

Simplicity Bias in Transformers and their Ability to Learn Sparse Boolean Functions

Satwik Bhattamishra[♠] Arkil Patel[◇] Varun Kanade[♠] Phil Blunsom^{♠♣}

[♠]University of Oxford [◇]Mila and McGill University [♣]Cohere

{satwik.bmishra, varun.kanade, phil.blunsom}@cs.ox.ac.uk arkil.patel@mila.quebec

Abstract

Despite the widespread success of Transformers on NLP tasks, recent works have found that they struggle to model several formal languages when compared to recurrent models. This raises the question of why Transformers perform well in practice and whether they have any properties that enable them to generalize better than recurrent models. In this work, we conduct an extensive empirical study on Boolean functions to demonstrate the following: (i) Random Transformers are relatively more biased towards functions of low sensitivity. (ii) When trained on Boolean functions, both Transformers and LSTMs prioritize learning functions of low sensitivity, with Transformers ultimately converging to functions of lower sensitivity. (iii) On sparse Boolean functions which have low sensitivity, we find that Transformers generalize near perfectly even in the presence of noisy labels whereas LSTMs overfit and achieve poor generalization accuracy. Overall, our results provide strong quantifiable evidence that suggests differences in the inductive biases of Transformers and recurrent models which may help explain Transformer’s effective generalization performance despite relatively limited expressiveness.

1 Introduction

Transformers (Vaswani et al., 2017) have supplanted recurrent models across a range of NLP tasks (Liu et al., 2019; Brown et al., 2020). In particular, effective large-scale pretrained models have predominantly been Transformer-based models and have found application in other areas such as computer vision and protein folding. Given the irrefutable importance of understanding these architectures, a significant effort has been devoted to analyze the inner workings of large-scale pretrained Transformers. However, the cause behind the difference in performance between Transformers and recurrent models has largely been unclear.

A line of work has attempted to understand neural sequence models through the lens of formal language theory. These works have sought to understand the expressive power of these architectures and identify differences in their ability to generalize across various formal languages. A notable result by Hahn (2020) showed that Transformers are limited in their ability to express the PARITY language¹ while it is well known that small-sized RNNs can express such languages. Across empirical studies, Transformers have been found to perform worse or comparably to LSTMs in almost all formal languages previously considered in the literature (Bhattamishra et al., 2020a; Delétang et al., 2022; Chiang and Cholak, 2022). In particular, Transformers have been shown to struggle with the PARITY language and certain other regular languages. This leads to a natural question: Why do Transformers perform so well in practice if they are arguably less expressive and perform worse than LSTMs across certain formal languages?

Although recurrent models such as LSTMs have been shown to perform better on formal languages such as PARITY, we find that they struggle to generalize well on several sparse Boolean functions such as SPARSE PARITIES. We find a clear contrast between the generalization abilities of Transformers and LSTMs on various k -SPARSE Boolean functions which have low sensitivity. Additionally, through extensive empirical analysis, we provide strong evidence to suggest differences in the bias towards low complexity functions between Transformers and recurrent models. Based on our results, we hypothesize that one of the reasons behind Transformer’s practical effectiveness could be that they are more biased towards simple functions in comparison to recurrent models which may lead to better generalization.

In particular, we focus on a complexity measure

¹Computing whether a bit string has odd or even number of ones.

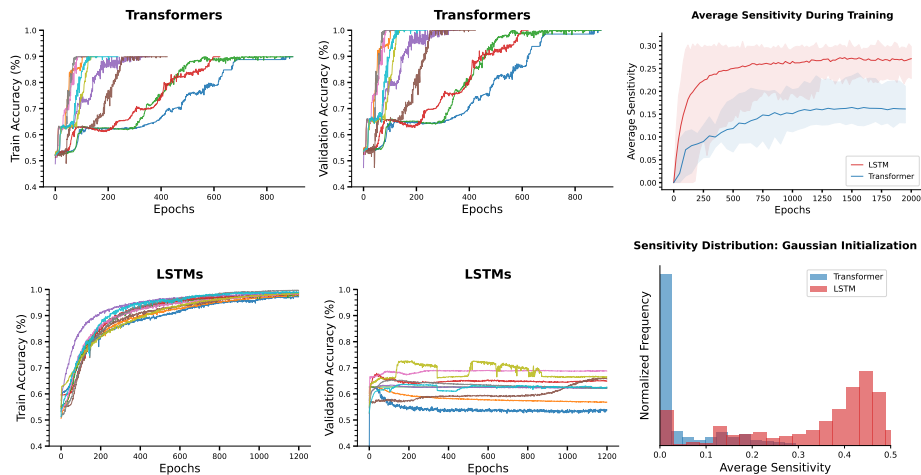


Figure 1: Overview of the main findings of the paper. Training and validation curves for Transformers (top left) and LSTMs (bottom left) on 10 different k -SPARSE functions in the presence of 10% noise. Sensitivity of Transformers and LSTMs during training on random Boolean functions (top right). Distribution of sensitivity of functions represented by random models with weights initialized according to normal distribution ($\mu=0, \sigma=10$) (bottom right).

called sensitivity (Kahn et al., 1989), which measures how likely it is that a function value changes due to a ‘small’ change in input. Sensitivity is related to several other complexity measures; functions with low sensitivity have low Kolmogorov complexity,² simpler Fourier spectra, and can be represented by decision trees of small depths. The relationship between sensitivity and generalization has also been previously studied in the literature (Novak et al., 2018; Franco, 2006).³ While measures such as Kolmogorov complexity are uncomputable, sensitivity can be tractably estimated and extensions of sensitivity can be used to estimate the complexity of functions in more realistic NLP tasks (Hahn et al., 2021).

Our Contributions. We investigate the bias in (a) *parameter space* by analyzing randomly initialized models, and (b) *learning procedure* by examining the sensitivity of models during the training process. Motivated by our findings indicating differences between the biases of Transformers and LSTMs, we evaluate their performance on functions of low sensitivity.

(i) We demonstrate that random Transformers are significantly more likely to represent functions of lower sensitivity than recurrent models when the weights are sampled uniformly or according to Normal distribution (see Figure 1, bottom right). When the weights are initialized following practical strategies (such as Xavier normal), then both

architectures are likely to have low sensitivity with Transformers having relatively lower sensitivity.

(ii) We show that both Transformers and LSTMs learn functions of increasing sensitivity when trained on a set of Boolean functions as well as practical datasets such as sentiment classification (see Figure 1, top right). For Boolean functions, Transformers converge to functions of lower sensitivity in comparison to LSTMs when both models achieve near-zero training error.

(iii) On various k -SPARSE Boolean functions, we find that Transformers generalize near-perfectly even in the presence of noise in the training data whereas LSTMs severely overfit and obtain poor generalization performance (see Figure 1, left).

Auxiliary Results. Although not the primary focus of the paper, we explore relations between sensitivity and generalization in Appendix D. In particular, we show how sensitivity can be used as a capacity measure to derive generalization bounds. Additionally, we explore the correlation between sensitivity and generalization gap for LSTMs and Transformer-based models on sentiment classification tasks. We also conduct experiments with three other complexity measures in Appendix B.1.

2 Related Work

Random Neural Networks. One approach to explaining deep learning’s unexpected generalization performance has been to study the inductive biases of random neural networks. Several prior works have shown theoretically (Palma et al., 2019) and

²The converse is not always true.

³A more thorough discussion is presented in Appendix B.2

empirically (Valle-Perez et al., 2019) that random untrained feedforward networks are biased towards ‘simple’ functions. Valle-Perez et al. (2019) showed that considering the distribution over functions generated via random neural networks as a prior leads to better PAC-Bayesian generalization bounds than traditional ones. Several works (Mingard et al., 2019; Wilson and Izmailov, 2020; Lee et al., 2017) have argued using heuristic methods that the inductive biases in random neural networks can be used to understand the properties of trained networks. Additionally, there is empirical and theoretical evidence (Oymak and Soltanolkotabi, 2019) that neural networks trained with SGD usually converge close to the initialization point. Hence, understanding the properties of random neural networks is imperative to understand their generalization abilities.⁴ In Section 4.1, we study the complexities of random Transformers and recurrent models and investigate the differences between them.

Formal Languages and Sequence Models. In the past few years, a strand of work⁵ primarily in the NLP community has attempted to understand neural sequence models’ capabilities and inner workings by analyzing them on formal languages, e.g. (Suzgun et al., 2019b; Sennhauser and Berwick, 2018). Given the recent success of Transformers, several works have sought to investigate them via the lens of formal languages. Hahn (2020) theoretically showed the limitations of Transformers in recognizing languages like Parity and Dyck-2. While Transformers are expressive enough to represent the PARITY language for bounded lengths (Chiang and Cholak, 2022), multiple works have observed that they struggle to generalize well on Parity and other regular languages when tested empirically (Bhattamishra et al., 2020a; Chiang and Cholak, 2022; Delétang et al., 2022). In contrast to this, we show that when evaluated on some simpler variants of these formal languages, Transformers generalize near perfectly whereas LSTMs achieve poor generalization performance.

3 Background and Preliminaries

3.1 Sensitivity of Boolean Functions

We will work with a complexity measure called Boolean Sensitivity which has been widely studied in computational complexity (Kahn et al., 1989;

⁴There is an enormous literature on this topic; refer to Mingard et al. (2021) for more references and discussion.

⁵See Appendix H for a more comprehensive version.

Ambainis et al., 2014). Sensitivity can be seen as a discrete analog (Gopalan et al., 2016) of the ‘smoothness’ of a continuous function which measures how gradually a function changes locally. For Boolean functions defined over the Hamming cube, sensitivity captures how many neighbours of a particular input have different outputs. Formally, the sensitivity of a Boolean function $f : \{0, 1\}^n \rightarrow \{\pm 1\}$ at input $x \in \{0, 1\}^n$ is defined as

$$s(f, x) = \sum_{i=1}^n \mathbb{I}[f(x) \neq f(x^{\oplus i})], \quad (1)$$

where \mathbb{I} denotes the indicator function and $x^{\oplus i} = (x_1, \dots, x_{i-1}, 1 - x_i, x_{i+1}, \dots, x_n)$ is the same as x at every coordinate or bit except the i -th one. The maximum sensitivity of a function f is defined as $ms(f) = \max_{x \in \{0, 1\}^n} s(f, x)$. The average sensitivity (also referred to as total influence) of a Boolean function measures the average of the sensitivity of the function across all inputs $x \in \{0, 1\}^n$ and is defined as

$$s(f) = \mathbb{E}_x[s(f, x)] = \frac{1}{2^n} \sum_{x \in \{0, 1\}^n} s(f, x). \quad (2)$$

See that $0 \leq s(f) \leq ms(f) \leq n$. To compare across inputs of different lengths, in our experiments we will normalize the average sensitivity across length $\mathcal{S}(f) = \frac{1}{n}s(f)$ which can also be interpreted as,

$$\mathcal{S}(f) = \Pr_{x \sim \{0, 1\}^n, i \sim [n]} [f(x) \neq f(x^{\oplus i})] \quad (3)$$

where $[n] = \{1, \dots, n\}$ and the sampling is over uniform distribution over the domains.

