal{A}, \mathbf{y})$ for every

$\mathbf{y} \in \Sigma^*$. We can then derive

$$H(\mathcal{A}) = - \sum_{\mathbf{y} \in \Sigma^*} \mathcal{A}(\mathbf{y}) \log \mathcal{A}(\mathbf{y}) \quad (14a)$$

$$= - \sum_{\boldsymbol{\pi} \in \Pi(\mathcal{A})} w(\boldsymbol{\pi}) \log w(\boldsymbol{\pi}). \quad (14b)$$

Now, let $\mathbf{M} \in \mathbb{R}^{|\mathcal{Q}| \times |\mathcal{Q}|}$ be the matrix of transition probabilities between each two states. That is, $M_{i,j}$ is the probability of transitioning from state $q_i \in \mathcal{Q}$ to state $q_j \in \mathcal{Q}$ in \mathcal{A} , and it is computed as

$$M_{i,j} \stackrel{\text{def}}{=} \sum_{y \in \Sigma} \mathbb{1} \left\{ q_i \xrightarrow{y/w} q_j \in \delta \right\} \cdot \text{softmax}(\mathbf{T}_{:,q_i})_y, \quad (15)$$

where $\mathbb{1} \left\{ q_i \xrightarrow{y/w} q_j \in \delta \right\}$ is the indicator function of whether the transition $q_i \xrightarrow{y/w} q_j$ exists in \mathcal{A} .⁷

Similarly, let $\boldsymbol{\alpha} \in \mathbb{R}^{|\mathcal{Q}|}$ be defined such that $\alpha_i \stackrel{\text{def}}{=} \lambda(q_i)$, i.e., the probability of starting a generation at state q_i . Finally, let $\boldsymbol{\xi} \in \mathbb{R}^{|\mathcal{Q}|}$ be defined as

$$\boldsymbol{\xi}_i \stackrel{\text{def}}{=} - \sum_{y \in \Sigma} \text{softmax}(\mathbf{T}_{:,q_i})_y \log \text{softmax}(\mathbf{T}_{:,q_i})_y. \quad (16)$$

This allows us to restate [Grenander \(1967, Thm. 4.2\)](#).

Theorem D.1 (Entropy of a DPFSA). *Given the definitions in Eqs. (15) and (16), $H(\mathcal{A})$ can be computed using the following expression*

$$H(\mathcal{A}) = \boldsymbol{\alpha}^\top (\mathbf{I} - \mathbf{M})^{-1} \boldsymbol{\xi}. \quad (17)$$

We use Eq. (17) to compute the entropy of all our regular LMs.

D.4.2 Expected String Length Under a DPFSA

To compute the expected length of a string under a DPFSA, we rely on the following lemma.

Lemma D.1 (Expected string length under a DPFSA). *The expected length of strings generated by a DPFSA $\mathcal{A} = (\Sigma, \mathcal{Q}, \delta, \lambda, \rho)$ can be computed using the following identity*

$$\mathbb{E}_{\boldsymbol{\pi} \in \Pi} [N(\boldsymbol{\pi})] = \sum_{q \in \mathcal{Q}} \left(\boldsymbol{\alpha}^\top (\mathbf{I} - \mathbf{M})^{-1} \right)_q - 1. \quad (18)$$

Proof. The following derivation proves the result

$$\sum_{q \in \mathcal{Q}} \left(\boldsymbol{\alpha}^\top (\mathbf{I} - \mathbf{M})^{-1} \right)_q = \sum_{q \in \mathcal{Q}} \left(\boldsymbol{\alpha}^\top \left(\sum_{n=0}^{\infty} \mathbf{M}^n \right) \right)_q \quad (19a)$$

$$= \sum_{q \in \mathcal{Q}} \left[\sum_{n=0}^{\infty} \left(\boldsymbol{\alpha}^\top \mathbf{M}^n \right)_q \right] \quad (19b)$$

$$= \sum_{n=0}^{\infty} \left[\sum_{q \in \mathcal{Q}} \left(\boldsymbol{\alpha}^\top \mathbf{M}^n \right)_q \right]. \quad (19c)$$

Notice that $(\boldsymbol{\alpha}^\top \mathbf{M}^n)_q$ is the total weight of paths in \mathcal{A} reaching the state q , but does not necessarily *terminate* there due to the absence of final weights ρ , after exactly n transitions, starting in initial states

