
Inductive Biases and Variable Creation in Self-Attention Mechanisms

Benjamin L. Edelman¹ Surbhi Goel² Sham Kakade¹ Cyril Zhang²

Abstract

Self-attention, an architectural motif designed to model long-range interactions in sequential data, has driven numerous recent breakthroughs in natural language processing and beyond. This work provides a theoretical analysis of the inductive biases of self-attention modules. Our focus is to rigorously establish which functions and long-range dependencies self-attention blocks prefer to represent. Our main result shows that bounded-norm Transformer networks “create sparse variables”: a single self-attention head can represent a sparse function of the input sequence, with sample complexity scaling only logarithmically with the context length. To support our analysis, we present synthetic experiments to probe the sample complexity of learning sparse Boolean functions with Transformers.

1. Introduction

Self-attention mechanisms have comprised an era-defining cornerstone of deep learning in recent years, appearing ubiquitously in empirical breakthroughs in generative sequence modeling and unsupervised representation learning. Starting with natural language (Vaswani et al., 2017), self-attention has enjoyed surprising empirical successes in numerous and diverse modalities of data. In many of these settings, self-attention has supplanted traditional recurrent and convolutional architectures, which are understood to incorporate inductive biases about temporal and translational invariances in the data. Self-attention models discard these functional forms, in favor of directly and globally modeling long-range interactions within the input context.

The proliferation of self-attention raises a fundamental question about its inductive biases: *which functions do self-attention networks prefer to represent?* Various intuitions

and empirics inform the design of these architectures, but formal statistical abstractions and analyses are missing in this space. To this end, this work initiates an analysis of the statistical foundations of self-attention.

We identify an inductive bias for self-attention, for which we coin the term *sparse variable creation*: a bounded-norm self-attention head learns a sparse function (which only depends on a small subset of input coordinates, such as a constant-fan-in gate in a Boolean circuit) of a length- T context, with sample complexity scaling as $\log(T)$. The main technical novelty in this work is a covering number-based capacity bound for attention mechanisms (including Transformer heads, as well as related and future architectures), implying norm-based generalization bounds. This is accompanied by a matching representational result, showing that bounded-norm self-attention heads are indeed capable of representing s -sparse functions with weight norms $2^{O(s)}$ (or $\text{poly}(s)$, for symmetric sparse functions). This provides a theoretical account for why attention models can learn long-range dependencies without overfitting.

Finally, we conduct synthetic experiments to probe the sample efficiency of learning sparse interactions with self-attention. We train Transformer models to identify sparse Boolean functions with randomly chosen indices, and corroborate the sample complexity scaling law predicted by the theory. A variant of this experiment (with i.i.d. samples) reveals a *computational* mystery, beyond the scope of our current statistical analysis: we find that Transformers can successfully learn the “hardest” (in the sense of SQ-dimension) s -sparse functions: the XOR (parity) functions.

1.1. Related work

The direct precursors to modern self-attention architectures were recurrent and convolutional networks augmented with attention mechanisms (Bahdanau et al., 2014; Luong et al., 2015; Xu et al., 2015). Landmark work by Vaswani et al. (2017) demonstrated significant improvements in machine translation via a pure self-attention architecture; autoregressive language models (Liu et al., 2018; Radford et al., 2018; 2019; Brown et al., 2020), and self-supervised representation learning via masked language modeling (Devlin et al., 2018) followed shortly.

¹Department of Computer Science, Harvard University, Cambridge, MA, USA ²Microsoft Research, New York, NY, USA. Correspondence to: Cyril Zhang <cyrilzhang@microsoft.com>.

Norm-based capacity bounds for neural nets. There is a vast body of literature dedicated to establishing statistical guarantees for neural networks, including VC-dimension and shattering bounds (dating back to (Anthony & Bartlett, 1999)). In recent years, classical norm-based generalization bounds have been established for various architectures (Bartlett et al., 2017; Neyshabur et al., 2015; 2017; Golowich et al., 2018; Long & Sedghi, 2019; Chen et al., 2019) using covering-based arguments. Jiang et al. (2019) provide an extensive empirical study of how well these bounds predict generalization in practice. Our work complements these results by establishing the first norm-based capacity analysis for attention models. Our main results rely on a novel reduction to the ℓ_∞ covering number bound for linear function classes given by (Zhang, 2002).

Other theoretical lenses on attention. Our work complements various existing theoretical perspectives on attention-based models. Vuckovic et al. (2020) formulate a dynamical system abstraction of attention layers, arriving at similar Lipschitz constant calculations to ours (which are coarser-grained, since they focus on contractivity and stability rather than finite-sample statistical guarantees). Zhang et al. (2019); Snell et al. (2021) study idealizations of the optimization problem of learning self-attention heads. Wei et al. (2021) propose a definition of *statistically meaningful approximation* of function classes that ties statistical learnability with expressivity, and show that Boolean circuits can be *SM-approximated* by Transformers with a sample complexity bound that depends mildly on circuit depth (rather than context size), using a margin amplification procedure. Kim et al. (2021) show that standard dot-product attention is not Lipschitz for an *unbounded* input domain, whereas our paper shows that norm-based generalization bounds are attainable with a $\|\cdot\|_{2,\infty}$ -bounded input domain.

See Appendix D for a broader survey of the literature on attention and self-attention networks.

2. Background and notation

Throughout this paper, the input $X := [x_1 x_2 \dots x_T]^\top \in \mathbb{R}^{T \times d}$ to an attention module (a.k.a. the context) will be a length- T sequence of embeddings $x_t \in \mathbb{R}^d$; m refers to the sample size (i.e. number of length- T sequences in a dataset). $\|\cdot\|_2$ denotes the spectral norm for matrices, and $\|\cdot\|_{p,q}$ denotes the (p,q) matrix norm where the p -norm is over columns and q -norm over rows. For vectors, $\|\cdot\|_p$ denotes the ℓ_p norm; we drop the subscript for the ℓ_2 norm. B is generally used to quantify bounds on norms of matrices and L for Lipschitz constants. Δ^{n-1} denotes the simplex in dimension n , that is, $\Delta^{n-1} := \{x \in \mathbb{R}^n : x \geq 0, \|x\|_1 = 1\}$.

Covering numbers. Our main technical contribution is a generalization bound arising from carefully counting the number of functions representable by a Transformer. The main technical ingredient is the notion of a covering number. We will use the following definition of ∞ -norm covering number adapted from (Zhang, 2002):

Definition 2.1 (Covering number). *For a given class of vector-valued functions \mathcal{F} , the covering number $\mathcal{N}_\infty(\mathcal{F}; \varepsilon; \{z^{(i)}\}_{i=1}^m; \|\cdot\|)$ is the smallest size of a collection (a cover) $\mathcal{C} \subset \mathcal{F}$ such that $\forall f \in \mathcal{F}, \exists \hat{f} \in \mathcal{C}$ satisfying*

$$\max_i \|f(z^{(i)}) - \hat{f}(z^{(i)})\| \leq \varepsilon.$$

Further, define

$$\mathcal{N}_\infty(\mathcal{F}, \varepsilon, m, \|\cdot\|) = \sup_{z^{(1)}, \dots, z^{(m)}} \mathcal{N}_\infty(\mathcal{F}; \varepsilon; z^{(1)}, \dots, z^{(m)}, \|\cdot\|).$$

If \mathcal{F} is real-valued (instead of vector-valued), we drop the norm from the notation. Furthermore for functions parameterized by a set of parameters Θ , we exploit the notation to replace \mathcal{F} by Θ .

Recall that for the class of linear functions,

$$\mathcal{F}_{\text{lin}} = \{x \mapsto w \cdot x : w \in \mathbb{R}^d, \|w\|_2 \leq B_W\},$$

we have the covering number bound (Zhang, 2002) of

$$\mathcal{N}_\infty(\mathcal{F}; \varepsilon; \{x^{(i)}\}_{i=1}^m) \leq O\left(\frac{B_X^2 B_W^2}{\varepsilon^2} \cdot \log\left(\frac{B_X B_W m}{\varepsilon}\right)\right),$$

where $\|x^{(i)}\| \leq B_X$ for $i \in [m]$. Importantly, note that the covering number has a mild dependence on m , only logarithmic; this logarithmic dependence on m will be helpful when we turn our analysis to the capacity of attention mechanisms.

Generalization bounds. This work focuses on providing log-covering number bounds, which imply uniform generalization bounds via standard arguments. The following lemma relates these quantities; we refer the reader to Appendix A.1 for a formal review.

Lemma 2.2 (Generalization bound via covering number; informal). *Suppose \mathcal{F} is a class of bounded functions, and $\log \mathcal{N}_\infty(\mathcal{F}; \varepsilon; x^{(1)}, \dots, x^{(m)}) \leq C_{\mathcal{F}}/\varepsilon^2$ for all $x^{(1)}, \dots, x^{(m)} \in \mathcal{X}^m$. Then for any $\delta > 0$, with probability at least $1 - \delta$, simultaneously for all $f \in \mathcal{F}$, the generalization error ε_{gen} satisfies*

$$\varepsilon_{\text{gen}}(f) \leq \tilde{O}\left(\sqrt{\frac{C_{\mathcal{F}}}{m}} + \sqrt{\frac{\log(1/\delta)}{m}}\right).$$




Figure 1. Diagrams of attention modules $f_{\text{tf-head}}$, $f_{\text{tf-layer}}$, $f_{\text{tf-scalar}}$ described in Section 3: alignment scores (grey edges) determine normalized attention weights (blue), which are used to mix the inputs $x_{1:T}$. *Left*: Attention with a general context z . *Center*: Self-attention layer, where both the input and the context come from $x_{1:T}$. *Right*: Auxiliary [CLS] token to extract a single scalar from a self-attention layer, providing a real-valued function class for classification or regression tasks.

3. Abstractions for (self-)attention

The precise definition of attention is less straightforward to define than for architectural components such as convolutions and residual connections. In this section, guided by the manifestations of attention discussed in (Luong et al., 2015), we present some notation and definitions which generalize attention mechanisms commonly seen in practice, including the Transformer. Intuitively, these definitions encompass neural network layers which induce context-dependent representation bottlenecks. Subsequently, we show how to represent the Transformer (the predominant attention-based architecture) as a special case of this formulation.

3.1. Attention

Intuitively, we would like to capture the notion that an output variable selects (“attends to”) a part of the input sequence on which it will depend, based on a learned function of global interactions (see Figure 1, left). To this end, we define an *attention head* as a function which maps a length- T input sequence (e.g. the tokens in a sentence, pixels in an image, or intermediate activations in a deep Transformer network) and an additional context $z \in \mathcal{Z}$ to an output $y \in \mathcal{Y}$. In this work, we will exclusively consider $\mathcal{X}, \mathcal{Y}, \mathcal{Z}$ to be \mathbb{R}^d . An attention head uses z to select the input coordinates in X to which the output y will “attend”, formalized below:

Definition 3.1 (Attention head). *An attention head is a function $f : \mathcal{X} \rightarrow \mathcal{Y}$, specified by an alignment score function $\text{Score} : \mathcal{X} \times \mathcal{Z} \rightarrow \mathbb{R}$ parameterized by $\theta_s \in \Theta_s$, normalization function $\text{Norm} : \mathbb{R}^T \rightarrow \Delta^{T-1}$, and position-wise maps $\phi_{\text{in}} : \mathcal{X} \rightarrow \mathcal{V}, \phi_{\text{out}} : \mathcal{V} \rightarrow \mathcal{Y}$ parameterized by $\theta_{\text{in}} \in \Theta_{\text{in}}$ and $\theta_{\text{out}} \in \Theta_{\text{out}}$. The output of an attention head on input $X \in \mathcal{X}^T, z \in \mathcal{Z}$ is*

$$y = \phi_{\text{out}} \left(\sum_{t=1}^T \left[\text{Norm} \left(\text{Score}(x_1, z; \theta_s), \dots, \text{Score}(x_T, z; \theta_s) \right) \right]_t \phi_{\text{in}}(x_t; \theta_{\text{in}}); \theta_{\text{out}} \right)$$

$$= \phi_{\text{out}} \left(\phi_{\text{in}}(X; \theta_{\text{in}})^\top \text{Norm} \left(\text{Score}(x_1, z; \theta_s), \dots, \text{Score}(x_T, z; \theta_s) \right); \theta_{\text{out}} \right)$$

where $\phi_{\text{in}}(X; \theta) = [\phi_{\text{in}}(x_1; \theta) \dots \phi_{\text{in}}(x_T; \theta)]^\top$ denotes the row-wise application of ϕ_{in} .

The above definition corresponds to the leftmost diagram in Figure 1. Here, \mathcal{V} is a vector space of input representations “mixed” by the normalized alignment scores; in this work, we will set $\mathcal{V} = \mathbb{R}^k$. A function class of attention heads is induced by specifying parameter classes for $\{\Theta_s, \Theta_{\text{in}}, \Theta_{\text{out}}\}$.

3.2. Self-attention and Transformers

A *self-attention head* is a special case of an attention head, in which the context z comes from one of the inputs x_t themselves: pairwise interactions among the elements in X are used to select the elements of X on which f depends. In this case, we use “input” and “context” interchangeably to refer to X . For example, a self-attention head which uses $z := x_t$ is defined by

$$y = \phi_{\text{out}} \left(\phi_{\text{in}}(X; \theta_{\text{in}})^\top \text{Norm}(\text{Score}(X, x_t; \theta_s)); \theta_{\text{out}} \right).$$

We now define the Transformer self-attention architecture as a special case of the above. Since a Transformer layer has shared parameters between multiple output heads, we will define all T outputs of the layer at once.

Definition 3.2 (Transformer layer). *A Transformer layer is a collection of T attention heads with outputs y_1, \dots, y_T , specified by the following choices of function classes (with shared parameters between the heads) where the context for y_τ is x_τ .*

- $\text{Score}(x, x_\tau; \{W_Q, W_K\}) := x_\tau^\top W_Q W_K^\top x$,
 $W_Q, W_K \in \mathbb{R}^{d \times k}$ (for output y_t)
- $\phi_{\text{in}}(x; W_V) := W_V^\top x$, $W_V \in \mathbb{R}^{d \times k}$
- $\phi_{\text{out}}(x; W_C) := W_C^\top \sigma(x)$, $W_C \in \mathbb{R}^{k \times d}$,

L_σ -Lipschitz activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ applied position-wise, with $\sigma(0) = 0$.

- $\text{Norm}(x) := \text{softmax}(x) = \frac{\exp(x)}{\mathbf{1}^\top \exp(x)}$

Defining $Y := [y_1 y_2 \dots y_T]^\top \in \mathbb{R}^{T \times d}$ and $[\text{RowSoftmax}(M)]_{t,:} := \text{softmax}(M_{t,:})$, we have

$$Y = \sigma(\text{RowSoftmax}(XW_Q(XW_K)^\top)XW_V)W_C.$$

Functions from the above class of Transformer layers map $\mathbb{R}^{T \times d}$ to itself, and can thus be iteratively composed. We discuss remaining discrepancies between Definition 3.2 and real Transformers (positional embeddings, position-wise feedforward networks, layer normalization, parallel heads, residual connections) in Section 4.3 and the appendix.

Extracting scalar outputs from a Transformer. Finally, we establish notation for a canonical way to extract a scalar prediction from the final layer of a Transformer. For a context of size T , a Transformer layer with $T + 1$ inputs is constructed, with a special index $[\text{CLS}]$.¹ The input at this position is a vector $x_{[\text{CLS}]} \in \mathbb{R}^d$ (which can be fixed or trainable); the output is a linear function $w^\top y_{[\text{CLS}]}$, for a trainable parameter $w \in \mathbb{R}^d$. This defines a class of functions mapping $\mathbb{R}^{T \times d} \rightarrow \mathbb{R}$, parameterized by a Transformer layer’s parameters and w , which we call the class of *scalar-output Transformers*. This is the setup used by the classification modules in BERT (Devlin et al., 2018) and all of its derivatives.

4. Capacity bounds for attention modules

In this section, we present covering number-based capacity bounds for generic attention heads and Transformers, along with overviews of their proofs. Section 4.1 bounds the capacity of a general attention head. Section 4.2 instantiates this bound for the case of a single Transformer self-attention head. Section 4.3 generalizes this bound for full depth- L Transformer networks. The sample complexity guarantees for Transformers scale only *logarithmically* in the context length T , providing rigorous grounding for the intuition that the architecture’s inductive bias selects sparse functions of the context.

Note: Throughout this section, assume that $\|x_t\|_2 \leq B_X$ for all $t \in [T]$ (i.e. $\|X\|_{2,\infty} \leq B_X$). Note that this allows for the Frobenius norm $\|X\|_F$ to scale with \sqrt{T} . The key challenge throughout our analysis is to avoid incurring factors of norms which take a sum over the t dimension, by constructing covers appropriately.

¹ $[\text{CLS}]$ stands for “class”, as in “treat the output at this position as the classifier’s prediction”.

4.1. Capacity of a general attention head

Recall that the attention head architecture can be represented as a function $f_{\text{head}} : \mathbb{R}^{T \times d} \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ parameterized by $\theta_s, \theta_{\text{in}}, \theta_{\text{out}}$ as

$$f_{\text{head}}(X, z; \theta_s, \theta_{\text{in}}, \theta_{\text{out}}) = \phi_{\text{out}}(\phi_{\text{in}}(X; \theta_{\text{in}})^\top \text{Norm}(\text{Score}(X, z; \theta_s)); \theta_{\text{out}}).$$

Denote the corresponding function class by

$$\mathcal{F}_{\text{head}} := \{(X, z) \mapsto f_{\text{head}}(X, z; \theta_s, \theta_{\text{in}}, \theta_{\text{out}}) : \theta_s \in \Theta_s, \theta_{\text{in}} \in \Theta_{\text{in}}, \theta_{\text{out}} \in \Theta_{\text{out}}\}$$

To convert the vector-valued function class to a scalar output function class, we define $\mathcal{F}_{\text{scalar}} := \{(X, z) \mapsto w^\top f(X, z) : f \in \mathcal{F}_{\text{head}}, w \in \mathbb{R}^d, \|w\| \leq B_w\}$.

For simplicity of presentation, we will focus only on the attention head, and assume that ϕ_{out} and w are fixed. We handle the general case of trainable downstream layers in the analysis of multi-layer Transformers in Appendix A.7.

Assumption 4.1. We make the following assumptions:

1. ϕ_{out} is L_{out} -Lipschitz in the ℓ_2 -norm, that is,

$$\forall a, b \in \mathbb{R}^k : \|\phi_{\text{out}}(a) - \phi_{\text{out}}(b)\| \leq L_{\text{out}}\|a - b\|.$$

2. ϕ_{in} is B_{in} -bounded in ℓ_2 -norm, that is,

$$\forall a \in \mathbb{R}^d, \theta_{\text{in}} \in \Theta_{\text{in}} : \|\phi_{\text{in}}(a; \theta_{\text{in}})\| \leq B_{\text{in}}\|a\|.$$

3. Norm is continuously differentiable and its Jacobian satisfies

$$\forall \theta \in \mathbb{R}^T, \|J \text{Norm}(\theta)\|_{1,1} \leq C_{\text{Norm}}.$$

Note that softmax (the most commonly used Norm function) satisfies the Jacobian assumption with $C_{\text{softmax}} = 2$ (see Corollary A.7).

We prove the following bound on the covering number of $\mathcal{F}_{\text{head}}$ for m samples,

Theorem 4.2 (Attention head capacity). *Under Assumptions 4.1, $\forall \alpha \in [0, 1]$ the covering number of $\mathcal{F}_{\text{head}}$ satisfies*

$$\begin{aligned} & \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{head}}; \varepsilon; \left\{ (X^{(i)}, z^{(i)}) \right\}_{i=1}^m; \|\cdot\|_2 \right) \leq \\ & \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{Score}}; \frac{\alpha \varepsilon}{C_{\text{Norm}} L_{\text{out}} B_{\text{in}} B_X}; \left\{ (x_t^{(i)}, z^{(i)}) \right\}_{t \in [T]}^{i \in [m]} \right) \\ & + \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{in}}; \frac{(1-\alpha)\varepsilon}{L_{\text{out}}}; \left\{ x_t^{(i)} \right\}_{t \in [T]}^{i \in [m]}; \|\cdot\|_2 \right), \end{aligned}$$

where $\mathcal{F}_{\text{Score}} = \{(x, z) \mapsto \text{Score}(x, z; \theta_s) : \theta_s \in \Theta_s\}$, and $\mathcal{F}_{\text{in}} = \{x \mapsto \phi_{\text{in}}(x; \theta_{\text{in}}) : \theta_{\text{in}} \in \Theta_{\text{in}}\}$.