Parity. The Parity function over $\{0, 1\}^n$ is defined as $\text{PARITY}(x) := (-1)^{\sum_{i=1}^n x_i}$. For any input $x \in \{0, 1\}^n$, the function PARITY has value +1 if the number of ones in the input is even and has value -1 otherwise. The sensitivity of the Parity function is the maximum among all functions since changing any bit of any input changes the function value. Hence, for PARITY over $\{0, 1\}^n$, $s(\text{PARITY}) = n$ and $\mathcal{S}(\text{PARITY}) = 1$.

Sparse Boolean functions. Another class of functions are the k -SPARSE functions (also referred to as k -juntas) where the function value depends on at most k coordinates of the input. More formally, a function $f : \{0, 1\}^n \rightarrow \{\pm 1\}$ is k -SPARSE if

there exist indices $1 \leq i_1 < i_2 < \dots < i_k \leq n$ and a function $g : \{0, 1\}^k \rightarrow \{\pm 1\}$, such that for every $x \in \{0, 1\}^n$, $f(x_1, x_2, \dots, x_n) = g(x_{i_1}, x_{i_2}, \dots, x_{i_k})$. Let $\text{SPARSE-}(k, n)$ be the class of k -SPARSE functions on inputs of length n that depend on at most k bits. It is easy to see that, for any $f \in \text{SPARSE-}(k, n)$, the average sensitivity $s(f) \leq k$ (and hence $\mathcal{S}(f) \leq \frac{k}{n}$).

When $k \ll n$, $\text{SPARSE-}(k, n)$ can be seen as a subclass of all Boolean functions with low average sensitivity. Other functions with low average sensitivity can also be approximated with k -SPARSE functions using Friedgut’s Junta Theorem (O’Donnell (2021), Page 269). The maximum average sensitivity $s(f) = k$ is attained by SPARSE PARITIES denoted f_{parity_k} which is the Parity over a subset of k coordinates. A sparse parity function f_{parity_k} over $S \subseteq [n]$, s.t. $|S| = k$ is $+1$ if the number of ones in the coordinates S is odd and -1 otherwise. Other Boolean functions such as sparse majority can be defined similarly. The majority function f_{maj} over $\{0, 1\}^n$ is $+1$ if the number of ones in the input is greater than the number of zeros and is -1 otherwise. Similarly, the sparse majority function f_{maj_k} is the majority function over coordinates $S \subseteq [n]$, s.t. $|S| = k$. Parities (and Sparse Parities) are an important class of Boolean functions since any Boolean function can be represented as a linear combination of a set of Parity functions.

4 Sensitivity Experiments

In this section, we conduct various experiments to investigate the differences in the bias of Transformers and RNNs towards functions of low sensitivity.⁶ From here onward, whenever sensitivity is mentioned, we will refer to the length normalized version of average sensitivity \mathcal{S} defined in Eq. (3). The first part of this section deals with analyzing the sensitivity of random Transformers and RNNs while the second part investigates the sensitivity of models trained to fit random Boolean functions.

4.1 Sensitivity of Randomly Initialized Models

We seek to understand the landscape of the complexity of functions in the parameter space of Transformers and RNNs. Let us assume that the parameter space Θ of our models is bounded, i.e. all the parameters (weights) take some value within some

⁶We have made our source code available at <https://github.com/satwik77/Transformer-Simplicity>.

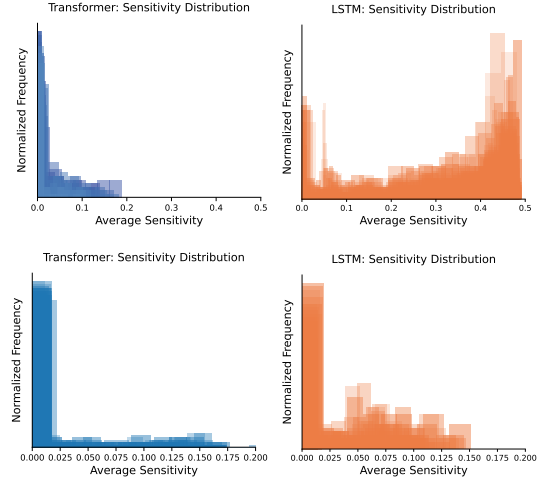


Figure 2: Distribution of sensitivity of different randomly initialized Transformers and LSTMs. *Top row:* Transformers (left) and LSTMs (right) with uniformly sampled weights across various hyperparameters. *Bottom row:* Analogous to top row but with Xavier normal initialization. Refer to Section 4.1 for details.

bounded range $[-B, B]$. A particular realization of the parameters with values in $[-B, B]$ leads to the model being a function from $\{0, 1\}^n \rightarrow \{0, 1\}$. We begin with a simple question: Out of all the parameterizations in the parameter space of Transformers (or RNNs), if we select one uniformly at random, then how likely is it to have low sensitivity?

Setup. In all our experiments, we consider binary classifiers with Transformers and RNN-based architectures. By Transformer, we refer to the encoder-only version of the original Transformer architecture (Vaswani et al., 2017) as used in models such as BERT (Devlin et al., 2019). The model takes a sequence of tokens along with a [CLF] token as input. The final classification is done based on the output vector of the [CLF] token. For recurrent models, we consider LSTMs (Hochreiter and Schmidhuber, 1997), GRUs, and RNNs with tanh activation. Most of the results in the main paper pertaining to recurrent models are based on experiments with LSTMs and we discuss when the behaviour is different for other recurrent models.

In our experiments, we explore four strategies to sample random networks: Uniform, Gaussian, Xavier uniform, and Xavier normal initialization. In uniform sampling, each parameter (weights and biases) is assigned a value by uniformly sampling in $[-10, 10]$. Similarly, for Gaussian initialization, each parameter is assigned by sampling from $\mathcal{N}(0, \sigma^2)$ where we set sigma as 10. Xavier normal (Glorot and Bengio, 2010) initialization is the one




Figure 3: Distribution of sensitivity of randomly initialized Transformers and LSTMs for a fixed hyperparameter (layers=2, width=256).

that is more commonly used in practice to train these models. All the weights are initialized with $\mathcal{N}(0, \sigma^2)$ where the standard deviation $\sigma = d^{-1/2}$ where d is the number of hidden units. All the input embedding vectors and positional embedding vectors are initialized with $\mathcal{N}(0, 1)$ which is the default scheme in PyTorch (Paszke et al., 2019). For input lengths greater than 10, we estimate the sensitivity of each model by computing the average over a sampled set of bit strings. We sample 10k bit strings and compute the average sensitivity across the samples. For each hyperparameter configuration, we sample 75-1000 different models to estimate their sensitivity depending on the computational costs associated with it. For most of the results reported in the main paper, we consider bit strings of length 20. But we also experiment with lengths $\in \{5, 7, 10, 15, 20, 50, 100, 200\}$.

Results. Figure 2 (upper row) and Figure 3 (left) shows the distribution of sensitivity for uniformly initialized Transformers and LSTMs. The distribution for Transformers is heavily skewed towards functions of very low sensitivity in comparison to LSTMs. The pattern holds across Gaussian initialization as well (see Figure 1, bottom right). For initialization strategies used in practice such as Xavier normal and Xavier uniform, we find that both Transformers and LSTMs have low sensitivity (see Figure 2, lower row and Figure 3, right) with Transformers having relatively lower average sensitivity. Refer to Section B.3 in the Appendix for results with Xavier uniform initialization.

Although we primarily discuss results with sensitivity in the main paper, similar experiments with other complexity measures are presented in Appendix B.1. Further experiments exploring the change in distribution across the number of layers, width, and lengths for both architectures are presented in Appendix B.3.

Discussion. These results imply that lower sen-




Figure 4: Distribution of sensitivity of Transformers and LSTMs trained on Boolean strings with random labels. Refer to Section 4.2 for details.

sitivity functions are *over-represented* in the parameter space of Transformers. If every Boolean function $f : \{0, 1\}^n \rightarrow \{0, 1\}$ would have had equal representation in the parameter space of the model, then the distribution would have concentrated around $1/2$. A learning algorithm based on a random search over the parameter space is more likely to encounter functions of low sensitivity. Note that, while recurrent models have higher sensitivity than Transformers, they are still lower than randomly sampling a Boolean function.

Why randomly initialized models? Each randomly initialized Transformer or RNN when restricted to domain $\{0, 1\}^n$ represents one of the 2^{2^n} Boolean functions f in \mathcal{F} . The distribution over \mathcal{F} induced by randomly initialized models can be seen as their prior $P(f)$. Given a set of (training) examples $S = \{(x_1, y_1), \dots, (x_m, y_m)\}$, let $P_B(f|S)$ denote the probability of sampling f conditioned on the event that the sampled function is consistent with S (matches all the input-output mappings in S). We can apply Bayes' rule to calculate the posterior $P_B(f|S) = P(S|f)P(f)/P(S)$ using the prior $P(f)$, the likelihood $P(S|f)$, and the marginal likelihood $P(S)$. Since we condition on f being consistent with S (zero training error), the likelihood $P(S|f) = 1$ if $\forall x_i \in S, f(x_i) = y_i$ and 0 otherwise. Let $U(S)$ denote the set of all functions $f \in \mathcal{F}$ which are consistent with S . Note that, given a fixed training set S , since $P(S)$ is constant, the probability over the choice of $f \in U(S)$ ultimately depends on the prior $P(f)$. In practice, we do not fit the training set by sampling models to find one that is consistent with the training set. However, recent work (Mingard et al., 2021) has shown that $P_B(f|S) \approx P_{SGD}(f|S)$ across a range of neural architectures and data sets. The SGD-based posterior $P_{SGD}(f|S)$ denotes the probabil-

ity that a neural network converges on function f when trained to fit S . Hence, our results suggest that for Transformers, $P_B(f|S)$ would be concentrated on low-sensitivity functions and consequently, $P_{SGD}(f|S)$ could be biased towards low-sensitivity functions as well.