⁷ \mathbf{M} is in fact the *stochastic matrix* describing \mathcal{A} .

according to their initial weights (Malagutti et al., 2024). It follows that $\sum_{i \in Q} (\boldsymbol{\alpha}^\top \mathbf{M}^n)_i$ equals the weights of all paths (visiting in any state $\in Q$) that generate a string of length of *at least* n :

$$\sum_{i \in Q} (\boldsymbol{\alpha}^\top \mathbf{M}^n)_i = \mathcal{A}(\{\mathbf{y} \mid |\mathbf{y}| \geq n\}). \quad (20)$$

Let $p(N(\boldsymbol{\pi}) = k \mid \mathcal{A})$ be the probability of the \mathcal{A} generating a string of length *exactly* k , terminating the generation after k transitions. In other words, $p(N(\boldsymbol{\pi}) = k \mid \mathcal{A})$ corresponds to the probability mass of the set $\{\mathbf{y} \mid |\mathbf{y}| = k\}$ under the DPFSA. By the reasoning above,

$$\sum_{q \in Q} (\boldsymbol{\alpha}^\top \mathbf{M}^n)_q = \sum_{k=n}^{\infty} p(N(\boldsymbol{\pi}) = k \mid \mathcal{A}), \quad (21)$$

and

$$\sum_{n=0}^{\infty} \left[\sum_{q \in Q} (\boldsymbol{\alpha}^\top \mathbf{M}^n)_q \right] \quad (22a)$$

$$= \sum_{n=0}^{\infty} \sum_{k=n}^{\infty} p(N(\boldsymbol{\pi}) = k \mid \mathcal{A}) \quad (22b)$$

$$= \sum_{k=0}^{\infty} p(N(\boldsymbol{\pi}) = k \mid \mathcal{A}) + \sum_{k=1}^{\infty} p(N(\boldsymbol{\pi}) = k \mid \mathcal{A}) + \sum_{k=2}^{\infty} p(N(\boldsymbol{\pi}) = k \mid \mathcal{A}) + \dots \quad (22c)$$

$$\begin{aligned} &= p(N(\boldsymbol{\pi}) = 0 \mid \mathcal{A}) + p(N(\boldsymbol{\pi}) = 1 \mid \mathcal{A}) + p(N(\boldsymbol{\pi}) = 2 \mid \mathcal{A}) + p(N(\boldsymbol{\pi}) = 3 \mid \mathcal{A}) + \dots \quad (22d) \\ &\quad + p(N(\boldsymbol{\pi}) = 1 \mid \mathcal{A}) + p(N(\boldsymbol{\pi}) = 2 \mid \mathcal{A}) + p(N(\boldsymbol{\pi}) = 3 \mid \mathcal{A}) + \dots \\ &\quad + p(N(\boldsymbol{\pi}) = 2 \mid \mathcal{A}) + p(N(\boldsymbol{\pi}) = 3 \mid \mathcal{A}) + \dots \\ &\quad + p(N(\boldsymbol{\pi}) = 3 \mid \mathcal{A}) + \dots \\ &\quad + \dots \end{aligned}$$

$$= \sum_{n=0}^{\infty} p(N(\boldsymbol{\pi}) = n \mid \mathcal{A}) + \sum_{n=0}^{\infty} n \cdot p(N(\boldsymbol{\pi}) = n \mid \mathcal{A}) \quad (22e)$$

$$= 1 + \sum_{n=0}^{\infty} n \cdot p(N(\boldsymbol{\pi}) = n \mid \mathcal{A}) \quad (22f)$$

$$= 1 + \mathbb{E}_{\boldsymbol{\pi} \in \Pi} [N(\boldsymbol{\pi})]. \quad (22g)$$

■

D.4.3 Cross Entropy Between a DPFSA and a Neural LM

To stochastically approximate the cross entropy between a ground-truth DPFSA and a neural LM, we compute

$$\hat{\mathbb{H}}(p(\mathcal{D}_{\text{test}}), q(\mathcal{D}_{\text{test}})) = -\frac{1}{N} \sum_{n=1}^N \left(\log q(\text{EOS} \mid \mathbf{y}^{(n)}) + \sum_{t=1}^{|\mathbf{y}_n|} \log q(y_t^{(n)} \mid \mathbf{y}_{<t}^{(n)}) \right), \quad (23)$$

where $q(\bar{y}_t \mid \mathbf{y}_{<t}^{(n)}) \stackrel{\text{def}}{=} \text{softmax}(\mathbf{E}h_{t-1})_{\bar{y}_t}$ is given by the neural LM. The strings $\mathbf{y}^{(n)}$ are sampled i.i.d. from p , language model induced by the DPFSA.