Note that the bound is in terms of the \mathcal{N}_∞ covering number of functions that dependent on dimensions d or k and not T . The effect of T only shows up in the number of samples to cover. Crucially, for some function classes (e.g. linear functions (Zhang, 2002)), \mathcal{N}_∞ scales only logarithmically with the number of samples. This is exactly what allows us to obtain our log T capacity bounds.

Since w is fixed, an ε -covering of $\mathcal{F}_{\text{head}}$ directly gives us an εB_w -covering for $\mathcal{F}_{\text{scalar}}$, implying

$$\begin{aligned} & \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{scalar}}; \varepsilon; \left\{ (X^{(i)}, z^{(i)}) \right\}_{i=1}^m \right) \\ & \leq \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{head}}; \varepsilon / B_w; \left\{ (X^{(i)}, z^{(i)}) \right\}_{i=1}^m, \|\cdot\|_2 \right). \end{aligned}$$

Proof overview. In order to prove the bound, we first show a Lipschitzness property of f_{head} . This property allows us to construct a cover by using covers for $\mathcal{F}_{\text{Score}}$ and \mathcal{F}_{in} .

Lemma 4.3 (ℓ_∞ -Lipschitzness of f_{head}). *For any $\theta_s, \hat{\theta}_s \in \Theta_s, \theta_{\text{in}}, \hat{\theta}_{\text{in}} \in \Theta_{\text{in}}$; for all $X \in \mathbb{R}^{T \times d}$, such that $\|X^\top\|_{2, \infty} \leq B_X$,*

$$\begin{aligned} & \left\| f_{\text{head}}(X, z; \theta_s, \theta_{\text{in}}, w) - f_{\text{head}}(X, z; \hat{\theta}_s, \hat{\theta}_{\text{in}}, w) \right\| \leq \\ & C_{\text{Norm}} L_{\text{out}} B_{\text{in}} B_X \left\| \text{Score}(X, z; \theta_s) - \text{Score}(X, z; \hat{\theta}_s) \right\|_\infty \\ & + L_{\text{out}} \left\| \phi_{\text{in}}(X; \theta_{\text{in}}) - \phi_{\text{in}}(X; \hat{\theta}_{\text{in}}) \right\|_{2, \infty}. \end{aligned}$$

The most crucial aspect of this proof is to avoid a spurious T dependence when accounting for the attention mechanism. The key observation here is that the attention part of the network is computed using Norm, whose Jacobian norm is bounded. This allows us to use the mean-value theorem to move to the maximum (ℓ_∞) error over T tokens instead of sum (ℓ_1), which could potentially incur a T factor. Furthermore, this allows us to combine all samples and tokens and construct an ℓ_∞ -cover directly for mT samples.

4.2. Capacity of a Transformer head

Let us now look at the case of a Transformer self-attention head and instantiate the covering bound. For ease of presentation and to focus on the self-attention part, we collapse $W_Q W_K^\top$ to a single matrix, set $k = d$ and remove the linear layer W_C^2 . Then the Transformer self-attention head (for any fixed $\tau \in [T]$) can be described as

$$\begin{aligned} f_{\text{tf-head}}(X; W_V, W_{QK}) := \\ \sigma \left(W_V^\top X^\top \text{softmax} \left(X W_{QK}^\top x_\tau \right) \right) \end{aligned}$$

which is obtained from the general formulation by setting the context to be x_τ , $\text{Score}(X, x_\tau; W_{QK}) = X W_{QK}^\top x_\tau$, $\text{Norm} = \text{softmax}$ and $\phi_{\text{out}} = \sigma$.

²See Appendix A.7 for an analysis of general deep Transformer models.

Because the number of parameters in a Transformer self-attention head is $O(d^2)$, with no dependence on T , one might presume by simple parameter counting that the capacity of the class of these heads does not grow as the context length T grows. But capacity is not solely governed by the number of parameters—for example, the class $\{x \in \mathbb{R} \mapsto \text{sign}(\sin(\alpha x))\}_{\alpha \in \mathbb{R}}$ has a single parameter but infinite VC-dimension. One might still hope to prove, for the special case of Transformer heads, a T -independent upper bound on the VC-dimension (or rather, its analog for real-valued functions, the *pseudo-dimension*). We observe that, in fact, the pseudo-dimension of this class does grow with T .

Proposition 4.4. *When the embedding dimension is $d = 3$, the class $\mathcal{F}_{\text{tf-head-unbounded}}$ ($:= \{X \mapsto f_{\text{tf-head}}(X; W_V, W_{QK}) : W_V, W_{QK} \in \mathbb{R}^{d \times d}\}$) of Transformer self-attention heads with unbounded norm has pseudo-dimension $\geq \lfloor \log T \rfloor$.*

The proofs for this subsection can be found in Appendix A.

Let us now define the function class of self-attention heads with bounded weight norms:

$$\begin{aligned} \mathcal{F}_{\text{tf-head}} := \{X \mapsto f_{\text{tf-head}}(X; W_V, W_{QK}) : \\ \|W_V\|_{2,1} \leq B_V^{2,1}, \|W_V\| \leq B_V, \|W_{QK}^\top\|_{2,1} \leq B_{QK}^{2,1}\}. \end{aligned}$$

Since W_V, W_{QK} have dimensions dependent on d and k , bounding their norms does not hide a T dependence. As before, to convert this vector-valued function class to a scalar output function class, we define

$$\begin{aligned} \mathcal{F}_{\text{tf-scalar}} := \\ \{X \mapsto w^\top f(X) : f \in \mathcal{F}_{\text{tf-head}}, w \in \mathbb{R}^d, \|w\| \leq B_w\}. \end{aligned}$$

We obtain the following bound on the covering number of $\mathcal{F}_{\text{tf-head}}$ as a corollary of Theorem 4.2:

Corollary 4.5. *For any $\varepsilon > 0$ and $X^{(1)}, \dots, X^{(m)} \in \mathbb{R}^{T \times d}$ such that $\|X^{(i)}\|_{2, \infty} \leq B_X$ for all $i \in [m]$, the covering number of $\mathcal{F}_{\text{tf-head}}$ satisfies*

$$\begin{aligned} & \log \mathcal{N}_\infty(\mathcal{F}_{\text{tf-head}}; \varepsilon; X^{(1)}, \dots, X^{(m)}, \|\cdot\|_2) \\ & \lesssim (L_\sigma B_X)^2 \cdot \frac{\left((B_V^{2,1})^{\frac{2}{3}} + (B_{QK}^{2,1} B_V)^{\frac{2}{3}} \right)^3}{\varepsilon^2} \cdot \log(mT) \end{aligned}$$

Here \lesssim hides logarithmic dependencies on quantities besides m and T .

Proof overview. The above result follows from bounding the covering numbers of

$$\begin{aligned} \mathcal{F}_{QK} := \{z \mapsto x_\tau^\top W_{QK} z : \|W_{QK}^\top\|_{2,1} \leq B_{QK}^{2,1}\}, \text{ and} \\ \mathcal{F}_V := \{z \mapsto W_V^\top z : \|W_V\|_{2,1} \leq B_V^{2,1}, \|W_V\| \leq B_V\}. \end{aligned}$$

Note that $|x_\tau^\top W_{QK}z - x_\tau^\top W_{QK}\hat{z}| \leq \|W_{QK}z - W_{QK}\hat{z}\|$ since $\|x_\tau\| \leq 1$, so the covering number of \mathcal{F}_{QK} is at most the covering number of the class of functions of the form $z \mapsto W_{QK}z$. Therefore, a covering number bound for the vector-valued linear function class suffices to handle both covering numbers:

Lemma 4.6. *Let $\mathcal{W} : \{W \in \mathbb{R}^{d_1 \times d_2} : \|W^\top\|_{2,1} \leq B_W\}$, and consider the function class $\mathcal{F} : \{x \mapsto Wx : W \in \mathcal{W}\}$. For any $\varepsilon > 0$ and $x^{(1)}, \dots, x^{(N)} \in \mathbb{R}^{d_1}$ satisfying $\forall i \in [N], \|x^{(i)}\| \leq B_X$,*

$$\begin{aligned} & \log \mathcal{N}_\infty(\mathcal{F}; \varepsilon; x^{(1)}, \dots, x^{(N)}; \|\cdot\|_2) \\ & \lesssim \frac{(B_X B_W)^2}{\varepsilon^2} \log(d_1 N). \end{aligned}$$

Note that this bound only depends logarithmically on the context length, as desired. The proof can be found in Appendix A.

Finally, our analysis is compatible with the following additional components:

Positional embeddings. In practice, the permutation-invariant symmetry of a Transformer network is broken by adding a *positional embedding* matrix $P \in \mathbb{R}^{T \times d}$ to the input X at the first layer. In practice, the embedding matrix is often fixed (not trainable). Our results extend to this setting in a straightforward way; see Appendix A.5. If these matrices are to be trained from a sufficiently large class (say, $\|P\|_{2,\infty} \leq 1$), the dependence of the log-covering number on T could become linear.

Residual connections. Including residual connections (e.g. redefining $f_{\text{tf-head}}(X)$ as $x_t + f_{\text{tf-head}}(X)$ for some index $t \in [T]$) simply increases the Lipschitz constant of each layer (w.r.t. the input) by at most 1. As long as $B_V = \Omega(1)$, this only changes our covering number bounds by a constant factor.

Multi-head self-attention. In almost all applications of Transformers, multiple parallel self-attention heads are used, and their outputs aggregated, to allow for a richer representation. Our analysis directly extends to this setting; see Appendix A.6 for details. When a single attention head is replaced with the sum of H parallel heads, the log-covering number scales up by a factor of $\text{poly}(H)$.

Layer normalization. State-of-the-art Transformer networks are trained with layer normalization modules (Ba et al., 2016), which is generally understood to aid optimization. We keep a variant of layer normalization in the covering number analysis—it proves to be useful in the analysis of full attention blocks (see Appendix A.7), as it keeps the norm of the embedding of each token bounded. Removing these layers would lead to a worse dependence on the spectral norm of the matrices.

4.3. Capacity bounds for multi-layer Transformers

In this section, we will extend our results for L -layer Transformer blocks. Denote the weights of layer i by $W^{(i)} := \{W_Q^{(i)}, W_K^{(i)}, W_V^{(i)}, W_C^{(i)}\}$. Further denote the set of weights up to layer i by $W^{1:i} = (W^{(1)}, \dots, W^{(i-1)})$. Denote the input representation of layer i by $g_{\text{tf-block}}^{(i)}(X; W^{1:i})$. We inductively define $g_{\text{tf-block}}^{(i)} : \mathbb{R}^{T \times d} \rightarrow \mathbb{R}^{T \times d}$ starting with $g_{\text{tf-block}}^{(1)}(X; W^{1:1}) = X$ (the input):

$$\begin{aligned} & g_{\text{tf-block}}^{(i+1)}(X; W^{1:i+1}) := \\ & \Pi_{\text{norm}} \left(\sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-block}}^{(i)}(X; W^{1:i}); W^{(i)} \right) \right) \right) W_C^{(i)} \right) \\ & \text{with } f(Z; \{W_Q, W_K, W_V, \cdot\}) := \\ & \text{RowSoftmax} \left(ZW_Q (ZW_K)^\top \right) ZW_V, \end{aligned}$$

where Π_{norm} denotes layer normalization³ applied to each row. We use a slightly modified version of LayerNorm where instead of normalizing to norm 1, we project it to the unit ball. Let the class of depth- L transformer blocks be

$$\begin{aligned} \mathcal{F}_{\text{tf-block}}^{(L)} := & \\ & \left\{ X \rightarrow g_{\text{tf-block}}^{(L+1)}(X; W^{1:L+1}) : \forall i \in [L], \right. \\ & \left. \left\| W_V^{(i)} \right\|_2, \left\| W_K^{(i)} W_Q^{(i)\top} \right\|_2, \left\| W_C^{(i)} \right\|_2 \leq C_2, \right. \\ & \left. \left\| W_V^{(i)} \right\|_{2,1}, \left\| W_K^{(i)\top} W_Q^{(i)} \right\|_{2,1}, \left\| W_C^{(i)} \right\|_{2,1} \leq C_{2,1} \right\}. \end{aligned}$$

To obtain a final scalar output, we use a linear function of the [CLS] output:

$$g_{\text{tf-scalar}}(X; W^{1:L+1}, w) = w^\top [g(X; W^{1:L+1})]_{[\text{CLS}]},$$

Let the scalar output function class be $\mathcal{F}_{\text{tf-scalar}}^{(L)} := \{X \rightarrow w^\top f(X)_{[\text{CLS}]} : f \in \mathcal{F}_{\text{tf-block}}^{(L)}, w \in \mathbb{R}^d, \|w\| \leq B_w\}$.

Theorem 4.7 (Theorem A.17 (simplified)). *Suppose $\forall i \in [m], \|X^{(i)}\|_{2,\infty} \leq B_X$, then we have*

$$\begin{aligned} & \log \mathcal{N}_\infty(\mathcal{F}_{\text{tf-block}}^{(L)}; \varepsilon; X^{(1)}, \dots, X^{(m)}) \\ & \lesssim (C_2 L \sigma)^{O(L)} \cdot \frac{B_X^2 B_w^2 C_{2,1}^2}{\varepsilon^2} \cdot \log(dmT). \end{aligned}$$

Note that the dependence on d and T is only logarithmic even for deeper networks. The dependence on $(2, 1)$ -norms

³Layer normalization allows for the norms of the outputs of each token in each layer to remain bounded by 1. Note that the norm of the entire input can still have a dependence on T . Our results would go through with a worse dependence on the spectral norms if we were to remove layer norm.

of the weight matrices is quadratic. As long as the spectral norms of the matrices are bounded by 1 and σ is 1-Lipschitz (which holds for sigmoids and ReLUs), the exponential dependence on L can be avoided.

5. Attention approximates sparse functions

The results in Section 4 show that function classes bottlenecked by self-attention mechanisms have “small” statistical capacity in terms of the context size. In this section, we answer the converse question: *which functions of interest are in these classes?* We show that Transformers are able to represent sparse interactions in the context with bounded weight norms, and can thus learn them sample-efficiently.

Consider the class of Boolean functions $f : \{0, 1\}^T \rightarrow \mathbb{R}$ which are s -sparse: they only depend on $s \ll T$ of their inputs. We will construct mappings from such functions to parameters of a self-attention head $f_{\text{tf-head}}$ composed with a feedforward network f_{mlp} ; note that $f_{\text{tf-head}} \circ f_{\text{mlp}}$ is the standard Transformer block. Intuitively, $f_{\text{tf-head}}$ is constructed to “keep” the correct s -dimensional subset of inputs and “forget” the rest, while f_{mlp} “memorizes” the values of f on these s inputs, using $2^{O(s)}$ parameters.

Setup. We consider the classes of Boolean functions $f : \{0, 1\}^T \rightarrow \mathbb{R}$ representable by bounded-norm scalar-output Transformer heads $f_{\text{tf-scalar}} : \mathbb{R}^{T \times d} \rightarrow \mathbb{R}$. To do this, we must first fix a mapping from $\{0, 1\}^T$ to $\mathbb{R}^{T \times d}$; we discuss several natural choices in Appendix B.1. The simplest of these uses a sum of token and positional embeddings $X(b)_{t,:} := e_{b_t} + v_t$, for a set of approximately orthogonal unit vectors $\{e_0, e_1\} \cup \{v_1, \dots, v_T\}$ of dimension $d = \Theta(\log T)$. After choosing a mapping $X(b)$, the setup of the representation problem is as follows: given $f(b)$, find Transformer weights $\theta_{\text{tf-head}}$ and feedforward network weights θ_{mlp} such that

$$f_{\text{tf+mlp}}(X(b); \theta_{\text{tf-head}}, \theta_{\text{mlp}}) := f_{\text{mlp}}(f_{\text{tf-head}}(X(b); \theta_{\text{tf-head}}); \theta_{\text{mlp}}) \approx f(b), \forall b \in \{0, 1\}^T.$$

Main representational results. For any size- s subset of indices $\mathcal{I} \subseteq [T]$, we show that Transformer blocks can represent all \mathcal{I} -sparse Boolean functions, whose values only depend on the inputs at the coordinates in \mathcal{I} . We give informal statements of these approximation results below, and present the precise statements in Appendix B.2.

Proposition 5.1 (Sparse variable creation via Transformers; informal). *Under any of the input mappings $X(b)$, we have the following guarantees:*

- $f_{\text{tf-scalar}}$ can approximate a particular monotone symmetric s -sparse Boolean function, with norms $\|W_Q\|_F \leq O(\log(Ts)); \|W_K\|_F, \|W_V\|_F, \|W_C\|_F \leq O(s)$.
- $f_{\text{tf+mlp}}$ can exactly represent symmetric s -sparse func-

tions, with the same Transformer weight norms as above; the feedforward network weights satisfy $\|W_1\|_F, \|W_2\|_F, \|w\|_F \leq O(\text{poly}(s))$.

- $f_{\text{tf+mlp}}$ can exactly represent general s -sparse functions, with the same Transformer weight norms as above; the feedforward network weights satisfy $\|W_1\|_F, \|W_2\|_F, \|w\|_F \leq O(2^s \cdot \text{poly}(s))$.

These results and the capacity bounds from Section 4 are simultaneously meaningful in the regime of $s \ll \log T$. An appealing interpretation for the $s = 2$ case is that a single Transformer head can learn a single logical gate (i.e. AND, OR, NAND, ...) in a Boolean circuit, with d and weight norms scaling as $\log T$.

Proof ideas. Each construction uses the same basic idea: select W_Q, W_K so that the attention mixture weights approximate the uniform distribution over the relevant positions, then use the ReLU network to memorize all distinct values of f . Full proofs are given in Appendix B.5.

Other realizable functions. Since there are $\binom{T}{s}$ s -sparse subsets of input indices, the sample complexity of learning a sparse Boolean function must scale at least as $\Omega(s \log T)$, matching the capacity bounds in terms of the $\log T$ dependence. However, sparse functions are not the *only* potentially useful functions realizable by bounded-norm Transformers. For instance, with $W_Q = W_K = 0$, so that all scores are zero, a Transformer head can take an average of T embeddings $X \rightarrow \frac{1}{T} \mathbf{1}^\top X W_V$. More generally, departing from the “orthogonal context vectors” embedding of Boolean inputs but using the same constructions as in this section, it is straightforward to conclude that bounded-norm Transformers can compute global averages of tokens whose $X W_K$ embeddings lie in an s -dimensional subspaces. This is why our results do not contradict the empirical finding of Clark et al. (2019) that some attention heads in trained Transformer models attend broadly. It is also straightforward to extend some of these results beyond Boolean domains; see Section B.4 for a sketch.

Bypassing theoretical limitations. Hahn (2020) points out that with constant weight norms, a Transformer’s ability to express global dependencies degrades with context length: as $T \rightarrow \infty$, the maximum change in output caused by altering a single input token approaches 0, and thus various interesting formal languages cannot be modeled by a Transformer in this particular limit. The constructions in this section show that this can be circumvented by allowing d and the weight norms to scale as $\log(T)$.

6. Experiments

Sections 4 and 5 show theoretically that Transformers can learn sparse Boolean functions, with sparse regression-like

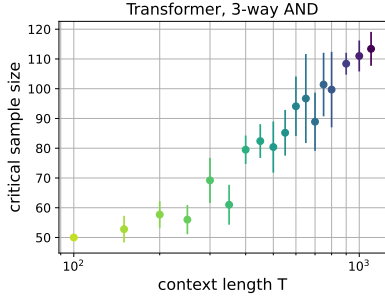


Figure 2. Main experimental finding: the sample complexity of learning a 3-sparse AND function of T input bits with Transformers. For each T , we measure the smallest sample size m necessary to reach 100% validation accuracy on at least 80% of random trials, finding that this threshold scales logarithmically with T .

sample complexity (in terms of the $\log T$ dependence). In this section, we present an empirical study which probes the end-to-end sample efficiency of Transformer architectures with standard training and architecture hyperparameters, and how it scales with the context length T .