4.2 Models learn functions of increasing sensitivity

In this section, we investigate the sensitivity of functions learned during the training process when Transformers and LSTMs are trained to fit datasets of Boolean strings with random labels.

Setup. We create datasets of size 1k each by uniformly sampling bit strings of length 40. The label for each input string is assigned randomly (+1 or -1 with probability 1/2). All the weights of the models are initialized with Xavier normal initialization and the biases are initialized with zero vectors. We consider Transformers and LSTMs across various hyperparameter configurations with a similar number of parameters. We train the models until they reach zero training error and estimate the sensitivity of the models at every epoch. We conduct the experiments over 20 different datasets with 100 runs for Transformers and LSTMs each.

Sensitivity during training. We find that both Transformers and LSTMs gradually learn functions of increasing sensitivity with Transformers converging to functions of much lower sensitivity than LSTMs (refer to Figure 1, top right). We observe similar behavior when the models are trained on various sparse Boolean functions including sparse parities. Even though sensitivity is defined over Boolean functions, we explore a few natural extensions to estimate the sensitivity of models trained on real datasets such as sentiment classification. On two sentiment classification datasets, namely SST and IMDB, we found similar observations where both Transformers and LSTMs seem to incrementally learn functions of increasing sensitivity. See Appendix C for more details.

Discussion. Even if sequence models such as Transformers or LSTMs are capable of representing arbitrary functions, our results suggest that they prioritize learning simpler patterns first. These results echo prior observations that indicate feedforward-like neural networks trained with SGD learn functions of increasing complexity (Nakkiran et al., 2019; Arpit et al., 2017). Rahaman et al. (2019) find that ReLU neural networks learn func-

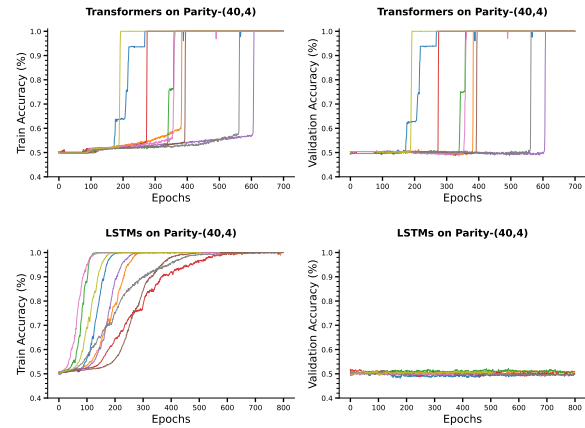


Figure 5: Training and validation curves for Transformers and LSTMs trained on SPARSE PARITIES of length $n = 40$ and $k = 4$.

tions of lower frequency modes first. Functions with lower average sensitivity also have a lower frequency and hence these observations are closely connected. More importantly, average sensitivity can be naturally extended to real data which allows us to empirically explore this for text data.

Sensitivity upon convergence. For Transformers and LSTMs trained until 0% training error, we estimate the sensitivity of functions learned by the models. We create 15 datasets and for each dataset, we compute the sensitivity of 100 trained models. The combined distribution of the sensitivity of the models across all datasets is shown in Figure 4. We observe that Transformers consistently learn functions of lower sensitivity in comparison to LSTMs. This supports our hypothesis that for Transformers the parameter search via algorithms such as Adam is more likely to find functions of lower sensitivity that fit the training set as opposed to LSTMs.

5 Experiments on Sparse Boolean Functions

Our results in the previous section indicate that relative to LSTMs, random Transformers are biased towards low-sensitivity functions and Transformers are biased towards learning Boolean functions of low sensitivity. Motivated by this difference in bias, we conduct experiments geared towards answering the following question: Is there any difference between the ability of Transformers and LSTMs to learn sparse Boolean functions which have low sensitivity?




Figure 6: Training and validation curves for Transformers trained on PARITY-(40,4) with 10% noise in the training data. See Section 5 for details.

5.1 Setup.

Boolean Functions. We focus on k -SPARSE Boolean functions which have low sensitivity when $k \ll n$ (refer to Section 3 for definition). We first consider certain Boolean functions which are widely studied in the analysis of Boolean functions. The first one is SPARSE PARITIES which can be interpreted as the k -SPARSE variation of standard parity. We denote an instance of SPARSE PARITIES as PARITY- (n, k) where n denotes the length of the input string and k denotes the number of relevant bits. We denote an instance of standard PARITY as PARITY- n where n denotes the length of the input string and the output is computed based on the number of ones in all indices. Learning PARITY- (n, k) with gradient-based methods has well-known hardness results – requiring at least $n^{\Omega(k)}$ computational steps to find the correct target function (Kearns, 1998). The other two Boolean functions we consider are sparse majorities (denoted by MAJ- (n, k)) and the dictator function (denoted by DICT- n). The output of the dictator function depends only on a single input bit, making it arguably one of the simplest Boolean functions with very low sensitivity. In MAJ- (n, k) , the output for a string of length n is determined by whether the number of ones is greater than the number of zeros in the k relevant indices.

The second set of Boolean functions we consider is random k -SPARSE functions (denoted by JUNTAS- (n, k)). For each instance of JUNTAS- (n, k) , the function is determined by randomly choosing k indices and assigning labels to each of the 2^k distinct inputs randomly.⁷

Noisy Labels. We also conduct experiments to examine the ability of the models to learn in the presence of noise. In these experiments, labels of

⁷Note that this is separate from the random Boolean functions in Section 4.2 where labels of unseen data is also random. In this scenario, the labels of the unseen inputs can be determined by learning the correct function from the training set.

training data are flipped with a certain probability η . Thus, about $1 - \eta$ fraction of the training data is clean and η fraction of the training data has incorrect labels. The validation set is clean without any modifications. The goal is to investigate whether a model is robust to noise during the training process.

Training Details. The training and validation sets are created by uniformly sampling bit strings over $\{0, 1\}^n$. In our experiments, we consider Transformers with 1-6 layers, 4-8 heads and width (usually referred to as `d_model`) within 8-128. We consider Transformers with both learnable and absolute positional encodings. For LSTMs, we consider up to 6 layers and widths (also referred to as `hidden_size`) within 8-256. The size of the token embeddings is kept the same as the width. We also consider the presence of learnable positional embeddings as a hyperparameter. We use batch sizes of 100 and 500 in all our experiments and tune across learning rates $\in \{1e-1, 5e-2, \dots, 1e-6\}$. For each dataset, we extensively tune the models across various hyperparameters, details of which are provided in Appendix G.

5.2 Experiments

Parities. For PARITY-(40, 4) and PARITY-40, we create 5 different datasets and report the results based on the maximum accuracy achieved on unseen test data. The train set consists of 30k samples and the validation sets contain 10k samples.

We observe a stark contrast between the performance of Transformers and LSTMs on different forms of parity tasks. We find that Transformers struggle to fit and generalize on PARITY-40 while LSTMs easily (across a range of hyperparameters) generalize well on them. On the other hand, perhaps surprisingly, on PARITY-(40, 4), we find that while Transformers generalize well, LSTMs severely overfit and achieve poor validation accuracy. Although LSTMs achieve 100% training accuracy over the training data, their validation accuracy does not move far beyond the chance level (50%). Figure 5 depicts the training and validation accuracy curves for Transformers and LSTMs on PARITY-(40, 4) task. We find similar behaviour for LSTMs even with learnable positional embeddings.

Robustness to noise. On SPARSE PARITIES datasets, we find that Transformers are surprisingly robust to noise. When the training data contains 5%-20% noise (η), Transformers achieve perfect generalization accuracy with training accuracy con-




Figure 7: Validation accuracies of LSTMs on sparse Majority and Dictator datasets averaged over 5 runs for best performing hyperparameter.

verging at $1 - \eta$. In some cases, after training for a large number of iterations post-convergence, Transformers begin to overfit on the noise. This observation echoes a similar finding in Tanzer et al. (2022) where they observed such behaviour while finetuning large pretrained models for sequence tagging tasks in the presence of noise. The training and validation accuracy curves are provided in Figure 6. The behaviour of recurrent models is the same as in the previous scenario with clean data: they overfit on the training data while achieving chance level validation accuracy. Additional results on PARITY- (n, k) across different dataset sizes, task variations, as well as exploring phenomena such as phase transitions and grokking are provided in Appendix E.

We observe this pattern across other sparse Boolean functions such as sparse majority and dictator functions as well. For sparse majority datasets MAJ- $(n, 5)$, we consider lengths $n \in \{50, 75, 100, 200\}$ and for dictator functions DICT- n , we consider lengths $n \in \{100, 200, 300, 500, 700\}$. We experiment with various rates of noise (10 - 30%). While LSTMs do generalize well up to certain lengths, they achieve poor validation accuracy ($<75\%$) as the lengths go higher. At the same time, they obtain 100% training accuracy on all the datasets. The validation accuracies of LSTMs are reported in Figure 7. In contrast, Transformers achieve near-perfect generalization even in the presence of significant noise.

Random k -sparse functions. For JUNTAS- (n, k) , we experiment with various datasets for JUNTAS- $(n, 5)$ with $n \in \{30, 50, 80, 150, 200\}$. For lengths $n < 150$, we find that LSTMs generalize well on some of the JUNTAS- $(n, 5)$ functions. However, in the presence of 10% noise (i.e., $\eta = 0.1$), their

performance degrades sharply. We create 10 datasets for JUNTAS- $(50, 5)$ with $\eta = 0.1$, and similar to previous scenarios, LSTMs struggle to generalize well ($>75\%$) whereas Transformers are able to generalize perfectly on all the datasets (see Figure 1, top middle). However, even when the validation accuracies of LSTMs were below 75%, their training accuracy reached 100% indicating that they overfit on the training data. Figure 1 (bottom left) shows the training and validation curves of LSTMs on the 10 datasets.

Sensitivity During Training. We observe that on k -SPARSE functions, both Transformers and LSTMs learn functions of increasing sensitivity. However, when LSTMs overfit and reach zero training error, they converge to functions of much higher sensitivity than that of the target function (see Figure 16. Since Transformers generalize perfectly, their sensitivity matches that of the target function.