E Additional Results

This section includes figures presenting the results of additional experiments augmenting and supporting the claims in the main paper.

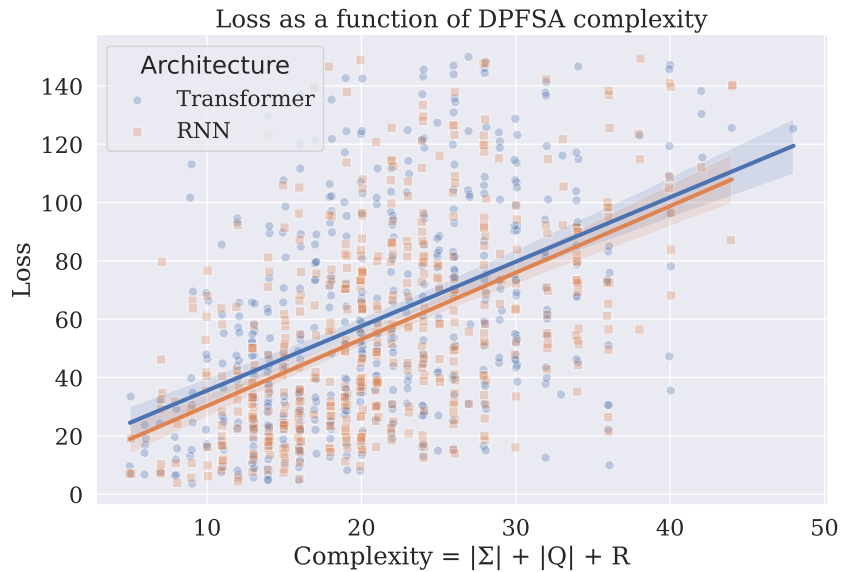


Figure 6: Validation performance of RNNs and Transformers as a function of the PFSA’s complexity, computed as $|\Sigma| + |Q| + R$. We compute the loss by summing it over symbols and dividing this sum by the number of strings in the test set. D_{KL} as a function of the expected sequence length of the DPFSA

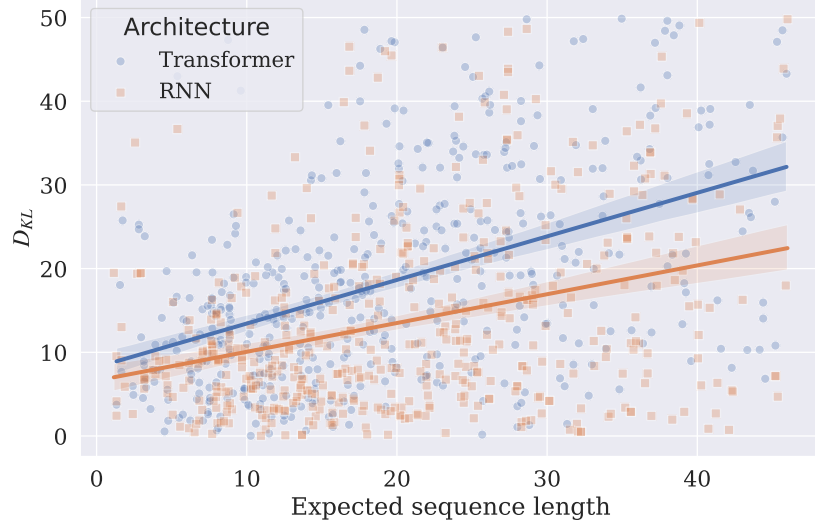


Figure 7: D_{KL} of RNNs and Transformers as a function of the average string length of the DPFSA.

Fig. 6. Overall performance of the neural LM models as a function of the “total complexity” of the PFSA they were optimised for, which we define as the sum of $|\Sigma|$, $|Q|$, and R . Performance is measured as the cross-entropy of the neural model on a held-out test set. We compute the loss by summing it over symbols and dividing this sum by the number of strings in the test set. We can see that RNNs tend to overperform Transformers, especially for more complex PFASAs.

Fig. 7. The D_{KL} of RNNs and Transformers as a function of the average string length of the DPFSA. Similarly to Tab. 4, we see that Transformers are much more sensitive to string length compared to RNNs.