Setup. We introduce a synthetic benchmark to support our analysis, in which we measure the statistical limit for learning sparse Boolean functions with Transformers. We choose a distribution \mathcal{D} on $\{0, 1\}^T$, and a family of distinct functions $\{f_i : \{0, 1\}^T \rightarrow \{0, 1\}\}_{i \in [N]}$, where N grows with T . Then, we choose an $i^* \in [N]$ uniformly at random, and train a Transformer binary classifier on m samples from \mathcal{D} , with labels given by f_{i^*} , evaluating generalization error via holdout samples. Then, for any learner to reach 100% accuracy on this sample, $m \geq \Omega(\log N)$ samples are required (one sample reveals at most one bit of information about i^*). We can then measure the empirical scaling of the sufficient sample size m to solve this problem, in terms of N (and thus T).

Learning sparse conjunctions. Concretely, we can choose f_i be the set of all $\binom{T}{s}$ conjunctions of s inputs (e.g. $y = x_2 \wedge x_3 \wedge x_{10}$), fixing the input distribution \mathcal{D} to be i.i.d. Bernoulli (we choose the bias to balance the labels). The model must learn which subset of s features are relevant, out of $\binom{T}{s}$ possibilities; this requires at least $m \geq \Omega(s \log T)$ samples. The theoretical analysis predicts that the sample complexity of learning any function realizable by a bounded-norm Transformer should asymptotically have the same $\log T$ scaling. We choose a fixed sparsity parameter $s = 3$, and measure how the empirical sample complexity (the smallest sample size $m(T)$ at which model training succeeds with non-negligible probability) scales with T .

Results. With architecture and training hyperparameters typical of real Transformer setups (except the number of

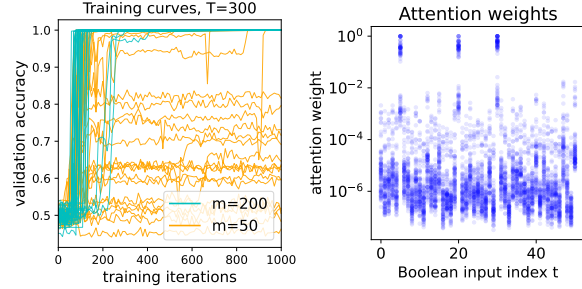


Figure 3. Additional visualizations for the sparse function learning experiments. *Left:* Examples of validation accuracy curves on the same problem instance ($T = 300$), with sample sizes above ($m = 200$) and below ($m = 50$) the threshold (≈ 70 from Figure 6). Training accuracy goes to 100% in both cases, but the Transformer overfits (orange curves) when m is too small. *Right:* Per-example attention weights for a successfully trained model ($T = 50$, $m = 300$, $\mathcal{I} = \{5, 20, 30\}$). The input-dependent attention weights approximately zero out the irrelevant bits.

layers, which we set to 1), we indeed observe that the empirical sample complexity appears to scale as $\log T$; see Figure 6. Despite the exponentially large support of input bit strings x and large total parameter count ($\sim 10^5$), the attention weights vanish on the $T - s$ irrelevant coordinates, and the model converges to sparse solutions; this is visualized in Figure 3 (*right*). Details are provided in Appendix C.1; in particular, model training near the statistical threshold is extremely unstable, and extensive variance reduction (best of 5 random restarts; 40 replicates; a total of $\sim 10^4$ training runs across each T, m) was necessary to produce these scaling plots.

Beyond our analysis: sparse parities. When choosing the family of sparse functions $\{f_i\}$, we can replace the AND operation with XOR: the label is the parity of a randomly chosen subset of i.i.d. uniform input bits. In this setting, unlike the AND case, there is a computational-statistical gap: $\Theta(s \log T)$ samples suffice to identify, but the fastest known algorithms for learning parities with noise require $T^{\Omega(s)}$ time. In the statistical query model, $\Omega(T^s)$ iterations of noisy batch gradient descent are necessary (Kearns, 1998). Figure 4 (with details in Appendix C.2) shows that when trained with i.i.d. samples, Transformer models can learn sparse parities. This raises an intriguing question, which is the *computational* analogue of the current work’s *statistical* line of inquiry: how does local search (i.e. gradient-based training) succeed at finding solutions that correspond to sparse discrete functions? The present work merely shows that these solutions exist; we intend to address the computational mystery in future work.

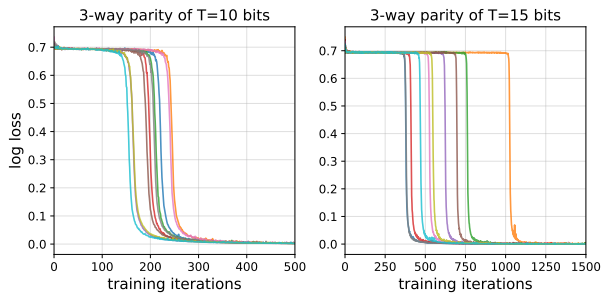


Figure 4. A curious empirical finding: Transformers can learn sparse parities. Loss curves (across 10 random seeds for initialization and SGD samples) are shown for this setup with $s = 3$, $T \in \{10, 15\}$, exhibiting phase transitions from random guessing to 100% accuracy. See Appendix C.2 for details.

7. Conclusion and future work

This work establishes a statistical analysis of attention and self-attention modules in neural networks. In particular, we identify an inductive bias we call *sparse variable creation*, consisting of (1) covering number-based capacity bounds which scale as $\log T$, and (2) constructions which show that self-attention models with small weight norms can represent sparse functions. This analysis is supported by an empirical study on learning sparse Boolean functions with Transformers. We hope that these rigorous connections between attention and sparsity, as well as the proposed experimental protocols, will inform the practice of training and regularizing these models, and the design of future attention-based architectures.

We believe that it is possible to refine the covering number bounds (where we have only sought to obtain optimal dependences on T) as well as the representation results (where we have not used the structure of the MLP, beyond its capacity for exhaustive memorization). Significant challenges (which are not specific to attention) remain in closing the theory-practice gap: precisely understanding the role of depth, as well as the trajectory of the optimization algorithm.

An exciting line of empirical work has made progress on understanding and interpreting state-of-the-art Transformer language models by examining the activations of their attention mechanisms (Clark et al., 2019; Tenney et al., 2019; Rogers et al., 2020). In some cases, these works have found instances in which Transformers seem to have learned features that are reminiscent of (sparse) hand-crafted features used in natural language processing. Reconciling our theoretical foundations work with this area of *BERTology* is an avenue for future synthesis.

Acknowledgements

Sham Kakade acknowledges funding from the Office of Naval Research under award N00014-22-1-2377 and the National Science Foundation Grant under award #CCF-1703574. We thank Nati Srebro for his questions regarding the removal of the dimension factor in an earlier version of this manuscript.

References

- Anthony, M. and Bartlett, P. L. *Neural network learning: Theoretical foundations*, volume 9. cambridge university press Cambridge, 1999.
- Ba, J. L., Kiros, J. R., and Hinton, G. E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bahdanau, D., Cho, K., and Bengio, Y. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- Bartlett, P., Foster, D. J., and Telgarsky, M. Spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1706.08498*, 2017.
- Bartlett, P. L. and Mendelson, S. Rademacher and gaussian complexities: Risk bounds and structural results. *Journal of Machine Learning Research*, 3:463–482, 2002. URL <http://www.jmlr.org/papers/v3/bartlett02a.html>.
- Bhattacharya, S., Ahuja, K., and Goyal, N. On the ability and limitations of transformers to recognize formal languages. *arXiv preprint arXiv:2009.11264*, 2020a.
- Bhattacharya, S., Patel, A., and Goyal, N. On the computational power of transformers and its implications in sequence modeling. *arXiv preprint arXiv:2006.09286*, 2020b.
- Brown, T. B., Mann, B., Ryder, N., Subbiah, M., Kaplan, J., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*, 2020.
- Chen, L., Lu, K., Rajeswaran, A., Lee, K., Grover, A., Laskin, M., Abbeel, P., Srinivas, A., and Mordatch, I. Decision transformer: Reinforcement learning via sequence modeling. *arXiv preprint arXiv:2106.01345*, 2021a.
- Chen, M., Li, X., and Zhao, T. On generalization bounds of a family of recurrent neural networks. *arXiv preprint arXiv:1910.12947*, 2019.
- Chen, M., Tworek, J., Jun, H., Yuan, Q., Ponde, H., Kaplan, J., Edwards, H., Burda, Y., Joseph, N., Brockman, G.,

- et al. Evaluating large language models trained on code. *arXiv preprint arXiv:2107.03374*, 2021b.
- Choromanski, K., Likhoshesterov, V., Dohan, D., Song, X., Gane, A., Sarlos, T., Hawkins, P., Davis, J., Mohiuddin, A., Kaiser, L., et al. Rethinking attention with performers. *arXiv preprint arXiv:2009.14794*, 2020.
- Clark, K., Khandelwal, U., Levy, O., and Manning, C. D. What does BERT look at? an analysis of BERT’s attention. *arXiv preprint arXiv:1906.04341*, 2019.
- Cybenko, G. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- d’Ascoli, S., Touvron, H., Leavitt, M., Morcos, A., Biroli, G., and Sagun, L. Convit: Improving vision transformers with soft convolutional inductive biases. *arXiv preprint arXiv:2103.10697*, 2021.
- Dehghani, M., Gouws, S., Vinyals, O., Uszkoreit, J., and Kaiser, Ł. Universal transformers. *arXiv preprint arXiv:1807.03819*, 2018.
- Devlin, J., Chang, M.-W., Lee, K., and Toutanova, K. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2018.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*, 2020.
- Dudley, R. M. The sizes of compact subsets of hilbert space and continuity of gaussian processes. *Journal of Functional Analysis*, 1(3):290–330, 1967.
- Elhage, N., Nanda, N., Olsson, C., Henighan, T., Joseph, N., Mann, B., Askell, A., Bai, Y., Chen, A., Conerly, T., DasSarma, N., Drain, D., Ganguli, D., Hatfield-Dodds, Z., Hernandez, D., Jones, A., Kernion, J., Lovitt, L., Ndousse, K., Amodei, D., Brown, T., Clark, J., Kaplan, J., McCandlish, S., and Olah, C. A mathematical framework for transformer circuits. *Transformer Circuits Thread*, 2021. <https://transformer-circuits.pub/2021/framework/index.html>.
- Golowich, N., Rakhlin, A., and Shamir, O. Size-independent sample complexity of neural networks. In *Conference On Learning Theory*, pp. 297–299. PMLR, 2018.
- Goyal, A., Lamb, A., Hoffmann, J., Sodhani, S., Levine, S., Bengio, Y., and Schölkopf, B. Recurrent independent mechanisms. In *International Conference on Learning Representations*, 2020.
- Goyal, A., Didolkar, A., Ke, N. R., Blundell, C., Beaudoin, P., Heess, N., Mozer, M. C., and Bengio, Y. Neural production systems. *Advances in Neural Information Processing Systems*, 34, 2021.
- Hahn, M. Theoretical limitations of self-attention in neural sequence models. *Transactions of the Association for Computational Linguistics*, 8:156–171, 2020.
- Hornik, K., Stinchcombe, M., and White, H. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- Hron, J., Bahri, Y., Sohl-Dickstein, J., and Novak, R. Infinite attention: Nngp and ntk for deep attention networks. In *International Conference on Machine Learning*, pp. 4376–4386. PMLR, 2020.
- Jaegle, A., Borgeaud, S., Alayrac, J.-B., Doersch, C., Ionescu, C., Ding, D., Koppula, S., Zoran, D., Brock, A., Shelhamer, E., et al. Perceiver io: A general architecture for structured inputs & outputs. *arXiv preprint arXiv:2107.14795*, 2021a.
- Jaegle, A., Gimeno, F., Brock, A., Zisserman, A., Vinyals, O., and Carreira, J. Perceiver: General perception with iterative attention. *arXiv preprint arXiv:2103.03206*, 2021b.
- Janner, M., Li, Q., and Levine, S. Reinforcement learning as one big sequence modeling problem. *arXiv preprint arXiv:2106.02039*, 2021.
- Jiang, Y., Neyshabur, B., Mobahi, H., Krishnan, D., and Bengio, S. Fantastic generalization measures and where to find them. *arXiv preprint arXiv:1912.02178*, 2019.
- Johnson, W. B., Lindenstrauss, J., and Schechtman, G. Extensions of lipschitz maps into banach spaces. *Israel Journal of Mathematics*, 54(2):129–138, 1986.
- Jumper, J., Evans, R., Pritzel, A., Green, T., Figurnov, M., Ronneberger, O., Tunyasuvunakool, K., Bates, R., Žídek, A., Potapenko, A., et al. Highly accurate protein structure prediction with alphafold. *Nature*, 596(7873): 583–589, 2021.
- Kearns, M. Efficient noise-tolerant learning from statistical queries. *Journal of the ACM (JACM)*, 45(6):983–1006, 1998.
- Kerg, G., Kanuparthi, B., Goyal, A., Goyette, K., Bengio, Y., and Lajoie, G. Untangling tradeoffs between recurrence and self-attention in artificial neural networks. *Advances in Neural Information Processing Systems*, 33, 2020.

- Kim, H., Papamakarios, G., and Mnih, A. The lipschitz constant of self-attention. In *International Conference on Machine Learning*, pp. 5562–5571. PMLR, 2021.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- Lee-Thorp, J., Ainslie, J., Eckstein, I., and Ontanon, S. Fnet: Mixing tokens with fourier transforms. *arXiv preprint arXiv:2105.03824*, 2021.
- Likhoshesterov, V., Choromanski, K., and Weller, A. On the expressive power of self-attention matrices. *arXiv preprint arXiv:2106.03764*, 2021.
- Liu, P. J., Saleh, M., Pot, E., Goodrich, B., Sepassi, R., Kaiser, L., and Shazeer, N. Generating wikipedia by summarizing long sequences. *arXiv preprint arXiv:1801.10198*, 2018.
- Long, P. M. and Sedghi, H. Generalization bounds for deep convolutional neural networks. *arXiv preprint arXiv:1905.12600*, 2019.
- Lu, K., Grover, A., Abbeel, P., and Mordatch, I. Pretrained transformers as universal computation engines. *arXiv preprint arXiv:2103.05247*, 2021.
- Luong, M.-T., Pham, H., and Manning, C. D. Effective approaches to attention-based neural machine translation. *arXiv preprint arXiv:1508.04025*, 2015.
- Neyshabur, B., Tomioka, R., and Srebro, N. Norm-based capacity control in neural networks. In *Conference on Learning Theory*, pp. 1376–1401. PMLR, 2015.
- Neyshabur, B., Bhojanapalli, S., and Srebro, N. A pac-bayesian approach to spectrally-normalized margin bounds for neural networks. *arXiv preprint arXiv:1707.09564*, 2017.
- Nielsen, M. A. *Neural networks and deep learning*, volume 25. Determination press San Francisco, CA, 2015.
- Olsson, C., Elhage, N., Nanda, N., Joseph, N., DasSarma, N., Henighan, T., Mann, B., Askell, A., Bai, Y., Chen, A., Conerly, T., Drain, D., Ganguli, D., Hatfield-Dodds, Z., Hernandez, D., Johnston, S., Jones, A., Kernion, J., Lovitt, L., Ndousse, K., Amodei, D., Brown, T., Clark, J., Kaplan, J., McCandlish, S., and Olah, C. In-context learning and induction heads. *Transformer Circuits Thread*, 2022. <https://transformer-circuits.pub/2022/in-context-learning-and-induction-heads/index.html>.
- Polu, S. and Sutskever, I. Generative language modeling for automated theorem proving. *arXiv preprint arXiv:2009.03393*, 2020.
- Power, A., Burda, Y., Edwards, H., Babuschkin, I., and Misra, V. Grokking: Generalization beyond overfitting on small algorithmic datasets. In *ICLR MATH-AI Workshop*, 2021.
- Radford, A., Narasimhan, K., Salimans, T., and Sutskever, I. Improving language understanding by generative pre-training. preprint, available at https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf, 2018.
- Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Rogers, A., Kovaleva, O., and Rumshisky, A. A primer in bertology: What we know about how bert works. *Transactions of the Association for Computational Linguistics*, 8:842–866, 2020.
- Siegelmann, H. T. and Sontag, E. D. On the computational power of neural nets. *Journal of computer and system sciences*, 50(1):132–150, 1995.
- Snell, C., Zhong, R., Klein, D., and Steinhardt, J. Approximating how single head attention learns. *arXiv preprint arXiv:2103.07601*, 2021.
- Tay, Y., Deghani, M., Abnar, S., Shen, Y., Bahri, D., Pham, P., Rao, J., Yang, L., Ruder, S., and Metzler, D. Long range arena: A benchmark for efficient transformers. *arXiv preprint arXiv:2011.04006*, 2020.
- Tenney, I., Das, D., and Pavlick, E. Bert rediscovers the classical nlp pipeline. *arXiv preprint arXiv:1905.05950*, 2019.
- Tolstikhin, I., Houlsby, N., Kolesnikov, A., Beyer, L., Zhai, X., Unterthiner, T., Yung, J., Keysers, D., Uszkoreit, J., Lucic, M., et al. Mlp-mixer: An all-mlp architecture for vision. *arXiv preprint arXiv:2105.01601*, 2021.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. In *Advances in neural information processing systems*, pp. 5998–6008, 2017.
- Vuckovic, J., Baratin, A., and Combes, R. T. d. A mathematical theory of attention. *arXiv preprint arXiv:2007.02876*, 2020.
- Wei, C., Chen, Y., and Ma, T. Statistically meaningful approximation: a case study on approximating turing machines with transformers. *arXiv preprint arXiv:2107.13163*, 2021.

Xu, K., Ba, J., Kiros, R., Cho, K., Courville, A., Salakhudinov, R., Zemel, R., and Bengio, Y. Show, attend and tell: Neural image caption generation with visual attention. In *International conference on machine learning*, pp. 2048–2057. PMLR, 2015.

Yang, G. Tensor programs ii: Neural tangent kernel for any architecture. *arXiv preprint arXiv:2006.14548*, 2020.

Yun, C., Bhojanapalli, S., Rawat, A. S., Reddi, S. J., and Kumar, S. Are transformers universal approximators of sequence-to-sequence functions? *arXiv preprint arXiv:1912.10077*, 2019.

Zhang, J., Karimireddy, S. P., Veit, A., Kim, S., Reddi, S. J., Kumar, S., and Sra, S. Why are adaptive methods good for attention models? *arXiv preprint arXiv:1912.03194*, 2019.

Zhang, T. Covering number bounds of certain regularized linear function classes. *Journal of Machine Learning Research*, 2(Mar):527–550, 2002.

A. Proofs of capacity bounds

In this section we present the full proofs (including the omitted proofs) of our capacity bounds. We also cover relevant background and useful technical lemmas.

A.1. Rademacher complexity and generalization bounds

Here we briefly review Rademacher complexity and its relationship to covering numbers and generalization bounds. We refer the reader to (Bartlett & Mendelson, 2002) for a more detailed exposition.

Definition A.1 (Empirical Rademacher complexity). *For a given class of functions $\mathcal{F} = \{f : \mathcal{X} \rightarrow \mathbb{R}\}$ and $\{z^{(i)} \in \mathcal{X}\}_{i=1}^m$, the empirical Rademacher complexity $\widehat{\mathcal{R}}(\mathcal{F}; z^{(1)}, \dots, z^{(m)})$ is defined as*

$$\widehat{\mathcal{R}}(\mathcal{F}; z^{(1)}, \dots, z^{(m)}) = \frac{1}{m} \mathbb{E}_\varepsilon \left[\sup_{f \in \mathcal{F}} \sum_{i=1}^m \varepsilon_i f(z^{(i)}) \right],$$

where ε is a vector of m i.i.d. Rademacher random variables ($\Pr[\varepsilon_i = 1] = \Pr[\varepsilon_i = -1] = 1/2$).

In order to relate the Rademacher complexity and ℓ_∞ -covering numbers, we use a modified version of Dudley's metric entropy.

Lemma A.2 (Dudley (1967); modified). *Consider a real-valued function class \mathcal{F} such that $|f| \leq A$ for all $f \in \mathcal{F}$. Then*

$$\widehat{\mathcal{R}}(\mathcal{F}; z^{(1)}, \dots, z^{(m)}) \leq c \cdot \inf_{\delta \geq 0} \left(\delta + \int_\delta^A \sqrt{\frac{\log \mathcal{N}_\infty(\mathcal{F}; \varepsilon; z^{(1)}, \dots, z^{(m)})}{m}} d\varepsilon \right)$$

for some constant $c > 0$.