6 Clarifications

(1) *Do our results imply that Transformers can learn any k -SPARSE functions with small (practical) number of examples?* No. For small lengths ($n = 50$) and $k = 3$, we could enumerate and verify that they are able to learn all functions in the presence of 10% noise. However, as the length n and the number of relevant bits k grow, Transformers struggle to perform well. Given the computational hardness associated with learning SPARSE PARITIES, the task becomes much more difficult with the increase in n and k . For $n = 100$ and $k = 5$, we were not able to obtain good generalization performance with Transformers.

(2) *Do Transformers never overfit on k -SPARSE functions?* They do overfit when the size of the training data is very small. For SPARSE PARITIES with $n = 40$ and $k = 4$, it is perhaps surprising that Transformers learn the correct function even with as little as 2500 training examples in less than 10000 computational steps. However, for training sets of size 1000, Transformers overfit across all runs. Additionally, with training sets of size 5000 - 10000, Transformers with higher depths overfit in some cases. See Appendix E for more details.

(3) *Does the low sensitivity bias of Transformer (Section 4) explain their good generalization performance on k -SPARSE functions such as SPARSE PARITIES?* No. Our findings in Section 4 motivated us to compare the performance of Transformers and LSTMs on functions of low sensitiv-

ity such as k -SPARSE functions. While the bias towards low sensitivity functions and strong performance on various k -SPARSE functions could be related, it is not a direct explanation for their performance on k -SPARSE. For SPARSE PARITIES, it is natural to expect Transformers to follow some mechanism along the lines presented in Barak et al. (2022) for FFNs trained with SGD. However, the exact details are unclear, and more importantly, why and how LSTMs overfit is unclear as well.

(4) *Are Transformers performing better than LSTMs because of learnable positional embeddings?* This seems unlikely since we found that Transformers with absolute positional encoding also generalize well on sparse parities (see Figure 18). Moreover, we found that LSTMs with learnable positional embeddings also fail to generalize on sparse parities and behave similarly to Figure 5.

(5) *Do LSTMs never succeed in learning SPARSE PARITIES from data?* They do succeed for smaller lengths. For lengths up to 20, we find that both Transformers and LSTMs are able to learn PARITY and SPARSE PARITIES. However, for higher lengths, Transformers struggle to fit PARITY and LSTMs begin to overfit on SPARSE PARITIES. For length $n = 20$ and $k = 4$, we could robustly find that even LSTMs without positional embeddings succeeded in learning sparse parities. On the other hand, for $n = 40$ and $k = 4$, we robustly found that LSTMs with learnable positional embeddings overfit and achieve poor generalization performance. Transformers were able to generalize well in the presence of noise across various hyperparameters for SPARSE PARITIES with $n = 40$ and $k = 4$. Our goal is not to identify the exact class of functions that Transformers can learn in practice. The key result is the juxtaposition of the performance between Transformer and LSTMs across various k -SPARSE functions.

(6) *Do Transformers work effectively in practice primarily due to their simplicity bias?* It is hard to answer this question. In our work, we try to highlight concrete differences between Transformers and LSTMs with respect to certain properties which have close connections to generalization. While these properties could partially be the reason behind their good generalization performance, it is also possible that they are ubiquitous in practice because they effectively model long-distance dependencies and can be trained efficiently.

7 Discussion and Final Remarks

A natural question that arises from our results is whether Transformers are performing better because the tasks are more suited to their architecture. Perhaps yes. One could argue that a number of regular languages that Transformers struggle to learn (Bhattamishra et al., 2020a; Delétang et al., 2022) are more suited to recurrent architecture. Transformers have been shown to perform poorly on languages that require modular counting. DFAs, which are often considered to be formal abstractions of recurrent models, can represent these more efficiently. For instance, languages like standard parity can be represented with a two-state DFA while representing sparse parities would require a larger number of states. In contrast, for circuits that have recently been related to Transformers (Hao et al., 2022; Merrill et al., 2022), representing sparse parities would be easier than representing standard parity. Our results indicate that previous works might have overestimated the performance of LSTMs by considering regular languages which are more suited for autoregressive architectures.

The question of which formal languages are more closely associated with practical tasks is not entirely clear. Prior works on analysis with formal languages have primarily followed Chomsky hierarchy owing to the conjecture that natural languages are mildly context-sensitive. While regular languages such as PARITY have high sensitivity ($\mathcal{S} = 1$), practical tasks are often structured and have typically much lower sensitivity (Hahn et al., 2021). In tasks such as sentiment analysis, the label often depends on a sparse subset of input tokens. When practical text datasets such as SST are labelled with random noise, then it can be shown that their sensitivity would be concentrated around $1/2$. As shown in Fig. 23, models take much longer to fit such datasets whereas, in the case of the true labels, they only need a few epochs to fit the dataset.

Our results indicate that while Transformers perform poorly on certain regular languages, they generalize more effectively than recurrent models on various sparse Boolean functions. Moreover, we showed that random Transformers as well as those trained with gradient-based algorithms are biased towards functions of low sensitivity. Our results add to the body of work that suggests that there is a form of implicit regularization in the procedure used to train neural models which prevent them from overfitting despite their incredible capacity.

Acknowledgments

We would like to thank Michael Hahn, Ard Louis, Kabir Ahuja, anonymous reviewers, and our colleagues at the University of Oxford for helpful discussions and for providing valuable feedback.

Limitations

A general limitation of this line of work is that most of the results are primarily confined to artificial datasets. Although such formal languages provide us with a controlled setting and clarity regarding the precise nature of the problem, the relation to practical tasks remains unclear. Hence, while our results highlight the contrast in the performance between the two types of architectures, its precise implications on real-world tasks remain unclear.

There are two negative results that do not support our hypothesis. (a) All the experiments discussed in the main paper are on strings of fixed lengths. We conducted some experiments on tasks with variable length sequences which in some sense have low sensitivity. The tasks can be seen as a variable length extension of sparse parities and sparse majorities. Unlike the fixed length setting, we found both LSTMs and Transformers perform similarly on those tasks. See Section E.1 in the Appendix for more details. (b) Although we found Transformers to consistently converge to low sensitivity functions in the case of Boolean functions, we did not find similar behaviour on sentiment classification datasets such as SST and IMDB (see Section C).

A caveat with empirical studies such as this is that the results depend on the hyperparameters and other aspects of the experimental setup. While we have tried to be as thorough as possible with hyperparameter tuning, there is always a chance that the results or behaviour could differ for some hyperparameter.

Ethics Statement

We have extensively discussed the limitations of our work in the previous section. We use two existing datasets, SST (Socher et al., 2013) and IMDB (Maas et al., 2011), which are publicly available and commonly used in NLP research. We synthetically generate datasets of formal languages which does not require ethical consideration. We have discussed the experimental details and computational budget in detail in Appendix G. The research presented in this paper focuses on analysing the inductive biases of Transformers and LSTMs based

on experiments on formal languages and subsequently we believe that our work does not raise any ethical concerns.

References

- Andris Ambainis, Mohammad Bavarian, Yihan Gao, Jieming Mao, Xiaoming Sun, and Song Zuo. 2014. Tighter relations between sensitivity and other complexity measures. In *International Colloquium on Automata, Languages, and Programming*, pages 101–113. Springer.
- Cem Anil, Yuhuai Wu, Anders Andreassen, Aitor Lewkowycz, Vedant Misra, Vinay Ramasesh, Ambrose Slone, Guy Gur-Ari, Ethan Dyer, and Behnam Neyshabur. 2022. Exploring length generalization in large language models. *arXiv preprint arXiv:2207.04901*.
- Devansh Arpit, Stanisław Jastrzundzki, Nicolas Ballas, David Krueger, Emmanuel Bengio, Maxinder S. Kanwal, Tegan Maharaj, Asja Fischer, Aaron Courville, Yoshua Bengio, and Simon Lacoste-Julien. 2017. A closer look at memorization in deep networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 233–242. JMLR.org.
- Boaz Barak, Benjamin L. Edelman, Surbhi Goel, Sham Kakade, Eran Malach, and Cyril Zhang. 2022. [Hidden progress in deep learning: Sgd learns parities near the computational limit](#).
- Satwik Bhattamishra, Kabir Ahuja, and Navin Goyal. 2020a. [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.
- Satwik Bhattamishra, Kabir Ahuja, and Navin Goyal. 2020b. [On the practical ability of recurrent neural networks to recognize hierarchical languages](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1481–1494, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