Proof sketch. The original statement is for 2-norm covering number, but the ∞ -norm case reduces to the 2-norm case because $N_2(\cdot) \leq N_\infty(\cdot)$. The original statement also fixes $\delta = 0$ rather than taking an infimum. Also, the standard statement has the integral going from 0 to ∞ , but these are easily replaced with δ and A . \square

For our paper, we will instantiate the above lemma for log covering numbers scaling as $1/\varepsilon^2$.

Corollary A.3 (Rademacher complexity via covering number). *Consider a real-valued function class \mathcal{F} such that $|f| \leq A$ for all $f \in \mathcal{F}$. Suppose $\log \mathcal{N}_\infty(\mathcal{F}; \varepsilon; z^{(1)}, \dots, z^{(m)}) \leq C_{\mathcal{F}}/\varepsilon^2$, then*

$$\widehat{\mathcal{R}}(\mathcal{F}; z^{(1)}, \dots, z^{(m)}) \leq c \cdot \sqrt{\frac{C_{\mathcal{F}}}{m}} \cdot \left(1 + \log \left(A\sqrt{m/C_{\mathcal{F}}} \right) \right)$$

for some constant $c > 0$.

Proof. Using Lemma A.2, we have for some constant $c > 0$,

$$\begin{aligned} \widehat{\mathcal{R}}(\mathcal{F}; z^{(1)}, \dots, z^{(m)}) &\leq c \inf_{\delta \geq 0} \left(\delta + \int_\delta^A \sqrt{\frac{\log \mathcal{N}_\infty(\mathcal{F}; \varepsilon; z^{(1)}, \dots, z^{(m)})}{m}} d\varepsilon \right) \\ &\leq c \inf_{\delta \geq 0} \left(\delta + \int_\delta^A \sqrt{\frac{C_{\mathcal{F}}}{\varepsilon^2 m}} d\varepsilon \right) \\ &= c \inf_{\delta \geq 0} \left(\delta + \sqrt{\frac{C_{\mathcal{F}}}{m}} \int_\delta^A \frac{1}{\varepsilon} d\varepsilon \right) \\ &= c \inf_{\delta \geq 0} \left(\delta + \sqrt{\frac{C_{\mathcal{F}}}{m}} \log(A/\delta) \right) \\ &= c \sqrt{\frac{C_{\mathcal{F}}}{m}} \left(1 + \log \left(A\sqrt{m/C_{\mathcal{F}}} \right) \right). \end{aligned}$$

\square

We can now obtain a generalization guarantee from the Rademacher complexity of a function class:

Theorem A.4 (Bartlett & Mendelson (2002)). *Let \mathcal{D} be a distribution over $\mathcal{X} \times \mathbb{R}$ and let $\ell : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ be a b -bounded loss function that is L -Lipschitz in its first argument. For a given function class \mathcal{F} and $f \in \mathcal{F}$, let $\text{risk}(f; \mathcal{D}) := \mathbb{E}_{(x,y) \sim \mathcal{D}}[\ell(f(x), y)]$ and $\widehat{\text{risk}}(f; (z^{(i)}, y^{(i)})_{i=1}^m) := \frac{1}{m} \sum_{i=1}^m \ell(f(z^{(i)}), y^{(i)})$. Then for any $\delta > 0$, with probability at least $1 - \delta$, simultaneously for all $f \in \mathcal{F}$,*

$$\left| \text{risk}(f; \mathcal{D}) - \widehat{\text{risk}}\left(f; (z^{(i)}, y^{(i)})_{i=1}^m\right) \right| \leq 4L \widehat{\mathcal{R}}\left(\mathcal{F}; z^{(1)}, \dots, z^{(m)}\right) + 2b \sqrt{\frac{\log(1/\delta)}{2m}}.$$

Combining the above, we get:

Lemma A.5 (Lemma 2.2 (restated)). *Consider a function class \mathcal{F} such that $|f| \leq A$ for all $f \in \mathcal{F}$ and $\log \mathcal{N}_\infty(\mathcal{F}; \varepsilon; x^{(1)}, \dots, x^{(m)}) \leq C_{\mathcal{F}}/\varepsilon^2$ for all $x^{(1)}, \dots, x^{(m)} \in \mathcal{X}^m$. Then for any $\delta > 0$, with probability at least $1 - \delta$, simultaneously for all $f \in \mathcal{F}$,*

$$\left| \text{risk}(f; \mathcal{D}) - \widehat{\text{risk}}\left(f; (x^{(i)}, y^{(i)})_{i=1}^m\right) \right| \leq 4cL \sqrt{\frac{C_{\mathcal{F}}}{m}} \left(1 + \log\left(A \sqrt{m/C_{\mathcal{F}}}\right)\right) + 2b \sqrt{\frac{\log(1/\delta)}{2m}},$$

for some constant $c > 0$.

A.2. Useful lemmas

Lemma A.6. *Consider function $f : \mathbb{R}^d \rightarrow \Delta^{d-1}$ such that the Jacobian of the function satisfies $\|J f(\theta)\|_{1,1} \leq c_f$ for all $\theta \in \mathbb{R}^d$, then for any vectors $\theta_1, \theta_2 \in \mathbb{R}^d$,*

$$\|f(\theta_1) - f(\theta_2)\|_1 \leq c_f \|\theta_1 - \theta_2\|_\infty.$$

Proof. By the fundamental theorem of calculus applied to $g(t) = f(t\theta_1 + (1-t)\theta_2)$, followed by a change of variables:

$$f(\theta_1) - f(\theta_2) = \left(\int_0^1 J(t\theta_1 + (1-t)\theta_2) dt \right) (\theta_1 - \theta_2),$$

We have

$$\|f(\theta_1) - f(\theta_2)\|_1 = \left\| \int_0^1 J(t\theta_1 + (1-t)\theta_2) (\theta_1 - \theta_2) dt \right\|_1$$

By Jensen's inequality:

$$\leq \int_0^1 \|J(t\theta_1 + (1-t)\theta_2) (\theta_1 - \theta_2)\|_1 dt$$

Using $\|Ax\|_1 \leq \|A\|_{1,1} \|x\|_\infty$:

$$\leq \int_0^1 \|J(t\theta_1 + (1-t)\theta_2)\|_{1,1} \|\theta_1 - \theta_2\|_\infty dt$$

By assumption on the Jacobian:

$$\leq c_f \|\theta_1 - \theta_2\|_\infty.$$

□

Corollary A.7. *For vectors $\theta_1, \theta_2 \in \mathbb{R}^p$, $\|\text{softmax}(\theta_1) - \text{softmax}(\theta_2)\|_1 \leq 2\|\theta_1 - \theta_2\|_\infty$.*

Proof. Observe that for softmax, the Jacobian satisfies:

$$J(\theta) = \text{diag}(\text{softmax}(\theta)) - \text{softmax}(\theta)\text{softmax}(\theta)^\top.$$

We have for all θ, h ,

$$\|J(\theta)\|_{1,1} = \sum_{i=1}^p \sum_{j=1}^p |\text{softmax}(\theta)_i (\mathbb{1}[i=j] - \text{softmax}(\theta)_j)|$$

$$\begin{aligned}
 &= \sum_{i=1}^p \text{softmax}(\theta)_i \left(1 - \text{softmax}(\theta)_i + \sum_{j \neq i} \text{softmax}(\theta)_j \right) \\
 &= 2 \sum_{i=1}^p \text{softmax}(\theta)_i (1 - \text{softmax}(\theta)_i) \\
 &\leq 2.
 \end{aligned}$$

Combining the above with Lemma A.6 gives the desired result. \square

Lemma A.8. For $\alpha_i, \beta_i \geq 0$, the solution to the following optimization

$$\begin{aligned}
 &\min_{x_1, \dots, x_n} \sum_{i=1}^n \frac{\alpha_i}{x_i^2} \\
 &\text{subject to } \sum_{i=1}^n \beta_i x_i = C
 \end{aligned}$$

is $\frac{\gamma^3}{C^2}$ and is achieved at $x_i = \frac{C}{\gamma} \left(\frac{\alpha_i}{\beta_i} \right)^{1/3}$ where $\gamma = \sum_{i=1}^n \alpha_i^{1/3} \beta_i^{2/3}$.

Proof. The proof follows by a standard Lagrangian analysis. \square

Lemma A.9 (Contractivity of Π_{norm}). Let Π_{norm} be the projection operator onto the unit norm ball. For any vectors u, v , we have $\|\Pi_{\text{norm}}(u) - \Pi_{\text{norm}}(v)\| \leq \|u - v\|$.

Proof. If u, v are both in the unit ball then this follows trivially. Let us assume that $\|u\| \geq \|v\|$ and $\|u\| \geq 1$ WLOG. First suppose $\|v\| \leq 1$. Let $B_V^{(1)} = \alpha u$ be the projection of v in the direction of u , and let $B_V^2 = v - B_V^{(1)}$. Then

$$\begin{aligned}
 \|\Pi_{\text{norm}}(u) - \Pi_{\text{norm}}(v)\|^2 &= \|u/\|u\| - v\|^2 \\
 &= \|u/\|u\| - (\alpha u + B_V^2)\|^2 \\
 &= \|(\|u\|^{-1} - \alpha)u - B_V^2\|^2 \\
 &= (\|u\|^{-1} - \alpha)^2 \|u\|^2 + \|B_V^2\|^2 \\
 &\leq (1 - \alpha^2) \|u\|^2 + \|B_V^2\|^2 && \text{since } \|u\|^{-1} < \alpha < 1 \\
 &= \|u - (\alpha u + B_V^2)\|^2 \\
 &= \|u - v\|^2
 \end{aligned}$$

If $\|v\| > 1$, then

$$\|\Pi_{\text{norm}}(u) - \Pi_{\text{norm}}(v)\| = \|\Pi_{\text{norm}}(u/\|v\|) - \Pi_{\text{norm}}(v/\|v\|)\| \leq \|u/\|v\| - v/\|v\|\| < \|u - v\|.$$

where the second-to-last inequality follows from the $\|v\| < 1$ case. \square

Lemma A.10 (Zhang (2002), Theorem 4). Let $\mathcal{V} : \{v : v \in \mathbb{R}^{d_1}, \|v\| \leq B_1\}$ and $\mathcal{F}_{\text{linear}} = \{x \mapsto v^\top x : v \in \mathcal{V}\}$. For any $\delta > 0$ and $x^{(1)}, \dots, x^{(N)}$ satisfying $\|x^{(i)}\| \leq B_2 \forall i$,

$$\log \mathcal{N}_\infty(\mathcal{F}_{\text{linear}}; \varepsilon; x^{(1)}, \dots, x^{(N)}) \leq 36 \frac{B_1^2 B_2^2}{\varepsilon^2} \log(2[4B_1 B_2 / \varepsilon + 2]N + 1).$$

A.3. Pseudo-dimension lower bound

Because the number of parameters in a Transformer self-attention head is $O(d^2)$, with no dependence on T , one might guess that the capacity of the class of these heads does not need to grow with the context length T . But parameter counting can be misleading—for example, the class $\{x \in \mathbb{R} \mapsto \text{sign}(\sin(\alpha x))\}_{\alpha \in \mathbb{R}}$ has a single parameter but infinite VC-dimension.

We observe that, when the weight norms are unbounded, the pseudo-dimension of the class of Transformer self-attention heads does grow with T , at least logarithmically, even when the embedding dimension is as small as 3.

Recall that a Transformer self-attention head is of the form

$$f_{\text{tf-head}}(X; W_V, W_{QK}) := \sigma(W_V^\top X^\top \text{softmax}(XW_{QK}^\top x_\tau))$$

Let $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ be any activation function that satisfies the following condition: there is a constant $c \in \mathbb{R}$ such that for $a > \frac{1}{2}$ and $b < \frac{1}{2}$, we have $\sigma(a) > c$, $\sigma(b) < c$.

The *pseudo-dimension* of a concept class \mathcal{F} is defined as

$$\text{Pdim}(\mathcal{F}) := \max_{c \in \mathbb{R}} \text{VCdim}(\{h(x) = \text{sign}(f(x) - c) : f \in \mathcal{F}\})$$

Proof of Proposition 4.4. For the purposes of this proof, we will treat x_τ as a fixed vector—so either it is not treated as part of the input, or it is set to the same value for all of the inputs. Thus, $W_{QK}^\top x_\tau$ is a fixed vector, which we call w_{QK} . Moreover, we will set W_V to be a $d \times 1$ matrix, so we will treat it as a vector $w_V \in \mathbb{R}^d$. Thus, we will be dealing with the following restricted concept class of unbounded Transformer self-attention heads:

$$f_{\text{tf-head}}(X; w_V, w_{QK}) := \sigma(w_V^\top X^\top \text{softmax}(Xw_{QK}))$$

for $w_V, w_{QK} \in \mathbb{R}^d$.

For simplicity, we consider the case where T is a power of 2. We will construct a set of $\log T$ inputs $\{X^{(i)}\}_{i=1}^{\log T}$ in \mathbb{R}^T that are shattered by $\mathcal{F}_{\text{tf-head-unbounded}}$.

In particular, indexing $t = 0, \dots, T - 1$, let

$$X_t^{(i)} := (\cos(2\pi t/T), \sin(2\pi t/T), \text{bin}(t)_i)$$

where $\text{bin}(t)_i$ is the i th bit of the binary expansion of the integer t (padded with 0s in the front such that the expansion is length $\log T$). We can think of the first two coordinates as a (fixed) “positional encoding” which allows for the attention mechanism to select a single position, and the third coordinate as the “token embedding”. The token embedding is designed such that for each binary vector of length $\log T$, there is a position t such that the vector of values at that position $(X_t^{(i)}[3])_{i=1}^{\log T}$ corresponds to the vector, thus inducing a shattering.

For $s = 0, \dots, T - 1$, let

$$w_{QK}^{(s)} := (T^2 \cos(2\pi s/T), T^2 \sin(2\pi s/T), 0).$$

These attention weights are of sufficiently large magnitude that they “pick out” a single coordinate. Also, let

$$w_V^{(s)} := (0, 0, 1).$$

Then we claim that the set

Observe that

$$w_{QK}^{(s)\top} x_t^{(i)} = T^2 \cos(2\pi(s - t)/T),$$

so we have

$$\text{softmax}(Xw_{QK}^{(s)})[t] = \frac{\exp(T^2 \cos(2\pi(s - t)/T))}{\sum_{\tau=0}^{T-1} \exp(T^2 \cos(2\pi(s - \tau)/T))} = \frac{\exp(T^2 \cos(2\pi(s - t)/T))}{\sum_{\tau=0}^{T-1} \exp(T^2 \cos(2\pi\tau/T))}$$

Let us bound the denominator using the fact that $\cos(\theta) \leq 1 - \frac{2}{\pi^2}\theta^2$ for $\theta \in [0, \pi]$:

$$\sum_{\tau=0}^{T-1} \exp(T^2 \cos(2\pi\tau/T)) \leq \exp(T^2) + 2 \sum_{\tau=1}^{\lfloor (T-1)/2 \rfloor} \exp(T^2 \cos(2\pi\tau/T))$$

$$\begin{aligned}
 &\leq \exp(T^2) + 2 \sum_{\tau=0}^{\lfloor (T-1)/2 \rfloor} \exp\left(T^2 \left(1 - \frac{2}{\pi^2} (2\pi\tau/T)^2\right)\right) \\
 &= e^{T^2} + 2e^{T^2} \sum_{\tau=1}^{\lfloor (T-1)/2 \rfloor} e^{-8\tau^2} \\
 &\leq e^{T^2} + 2e^{T^2} \int_{\rho=0}^{\lfloor (T-1)/2 \rfloor} e^{-8\rho^2} d\rho \\
 &\leq e^{T^2} + 2e^{T^2} \int_{\rho=0}^{\infty} e^{-8\rho^2} d\rho \\
 &= e^{T^2} + \sqrt{\frac{\pi}{8}} e^{T^2}
 \end{aligned}$$

Hence, for each s ,

$$\text{softmax}(X w_{QK}^{(s)})[s] \geq \frac{1}{1 + \sqrt{\frac{\pi}{8}}} > \frac{1}{2}$$

and

$$\sum_{t \neq s} \text{softmax}(X w_{QK}^{(s)})[t] \leq 1 - \frac{1}{1 + \sqrt{\frac{\pi}{8}}} < \frac{1}{2}$$

Then we claim that the set $\{f_{\text{tf-head}}(\cdot; w_V^{(s)}, w_{QK}^{(s)})\}_{s=0}^{T-1}$ shatters $\{X^{(i)}\}_{i=1}^{\log T}$:

$$\begin{aligned}
 f_{\text{tf-head}}(X^{(i)}; w_V^{(s)}, w_{QK}^{(s)}) &= \sigma\left(w_V^{(s)\top} X^\top \text{softmax}\left(X^{(i)} w_{QK}^{(s)}\right)\right) \\
 &= \sigma\left(w_V^{(s)\top} \sum_{t=0}^{T-1} \text{softmax}\left(X^{(i)} w_{QK}^{(s)}\right)[t] \cdot \text{bin}(t)_i\right)
 \end{aligned}$$

which is greater than c if $\text{bin}(s)_i = 1$ and is less than c if $\text{bin}(s)_i = 0$, since the $t = s$ term in the sum dominates all the other terms. Thus, the different choices of s induce a shattering. \square

A.4. Covering number upper bounds

Proof of Lemma 4.3. Observe that,

$$\begin{aligned}
 &\left\| f_{\text{head}}(X, z; \theta_s, \theta_{\text{in}}) - f_{\text{head}}(X, z; \hat{\theta}_s, \hat{\theta}_{\text{in}}) \right\| \\
 &= \left\| \phi_{\text{out}}\left(\phi_{\text{in}}(X; \theta_{\text{in}})^\top \text{Norm}(\text{Score}(X, z; \theta_s))\right) - \phi_{\text{out}}\left(\phi_{\text{in}}(X; \hat{\theta}_{\text{in}})^\top \text{Norm}(\text{Score}(X, z; \hat{\theta}_s))\right) \right\|
 \end{aligned}$$

By L_{out} -Lipschitzness of ϕ_{out} and bound on $\|w\|$:

$$\leq L_{\text{out}} \left\| \phi_{\text{in}}(X; \theta_{\text{in}})^\top \text{Norm}(\text{Score}(X, z; \theta_s)) - \phi_{\text{in}}(X; \hat{\theta}_{\text{in}})^\top \text{Norm}(\text{Score}(X, z; \hat{\theta}_s)) \right\|$$

By triangle inequality:

$$\begin{aligned}
 &\leq L_{\text{out}} \left\| \phi_{\text{in}}(X; \theta_{\text{in}})^\top \left(\text{Norm}(\text{Score}(X, z; \theta_s)) - \text{Norm}(\text{Score}(X, z; \hat{\theta}_s)) \right) \right\| \\
 &\quad + L_{\text{out}} \left\| \left(\phi_{\text{in}}(X; \theta_{\text{in}}) - \phi_{\text{in}}(X; \hat{\theta}_{\text{in}}) \right)^\top \text{Norm}(\text{Score}(X, z; \hat{\theta}_s)) \right\|
 \end{aligned}$$

Using $\|Pv\| \leq \|P\|_{2,\infty} \|v\|_1$ and B_{in} -boundedness of ϕ_{in} :

$$\leq L_{\text{out}} B_{\text{in}} \left\| \text{Norm}(\text{Score}(X, z; \theta_s)) - \text{Norm}(\text{Score}(X, z; \hat{\theta}_s)) \right\|_1$$

$$+ L_{\text{out}} \left\| \left(\phi_{\text{in}}(X; \theta_{\text{in}}) - \phi_{\text{in}}(X; \widehat{\theta}_{\text{in}}) \right)^\top \right\|_{2, \infty} \left\| \text{Norm}(\text{Score}(X, z; \widehat{\theta}_s)) \right\|_1$$

By Lemma A.6 and the assumption on Norm:

$$\leq L_{\text{out}} C_{\text{Norm}} \left\| \phi_{\text{in}}(X; \theta_{\text{in}})^\top \right\|_{2, \infty} \left\| \text{Score}(X, z; \theta_s) - \text{Score}(X, z; \widehat{\theta}_s) \right\|_\infty + L_{\text{out}} \left\| \left(\phi_{\text{in}}(X; \theta_{\text{in}}) - \phi_{\text{in}}(X; \widehat{\theta}_{\text{in}}) \right)^\top \right\|_{2, \infty}$$

By boundedness of ϕ_{in} and $\|X^\top\|_{2, \infty} \leq B_X$:

$$\leq L_{\text{out}} C_{\text{Norm}} B_{\text{in}} B_X \left\| \text{Score}(X, z; \theta_s) - \text{Score}(X, z; \widehat{\theta}_s) \right\|_\infty + L_{\text{out}} \left\| \left(\phi_{\text{in}}(X; \theta_{\text{in}}) - \phi_{\text{in}}(X; \widehat{\theta}_{\text{in}}) \right)^\top \right\|_{2, \infty}.$$

□

Proof of Theorem 4.2. Our goal is to show that for every $\varepsilon > 0$, collection of inputs $(X^{(1)}, z^{(1)}), \dots, (X^{(m)}, z^{(m)})$, there is a cover $\mathcal{C}_{\text{head}}$ such that for all $\theta_s \in \Theta_s, \theta_{\text{in}} \in \Theta_{\text{in}}$, there is some $(\widehat{\theta}_s, \widehat{\theta}_{\text{in}}) \in \mathcal{C}_{\text{head}}$ such that $\max_i \left\| f_{\text{head}}(X^{(i)}, z^{(i)}; \theta_s, \theta_{\text{in}}) - f_{\text{head}}(X^{(i)}, z^{(i)}; \widehat{\theta}_s, \widehat{\theta}_{\text{in}}) \right\| \leq \varepsilon$.

Observe that for all $\theta_s, \widehat{\theta}_s$,

$$\max_{i \in [m]} \left\| \text{Score}(X^{(i)}, z^{(i)}; \theta_s) - \text{Score}(X^{(i)}, z^{(i)}; \widehat{\theta}_s) \right\|_\infty = \max_{i \in [m], t \in [T]} \left| \text{Score}(x_t^{(i)}, z^{(i)}; \theta_s) - \text{Score}(x_t^{(i)}, z^{(i)}; \widehat{\theta}_s) \right|.$$

Similarly, for all $\theta_{\text{in}}, \widehat{\theta}_{\text{in}}$,

$$\max_{i \in [m]} \left\| \left(\phi_{\text{in}}(X^{(i)}; \theta_{\text{in}}) - \phi_{\text{in}}(X^{(i)}; \widehat{\theta}_{\text{in}}) \right)^\top \right\|_{2, \infty} = \max_{i \in [m], t \in [T]} \left\| \phi_{\text{in}}(x_t^{(i)}; \theta_{\text{in}}) - \phi_{\text{in}}(x_t^{(i)}; \widehat{\theta}_{\text{in}}) \right\|.$$

This crucially allows us to aggregate over the i and t dimensions together.⁴ Therefore, we can consider \mathcal{N}_∞ covers for the above to bound the overall covering number.

Let $\mathcal{C}_{\text{Score}}$ be the $\varepsilon_{\text{Score}}$ -cover (∞) for $\mathcal{F}_{\text{Score}}$ over inputs $\{(x_t^{(i)}, z^{(i)})\}_{i \in [m], t \in [T]}$ of size

$$\mathcal{N}_\infty \left(\mathcal{F}_{\text{Score}}; \varepsilon_{\text{Score}}; \{(x_t^{(i)}, z^{(i)})\}_{i \in [m], t \in [T]} \right).$$

Also, Let \mathcal{C}_{in} be the ε_{in} -cover (∞) for \mathcal{F}_{in} over inputs $\{x_t^{(i)}\}_{i \in [m], t \in [T]}$ of size

$$\mathcal{N}_\infty \left(\mathcal{F}_{\text{in}}; \varepsilon_{\text{in}}; \{x_t^{(i)}\}_{i \in [m], t \in [T]}; \|\cdot\|_2 \right).$$

We are ready to construct the cover for $\mathcal{F}_{\text{head}}$. Set $\mathcal{C}_{\text{head}} = \{f_{\text{head}}(\cdot; \widehat{\theta}_s, \widehat{\theta}_{\text{in}})\}_{i \in [m]} : \widehat{\theta}_s \in \mathcal{C}_{\text{Score}}, \widehat{\theta}_{\text{in}} \in \mathcal{C}_{\text{in}}\}$. Then for any $\theta_s \in \Theta_s, \theta_{\text{in}} \in \Theta_{\text{in}}$, there exists $\widehat{\theta}_s, \widehat{\theta}_{\text{in}} \in \mathcal{C}_{\text{head}}$, such that for all $i \in [m]$, using Lemma 4.3:

$$\left\| f_{\text{head}}(X^{(i)}, z^{(i)}; \theta_s, \theta_{\text{in}}) - f_{\text{head}}(X^{(i)}, z^{(i)}; \widehat{\theta}_s, \widehat{\theta}_{\text{in}}) \right\| \leq C_{\text{Norm}} L_{\text{out}} B_{\text{in}} B_X \varepsilon_{\text{Score}} + L_{\text{out}} \varepsilon_{\text{in}}.$$

The size of the cover we have constructed is,

$$\begin{aligned} \log |\mathcal{C}_{\text{head}}| &= \log |\mathcal{C}_{\text{Score}}| + \log |\mathcal{C}_{\text{in}}| \\ &= \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{Score}}; \varepsilon_{\text{Score}}; \{(x_t^{(i)}, z^{(i)})\}_{i \in [m], t \in [T]} \right) + \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{in}}; \varepsilon_{\text{in}}; \{x_t^{(i)}\}_{i \in [m], t \in [T]}; \|\cdot\|_2 \right) \end{aligned}$$

and we are done. □

⁴In the case of the Transformer self-attention mechanism, we will obtain ∞ -norm covering numbers for Score and ϕ_{in} that have only logarithmic dependence on the number of examples. Because of this aggregation trick, the resulting covering number for the whole layer will have merely logarithmic dependence on the context length T .

Proof of Corollary 4.5. By Theorem 4.2, the covering number of $\mathcal{F}_{\text{tf-head}}$ satisfies

$$\begin{aligned} & \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{tf-head}}; \varepsilon; \left\{ (X^{(i)}, z^{(i)}) \right\}_{i=1}^m \right) \\ & \leq \inf_{\alpha \in [0,1]} \left[\log \mathcal{N}_\infty \left(\mathcal{F}_{QK}; \frac{\alpha \varepsilon}{2L_\sigma B_V B_X}; \{(x_t^{(i)}, z^{(i)})\}_{i \in [m], t \in [T]} \right) \right. \\ & \quad \left. + \log \mathcal{N}_\infty \left(\mathcal{F}_V; \frac{(1-\alpha)\varepsilon}{L_\sigma}; \{x_t^{(i)}\}_{i \in [m], t \in [T]}; \|\cdot\|_2 \right) \right]. \end{aligned}$$

where we have used the fact that for a scalar-output Transformer layer:

- softmax satisfies the Jacobian assumption with $C_{\text{softmax}} = 2$ using Corollary A.7.
- L_{out} is the Lipschitz constant of σ : L_σ .
- B_{in} is a bound on the norm of $W_V^\top x$ with respect to norm of x : B_V .

By Lemma 4.6, for any $\varepsilon_{QK}, \varepsilon_V > 0$:

$$\begin{aligned} \log \mathcal{N}_\infty \left(\mathcal{F}_{QK}; \varepsilon_{QK}; \{(x_t^{(i)}, z^{(i)})\}_{i \in [m], t \in [T]} \right) & \lesssim \frac{(B_{QK}^{2,1} B_X)^2 \log(dmT)}{\varepsilon_{QK}^2} \\ \log \mathcal{N}_\infty \left(\mathcal{F}_V; \varepsilon_V; \{(x_t^{(i)}, z^{(i)})\}_{i \in [m], t \in [T]}; \|\cdot\|_2 \right) & \lesssim \frac{(B_V^{2,1} B_X)^2 \log(dmT)}{\varepsilon_V^2} \end{aligned}$$

since $W_{QK}, W_V \in \mathbb{R}^{d \times d}$ ($k = d$). We want to choose ε_{QK} and ε_V to minimize the sum of the above two terms, subject to

$$2L_\sigma B_V B_X \varepsilon_{QK} + L_\sigma \varepsilon_V \leq \varepsilon.$$

By Lemma A.8, the solution to this optimization leads to an optimal bound of:

$$\log \mathcal{N}_\infty(\mathcal{F}_{\text{tf-head}}; \varepsilon; X^{(1)}, \dots, X^{(M)}) \lesssim (L_\sigma B_X)^2 \cdot \frac{\left((B_V^{2,1})^{\frac{2}{3}} + (B_{QK}^{2,1} B_V B_X)^{\frac{2}{3}} \right)^3}{\varepsilon^2} \cdot \log(dmT).$$

□

Proof of Lemma 4.6. Our approach will be to construct a cover by decomposing the problem into two separate cover problems, (1) ℓ_2 -cover over the possible norms of the rows of W , and (2) ℓ_∞ -cover of the set \mathcal{W} constrained to the norms dictated by the first cover. More formally, let us define:

$$\begin{aligned} \mathcal{W}_{:,2} & = \left\{ \left[\begin{array}{c} \|w_1\| \\ \vdots \\ \|w_{d_1}\| \end{array} \right] : W = \begin{bmatrix} w_1 \\ \vdots \\ w_{d_1} \end{bmatrix} \in \mathcal{W} \right\} = \{v \in \mathbb{R}^{d_1} : \|v\|_1 \leq B_W\}, \text{ and} \\ \mathcal{F}_v & = \left\{ x \rightarrow Wx : W = \begin{bmatrix} w_1 \\ \vdots \\ w_{d_1} \end{bmatrix} \in \mathbb{R}^{d_1 \times d_2}, \forall i \|w_i\| \leq v_i \right\}. \end{aligned}$$

Denote $\mathcal{C}_\mathcal{W}$ to be the ε_1 -cover (in terms of ℓ_2 -norm) for $\mathcal{W}_{:,2}$ ⁵, and for any $v \in \mathbb{R}^{d_1}$, denote $\mathcal{C}_v := \mathcal{N}_\infty(\mathcal{F}_v, \varepsilon_2; x^{(1)}, \dots, x^{(N)}; \|\cdot\|_2)$. We will set ε_1 and ε_2 later.

⁵Here the cover is for a set and not a function.

We will first show that $\mathcal{C} := \{f : v \in \mathcal{C}_W, f \in \mathcal{F}_v\}$ is an $(\varepsilon_2 + B_X \varepsilon_1)$ -cover (in terms of ℓ_∞) for \mathcal{F} . Consider $f \in \mathcal{F}$

parameterized by some $W = \begin{bmatrix} w_1 \\ \vdots \\ w_{d_1} \end{bmatrix}$. Since $W \in \mathcal{W}$, we know that there is a $\bar{v} \in \mathcal{C}_W$, such that,

$$\sqrt{\sum_{i=1}^{d_1} (\|w_i\| - \bar{v}_i)^2} \leq \varepsilon_1.$$

Define $\bar{W} = \begin{bmatrix} \bar{v}_1 \cdot \frac{w_1}{\|w_1\|} \\ \vdots \\ \bar{v}_{d_1} \cdot \frac{w_{d_1}}{\|w_{d_1}\|} \end{bmatrix}$. Then there exists $f \in \mathcal{F}_{\bar{v}} \subseteq \mathcal{C}$ such that

$$\max_{i \in [N]} \|\bar{W}x^{(i)} - f(x^{(i)})\| \leq \varepsilon_2.$$

Now we have, for all $i \in [N]$,

$$\begin{aligned} \|Wx^{(i)} - f(x^{(i)})\| &\leq \|\bar{W}x^{(i)} - f(x^{(i)})\| + \|(\bar{W} - W)x^{(i)}\| \\ &\leq \varepsilon_2 + \sqrt{\sum_{j=1}^{d_1} (\bar{v}_j - \|w_j\|)^2 \left(\frac{w_j \cdot x^{(i)}}{\|w_j\|}\right)^2} \\ &\leq \varepsilon_2 + B_X \sqrt{\sum_{j=1}^{d_1} (\bar{v}_j - \|w_j\|)^2} \\ &\leq \varepsilon_2 + B_X \varepsilon_1. \end{aligned}$$

Thus, we get the desired cover. Note that the size of the cover satisfies

$$\log |\mathcal{C}| \leq \log |\mathcal{C}_W| + \max_{v \in \mathcal{C}_W} \log |\mathcal{C}_v|.$$

Now we need to construct the sub-covers and bound the size of \mathcal{C} . Let us first construct the cover \mathcal{C}_W . Using Maurey's sparsification (see Theorem 3 in (Zhang, 2002)), we can find a proper cover of \mathcal{C}_W which satisfies,

$$\log |\mathcal{C}_W| \leq \frac{B_W^2}{\varepsilon_1^2} \log(2d_1 + 1).$$

Let us now construct \mathcal{C}_v for a fixed v . The approach will be to cover each of the rows of W independently, treating each as specifying a linear function from $\mathbb{R}^{d_2} \rightarrow \mathbb{R}$. By Lemma A.10, letting $\mathcal{Z}(b) := \{z : z \in \mathbb{R}^{d_2}, \|z\| \leq b\}$ and $\mathcal{F}_{\text{linear}}(b) = \{x \mapsto z^\top x : z \in \mathcal{Z}(b)\}$, for any $\delta > 0$

$$\log \mathcal{N}_\infty(\mathcal{F}_{\text{linear}}(b); \delta; x^{(1)}, \dots, x^{(N)}) \leq \frac{cb^2 B_X^2 \log((1 + bB_X/\delta)N)}{\delta^2}.$$

In fact the cover, which we denote by $\bar{\mathcal{F}}_{\text{linear}}(b; \delta)$, is proper: $\bar{\mathcal{F}}_{\text{linear}}(b; \delta) = \{x \mapsto \bar{z}^\top x : \bar{z} \in \bar{\mathcal{Z}}\}$ for some finite subset $\bar{\mathcal{Z}} \subset \mathcal{Z}(b)$. Then the cover for the matrix can be constructed as,

$$\mathcal{C}_v = \left\{ x \mapsto \bar{Z}x : \bar{Z} = \begin{bmatrix} \bar{z}_1 \\ \vdots \\ \bar{z}_{d_1} \end{bmatrix} : \bar{z}_i \in \bar{\mathcal{F}}_{\text{linear}}\left(v_i; \varepsilon_2 \sqrt{\frac{v_i}{\|v\|_1}}\right) \text{ for all } i \right\}.$$

Observe that this forms a cover for \mathcal{F}_v . For any $f \in \mathcal{F}_v$ parameterized by W , let \bar{w}_i be the closest element in the corresponding row covers, then we have

$$\|Wx^{(i)} - \bar{W}x^{(i)}\| = \sqrt{\sum_{j=1}^{d_2} (w_j^\top x^{(i)} - \bar{w}_j^\top x^{(i)})^2} \leq \sqrt{\sum_{j=1}^{d_2} \varepsilon_2^2 \cdot \frac{v_j}{\|v\|_1}} = \varepsilon_2.$$

□

Note that the size of \mathcal{C}_v can be bounded as,

$$\begin{aligned} \log |\mathcal{C}_v| &= \sum_{i=1}^{d_2} \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{linear}}(v_i); \varepsilon_2 \sqrt{\frac{v_i}{\|v\|_1}}; x^{(1)}, \dots, x^{(N)} \right) \\ &\leq \sum_{i=1}^{d_2} \frac{c v_i^2 \|v\|_1 B_X^2 \log((1 + \|v\|_1 B_X / \varepsilon_2) N)}{\varepsilon_2^2 v_i} \\ &= \frac{c \|v\|_1^2 B_X^2 \log((1 + \|v\|_1 B_X / \varepsilon_2) N)}{\varepsilon_2^2} \leq \frac{c B_W^2 B_X^2 \log((1 + B_W B_X / \varepsilon_2) N)}{\varepsilon_2^2}. \end{aligned}$$

Combining the above and setting $\varepsilon_1, \varepsilon_2$ appropriately, we get

$$\log |\mathcal{C}| \leq \frac{B_W^2 B_X^2}{\varepsilon^2} \log(d_1 N).$$

A.5. Capacity with positional embeddings

Since the Transformer architecture is permutation invariant for all $t \neq \tau$, positional embeddings (fixed or trainable) are typically added to the inputs to distinguish the different positions of the tokens. These positional embeddings are matrices $P \in \mathbb{R}^{T \times d}$ such that $P = [p_1 \dots p_T]^\top$ for $p_i \in \mathbb{R}^d$. Accounting for the positional embeddings as input, a single Transformer attention head can be expressed as:

$$f_{\text{tf-pos}}(X, P; W_V, W_{QK}) := \sigma \left(W_V^\top (X + P)^\top \text{softmax} \left((X + P) W_{QK}^\top (x_\tau + p_\tau) \right) \right).$$

For a fixed positional embedding P , let us define

$$\mathcal{F}_{\text{tf-pos}}(P) := \{X \rightarrow f_{\text{tf-pos}}(X, P; W_V, W_{QK}) : \|W_V\|_{2,1} \leq B_V^{2,1}, \|W_V\| \leq B_V, \|W_{QK}^\top\|_{2,1} \leq B_{QK}^{2,1}\}$$

. Position embedding just impacts the input into the covering bound argument which effects the bound in terms of the $\|P^\top\|_{2,\infty}$ as given below,

Lemma A.11. *For all $X^{(1)}, \dots, X^{(m)} \in \mathbb{R}^{T \times d}$ such that $\|X^{(i)\top}\|_{2,\infty} \leq B_X$ for all $i \in [m]$, and $P \in \mathbb{R}^{T \times d}$ such that $\|P^\top\|_{2,\infty} \leq B_P$, the covering number of $\mathcal{F}_{\text{tf-pos}}(P)$ satisfies*

$$\log \mathcal{N}_\infty(\mathcal{F}_{\text{tf-pos}}(P); \varepsilon; X^{(1)}, \dots, X^{(m)}, \|\cdot\|_2) \lesssim (L_\sigma(B_X + B_P))^2 \cdot \frac{\left((B_V^{2,1})^{\frac{2}{3}} + (2B_{QK}^{2,1} B_V (B_X + B_P))^{\frac{2}{3}} \right)^3}{\varepsilon^2} \cdot \log(dmT).$$

Proof. Observe that $f_{\text{tf-pos}}(X, P; W_V, W_{QK}) = f_{\text{tf-head}}(X + P; W_V, W_{QK})$. Thus we have,

$$\log \mathcal{N}_\infty \left(\mathcal{F}_{\text{tf-pos}}(P); \varepsilon; \left\{ X^{(i)} \right\}_{i=1}^m, \|\cdot\|_2 \right) = \log \mathcal{N}_\infty \left(\mathcal{F}_{\text{tf-head}}; \varepsilon; \left\{ X^{(i)} + P \right\}_{i=1}^m, \|\cdot\|_2 \right).$$

For all $i \in [m]$, $\|(X^{(i)} + P)^\top\|_{2,\infty} \leq \|X^{(i)\top}\|_{2,\infty} + \|P^\top\|_{2,\infty} \leq B_X + B_P$. Therefore, using Corollary 4.5, we get the desired result. \square

Therefore our bounds go through for fixed positional embeddings. If we were to train the embeddings, we would need a much finer cover on the embeddings which could incur a T dependence.