- David Chiang and Peter Cholak. 2022. [Overcoming a theoretical limitation of self-attention](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7654–7664, Dublin, Ireland. Association for Computational Linguistics.
- Grégoire Delétang, Anian Ruoss, Jordi Grau-Moya, Tim Genewein, Li Kevin Wenliang, Elliot Catt, Marcus Hutter, Shane Legg, and Pedro A Ortega. 2022. Neural networks and the chomsky hierarchy. *arXiv preprint arXiv:2207.02098*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). 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 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Javid Ebrahimi, Dhruv Gelda, and Wei Zhang. 2020. How can self-attention networks recognize dyck-n languages? *arXiv preprint arXiv:2010.04303*.
- Benjamin L Edelman, Surbhi Goel, Sham Kakade, and Cyril Zhang. 2022. Inductive biases and variable creation in self-attention mechanisms. In *International Conference on Machine Learning*, pages 5793–5831. PMLR.
- Leonardo Franco. 2006. Generalization ability of boolean functions implemented in feedforward neural networks. *Neurocomputing*, 70(1-3):351–361.
- Felix A Gers and E Schmidhuber. 2001. Lstm recurrent networks learn simple context-free and context-sensitive languages. *IEEE Transactions on Neural Networks*, 12(6):1333–1340.
- Xavier Glorot and Yoshua Bengio. 2010. Understanding the difficulty of training deep feedforward neural networks. In *Proceedings of the thirteenth international conference on artificial intelligence and statistics*, pages 249–256. JMLR Workshop and Conference Proceedings.
- Parikshit Gopalan, Noam Nisan, Rocco A Servedio, Kunal Talwar, and Avi Wigderson. 2016. Smooth boolean functions are easy: Efficient algorithms for low-sensitivity functions. In *Proceedings of the 2016 ACM Conference on Innovations in Theoretical Computer Science*, pages 59–70.
- Mark W Goudreau, C Lee Giles, Srimat T Chakradhar, and Dong Chen. 1994. First-order versus second-order single-layer recurrent neural networks. *IEEE Transactions on Neural Networks*, 5(3):511–513.
- Michael Hahn. 2020. [Theoretical limitations of self-attention in neural sequence models](#). *Transactions of the Association for Computational Linguistics*, 8:156–171.
- Michael Hahn, Dan Jurafsky, and Richard Futrell. 2021. [Sensitivity as a complexity measure for sequence classification tasks](#). *Transactions of the Association for Computational Linguistics*, 9:891–908.
- Yiding Hao, Dana Angluin, and Robert Frank. 2022. [Formal language recognition by hard attention transformers: Perspectives from circuit complexity](#). *Transactions of the Association for Computational Linguistics*, 10:800–810.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long short-term memory. *Neural computation*, 9(8):1735–1780.
- Jeff Kahn, Gil Kalai, and Nathan Linial. 1989. *The influence of variables on Boolean functions*. Citeseer.
- Michael Kearns. 1998. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- John F Kolen and Stefan C Kremer. 2001. *A field guide to dynamical recurrent networks*. John Wiley & Sons.
- Samuel A Korsky and Robert C Berwick. 2019. On the computational power of rnns. *arXiv preprint arXiv:1906.06349*.
- Jaehoon Lee, Yasaman Bahri, Roman Novak, Samuel S Schoenholz, Jeffrey Pennington, and Jascha Sohl-Dickstein. 2017. Deep neural networks as gaussian processes. *arXiv preprint arXiv:1711.00165*.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. Roberta: A robustly optimized bert pretraining approach. *arXiv preprint arXiv:1907.11692*.
- Andrew Maas, Raymond E Daly, Peter T Pham, Dan Huang, Andrew Y Ng, and Christopher Potts. 2011. Learning word vectors for sentiment analysis. In *Proceedings of the 49th annual meeting of the association for computational linguistics: Human language technologies*, pages 142–150.
- 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, Gail Weiss, Yoav Goldberg, Roy Schwartz, Noah A Smith, and Eran Yahav. 2020. A formal hierarchy of rnn architectures. *arXiv preprint arXiv:2004.08500*.
- Aaron Meurer, Christopher P. Smith, Mateusz Paprocki, Ondřej Čertík, Sergey B. Kirpichev, Matthew Rocklin, AMiT Kumar, Sergiu Ivanov, Jason K. Moore, Sartaj Singh, Thilina Rathnayake, Sean Vig,

- Brian E. Granger, Richard P. Muller, Francesco Bonazzi, Harsh Gupta, Shivam Vats, Fredrik Johansson, Fabian Pedregosa, Matthew J. Curry, Andy R. Terrel, Štěpán Roučka, Ashutosh Saboo, Isuru Fernando, Sumith Kulal, Robert Cimrman, and Anthony Scopatz. 2017. *Sympy: symbolic computing in python*. *PeerJ Computer Science*, 3:e103.
- Chris Mingard, Joar Skalse, Guillermo Valle-Pérez, David Martínez-Rubio, Vladimir Mikulik, and Ard A Louis. 2019. Neural networks are a priori biased towards boolean functions with low entropy. *arXiv preprint arXiv:1909.11522*.
- Chris Mingard, Guillermo Valle-Pérez, Joar Skalse, and Ard A Louis. 2021. Is sgd a bayesian sampler? well, almost. *Journal of Machine Learning Research*, 22.
- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. 2018. *Foundations of machine learning*. MIT press.
- Preetum Nakkiran, Gal Kaplun, Dimitris Kalimeris, Tristan Yang, Benjamin L. Edelman, Fred Zhang, and Boaz Barak. 2019. *SGD on Neural Networks Learns Functions of Increasing Complexity*. Curran Associates Inc., Red Hook, NY, USA.
- Roman Novak, Yasaman Bahri, Daniel A. Abolafia, Jeffrey Pennington, and Jascha Sohl-Dickstein. 2018. *Sensitivity and generalization in neural networks: an empirical study*. In *International Conference on Learning Representations*.
- Ryan O’Donnell. 2021. *Analysis of boolean functions*.
- Samet Oymak and Mahdi Soltanolkotabi. 2019. Over-parameterized nonlinear learning: Gradient descent takes the shortest path? In *International Conference on Machine Learning*, pages 4951–4960. PMLR.
- Giacomo De Palma, Bobak Toussi Kiani, and Seth Lloyd. 2019. *Random deep neural networks are biased towards simple functions*. In *Advances in Neural Information Processing Systems 32: Annual Conference on Neural Information Processing Systems 2019, NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 1962–1974.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. 2019. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32.
- Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. 2022. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*.
- Nasim Rahaman, Aristide Baratin, Devansh Arpit, Felix Draxler, Min Lin, Fred Hamprecht, Yoshua Bengio, and Aaron Courville. 2019. *On the spectral bias of neural networks*. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5301–5310. PMLR.
- Alexander Rush. 2018. *The annotated transformer*. In *Proceedings of Workshop for NLP Open Source Software (NLP-OSS)*, pages 52–60, Melbourne, Australia. Association for Computational Linguistics.
- Avi Schwarzschild, Eitan Borgnia, Arjun Gupta, Furong Huang, Uzi Vishkin, Micah Goldblum, and Tom Goldstein. 2021. Can you learn an algorithm? generalizing from easy to hard problems with recurrent networks. *Advances in Neural Information Processing Systems*, 34:6695–6706.
- Luzi Sennhauser and Robert Berwick. 2018. *Evaluating the ability of LSTMs to learn context-free grammars*. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 115–124, Brussels, Belgium. Association for Computational Linguistics.
- Natalia Skachkova, Thomas Trost, and Dietrich Klakow. 2018. *Closing brackets with recurrent neural networks*. In *Proceedings of the 2018 EMNLP Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 232–239, Brussels, Belgium. Association for Computational Linguistics.
- Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng, and Christopher Potts. 2013. Recursive deep models for semantic compositionality over a sentiment treebank. In *Proceedings of the 2013 conference on empirical methods in natural language processing*, pages 1631–1642.
- 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.
- Michael Tänzler, Sebastian Ruder, and Marek Rei. 2022. *Memorisation versus generalisation in pre-trained language models*. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7564–7578, Dublin, Ireland. Association for Computational Linguistics.
- Guillermo Valle-Perez, Chico Q. Camargo, and Ard A. Louis. 2019. *Deep learning generalizes because the parameter-function map is biased towards simple functions*. In *International Conference on Learning Representations*.

- 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*, pages 5998–6008.
- Qinglong Wang, Kaixuan Zhang, Alexander G. Ororbia II, Xinyu Xing, Xue Liu, and C. Lee Giles. 2018. [A comparative study of rule extraction for recurrent neural networks](#). *CoRR*, abs/1801.05420v2.
- Shunjie Wang. 2021. *Evaluating Transformer’s Ability to Learn Mildly Context-Sensitive Languages*. University of Washington.
- 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.
- Andrew G Wilson and Pavel Izmailov. 2020. Bayesian deep learning and a probabilistic perspective of generalization. *Advances in neural information processing systems*, 33:4697–4708.
- Shunyu Yao, Binghui Peng, Christos Papadimitriou, and Karthik Narasimhan. 2021. [Self-attention networks can process bounded hierarchical 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 3770–3785, Online. Association for Computational Linguistics.
- Chiyuan Zhang, Samy Bengio, Moritz Hardt, Benjamin Recht, and Oriol Vinyals. 2021. Understanding deep learning (still) requires rethinking generalization. *Communications of the ACM*, 64(3):107–115.

A Roadmap

The appendix is organized as follows.

- In Section B, we report and discuss additional results on the complexity of random models.
- In Section C, we investigate the sensitivity of models on real data. In particular, we demonstrate that models learn functions of increasing sensitivity on sentiment classification datasets such as SST and IMDB.
- In Section D, we discuss some additional results relating sensitivity and generalization.
- In Section E, we present additional experiments investigating the ability of Transformers and LSTMs to learn sparse boolean functions.
- In Section F, we present some experiments to show that both Transformers and LSTMs can easily fit practical datasets even when they are labelled randomly.
- In Section G, details of implementation and experimental setup are discussed which are relevant for the reproducibility of the results.
- In Section H, we discuss some additional works related to our paper.

B Complexity of Random Models

In this section, we discuss additional results related to the complexity of random Transformers and LSTMs. We present results with additional complexity measures, initialization strategies, and variations across hyperparameters.

B.1 Additional Measures

As discussed in Section 3, sensitivity is related to several other complexity measures. Since it is more tractable to estimate sensitivity as opposed to certain other measures, we primarily focused on estimating and comparing sensitivity in the main paper. We explore three other complexity measures which have been previously explored in the literature to compute the complexity of functions represented by neural models. The measures are defined as follows:

1. **SOP (Size of Boolean Expression)**: This measure computes the size of the smallest Boolean

expression in Sum-of-Product form that represents the function. In order to compute this for a neural network over $\{0, 1\}^n$, we compute the output of the model over all 2^n inputs and then use standard libraries (SymPy (Meurer et al., 2017)) to find the Boolean expression. The size indicates the number of operators and operands in the smallest expression. Since the problem of minimizing Boolean expressions is NP-complete, the runtime grows exponentially, and hence, we can only compute this up to length 10 for several samples of random models. This measure was explored in Valle-Perez et al. (2019).

2. **Entropy**: This measure takes the output labels for all 2^n inputs and simply computes the entropy over the labels. This is a weak measure and primarily indicates how imbalanced the label set is. This measure was explored in Mingard et al. (2019); Valle-Perez et al. (2019).
3. **CSR (Critical Sample Ratio)**: This measure computes the fraction of inputs for which the function label changes at a small fixed distance from the inputs (Arpit et al., 2017). For discrete inputs such as $\{0, 1\}^n$, CSR can be seen as the fraction of inputs for which the function label changes at a Hamming distance of 1. This was also explored in Valle-Perez et al. (2019).

Figure 9 shows the distribution of different complexity measures and scatter plots depicting relations among them. The measures are computed for random Transformers and LSTMs with weights sampled uniformly between -10 and 10. The measures are computed for sequences of length 7 with $200k$ samples of models. We take Transformers and LSTMs with depth $\in \{1, 2, 4, 8\}$ and width ($d_{\text{model}}/hidden_size$) $\in \{8, 32, 64, 256, 768\}$. We take an equal number of samples for each hyperparameter. Figure 8 shows the distribution of SOP based on 50k samples for a fixed hyperparameter configuration of Transformer and LSTM. It includes a 1-layer LSTM with width 64 and a 4-layer Transformer with width 64.