A.6. Capacity of multiple parallel heads

In virtually all practical applications of Transformers since their inception, instead of using one set of weights for an attention head, there are parallel attention heads, which have separate identically-shaped parameters; their outputs are concatenated. For the purposes of this analysis, suppose we have

$$f_{\text{tf-heads}} \left(X; \left\{ W_V^{[h]}, W_{QK}^{[h]} \right\}_{h=1}^H \right) := \sum_{h=1}^H f_{\text{tf-head}} \left(X; W_V^{[h]}, W_{QK}^{[h]} \right).$$

Let us define the class of multi-head self-attention with H heads as

$$\mathcal{F}_{\text{tf-heads}} := \left\{ X \mapsto f_{\text{tf-heads}} \left(X; \left\{ W_V^{[h]}, W_{QK}^{[h]} \right\}_{h=1}^H \right) : \right. \\ \left. \forall h \in [H], \left\| W_V^{[h]} \right\|_{2,1} \leq B_V^{2,1[h]}, \left\| W_V^{[h]} \right\| \leq B_V^{[h]}, \left\| W_{QK}^{[h]} \right\|_{2,1} \leq B_{QK}^{2,1[h]} \right\}.$$

Lemma A.12. For all $X^{(1)}, \dots, X^{(m)} \in \mathbb{R}^{T \times d}$ such that $\left\| X^{(i)} \right\|_{2,\infty} \leq B_X$ for all $i \in [m]$, the covering number of $\mathcal{F}_{\text{tf-heads}}$ satisfies

$$\log \mathcal{N}_\infty(\mathcal{F}_{\text{tf-heads}}; \varepsilon; X^{(1)}, \dots, X^{(m)}, \|\cdot\|_2) \lesssim (L_\sigma B_X)^2 \cdot \frac{\left(\sum_{h=1}^H (B_V^{2,1[h]})^{\frac{2}{3}} + (2B_{QK}^{2,1[h]} B_V^{[h]})^{\frac{2}{3}} \right)^3}{\varepsilon^2} \cdot \log(dmT).$$

Proof. For all $h \in [H]$, let \mathcal{C}_h be an ε_h -covering of $\mathcal{F}_{\text{tf-head}}$ with weight bounds corresponding to head h . Since $f_{\text{tf-heads}} \left(X; \left\{ W_V^{[h]}, W_{QK}^{[h]} \right\}_{h=1}^H \right) = \sum_{h=1}^H f_{\text{tf-head}} \left(X; W_V^{[h]}, W_{QK}^{[h]} \right)$, we have $\mathcal{C} := \mathcal{C}_1 \times \dots \times \mathcal{C}_H$ ⁶ is an $\left(\sum_{h=1}^H \varepsilon_h \right)$ -covering for $\mathcal{F}_{\text{tf-heads}}$. Using Corollary 4.5 (and optimizing for ε_h using Lemma A.8, by breaking them into individual errors for each head), we have

$$\log |\mathcal{C}| = \sum_{h=1}^H \log |\mathcal{C}_h| \leq \sum_{h=1}^H \leq (L_\sigma B_X)^2 \cdot \frac{\left(\sum_{h=1}^H (B_V^{2,1[h]})^{\frac{2}{3}} + (2B_{QK}^{2,1[h]} B_V^{[h]})^{\frac{2}{3}} \right)^3}{\varepsilon^2} \cdot \log(dmT).$$

□

To see the dependence on H , consider the setting where the weight bounds are the same for each head (dropping the $[h]$ subscript), then we get,

$$\log \mathcal{N}_\infty(\mathcal{F}_{\text{tf-heads}}; \varepsilon; X^{(1)}, \dots, X^{(m)}, \|\cdot\|_2) \lesssim (L_\sigma B_X)^2 \cdot H^3 \cdot \frac{\left((B_V^{2,1})^{\frac{2}{3}} + (2B_{QK}^{2,1} B_V)^{\frac{2}{3}} \right)^3}{\varepsilon^2} \cdot \log(dmT).$$

A.7. Capacity of multi-layer Transformers

This section analyzes the capacity of an L -layer Transformer. Let us denote the weights of layer i by $W^{(i)} := \{W_Q^{(i)}, W_K^{(i)}, W_V^{(i)}, W_C^{(i)}\}$ such that $\left\| W_K^{(i)} W_Q^{(i)\top} \right\|_2 \leq B_{QK}^{(i)}$, $\left\| W_V^{(i)} \right\|_2 \leq B_V^{(i)}$, $\left\| W_C^{(i)} \right\|_2 \leq B_C^{(i)}$ and $\left\| W_K^{(i)\top} W_Q^{(i)} \right\|_{2,1} \leq B_{QK}^{2,1(i)}$, $\left\| W_V^{(i)} \right\|_{2,1} \leq B_V^{2,1(i)}$ and $\left\| W_C^{(i)} \right\|_{2,1} \leq B_C^{2,1(i)}$. Let us further denote the set of weights up to layer i by $W^{1:i} = (W^{(1)}, \dots, W^{(i)})$. Let the input representation of layer i be $g_{\text{tf-head}}^{(i)}(X; W^{1:i})$. We inductively define g with $g_{\text{tf-head}}^{(1)}(X; W^{1:1}) = X$

$$g_{\text{tf-head}}^{(i+1)}(X; W^{1:i+1}) = \Pi_{\text{norm}} \left(\sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; W^{1:i}); W^{(i)} \right) \right) \right) W_C^{(i)} \right) \text{ with} \\ f(Z; \{W_Q, W_K, W_V, W_C\}) = \text{RowSoftmax} \left(ZW_Q (ZW_K)^\top \right) ZW_V,$$

where Π_{norm} is applied row-wise. Our final output is $g_{\text{tf-scalar}}(X; W^{1:L+1}, w) = w^\top g_{\text{tf-head}}^{(L)}(X; W^{1:L+1})$ [CLS] for $\|w\| \leq B_w$.

In order to construct a cover, we will first bound the distance between the function g with different weight parameters $W^{1:L+1}$ and $\widehat{W}^{1:L+1}$. This bound will depend on the closeness of the parameters which will allow us to construct a cover of the network in an iterative fashion by constructing covers of each layer.

⁶Here, \times denotes the Cartesian product: the functions obtained by using the every combination of parameters of each individual cover.

A.7.1. LIPSCHITZNESS OF THE NETWORK

To bound the Lipschitzness of the network, we will first bound the distance between f with different weights and inputs.

Lemma A.13 (Instantiation of Lemma 4.3). *For any $W_K, \widehat{W}_K, W_V, \widehat{W}_V, W_Q, \widehat{W}_Q \in \mathbb{R}^{d \times k}$, for all $Z \in \mathbb{R}^{T \times d}$ such that $\|Z^\top\|_{2,\infty} \leq 1$,*

$$\begin{aligned} & \left\| \left(f(Z; \{W_Q, W_K, W_V, \cdot\}) - f(Z; \{\widehat{W}_Q, \widehat{W}_K, \widehat{W}_V, \cdot\}) \right)^\top \right\|_{2,\infty} \\ & \leq 2 \|W_V\|_2 \left\| \left(W_Q W_K^\top - \widehat{W}_Q \widehat{W}_K^\top \right) Z^\top \right\|_{2,\infty} + \left\| (W_V - \widehat{W}_V)^\top Z^\top \right\|_{2,\infty} \end{aligned}$$

Proof. Consider a fixed row τ of the output of the functions,

$$\begin{aligned} & \left\| f(Z; \{W_Q, W_K, W_V, \cdot\})[\tau] - f(Z; \{\widehat{W}_Q, \widehat{W}_K, \widehat{W}_V, \cdot\})[\tau] \right\| \\ & = \left\| W_V^\top Z^\top \text{softmax}(Z W_K W_Q^\top z_\tau) - \widehat{W}_V^\top Z^\top \text{softmax}(Z \widehat{W}_K \widehat{W}_Q^\top z_\tau) \right\| \end{aligned}$$

By triangle inequality:

$$\begin{aligned} & \leq \left\| W_V^\top Z^\top \left(\text{softmax}(Z W_K W_Q^\top z_\tau) - \text{softmax}(Z \widehat{W}_K \widehat{W}_Q^\top z_\tau) \right) \right\| \\ & \quad + \left\| (W_V - \widehat{W}_V)^\top Z^\top \text{softmax}(Z \widehat{W}_K \widehat{W}_Q^\top z_\tau) \right\| \end{aligned}$$

Using $\|Pv\| \leq \|P\|_{2,\infty} \|v\|_1$:

$$\begin{aligned} & \leq \|W_V^\top Z^\top\|_{2,\infty} \left\| \text{softmax}(Z W_K W_Q^\top z_\tau) - \text{softmax}(Z \widehat{W}_K \widehat{W}_Q^\top z_\tau) \right\|_1 \\ & \quad + \left\| (W_V - \widehat{W}_V)^\top Z^\top \right\|_{2,\infty} \left\| \text{softmax}(Z \widehat{W}_K \widehat{W}_Q^\top z_\tau) \right\|_1 \end{aligned}$$

By Corollary A.7, $\|Z^\top\|_{2,\infty} \leq 1$, $\|PQ\|_{2,\infty} \leq \|P\|_2 \|Q\|_{2,\infty}$, and $\|P^\top\|_2 = \|P\|_2$:

$$\begin{aligned} & \leq 2 \|W_V\|_2 \left\| Z W_K W_Q^\top z_\tau - Z \widehat{W}_K \widehat{W}_Q^\top z_\tau \right\|_\infty + \left\| (W_V - \widehat{W}_V)^\top Z^\top \right\|_{2,\infty} \\ & \leq 2 \|W_V\|_2 \left\| \left(W_Q W_K^\top - \widehat{W}_Q \widehat{W}_K^\top \right) Z^\top \right\|_{2,\infty} + \left\| (W_V - \widehat{W}_V)^\top Z^\top \right\|_{2,\infty}. \end{aligned}$$

□

Lemma A.14. *For any $W_K, W_V, W_Q \in \mathbb{R}^{d \times k}$, for all $Z, \widehat{Z} \in \mathbb{R}^{T \times d}$ such that $\|Z^\top\|_{2,\infty} \leq 1$, $\|\widehat{Z}^\top\|_{2,\infty} \leq 1$,*

$$\begin{aligned} & \left\| \left(f(Z; \{W_Q, W_K, W_V, \cdot\}) - f(\widehat{Z}; \{W_Q, W_K, W_V, \cdot\}) \right)^\top \right\|_{2,\infty} \\ & \leq \|W_V\|_2 \left(1 + 4 \|W_K W_Q^\top\|_2 \right) \left\| (Z - \widehat{Z})^\top \right\|_{2,\infty}. \end{aligned}$$

Proof. Consider a fixed row τ of the output of the functions,

$$\begin{aligned} & \left\| f(Z; \{W_Q, W_K, W_V, \cdot\})[\tau] - f(\widehat{Z}; \{W_Q, W_K, W_V, \cdot\})[\tau] \right\| \\ & = \left\| W_V^\top Z^\top \text{softmax}(Z W_K W_Q^\top z_\tau) - W_V^\top \widehat{Z}^\top \text{softmax}(\widehat{Z} W_K W_Q^\top \widehat{z}_\tau) \right\| \end{aligned}$$

By triangle inequality:

$$\leq \left\| W_V^\top (Z - \widehat{Z})^\top \text{softmax}(Z W_K W_Q^\top z_\tau) \right\| + \left\| W_V^\top \widehat{Z}^\top \left(\text{softmax}(Z W_K W_Q^\top z_\tau) - \text{softmax}(\widehat{Z} W_K W_Q^\top \widehat{z}_\tau) \right) \right\|$$

Using $\|Pv\| \leq \|P\|_{2,\infty} \|v\|_1$:

$$\begin{aligned} & \leq \left\| W_V^\top (Z - \widehat{Z})^\top \right\|_{2,\infty} \left\| \text{softmax}(Z W_K W_Q^\top z_\tau) \right\|_1 \\ & \quad + \left\| W_V^\top \widehat{Z}^\top \right\|_{2,\infty} \left\| \text{softmax}(Z W_K W_Q^\top z_\tau) - \text{softmax}(\widehat{Z} W_K W_Q^\top \widehat{z}_\tau) \right\|_1 \end{aligned}$$

By Corollary A.7, $\|\widehat{Z}^\top\|_{2,\infty} \leq 1$ and $\|PQ\|_{2,\infty} \leq \|P\|_2\|Q\|_{2,\infty}$:

$$\leq \|W_V\|_2 \left\| (Z - \widehat{Z})^\top \right\|_{2,\infty} + 2 \|W_V\|_2 \left\| ZW_KW_Q^\top z_\tau - \widehat{Z}W_KW_Q^\top \widehat{z}_\tau \right\|_\infty$$

By triangle inequality:

$$\leq \|W_V\|_2 \left\| (Z - \widehat{Z})^\top \right\|_{2,\infty} + 2 \|W_V\|_2 \left(\left\| (Z - \widehat{Z})W_KW_Q^\top z_\tau \right\|_\infty + \left\| \widehat{Z}W_KW_Q^\top (z_\tau - \widehat{z}_\tau) \right\|_\infty \right)$$

Since $\|\widehat{Z}^\top\|_{2,\infty} \leq 1$ and $\|Pv\|_\infty \leq \|P^\top\|_{2,\infty}\|v\|$:

$$\leq \|W_V\|_2 \left(1 + 4 \|W_KW_Q^\top\|_2 \right) \left\| (Z - \widehat{Z})^\top \right\|_{2,\infty}.$$

□

With the above lemmas, we are ready to prove the effect of change of weights on g .

Lemma A.15. For any $W_1^{i+1}, \widehat{W}_1^{i+1}$ satisfying the norm constraints,

$$\begin{aligned} & \left\| \left(g_{\text{tf-block}}^{(i+1)}(X; W^{1:i+1}) - g_{\text{tf-block}}^{(i+1)}(X; \widehat{W}^{1:i+1}) \right)^\top \right\|_{2,\infty} \\ & \leq \left\| \left(W_C^{(i)} - \widehat{W}_C^{(i)} \right)^\top \sigma \left(\Pi_{\text{norm}} \left(f \left((X; \widehat{W}^{1:i}); \widehat{W}^{(i)} \right) \right) \right)^\top \right\|_{2,\infty} \\ & \quad + L_\sigma B_C^{(i)} B_V^{(i)} \left(1 + 4B_{QK}^{(i)} \right) \left\| \left(g_{\text{tf-block}}^{(i)}(X; W^{1:i}) - g_{\text{tf-block}}^{(i)}(X; \widehat{W}^{1:i}) \right)^\top \right\|_{2,\infty} \\ & \quad + 2L_\sigma B_C^{(i)} B_V^{(i)} \left\| \left(W_Q^{(i)} W_K^{(i)\top} - \widehat{W}_Q^{(i)} \widehat{W}_K^{(i)\top} \right) g_{\text{tf-block}}^{(i)}(X; \widehat{W}^{1:i})^\top \right\|_{2,\infty} \\ & \quad + L_\sigma B_C^{(i)} \left\| (W_V - \widehat{W}_V)^\top g_{\text{tf-block}}^{(i)}(X; \widehat{W}^{1:i})^\top \right\|_{2,\infty}. \end{aligned}$$

Proof. Unrolling one layer, we have

$$\begin{aligned} & \left\| \left(g_{\text{tf-head}}^{(i+1)}(X; W^{1:i+1}) - g_{\text{tf-head}}^{(i+1)}(X; \widehat{W}^{1:i+1}) \right)^\top \right\|_{2,\infty} \\ & = \left\| \left(\Pi_{\text{norm}} \left(\sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; W^{1:i}); W^{(i)} \right) \right) \right) \right) W_C^{(i)} \right) \right. \\ & \quad \left. - \Pi_{\text{norm}} \left(\sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; \widehat{W}^{1:i}); \widehat{W}^{(i)} \right) \right) \right) \widehat{W}_C^{(i)} \right) \right)^\top \right\|_{2,\infty} \end{aligned}$$

Using Lemma A.9 for each row:

$$\leq \left\| W_C^{(i)\top} \sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; W^{1:i}); W^{(i)} \right) \right) \right)^\top - \widehat{W}_C^{(i)\top} \sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; \widehat{W}^{1:i}); \widehat{W}^{(i)} \right) \right) \right)^\top \right\|_{2,\infty}$$

By triangle inequality for each row:

$$\begin{aligned} & \leq \underbrace{\left\| W_C^{(i)\top} \left(\sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; W^{1:i}); W^{(i)} \right) \right) \right) - \sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; \widehat{W}^{1:i}); \widehat{W}^{(i)} \right) \right) \right) \right)^\top \right\|_{2,\infty}}_{(A)} \\ & \quad + \left\| \left(W_C^{(i)} - \widehat{W}_C^{(i)} \right)^\top \sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; \widehat{W}^{1:i}); \widehat{W}^{(i)} \right) \right) \right)^\top \right\|_{2,\infty}. \end{aligned}$$

Let us focus on term (A).

Bounding the norm per row:

$$(A) \leq \left\| W_C^{(i)} \right\|_2 \left\| \sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; W^{1:i}); W^{(i)} \right) \right) \right)^\top - \sigma \left(\Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)}(X; \widehat{W}^{1:i}); \widehat{W}^{(i)} \right) \right) \right)^\top \right\|_{2,\infty}$$

Since σ is L_σ -Lipschitz and $\|W_C^{(i)}\|_2 \leq B_C^{(i)}$, for each row:

$$\leq L_\sigma B_C^{(i)} \left\| \Pi_{\text{norm}} \left(f \left(g_{\text{tf-head}}^{(i)} \left(X; W^{1:i}; W^{(i)} \right) \right)^\top - \Pi_{\text{norm}} \left(f \left(g \left(X; \widehat{W}^{1:i}; \widehat{W}^{(i)} \right) \right)^\top \right) \right\|_{2,\infty}$$

Using Lemma A.9 for each row:

$$\leq L_\sigma B_C^{(i)} \left\| f \left(g_{\text{tf-head}}^{(i)} \left(X; W^{1:i}; W^{(i)} \right) \right)^\top - f \left(g_{\text{tf-head}}^{(i)} \left(X; \widehat{W}^{1:i}; \widehat{W}^{(i)} \right) \right)^\top \right\|_{2,\infty}$$

By triangle inequality:

$$\begin{aligned} &\leq L_\sigma B_C^{(i)} \left\| f \left(g_{\text{tf-head}}^{(i)} \left(X; W^{1:i}; W^{(i)} \right) \right)^\top - f \left(g_{\text{tf-head}}^{(i)} \left(X; \widehat{W}^{1:i}; W^{(i)} \right) \right)^\top \right\|_{2,\infty} \\ &\quad + L_\sigma B_C^{(i)} \left\| f \left(g_{\text{tf-head}}^{(i)} \left(X; \widehat{W}^{1:i}; W^{(i)} \right) \right)^\top - f \left(g_{\text{tf-head}}^{(i)} \left(X; \widehat{W}^{1:i}; \widehat{W}^{(i)} \right) \right)^\top \right\|_{2,\infty} \end{aligned}$$

By Lemma A.13 and A.14 and norm bounds on the matrices:

$$\begin{aligned} &\leq L_\sigma B_C^{(i)} B_V^{(i)} \left(1 + 4B_{QK}^{(i)} \right) \left\| g_{\text{tf-head}}^{(i)} \left(X; W^{1:i} \right)^\top - g \left(X; \widehat{W}^{1:i} \right)^\top \right\|_{2,\infty} \\ &\quad + 2L_\sigma B_C^{(i)} B_V^{(i)} \left\| \left(W_Q^{(i)} W_K^{(i)\top} - \widehat{W}_Q^{(i)} \widehat{W}_K^{(i)\top} \right) g_{\text{tf-head}}^{(i)} \left(X; \widehat{W}^{1:i} \right)^\top \right\|_{2,\infty} \\ &\quad + L_\sigma B_C^{(i)} \left\| \left(W_V - \widehat{W}_V \right)^\top g_{\text{tf-head}}^{(i)} \left(X; \widehat{W}^{1:i} \right)^\top \right\|_{2,\infty}. \end{aligned}$$

Combining the above gives us the desired result. \square

Lastly, we take account of the last linear weight and observe that,

Lemma A.16. For any $W^{1:L+1}, \widehat{W}^{1:L+1}$ and w, \widehat{w} ,

$$\begin{aligned} &\left| g_{\text{tf-scalar}} \left(X; W^{1:L+1}, w \right) - g_{\text{tf-scalar}} \left(X; \widehat{W}^{1:L+1}, \widehat{w} \right) \right| \\ &\leq \|w\| \left\| g_{\text{tf-block}}^{(L+1)} \left(X; W^{1:L+1} \right)_{[\text{CLS}]} - g_{\text{tf-block}}^{(L+1)} \left(X; \widehat{W}^{1:L+1} \right)_{[\text{CLS}]} \right\| + \left| (w - \widehat{w})^\top g_{\text{tf-block}}^{(L+1)} \left(X; \widehat{W}^{1:L+1} \right)_{[\text{CLS}]} \right|. \end{aligned}$$

Proof. Observe that,

$$\begin{aligned} &\left| g_{\text{tf-scalar}} \left(X; W^{1:L+1}, w \right) - g_{\text{tf-scalar}} \left(X; \widehat{W}^{1:L+1}, \widehat{w} \right) \right| \\ &= \left| w^\top g_{\text{tf-block}}^{(L+1)} \left(X; W^{1:L+1} \right)_{[\text{CLS}]} - \widehat{w}^\top g_{\text{tf-block}}^{(L+1)} \left(X; \widehat{W}^{1:L+1} \right)_{[\text{CLS}]} \right| \end{aligned}$$

By triangle inequality:

$$\leq \left| w^\top \left(g_{\text{tf-block}}^{(L+1)} \left(X; W^{1:L+1} \right)_{[\text{CLS}]} - g_{\text{tf-block}}^{(L+1)} \left(X; \widehat{W}^{1:L+1} \right)_{[\text{CLS}]} \right) \right| + \left| (w - \widehat{w})^\top g_{\text{tf-block}}^{(L+1)} \left(X; \widehat{W}^{1:L+1} \right)_{[\text{CLS}]} \right|$$

Bounding the inner product by norms:

$$\leq \|w\| \left\| g_{\text{tf-block}}^{(L+1)} \left(X; W^{1:L+1} \right)_{[\text{CLS}]} - g_{\text{tf-block}}^{(L+1)} \left(X; \widehat{W}^{1:L+1} \right)_{[\text{CLS}]} \right\| + \left| (w - \widehat{w})^\top g_{\text{tf-block}}^{(L+1)} \left(X; \widehat{W}^{1:L+1} \right)_{[\text{CLS}]} \right|.$$

\square

A.7.2. CONSTRUCTING THE COVER

The cover construction follows the standard recipe of composing covers per layer (as in (Bartlett et al., 2017)).

Theorem A.17. Let $\mathcal{F}_{\text{tf-scalar}}^{(L)}$ represent the class of functions of L -layer Transformer blocks satisfying the norm bounds (specified before) followed by linear layer on the [CLS] token. Then, for all $X^{(i)}$

$$\log \mathcal{N}_{\infty}(\mathcal{F}_{\text{tf-scalar}}^{(L)}; \varepsilon; X^{(1)}, \dots, X^{(m)}, \|\cdot\|_2) \lesssim \frac{\log(dmT)}{\varepsilon^2} \times \left(B_w^{\frac{2}{3}} + \sum_{i=1}^L \alpha_i^{\frac{2}{3}} \left(B_C^{2,1(i)\frac{2}{3}} + d^{\frac{2}{3}} \left(2L_{\sigma} B_C^{(i)} B_V^{(i)} B_{QK}^{2,1(i)\frac{2}{3}} + k^{\frac{2}{3}} \left(L_{\sigma} B_C^{(i)} B_V^{2,1(i)\frac{2}{3}} \right) \right) \right) \right)^3$$

where $\alpha_i = \prod_{j < i} L_{\sigma} B_C^{(j)} B_V^{(j)} (1 + 4B_{QK}^{(j)})$.

Proof. Our goal is to show that for every $\varepsilon > 0$, and collection of inputs $X^{(1)}, \dots, X^{(m)}$, there is a cover \mathcal{C} of vectors in $\mathbb{R}^{(m)}$ such that for all $W^{1:L+1}$ and w satisfying the norm bounds, there is some $v \in \mathcal{C}$ such that $\max_i |g_{\text{tf-scalar}}(X^{(i)}; W^{1:L+1}, w) - v| \leq \varepsilon$.

In each layer of the transformer, $W_Q^{(i)}$ and $W_K^{(i)}$ always appear together in the form $W_K^{(i)} W_Q^{(i)\top}$. Therefore, we will overload notation and define $W_{QK}^{(i)} : W_K^{(i)} W_Q^{(i)\top}$. Our cover \mathcal{C} will be proper, consisting of vectors of the form $(g_{\text{tf-scalar}}(X^{(i)}; \widehat{W}^{1:L+1}, \widehat{w}))_{i \in [m]}$. We will build the cover iteratively by finding finite collections of matrices $\widehat{W}^{1:i}$ for each layer.

First observe that for any collection of $Z^{(1)}, \dots, Z^{(m)} \in \mathbb{R}^{T \times d_1}$, and any $W, \widehat{W} \in \mathbb{R}^{d_1 \times d_2}$,

$$\max_{i \in [m]} \left\| W^{\top} Z^{(i)\top} - \widehat{W}^{\top} Z^{(i)\top} \right\|_{2, \infty} = \max_{i \in [m], t \in [T]} \left\| W^{\top} z_t^{(i)} - \widehat{W}^{\top} z_t^{(i)} \right\|.$$

This crucially allows us to aggregate over the samples and context length. In particular, we can apply Lemma 4.6 with the input vectors $(z_t^{(i)})_{i \in [m], t \in [T]}$; a total of mT input vectors. Specifically, for any ε and $\mathcal{W}(d_1, d_2, \alpha) := \{W \in \mathbb{R}^{d_1 \times d_2} \mid \|W\|_{2,1} \leq \alpha\}$ with fixed $Z^{(i)}$ satisfying $\|Z^{(i)\top}\|_{2, \infty} \leq 1$, Lemma 4.6 gives us such a cover.

First let us build a cover for one Transformer layer with inputs $Z^{(1)}, \dots, Z^{(m)}$. We will begin with creating an ε_V -cover $\widehat{\mathcal{W}}_V$ for the function class of linear transformations given by $\mathcal{W}_V : \{W \in \mathbb{R}^{d \times k}, \|W\|_{2,1} \leq \alpha, \|W\|_2 \leq s\}$ and ε_{QK} -cover $\widehat{\mathcal{W}}_{QK}$ for $\mathcal{W}_{QK} := \{W \in \mathbb{R}^{d \times d}, \|W^{\top}\|_{2,1} \leq \beta, \|W\|_2 \leq r\}$ and inputs $Z^{(1)}, \dots, Z^{(m)}$. For each pair of $\widehat{W}_V \in \widehat{\mathcal{W}}_V$ and $\widehat{W}_{QK} \in \widehat{\mathcal{W}}_{QK}$, we construct an ε_C -cover $\widehat{\mathcal{W}}_C(\widehat{W}_V, \widehat{W}_{QK})$ for $\mathcal{W}_C : \{W \in \mathbb{R}^{k \times d}, \|W\|_{2,1} \leq \gamma, \|W\|_2 \leq c\}$ and inputs $\left\{ \sigma \left(\Pi_{\text{norm}} \left(f \left(Z^{(i)}; \widehat{W}_V, \widehat{W}_{QK} \right) \right) \right) \right\}_{i=1}^m$. Our final cover is

$$\widehat{\mathcal{W}} := \left\{ (\widehat{W}_V, \widehat{W}_{QK}, \widehat{W}_C) : \widehat{W}_V \in \widehat{\mathcal{W}}_V, \widehat{W}_{QK} \in \widehat{\mathcal{W}}_{QK}, \widehat{W}_C \in \widehat{\mathcal{W}}_C(\widehat{W}_V, \widehat{W}_{QK}) \right\}.$$

Using Lemma A.15, we can show that $\widehat{\mathcal{W}}$ is an ε -cover for $g(\cdot; \{W_V, W_{QK}, W_C\})$ and inputs $Z^{(1)}, \dots, Z^{(m)}$ where

$$\varepsilon = \varepsilon_C + 2L_{\sigma} c s \varepsilon_{QK} + L_{\sigma} c \varepsilon_V.$$

Using Lemma 4.6, the size of the cover is

$$\begin{aligned} |\widehat{\mathcal{W}}| &\leq |\widehat{\mathcal{W}}_V| |\widehat{\mathcal{W}}_{QK}| \max_{\substack{\widehat{W}_V \in \widehat{\mathcal{W}}_V \\ \widehat{W}_{QK} \in \widehat{\mathcal{W}}_{QK}}} |\widehat{\mathcal{W}}_C(\widehat{W}_V, \widehat{W}_{QK})| \\ \implies \log |\widehat{\mathcal{W}}| &\lesssim \left(\frac{\alpha^2}{\varepsilon_V^2} + \frac{\beta^2}{\varepsilon_{QK}^2} + \frac{\gamma^2}{\varepsilon_C^2} \right) \log(dmT). \end{aligned}$$

We are now ready to inductively construct a cover for the deeper network. Suppose we have a $\varepsilon^{(i)}$ -cover $\widehat{\mathcal{W}}^{1:i}$ for $g(\cdot; W^{1:i})$ on $X^{(1)}, \dots, X^{(m)}$. We show how to construct an $\varepsilon^{(i+1)}$ -cover for $g(\cdot; W^{1:i+1})$. For every element $\widehat{W}^{1:i} \in \widehat{\mathcal{W}}^{1:i}$

we construct a $(\varepsilon_C^{(i)} + 2L_\sigma B_C^{(i)} B_V^{(i)} \varepsilon_{QK}^{(i)} + L_\sigma B_C^{(i)} \varepsilon_V^{(i)})$ -cover $\widehat{\mathcal{W}}_i(\widehat{W}^{1:i})$ for the transformer layer (as above) on inputs $\{g(X^{(j)}; \widehat{W}^{1:i})\}_{j=1}^m$. Consider the cover

$$\widehat{\mathcal{W}}^{1:i+1} := \left\{ (\widehat{W}^{1:i}, \widehat{W}^{(i)}) : \widehat{W}^{1:i} \in \widehat{\mathcal{W}}^{1:i}, \widehat{W}^{(i)} \in \widehat{\mathcal{W}}_i(\widehat{W}^{1:i}) \right\}.$$

By Lemma A.15, this gives,

$$\varepsilon^{(i+1)} = L_\sigma B_C^{(i)} B_V^{(i)} (1 + 4B_{QK}^{(i)}) \varepsilon_C^{(i)} + \varepsilon_C^{(i)} + 2L_\sigma B_C^{(i)} B_V^{(i)} \varepsilon_{QK}^{(i)} + L_\sigma B_C^{(i)} \varepsilon_V^{(i)}.$$

The size of the cover is

$$|\widehat{\mathcal{W}}^{1:i+1}| \leq |\widehat{\mathcal{W}}^{1:i}| \max_{\widehat{W}^{1:i} \in \widehat{\mathcal{W}}^{1:i}} |\widehat{\mathcal{W}}_i(\widehat{W}^{1:i})|.$$

Inductively applying this, we get

$$\begin{aligned} \varepsilon^{(L+1)} &= \sum_{i=1}^L \left(\prod_{j<i} L_\sigma B_C^{(j)} B_V^{(j)} (1 + 4B_{QK}^{(j)}) \right) \left(\varepsilon_C^{(i)} + 2L_\sigma B_C^{(i)} B_V^{(i)} \varepsilon_{QK}^{(i)} + L_\sigma B_C^{(i)} \varepsilon_V^{(i)} \right) \\ &= \sum_{i=1}^L \alpha_i \left(\varepsilon_C^{(i)} + 2L_\sigma B_C^{(i)} B_V^{(i)} \varepsilon_{QK}^{(i)} + L_\sigma B_C^{(i)} \varepsilon_V^{(i)} \right) \end{aligned}$$

where $\alpha_i = \prod_{j<i} L_\sigma B_C^{(j)} B_V^{(j)} (1 + 4B_{QK}^{(j)})$.

The size of the cover is

$$\log \left(|\widehat{\mathcal{W}}^{1:L+1}| \right) \leq \sum_{i=1}^L \left(\frac{B_V^{2,1(i)^2}}{\varepsilon_V^{(i)^2} + \frac{B_{QK}^{2,1(i)^2}}{\varepsilon_{QK}^{(i)^2} + \frac{B_C^{2,1(i)^2}}{\varepsilon_C^{(i)^2}} \right) \log(dmT).$$

Notice that the layer-norm maintains the norm bound on the inputs. Lastly, we need to cover the linear layer on the [CLS] token and compose it with the cover of $g^{1:L}$ (as before). Using Lemma A.10 and A.16, we can get the final ε -cover \mathcal{C} with

$$\varepsilon = B_w \sum_{i=1}^L \alpha_i \left(\varepsilon_C^{(i)} + 2L_\sigma B_C^{(i)} B_V^{(i)} \varepsilon_{QK}^{(i)} + L_\sigma B_C^{(i)} \varepsilon_V^{(i)} \right) + \varepsilon_w$$

and size

$$\log |\mathcal{C}| \lesssim \frac{B_w^2 \log(m)}{\varepsilon_w^2} + \sum_{i=1}^L \left(\frac{B_V^{2,1(i)^2}}{\varepsilon_V^{(i)^2} + \frac{B_{QK}^{2,1(i)^2}}{\varepsilon_{QK}^{(i)^2} + \frac{B_C^{2,1(i)^2}}{\varepsilon_C^{(i)^2}} \right) \log(dmT).$$

Using Lemma A.8, the size of the cover for fixed ε gives us the desired result. \square

B. Sparse function representation via bounded-norm Transformers

B.1. Setup

Reductions from Boolean functions to Transformers. In order to establish our function approximation results, we must first define a canonical mapping between length- T Boolean strings $b \in \{0, 1\}^T$ and Transformer inputs $X \in \mathbb{R}^{T \times d}$. The key point (which has also been considered since the inception of the Transformer (Vaswani et al., 2017), and continues to be a crucial consideration in practice (Dosovitskiy et al., 2020)) is that the network's permutation-equivariant symmetry needs to be broken by assigning different embeddings to different indices of b . There are several possible natural choices here, which are all of practical interest:

- *Deterministic positional embeddings.* Fix positional embedding matrices $P \in \mathbb{R}^{T \times d}$, $E \in \mathbb{R}^{\{0,1\} \times d}$, and a special direction $v_{[\text{CLS}]} \in \mathbb{R}^d$, such that the $T+3$ vectors $\{P_{t,:}\}_{t=1}^T \cup \{E_{j,:}\}_{j \in \{0,1\}} \cup \{v_{[\text{CLS}]}\}$ are an approximately orthonormal basis for \mathbb{R}^d (see below). The input to the Transformer is then $X = E_b + P$, where $E_b \in \mathbb{R}^{T \times d}$ such that $[E_b]_{t,:} = E_{b_t,:}$ for each $t \in [T]$. In the $f_{\text{tf-scalar}}$ formulation, we choose the auxiliary input $x_{[\text{CLS}]}$ to be the constant vector $v_{[\text{CLS}]}$. This closely matches applications of Transformers in NLP (Vaswani et al., 2017).

- *Trainable positional embeddings.* Like the above, but P is a trainable parameter; we still require approximate orthogonality of $\{E_{j,:}\}_{j \in 0,1} \cup \{v_{[\text{CLS}]}\}$. It is also possible to consider the case where E and $v_{[\text{CLS}]}$ are trainable (matching the way token embeddings are trained in practice). This becomes important in the regime of large vocabulary sizes that require embeddings to capture shared information between tokens; however, this is not necessary for our constructions, as we limit our consideration to binary tokens. This simplifies our constructions and improves statistical rates; additionally, it is a popular and well-studied alternative (Vaswani et al., 2017; Devlin et al., 2018; Radford et al., 2018; 2019; Brown et al., 2020).
- *Bag of vectors.* Fix a matrix $V \in \mathbb{R}^{T \times d}$ with approximately orthogonal rows (like the deterministic P), but choose the Transformer input

$$X := V \text{diag}(b).$$

This construction replaces positional embeddings with positional “indicator vectors” which can be swapped between any of the Transformer’s input positions. It has the advantage of being symmetric with respect to permutation of the Transformer’s input positions: it turns out that

$$f_{\text{tf-scalar}}(V \text{diag}(b)) = f_{\text{tf-scalar}}(V \Pi \text{diag}(b)),$$

for any $T \times T$ permutation matrix Π . It is also the most natural construction when considering the composition of sparse Boolean functions across multiple layers: a layer can output combinations of the basis rows v_i for further function composition, like Boolean gates.

Approximately orthonormal basis. Each of the Boolean function approximation constructions will rely on a basis set of vectors, which will be used as positional embeddings (or the variable indices in the bag-of-vectors construction). We will fix a set of approximately orthonormal vectors $\{v_i : \|v_i\| = 1\}_{i=1}^{T'}$ in \mathbb{R}^d : for each $i \neq j$, we have $|v_i^\top v_j| \leq \Delta$. When $\Delta = 0$, the maximal T' for which such a set exists is d ; for $\Delta \in (0, \frac{1}{2})$, the Johnson-Lindenstrauss lemma (Johnson et al., 1986) implies that the maximal set is of size $\exp(\Theta(d\Delta^2))$. For given choices of d, Δ and a maximal $\{v_1, \dots, v_{T'}\}$, our construction is valid for contexts of length $T \leq T'$. For the special vectors $e_0, e_1, v_{[\text{CLS}]}$, we will assume that these are exactly orthogonal to the v_i and each other, so that the v_i must be a basis in dimension at least $d - 3$. This is for clarity only—it reduces the number of error terms to propagate through the analysis.

Self-attention block. In each construction (which specifies an input $X \in \mathbb{R}^{T \times d}$, we will specify the parameters $W_Q, W_K, W_V, W_C, w = e_1$ of a scalar-output Transformer $f_{\text{tf-scalar}}$, which takes an input $X \in \mathbb{R}^{(T+1) \times d}$; the auxiliary token input will be the constant vector $x_{[\text{CLS}]} := v_{[\text{CLS}]} \in \mathbb{R}^d$. The internal activation function σ is chosen to be the identity. Summarizing, the functional form of $f_{\text{tf-scalar}} \in \mathcal{F}_{\text{tf-scalar}}$ in these constructions is

$$f_{\text{tf-scalar}}(X; W_Q, W_K, W_V, W_C, e_1) = \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) X W_V W_C e_1.$$

In the intermediate lemmas, it will also be useful to consider the corresponding attention head output

$$f_{\text{tf-head}}(X; W_Q, W_K, W_V, W_C) = \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) X W_V W_C,$$

and its projections $f_{\text{tf-head}} \circ \Pi_{d_{\text{proj}}}$ onto the first d_{proj} coordinates.

Feedforward networks. We establish some notation for feedforward networks. An L -layer feedforward network, with activation function $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ and dimensions d_1, \dots, d_{L+1} , is parameterized by weight matrices $W_i \in \mathbb{R}^{d_{i+1} \times d_i}$, and maps $x \in \mathbb{R}^{d_1}$ to $y \in \mathbb{R}^{d_{L+1}}$, by the iterative equations

$$\begin{aligned} y_1^\top &:= \sigma(x^\top W_1), \\ y_{i+1}^\top &:= \sigma(y_i^\top W_i), \quad i = 1, \dots, L-1, \\ f_{\text{mlp}}(x; W_1, \dots, W_L)^\top &= y^\top := y_L^\top W_L. \end{aligned}$$

When $d_{L+1} = 1$, we will use the notation w instead of W_L . It will be convenient to incorporate *bias* weights by introducing an extra input coordinate $W_i \in \mathbb{R}^{(d_{i+1}+1) \times d_i}$, and augmenting the linear function accordingly:

$$y_i^\top W_i \mapsto [y_i^\top \ 1] W_i.$$

Self-attention composed with a feedforward network. The full definition of the Transformer layer composes a self-attention layer ($f_{\text{tf-layer}} : \mathbb{R}^{T \times d} \rightarrow \mathbb{R}^{T \times d}$) with a position-wise feedforward network ($f_{\text{mlp}} : \mathbb{R}^d \rightarrow \mathbb{R}^d$). We will use this combination of modules to establish our function approximation results: $f_{\text{tf-layer}}$ acts as a sparse bottleneck, while f_{mlp} approximates an arbitrary function of the selected coordinates. For our single-layer constructions, it is most convenient to establish notation for a scalar-output Transformer with a feedforward network. To this end, define $\mathcal{F}_{\text{tf+mlp}}$ to be the function class with the same Score, Norm, ϕ_{in} functions as in $\mathcal{F}_{\text{tf-scalar}}$ (thus, the same parameters W_Q, W_K, W_V), with identity activation function, but a feedforward neural network replacing the linear ϕ_{out} and w . Concretely, with $L = 3$ and the ReLU activation function $(\cdot)_+$, $\mathcal{F}_{\text{tf+mlp}}$ contains functions of the form

$$f_{\text{tf+mlp}}(X; \theta) = ((y^\top W_1)_+ W_2)_+ w,$$

$$y = \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) X W_V W_C w,$$

with parameters $\theta := (W_Q, W_K, W_V, W_C, W_1, W_2, w)$.

Multiple self-attention heads. The final component we will need for the function approximation setup is multi-headed self-attention. We will extend the definition of the single-headed $f_{\text{tf-head}}$ to

$$f_{\text{tf-heads}} \left(X; \left\{ W_Q^{[h]}, W_K^{[h]}, W_V^{[h]}, W_C^{[h]} \right\}_{h=1}^H \right) := \sum_{h=1}^H f_{\text{tf-head}} \left(X; W_Q^{[h]}, W_K^{[h]}, W_V^{[h]}, W_C^{[h]} \right),$$

and substitute this definition into $f_{\text{tf+mlp}}$ when discussing a multi-headed construction.