As can be seen in Figure 9, there exists significant correlation between sensitivity and other measures. Note that, high sensitivity functions will always have high entropy and high CSR but the




Figure 8: Distribution of SOP measure based on 50k samples of 4-layer Transformers and 1-layer LSTMs with width 64. Refer to Section B.1 for details.

converse is not true. Functions with maximum entropy can also have low sensitivity. For instance, the dictator function has maximum entropy (since half the inputs have label 1 and the other half have label 0) while having very low sensitivity. Similarly, CSR can be seen as a weaker version of sensitivity.

B.2 Why Sensitivity?

Sensitivity can be seen as a discrete analog (Gopalan et al., 2016) of the ‘smoothness’ of a continuous function which measures how gradually a function changes locally. Functions of higher sensitivity can be considered more complex since the function value can be changed by changing any of a large subset of bits whereas functions of lower sensitivity depend on fewer bits and their function value can be determined based on a small number of input coordinates. Sensitivity measures are also polynomially related to several other notions of complexity such as the depth of a decision tree, certificate complexity, and the degree of the Fourier expansion of Boolean functions (see Ambainis et al. (2014) for more details). The correlation between generalization and a different notion of sensitivity has been demonstrated in Novak et al. (2018) for computer vision models. The relation between generalization and a variant of Boolean sensitivity has even been explored over a decade ago by Franco (2006). More recently, Hahn et al. (2021) extend the notion of block sensitivity to incorporate variable length sequences and propose it as a measure to estimate the difficulty of various NLP tasks.

B.3 Additional Sensitivity Results

The distribution of the sensitivity of Transformers and LSTMs initialized with Xavier uniform distribution are given in Figure 11 respectively. For

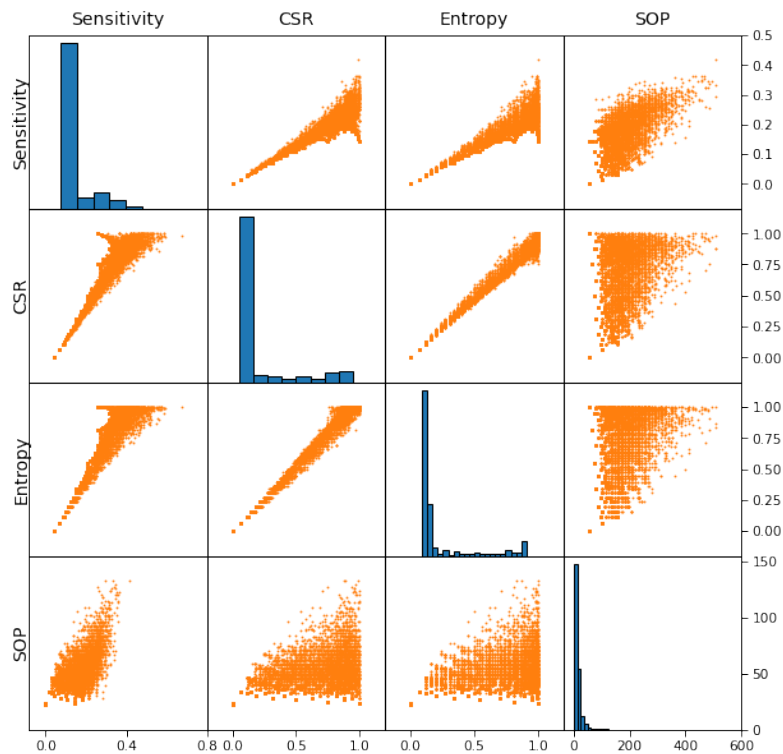
Gaussian initialization, the weights are sampled with mean 0 and $\sigma = 10$. For Xavier uniform initialization all the values in weight matrices are sampled uniformly between $-d^{1/2}$ and $d^{1/2}$ where d is the number of hidden units. The values in the bias vectors are set to zero and the ones in the input embedding vectors are sampled from $\mathcal{N}(0, 1)$.

Finding Parity. For strings of length 5, the total number of possible functions is 2^{2^5} . If Boolean functions are sampled uniformly, then the probability of picking the PARITY function is less than 1 in two billion. However, on uniformly sampling 10 million LSTMs of depth 2 and hidden size 8, we found that the probability of finding one that represents PARITY is 1 in 30,000. Hence, it is over **60,000** times more likely to find PARITY function by sampling LSTMs than randomly sampling Boolean functions. This indicates that the parameter space of recurrent models such as LSTMs has a significant representation of PARITY functions which might help explain why it is easier for them to learn PARITY. On the other hand, for Transformers, we did not find a single sample which represented PARITY based on 10 million samples.

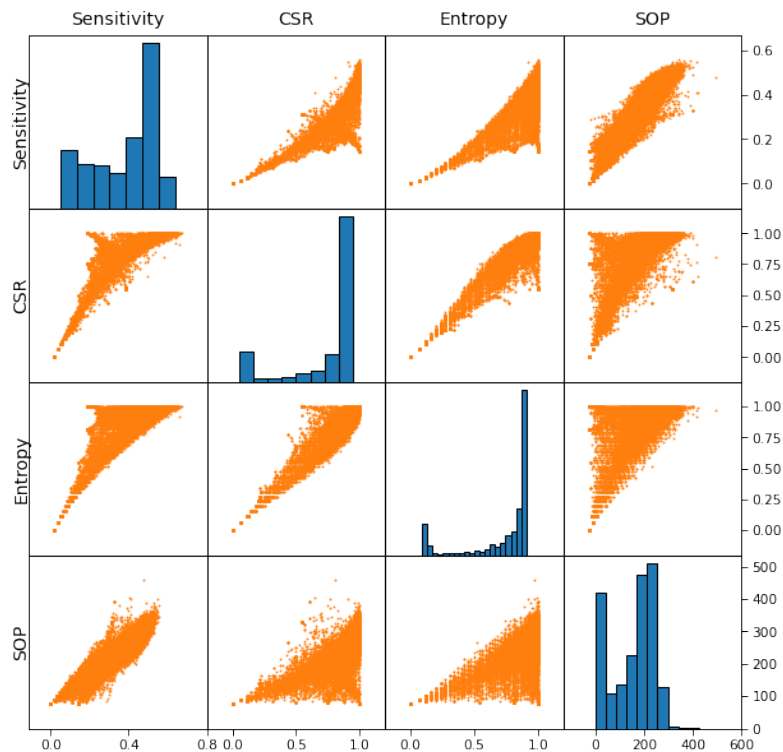
Change across hyperparameters. For uniform sampling, a general observation for both the architectures is that the likelihood of higher sensitivity functions increases with the number of layers (see Figure 15, left and Figure 10), however, even for Transformers with depth 12, the distribution is heavily skewed towards low sensitivity functions in comparison to recurrent models with depth 1. Unlike recurrent models, the sensitivity of Transformers decreases when the width of the model is increased (see Figure 15, middle). For Transformers, the average sensitivity decreases with the increase in the length of the strings (see Figure 15, right), whereas for LSTMs, it remains quite high even for lengths up to 200.

For LSTMs with uniform sampling, the change in sensitivity across different widths (hidden_size) and lengths is provided in Figure 10. As can be seen, the sensitivity of LSTMs does not significantly reduce across higher lengths and widths, unlike Transformers.

While it is not entirely clear why random Transformers are relatively more biased towards low complexity functions, we observe that they behave similar to hard-attention Transformers upon inspection of attention weights. Recent works (Hao et al., 2022; Hahn, 2020) have shown that hard-attention



(a) Transformers



(b) LSTMs

Figure 9: The distributions (diagonals) and comparisons (off-diagonals) of 4 different complexity measures for random (a) Transformers and (b) LSTMs with weights uniformly sampled in $[-10, 10]$. The measures are computed over all inputs in $\{0, 1\}^7$. CSR denotes critical sample ratio and SOP denotes the size of the smallest expression in sum-of-product form that represents the truth table. Refer to Section B.1 for details.




Figure 10: Mean sensitivity of Random LSTMs across different widths and lengths. Refer to Section B.3 for details.

Transformers can only represent functions in AC^0 (which contain functions that can be represented by constant depth AND/OR circuits). Since AC^0 circuits can only represent functions of low average sensitivity (O’Donnell, 2021), it might help explain why random Transformers have low sensitivity.

C Sensitivity During Learning Sentiment Classification

In this section, we discuss experiments on measuring the sensitivity of Transformers and LSTMs when trained on the sentiment classification task.

C.1 Experimental Setup

Datasets. We experiment with two sentiment classification datasets: SST (Socher et al., 2013) and IMDB (Maas et al., 2011). For SST, we train on the full train set of size 67349 examples and evaluate both sensitivity and validation accuracy on the validation set of size 872 examples. For IMDB, we preprocess the dataset to only include sentences of length up to 500. This leads to a train set of size 22156. The validation set consists of 8939 examples randomly sampled from the test set. Since the sentences in IMDB dataset are of much longer lengths, in order to save compute, we evaluate sensitivity of models on a dataset of size 894 examples randomly sampled from the test set.

Sensitivity Metrics. Boolean sensitivity as defined in Section 3 cannot be directly applied to sequences of variable length and larger vocabulary. As an alternative, we compute certain proxy metrics which measure how likely it is for the function value to change due to a change in one token of the input sequence. To that end, we design three simple metrics to measure the sensitivity of models trained on sentiment classification:

1. *Word Label-Sensitivity:* For each word in the sentence (one word at a time), we replace it n times with a word sampled randomly from the

vocabulary and measure the average (over n) number of times the predicted label changes. We sum this value for all the words in the sentence and normalize the value by its length.

2. *Word Softmax-Sensitivity:* For each word in the sentence (one word at a time), we replace it n times with a word sampled randomly from the vocabulary and measure the average (over n) L_2 -distance between the predicted softmax normalized output vector before and after the replacement. Again, we sum this value for all the words in the sentence and normalize by its length.
3. *Embedding Label-Sensitivity:* For each word in the sentence (one word at a time), we add Gaussian noise with mean 0 and variance σ^2 to its embedding n different times and measure the average (over n) number of times the predicted label changes. We sum this value for all the words in the sentence and normalize by its length.

For all metrics, the final score is obtained by averaging across all the examples in the dataset. In all our experiments, we set $n = 10$, and $\sigma^2 = 15$.

Hyperparameter Details. For both Transformers and LSTMs, we vary the number of layers $\in \{1, 2\}$, learning rate $\in \{0.0001, 0.0003, 0.0005\}$, and model width (d_model/hidden_size) $\{128, 256\}$. We set the batch size as 128 and the FFN size as twice the width. For LSTMs, we keep the embedding size the same as the hidden size. Both models are trained with Adam (Kingma and Ba, 2014) optimization and using Dropout regularization with probability 0.2.

Results. Figure 12 shows the word softmax-sensitivity for both models across different iterations of training for SST and IMDB datasets. The word label-sensitivity and embedding label-sensitivity for SST is provided in Figure 13.

We find that across all three measures, both Transformers and LSTMs learn functions of increasing sensitivity where they prioritize learning functions of lower sensitivity first. We found ‘word label-sensitivity’ and ‘word softmax-sensitivity’ to correlate well with *generalization gap* (i.e., the difference between train accuracy and test accuracy). Since the measures are very similar, there is a strong correlation between the two measures themselves. We did not find any non-trivial correlation




Figure 11: Distribution of sensitivity of different Transformers and LSTMs with Xavier uniform initialization. Sensitivity of Transformers (left) and LSTMs (middle) across various hyperparameters. Right: Comparison for a fixed hyperparameter. Refer to Section B.3 for details.




Figure 12: The ‘word softmax-sensitivity’ of Transformers and LSTMs at different stages of training on SST and IMDB datasets.

between ‘embedding label-sensitivity’ and generalization gap. Note that unlike random and sparse Boolean functions, on real datasets, we did not find Transformers converging to functions with lower sensitivity.

D Sensitivity and Generalization

D.1 Sensitivity as Capacity Measure

We show how maximum sensitivity can be used as a capacity measure to derive generalization bounds. Capacity measures such as the VC Dimension are a classical approach to derive sample complexities and probabilistic upper bounds for the test error of a classification model.

Let $\mathcal{F}_k : \{0, 1\}^n \rightarrow \{\pm 1\}$ be a class of functions such that the maximum sensitivity for any function $f \in \mathcal{F}_k$ is upper bounded by k where $0 \leq k \leq n$. Any function f with a maximum sensitivity k can be uniquely determined by its values on any Hamming ball of radius $2k$ in $\{0, 1\}^n$ (Gopalan et al., 2016). This can be used to upper bound the size of the function class $|\mathcal{F}_k| \leq 2^{\binom{n}{\leq 2k}}$. Since the VC Dimension (denoted VCD) of a class of functions \mathcal{F} is upper bounded by $\log |\mathcal{F}|$, we have that,

$$\text{VCD}(\mathcal{F}_k) \leq \binom{n}{\leq 2k} \leq \binom{n+2k}{2k} \quad (4)$$

$$\leq \left(\frac{e(n+2k)}{2k} \right)^{2k} = \mathcal{O}(n^{2k})$$

Let $f \in \mathcal{F}_k$ be a target function and $\hat{f} \in \mathcal{F}_k$ be a hypothesis produced by a learning algorithm. Let $L(\hat{f}, f) = \mathbb{E}_{x \sim \{0,1\}^n} [\mathbb{I}[\hat{f}(x) \neq f(x)]]$ be the true error between f and \hat{f} . Similarly, let $\hat{L}_S(\hat{f}, f) = \frac{1}{m} \sum_{i=1}^m \mathbb{I}[\hat{f}(x) \neq f(x)]$ be the empirical error on a sample set S . Then using Equation (4) and basic properties of VC dimension (Mohri et al. (2018)), we can upper bound the distance of the true error L from the sample error \hat{L} using maximum sensitivity.

Proposition D.1. *For any $\delta > 0$, with probability at least $1 - \delta$, the following holds for any function $f, \hat{f} \in \mathcal{F}_k$,*

$$L(\hat{f}, f) \leq \hat{L}_S(\hat{f}, f) + \sqrt{\frac{cn^{2k} \log \frac{2em}{cn^{2k}} + 8 \log \frac{4}{\delta}}{m}} \quad (5)$$

where $c > 0$ is some constant. Functions with low maximum sensitivity can be learned with better sample efficiency. Functions with low average sensitivity can also be learned efficiently when the data generating distribution is uniformly distributed over the input (O’Donnell (2021), Sec 3.4).

D.2 Sensitivity and Generalization Gap

The correlation between sensitivity and generalization has previously been studied for networks trained on Boolean functions (Franco, 2006) and image datasets (Novak et al., 2018). We examine the relation between simple variants of sensitivity described in Section C and generalization.

We train various models on SST dataset until convergence and then compare sensitivity with generalization gap. The generalization gap is simply the difference between the train error and test error; higher gap indicates overfitting. We plot the




Figure 13: The ‘word label-sensitivity’ and ‘embedding label-sensitivity’ of Transformers and LSTMs at different stages of training on SST dataset.




Figure 14: Sensitivity vs Generalization gap for Transformers, LSTMs, and RoBERTa LLM trained on SST.

word label-sensitivity and word softmax-sensitivity (defined in Section C) for Transformers, LSTMs, and a pretrained Large Language Model (RoBERTa (Liu et al., 2019)) against the generalization gap (see Figure 14). We observe positive correlation between the measures and generalization gap indicating that when sensitivity is higher, the models are more likely to overfit and achieve poorer generalization performance. Large language models such as RoBERTa have a lower sensitivity while achieving better test accuracies than Transformers and LSTMs trained from scratch.

E Additional Experiments on Sparse Boolean Functions

Standard Parity. The training curves for LSTMs on standard parity are provided in Figure 17. The models are trained on datasets of size 20k where the input strings are of length 30. Similar to Transformers on SPARSE PARITIES, we observe phase transitions for LSTMs on standard PARITY task.

Sparse Parities. The results on sparse parities with length $n=40$ and $k=4$ relevant bits for Transformers with absolute positional encodings are provided in Figure 18. We find that Transformers with absolute positional encodings are able to generalize well on SPARSE PARITIES task and exhibit grokking on relatively larger datasets (30k samples) in comparison to models with learnable positional embeddings. For Transformers trained with learnable encoding, we robustly observe grokking on small

datasets. Figure 19 depicts the training curves for Transformers trained on datasets of size 5k.

Overfitting. We found that Transformers overfit on training data when the sample size is too low (see Figure 20). Apart from that, for datasets of certain sizes, we find that while Transformers with depth up to 6 generalize well, those with much higher depths (> 8) overfit across several runs.

Effect of Regularization on LSTMs. We explore the effect of dropout and L2 regularization on training with LSTMs. While training on sparse parities, we find that increasing regularization increases the convergence time but the model still overfits and converges to a function with higher sensitivity than the target function. Upon further increasing regularization, the model fails to fit the training data.

Mixed Parity. To explore the difference in bias between Transformers and LSTMs, we conduct a simple experiment described as follows. We create a dataset of size 15k called ‘Mixed Parity’ where half of the examples are labelled as standard PARITY (label is determined by all bits) and the other half is labelled as SPARSE PARITIES with 4 relevant bits. The inputs are of length 30 and the first bit determines whether the input is labelled according to standard PARITY function (when the first bit is 1) or as a SPARSE PARITIES function (when the first bit is 0).

We train Transformers and LSTMs (with learnable positional encodings) of depth 2 and width 64 across various learning rates $\in [0.01, 0.00001]$ on the Mixed Parity dataset. We find that LSTMs obtain 100% training accuracy on the dataset (see Figure 21, right); LSTMs validation accuracy on the PARITY task is near 100% whereas it is 50% on the SPARSE PARITIES task. In contrast, the training accuracy of Transformers converges around 75% (see Figure 21, left); their validation accuracy on the PARITY task is 50% whereas on SPARSE PARITIES they achieve near 100% validation accuracy.

Convergence time vs Sample Size. For Transformers trained on SPARSE PARITIES, we conduct experiments to compare the number of computational steps required to successfully learn SPARSE PARITIES with the size of the dataset it is trained on. We consider length $n=40$ and $k=4$, and create datasets of five different sizes ($\{5k, 25k, 50k, 100k \text{ and } 250k\}$). For each dataset, we train a Transformer of depth 2 and width 128




Figure 15: Sensitivity of random Transformers across different depths, widths, and lengths. Refer to Section 4.1 for details.




Figure 16: The sensitivity of Transformers and LSTMs at different stages of training on various k -SPARSE functions of length 50. Refer to Section 5 for details.




Figure 17: Training and validation curves for LSTMs trained on standard parity (PARITY-30) dataset containing 30k samples. Similar to the behaviour of Transformers on SPARSE PARITIES, we observe abrupt phase transitions for LSTMs where there is a sharp increase in the train and validation accuracy.

across 100 different initializations with learning rate $\in \{0.0001, 0.0005\}$ and with batch size 500. We consider each iteration as a computational step. We report the median, minimum and maximum steps for each dataset in Figure 22. We find that neural networks such as Transformers can successfully learn SPARSE PARITIES with relatively small number of computational steps on small-sized training sets. It is perhaps surprising that for SPARSE PARITIES with $n=40$ and $k=4$, Transformers can successfully generalize well with less than 20000 computational steps on over 75 out of 100 runs.




Figure 18: Training and validation curves for Transformers with absolute positional encodings trained on sparse parity task (length $n=40$ and $k=4$). The training accuracy gradually starts to increase with no change in validation accuracy. After a point, both training and validation accuracies match.

Phase Transitions. As reported in Barak et al. (2022), we observe phase transitions on Parity tasks where the training and validation accuracies do not change for a large number of training iterations and then abruptly reach near-perfect accuracy in a few iterations (see Figure 6). This phenomenon was observed for feedforward networks (FFNs) and Transformers in Barak et al. (2022) and theoretically explained for ReLU FFNs trained with SGD. We observe another such behaviour for LSTMs on PARITY- n (see Figure 17 for training curves on PARITY-30). For both LSTMs and Transformers, we were unable to get them to generalize well with SGD on either SPARSE PARITIES or standard PARITY. Both the architectures seem to succeed with the Adam optimizer (Kingma and Ba, 2014).

Grokking. Another interesting phenomenon we observe in Transformers is that in some cases the training accuracies start increasing gradually with no change in the validation accuracy. After some iterations, the validation accuracy increases and matches the training accuracy. We reliably observed this phenomenon while training Transformers with absolute positional encodings across training sets of various sizes (see Figure 18) and while training with learnable encodings on small-sized

training sets (see Figure 19). Similar observations for grokking (Power et al., 2022) were made in Barak et al. (2022) for ReLU FFNs trained on small-sized training sets.

E.1 Experiments on Variable Length Inputs

We conducted some additional experiments on tasks with variable length inputs. These tasks are simple extensions of sparse parities and majorities to variable length input and have (in an informal sense) low sensitivity.

Task. Let $\text{VARPARITY-}(n, k)$ denote the extension of $\text{PARITY-}(n, k)$ to variable length sequences. A function in $\text{VARPARITY-}(n, k)$ is defined over sequence of $\{0, 1, 2\}$ where the total number of 0s and 1s are exactly n , along with k relevant indices which determine the label. The input distribution is such that there could be token 2 between any zeros and ones with some probability. The tokens 2 however do not influence the output of the function and are merely constructed to vary the input lengths. The label is determined by removing all the tokens 2 from the input and applying the regular $\text{PARITY-}(n, k)$ over the remaining string over $\{0, 1\}^n$.

For illustration, for $f_S \in \text{PARITY-}(4, 2)$, where $S = \{1, 3\}$, for an input ‘1001’, the function $f_S(\underline{1}00\underline{1}) = 1$ since the number of 1s in position 1 and 3 is odd. For a similar function $\hat{f}_S \in \text{VARPARITY-}(4, 2)$, here are some examples on various inputs: $\hat{f}_S(\underline{2}1\underline{0}22\underline{0}222\underline{1}) = \hat{f}_S(\underline{1}2\underline{0}0\underline{2}22\underline{1}2) = \hat{f}_S(\underline{1}22\underline{0}2\underline{0}21\underline{2}22) = f_S(\underline{1}00\underline{1}) = 1$. The function $\text{VARMAJ-}(n, k)$ is defined similarly, where it takes an input string over $\{0, 1, 2\}$ and the label is determined by removing all 2s and applying regular $\text{MAJ-}(n, k)$ on the remaining string over $\{0, 1\}^n$.

Results. Contrary to the fixed length setting, we observe that both Transformers and LSTMs perform similarly on these tasks. For both $\text{VARPARITY-}(n, k)$ and $\text{VARMAJ-}(n, k)$ we experiment with various mean lengths and variances with $k = 5$. The general behaviour is that both Transformers and LSTMs generalize well when the tasks over short sequences (< 40 for $\text{VARPARITY-}(n, k)$ and < 100 for $\text{VARMAJ-}(n, k)$). However as the lengths of the input go beyond that, both architectures do not generalize well.

In comparison to LSTMs, Transformers only performed better when the variance of the lengths of the inputs was very low. An interesting observation about Transformers is that they only seemed to gen-




Figure 19: Training and validation curves for Transformers trained on sparse parity dataset (length $n=40$ and $k=4$) with only 5k training examples. The training accuracy gradually starts to increase with no change in validation accuracy. After a point, both training and validation accuracies match.




Figure 20: Training and validation accuracy curves for Transformers trained on sparse parity dataset (length $n=40$ and $k=4$) with 1k training examples. The behaviour is similar to how LSTMs overfit for much larger training sets (40k samples). We observe similar behaviour for Transformer with large depths on datasets of size 5k.

eralize well with positional masking (also referred to as causal masking) along with positional encodings. Their performance was notably worse with only positional encodings (learnable or absolute).

These results do not support the hypothesis posed in Section 1 and we intend to explore this further in the future.

F Fitting Randomly Labelled Data

We conduct some experiments to examine the ability of LSTMs and Transformers to fit random noise. The capacity of a class of functions to fit random noise is often theoretically measured as its Rademacher complexity. Given the incredible expressive power of neural networks, measures such as Rademacher complexity lead to vacuous generalization bounds. One assumption was that, despite their capacity, deep neural networks trained with gradient-based methods can only learn a small subset of such functions. The work of Zhang et al. (2021) demonstrated that large feedforward-like networks trained with gradient-based methods are able to fit random noise on image datasets. We conduct similar experiments to evaluate the ability of sequence models to fit noise on text data. We con-




Figure 21: Training curves for Transformers and LSTMs trained on the Mixed Parity dataset (length $n=30$) with 15k training examples. Refer to Section E for details.




Figure 22: The number of computational steps till convergence vs size of training data for Transformers trained on PARITY- $(40, 4)$. The curve depicts the median number (along with min/max) of iterations with a batch size of 500 across 50-100 successful runs for each dataset size. Refer to Section E for details.


sider the SST dataset (Socher et al., 2013) as used in the GLUE benchmark. The training data contains approximately 65k samples and we label each sample either $+1$ or -1 randomly (with probability $1/2$ each).

Figure 23 depicts the training curves for Transformers and LSTMs. We find that both the models are able to conveniently fit the training set near-perfectly. For both the models, the training takes significantly more number of iterations/epochs in comparison to training on the original dataset with true labels which only takes a few epochs.

G Implementation Details

Our implementation of Transformer is based on Rush (2018). For various recurrent models such as RNNs, GRUs, and LSTMs, we use PyTorch’s standard implementation (Paszke et al., 2019).

For each dataset, we extensively tune across several hyperparameters and report the results based on the best-performing models. Table 1 lists the hyperparameters used for tuning the models for Boolean function experiments in Section 5. We use



(a) Transformers

(b) LSTMs

Figure 23: Training curves of Transformers and LSTMs when trained on SST dataset with randomly assigned labels. Both architectures conveniently fit the entire dataset although they take much longer in comparison to when trained with original labels.

a grid search procedure to tune the hyperparameters. The models were trained with cross-entropy loss. For all our results, we used Adam Optimizer and tuned the learning rates. We also tried SGD with weight decay but could not get either Transformers or LSTMs to perform well on parities, sparse parities, or random k -sparse functions.

Compute. All our experiments were conducted using 16 NVIDIA Tesla V100 GPUs each with 16GB memory. Since the datasets are synthetic and relatively smaller than practical datasets, most training runs took 10-30mins on a single GPU. The larger expenditure was on tuning LSTMs to find whether any of the hyperparameters succeed. Some experiments with LSTMs on PARITY- $(40, 4)$ for over 100k steps took 3 hours and across multiple hyperparameters took ≈ 400 GPU hours in total. The experiments conducted for Figure 22 took similar amount of GPU hours (≈ 300). The rest of the experiments took less than 10% of this time. The experiments conducted in Section 4 with random models were not as compute intensive.

H Additional Related Work

Formal Languages and Recurrent Models. For recurrent models, analysis on formal languages dates back to a few decades ago (see Kolen and Kremer (2001)). Several works have examined the ability of RNNs and LSTMs to recognize various context-free and counter languages (Gers and Schmidhuber, 2001; Weiss et al., 2018; Suzgun et al., 2019b), the most prominent one being the Dyck⁸ languages (Skachkova et al., 2018; Bhattamishra et al., 2020b). Connections between RNNs and finite state automata have been explored for a long time (Goudreau et al., 1994; Korsky and

⁸Checking whether a sequence of brackets is well-balanced.

Hyperparameter	Bounds Transformer LSTM
D_model/Hidden Size	[16, 128] [8, 256]
Heads	[4, 8]
Number of Layers	[1, 6] [1, 6]
Learning Rate	[1e-2, 1e-5]
Position Encoding Scheme	[Learnable, Absolute]

Table 1: Different hyperparameters and the values considered for each of them. Note that certain parameters like Heads and Position Encoding Scheme are only relevant for Transformer-based models and not for LSTMs.

Berwick, 2019). Prior works have also sought to extract finite state automata from recurrent models trained on regular languages (see Wang et al. (2018) for a survey). Connections between LSTMs and counter automata have also been established empirically (Suzgun et al., 2019a) and theoretically (Merrill et al., 2020). More recently, multiple works have investigated the ability of Transformers to recognize various regular, context-free (Ebrahimi et al., 2020; Yao et al., 2021; Bhattamishra et al., 2020b), and mildly context-sensitive languages (Wang, 2021).

Neural Networks and Parities. Prior works can be divided into two categories. One set of works focuses on the Parity language containing strings of arbitrary length that can be represented by a 2-state DFA. The other set examines the PARITY- (n, k) problem with strings over $\{0, 1\}^n$ where the output depends on a subset of bits. The PARITY- (n, k) problem has been widely studied in learning theory and has several well-understood properties. On the PARITY- (n, k) , Barak et al. (2022) theoretically analyze the ability of feedforward-like networks trained with SGD and conduct experiments with various architectures including Transformers. Some of our results corroborate their findings and we empirically explore the phenomenon further. Preliminary experiments along this direction were also explored in Edelman et al. (2022) where they showed that Transformers can efficiently express k -SPARSE functions.

Since variable length Parity can be represented by a 2-state DFA, small-sized RNNs can efficiently represent them and several works have found LSTMs to generalize well when tested empirically (Schwarzschild et al., 2021). On the other hand, Transformers are limited in their ability to express such a language (Hahn, 2020). While they can express them for bounded lengths (Chiang and Cholak, 2022), they have been found to strug-

gle when tested empirically (Bhattamishra et al., 2020a; Delétang et al., 2022; Chiang and Cholak, 2022). Anil et al. (2022) explore length generalization abilities of large language models on Parity. Finding PARITY and other languages by uniformly initializing weights was also explored in the 1990s (see Chap. 13 Kolen and Kremer (2001)) for older versions of recurrent architectures which are not used anymore.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations
- A2. Did you discuss any potential risks of your work?
Not applicable. Left blank.
- A3. Do the abstract and introduction summarize the paper’s main claims?
1
- A4. Have you used AI writing assistants when working on this paper?
Left blank.

B Did you use or create scientific artifacts?

4,5

- B1. Did you cite the creators of artifacts you used?
1,2,4
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Our use of existing artifacts pertains to usage of the Transformer and LSTM models and the SST and IMDB datasets which are all open-sourced and usage is allowed without restriction.
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Not applicable. Left blank.
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Not applicable. Left blank.
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Not applicable. Left blank.
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
see section 4,5

C Did you run computational experiments?

4,5

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
H

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

4,5,H

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

See result figures.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

H

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Left blank.

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

No response.

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

No response.

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

No response.

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

No response.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

No response.