Classes and properties of Boolean functions. We will call a Boolean function $f : \{0, 1\}^T \rightarrow \mathcal{Y}$ \mathcal{I} -sparse if it only depends on a fixed subset $\mathcal{I} \subseteq [T]$ of its inputs:

$$b_i = b'_i \quad \forall i \in \mathcal{I} \implies f(b) = f(b').$$

Overloading notation, if $\mathcal{I} = s$, we will also call f s -sparse. We will call an \mathcal{I} -sparse Boolean function f *symmetric* if its value is invariant under permutation of the indices in \mathcal{I} :

$$|\{i \in \mathcal{I} : b_i = 1\}| = |\{i \in \mathcal{I} : b'_i = 1\}| \implies f(b) = f(b').$$

Further, we will call an \mathcal{I} -sparse real-valued symmetric Boolean function $f : \{0, 1\}^T \rightarrow \mathcal{Y}$ *monotone* if $f(b)$ is monotonically increasing in $r := |\{i \in \mathcal{I} : b_i = 1\}|$. If, for some $\gamma > 0$, it holds that $f(r+1) \geq f(r) + \gamma$ for each $r = 0, \dots, s-1$, we call f γ -*strictly monotone*. A vector-valued \mathcal{I} -sparse Boolean function $f : \{0, 1\}^T \rightarrow \mathbb{R}^{d_f}$ is γ -*injective* if

$$\|f(b) - f(b')\|_\infty \geq \gamma$$

for each b, b' that differ at some position $i \in \mathcal{I}$; f is called B -bounded if $\|f(b)\|_\infty \leq B$ for all $b \in \{0, 1\}^T$.

Uniform approximation. For some $\varepsilon \geq 0$ and a function $f : \{0, 1\}^T \rightarrow \mathbb{R}^d$, we say that $\widehat{f} \in \mathcal{F}$ ε -uniformly approximates f under the mapping $b \mapsto X(b)$ if

$$\left\| \widehat{f}(X(b)) - f(b) \right\|_\infty \leq \varepsilon, \quad \forall b \in \{0, 1\}^T.$$

B.2. Results

We give an overview of the function approximation results under each input mapping $b \mapsto X(b)$, as a multi-part proposition:

Proposition B.1 (Sparse variable creation with Transformers). *The function classes $\mathcal{F}_{\text{tf-scalar}}, \mathcal{F}_{\text{tf+mlp}}$ contain the following classes of sparse Boolean functions:*

- **Deterministic positional embeddings:** For any \mathcal{I} , $\mathcal{F}_{\text{tf-scalar}}$ can approximate a particular monotone symmetric \mathcal{I} -sparse f , with Transformer weight norm bounds from the real-valued construction in Lemma B.2. $\mathcal{F}_{\text{tf+mlp}}$ with 1 head can exactly represent any symmetric s -sparse f , with the same bounds on Transformer weight norms, and feedforward network weight norms scaling as $O(\text{poly}(s))$. $\mathcal{F}_{\text{tf+mlp}}$ with s heads can exactly represent any s -sparse f , with Transformer weight norm bounds from the vector-valued construction in Lemma B.2, and feedforward network weight norms scaling as $O(\text{poly}(s) \cdot 2^s)$.

- *Trainable positional embeddings:* For any \mathcal{I} , $\mathcal{F}_{\text{tf-scalar}}$ can approximate a particular monotone symmetric \mathcal{I} -sparse f , with positional embedding and Transformer weight norm bounds from the real-valued construction in Lemma B.3. $\mathcal{F}_{\text{tf+mlp}}$ with 1 head can exactly represent any symmetric s -sparse f , with the same bounds on P and Transformer weight norms, and feedforward network weight norms scaling as $O(\text{poly}(s))$. $\mathcal{F}_{\text{tf+mlp}}$ with s heads can exactly represent any sparse f , with P and Transformer weight norm bounds from the vector-valued construction in Lemma B.3, and feedforward network weight norms scaling as $O(\text{poly}(s) \cdot 2^s)$.
- *Bag of vectors:* For any \mathcal{I} , $\mathcal{F}_{\text{tf-scalar}}$ can approximate a particular monotone symmetric \mathcal{I} -sparse f , with Transformer weight norms from Lemma B.4. $\mathcal{F}_{\text{tf+mlp}}$ with 1 head can represent any symmetric s -sparse f , with the same Transformer weight norm bounds, and feedforward network weight norms scaling as $O(\text{poly}(s))$. $\mathcal{F}_{\text{tf+mlp}}$ with 1 head can also exactly represent any s -sparse f , with the same bounds on Transformer weight norms, and feedforward network weight norms scaling as $O(\text{poly}(s) \cdot 2^s)$.

The formal statements are obtained by $(\gamma/4)$ -uniformly approximating a γ -strictly monotone or γ -injective function with self-attention alone (Lemmas B.2, B.3, B.4), then applying a robust universal function representation construction (Lemmas B.5, B.6) appropriately. They are organized as follows:

Lemma B.2 (Deterministic P , no MLP). *Suppose $X(b) = P + E_b$ with deterministic P . Let $\mathcal{I} \subseteq [T]$ such that $|\mathcal{I}| = s \leq d, k$, and $\Delta < 1/s$. Then, for all $0 < \gamma \leq 1$, there exists a 1-bounded, $(2/s)$ -strictly monotone symmetric \mathcal{I} -sparse Boolean function $g_{\mathcal{I}} : \{0, 1\}^T \rightarrow \mathbb{R}$ and Transformer head parameters such that $f_{\text{tf-scalar}}(X(b); W_Q, W_K, W_V, W_C, w = e_1)$ $(\gamma/4)$ -uniformly approximates $g_{\mathcal{I}}$. The norms satisfy*

$$\|W_Q\|_F \leq \frac{\log\left(\frac{8T}{\gamma}\right)}{1 - s\Delta}, \quad \|W_K\|_F \leq s, \quad \|W_V\|_F \leq 2, \quad \|W_C\|_F \leq 1.$$

Also, there exists a 1-bounded, 2-injective \mathcal{I} -sparse Boolean function $g'_{\mathcal{I}} : \{0, 1\}^T \rightarrow \mathbb{R}^s$ and s -headed Transformer parameters such that $f_{\text{tf-head}}\left(X(b); \left\{W_Q^{[h]}, W_K^{[h]}, W_V^{[h]}, W_C^{[h]}\right\}_{h=1}^s\right) \circ \Pi_s$ uniformly approximates $g'_{\mathcal{I}}$. The norms of each head satisfy

$$\|W_Q^{[h]}\|_F \leq \frac{\log\left(\frac{8T}{\gamma}\right)}{1 - s\Delta}, \quad \|W_K^{[h]}\|_F \leq 1, \quad \|W_V^{[h]}\|_F \leq 2, \quad \|W_C^{[h]}\|_F \leq 1.$$

Lemma B.3 (Trainable P , no MLP). *Suppose $X(b) = P + E_b$ with trainable P . Let $\mathcal{I} \subseteq [T]$ such that $|\mathcal{I}| = s \leq d, k$. Then, for any $0 < \gamma \leq 1$, and with the same $g_{\mathcal{I}}$ as in Lemma B.2, there exists P and Transformer head parameters such that $f_{\text{tf-scalar}}(X(b); W_Q, W_K, W_V, W_C, w = e_1)$ $(\gamma/4)$ -uniformly approximates $g_{\mathcal{I}}$. The norms satisfy*

$$\|P^\top\|_{2,1} \leq s, \quad \|W_Q\|_F \leq \log\left(\frac{8T}{\gamma}\right), \quad \|W_K\|_F \leq 1, \quad \|W_V\|_F \leq 2, \quad \|W_C\|_F \leq 1.$$

Also, for the same $g'_{\mathcal{I}}$ as in Lemma B.2, there exists P and s -headed Transformer parameters such that $f_{\text{tf-head}}\left(X(b); \left\{W_Q^{[h]}, W_K^{[h]}, W_V^{[h]}, W_C^{[h]}\right\}_{h=1}^s\right) \circ \Pi_s$ uniformly approximates $g'_{\mathcal{I}}$. The norms of each head satisfy

$$\|P^\top\|_{2,1} \leq s, \quad \|W_Q^{[h]}\|_F \leq \log\left(\frac{8T}{\gamma}\right), \quad \|W_K^{[h]}\|_F \leq 1, \quad \|W_V^{[h]}\|_F \leq 2, \quad \|W_C^{[h]}\|_F \leq 1.$$

Lemma B.4 (Bag of vectors, no MLP). *Suppose $X(b) = V + \text{diag}(b)$. Let $\mathcal{I} \subseteq [T]$ such that $|\mathcal{I}| = s \leq d, k$, and $\Delta < 1/s$. Then, for all $s\Delta < \gamma < 1$, there exists an s -bounded, $(1/s)$ -strictly monotone symmetric \mathcal{I} -sparse Boolean function $g_{\mathcal{I}} : \{0, 1\}^T \rightarrow \mathbb{R}$ and Transformer head parameters such that $f_{\text{tf-scalar}}(X(b); W_Q, W_K, W_V, W_C, w = e_1)$ $(\gamma/4)$ -uniformly approximates $g_{\mathcal{I}}$. The norms satisfy*

$$\|W_Q\|_F \leq \frac{\log\left(\frac{8Ts(1+\Delta)}{\gamma - s\Delta}\right)}{1 - s\Delta}, \quad \|W_K\|_F \leq s + 1, \quad \|W_V\|_F \leq 2s, \quad \|W_C\|_F \leq s.$$

Also, there exists a 1-bounded, $(1/s)$ -injective \mathcal{I} -sparse Boolean function $g'_{\mathcal{I}} : \{0, 1\}^T \rightarrow \mathbb{R}^s$ and Transformer head parameters such that $f_{\text{tf-head}}(X(b); W_Q, W_K, W_V, W_C) \circ \Pi_s$ uniformly approximates $g'_{\mathcal{I}}$. The norms satisfy the same bounds as above.

Lemma B.5 (Monotone to symmetric functions via MLP). *Let $f : \{0, 1\}^T \rightarrow \mathbb{R}$ be any real-valued symmetric s -sparse Boolean function with index set \mathcal{I} . Let W_Q, W_K, W_V, W_C be the parameters of a function*

$$f_{\text{tf-head}}(X; W_Q, W_K, W_V, W_C) := \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) X W_V W_C,$$

and let $\Pi_1 : \mathbb{R}^d \rightarrow \mathbb{R}$ be the projection onto the first coordinate. Suppose that under some mapping $b \mapsto X(b)$, $f_{\text{tf-head}} \circ \Pi_s$ ($\gamma/4$)-uniformly approximates a B -bounded γ -strictly monotone symmetric \mathcal{I} -sparse Boolean function $g : \{0, 1\}^T \rightarrow \mathbb{R}$, for some γ . Then, there exists a function $f_{\text{tf+mlp}} \in \mathcal{F}_{\text{tf+mlp}}$ with the same weights W_Q, W_K, W_V, W_C , and 3-layer feedforward network weights W_1, W_2, w , such that

$$f_{\text{tf+mlp}}(X(b)) = f(b), \quad \forall b \in \{0, 1\}^T,$$

with dimensions $(d_2, d_3) = (4(s+1), 2(s+1))$ and weight norms satisfying

$$\|W_1\|_\infty \leq \frac{8 \max(1, B)}{\gamma}, \quad \|W_2\|_\infty \leq \frac{8s}{\gamma}, \quad \|w\|_\infty \leq \max_{b \in \{0, 1\}^T} |f(b)|.$$

Lemma B.6 (Injective to arbitrary functions via MLP). *Let $f : \{0, 1\}^T \rightarrow \mathbb{R}$ be any real-valued s -sparse Boolean function with index set \mathcal{I} such that $|\mathcal{I}| = s \leq d$. Let W_Q, W_K, W_V, W_C be the parameters of a function*

$$f_{\text{tf-head}}(X; W_Q, W_K, W_V, W_C) := \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) X W_V W_C,$$

and let $\Pi_s : \mathbb{R}^d \rightarrow \mathbb{R}^s$ be the projection onto the first s coordinates. Suppose that under some mapping $b \mapsto X(b)$, $f_{\text{tf-head}} \circ \Pi_s$ ($\gamma/4$)-uniformly approximates a γ -injective function $g : \{0, 1\}^T \rightarrow \mathbb{R}^s$ satisfying $\|g(b)\|_\infty \leq B$. Then, there exists a function $f_{\text{tf+mlp}} \in \mathcal{F}_{\text{tf+mlp}}$ with the same weights W_Q, W_K, W_V, W_C , and 3-layer feedforward network weights W_1, W_2, w , such that

$$f_{\text{tf+mlp}}(X(b)) = f(b), \quad \forall b \in \{0, 1\}^T,$$

with dimensions $(d_2, d_3) = (4s2^s, 2 \cdot 2^s)$ and weight norms satisfying

$$\|W_1\|_\infty \leq \frac{8 \max(1, B)}{\gamma}, \quad \|W_2\|_\infty \leq \frac{8s}{\gamma}, \quad \|w\|_\infty \leq \max_{b \in \{0, 1\}^T} |f(b)|.$$

B.3. Useful lemmas

We will use a construction which approximates a ‘‘hard selection’’ of s indices using the softmax mixture; for this, we will need to quantify the approximation error when the inputs to the softmax function are bounded.

Lemma B.7 (Softmax truncation). *Let $z \in (\mathbb{R} \cup \{-\infty\})^T$ such that $z_t \geq R$ for each $1 \leq t \leq s$, and $z_t \leq 0$ for each $s+1 \leq t \leq T$. Define $z' \in (\mathbb{R} \cup \{-\infty\})^T$ so that $z'_t = z_t$ for $1 \leq t \leq s$, and $z'_t = -\infty$ for $s+1 \leq t \leq T$. Then, letting $e^{-\infty} = 0$ in the definition of $\text{softmax}(\cdot)$, we have*

$$\|\text{softmax}(z') - \text{softmax}(z)\|_1 \leq 2 \frac{T-s}{s \exp(R)} < \frac{2T}{\exp(R)}.$$

Proof. We have

$$\|\text{softmax}(z') - \text{softmax}(z)\|_1 = \sum_{t=1}^s \exp(z_t) \left(\frac{1}{\mathbf{1}^\top \exp(z')} - \frac{1}{\mathbf{1}^\top \exp(z)} \right) + \sum_{t=s+1}^T \frac{\exp(z_t)}{\mathbf{1}^\top \exp(z)}.$$

The first summation is equal to

$$1 - \frac{\mathbf{1}^\top \exp(z')}{\mathbf{1}^\top \exp(z)} \leq \frac{T-s}{s \exp(R)},$$

while the same upper bound holds for the second summation, since each term is at most $\frac{1}{s \exp(R)}$. \square

Our results on approximating arbitrary sparse Boolean functions will depend on a generic construction for robustly approximating an arbitrary function $f : \mathbb{R}^d \rightarrow \mathbb{R}$ with a feedforward neural network. For simplicity of presentation, we use a standard⁷ 3-layer ReLU network construction, which *exactly* represents a piecewise constant function in specified regions.

Lemma B.8 (Exact function representation with a 3-layer ReLU net). *Let $f : \mathbb{R}^{d_f} \rightarrow \mathbb{R}$, and let $x_1, \dots, x_n \in \mathbb{R}^{d_f}$ such that $\|x_i\|_\infty \leq B$ for each $i \in [n]$, $\|x_i - x_j\|_\infty \geq 4\delta$ for each $i \neq j \in [n]$. Then, there is a 3-layer feedforward network with ReLU activations, with parameters $W_1 \in \mathbb{R}^{(d_f+1) \times d_2}$, $W_2 \in \mathbb{R}^{(d_2+1) \times d_3}$, $w \in \mathbb{R}^{d_3}$, such that*

$$f_{\text{mlp}}(x_i + z) = f(x_i)$$

for all $i \in [n]$ and $\|z\|_\infty \leq \delta$, where $\text{ReLU}(x) := x_+ = \max(0, x)$ is applied entrywise, with

$$d_2 = 4nd_f, \quad d_3 = 2n,$$

$$\|W_1\|_\infty \leq \frac{\max(1, B)}{\delta}, \quad \|W_2\|_\infty \leq \frac{d_f}{\delta}, \quad \|w\|_\infty \leq \max_{i \in [n]} |f(x_i)|.$$

Proof. First, we construct a one-dimensional ‘‘bump’’ function basis, and propagate the Lipschitz constants. A threshold function with a linear ‘‘ramp’’ of width δ can be obtained from a linear combination of 2 ReLU functions:

$$\nu_\delta(x) := (x/\delta + 1)_+ - (x/\delta)_+ = \begin{cases} 0 & x \leq -\delta \\ x/\delta + 1 & -\delta \leq x \leq 0 \\ 1 & x \geq 0 \end{cases}.$$

Next, we construct the bump function

$$\psi_\delta(x) := \nu_\delta(x) - \nu_\delta(2\delta - x).$$

By this construction, we have $\psi_\delta(x) = 1$ for $0 \leq x \leq 2\delta$ and $\psi_\delta(x) = 0$ for $x \leq -\delta$ and $x \geq 3\delta$, interpolating linearly on $[-\delta, 0]$ and $[2\delta, 3\delta]$. Next, define

$$\begin{aligned} \psi_\delta(x; x_0) &:= \psi_\delta(x - x_0 + \delta) \\ &= \left(\frac{x - x_0}{\delta} + 2 \right)_+ - \left(\frac{x - x_0}{\delta} + 1 \right)_+ - \left(\frac{x_0 - x}{\delta} + 2 \right)_+ + \left(\frac{x_0 - x}{\delta} + 1 \right)_+ \end{aligned}$$

so that $\psi_\delta(x; x_0) = 1$ for $|x - x_0| \leq \delta$, $\psi_\delta(x; x_0) = 0$ for $|x - x_0| \geq 2\delta$.

We construct the first layer $W_1 \in \mathbb{R}^{(d+1) \times (4nd)}$ using these bump functions: indexing the $4nd$ dimension by $(h \in [4], i \in [n], j \in [d])$, we construct

$$\begin{aligned} [W_1]_{:, (1, i, :)} &:= \begin{bmatrix} \frac{1}{\delta} I \\ -\frac{x_i}{\delta} + 2 \cdot \mathbf{1}^\top \end{bmatrix}, & [W_1]_{:, (2, i, :)} &:= \begin{bmatrix} \frac{1}{\delta} I \\ -\frac{x_i}{\delta} + \mathbf{1}^\top \end{bmatrix}, \\ [W_1]_{:, (3, i, :)} &:= \begin{bmatrix} -\frac{1}{\delta} I \\ \frac{x_i}{\delta} + 2 \cdot \mathbf{1}^\top \end{bmatrix}, & [W_1]_{:, (4, i, :)} &:= \begin{bmatrix} -\frac{1}{\delta} I \\ \frac{x_i}{\delta} + \mathbf{1}^\top \end{bmatrix}, \end{aligned}$$

so that

$$([x \ 1]^\top [[W_1]_{j, (1, i, :)} [W_1]_{j, (2, i, :)} [W_1]_{j, (3, i, :)} [W_1]_{j, (4, i, :)}])_+ [1 \ -1 \ -1 \ 1]^\top = \psi_\delta(x; [x_i]_j).$$

The second layer is used to construct n activations which are indicators of whether x is in the neighborhood of each x_i . For each x_i , we will simply average the d_f one-dimensional indicators for each coordinate, and implement a threshold function $\nu_{\delta/d_f}(1 - x)$. We choose $W_2 \in \mathbb{R}^{(4nd_f+1) \times (2n)}$, with the $4nd_f + 1$ dimension indexed by (h, i, j) and an extra bias dimension \perp , and the $2n$ dimension indexed by $(h' \in \{1, 2\}, i' \in [n])$ so that

$$[W_2]_{(h, i, :), (h', i')} := [1 \ -1 \ -1 \ 1]_h \cdot \mathbb{1}[i = i'] \cdot \frac{1}{\delta} \cdot \mathbf{1},$$

⁷For example, this follows from the discussion in Chapter 4 of (Nielsen, 2015).

⁸Here, W_1, W_2 have bias terms; w does not.

$$[W_2]_{\perp,(1,i')} := 1 - \frac{d_f}{\delta}, \quad [W_2]_{\perp,(2,i')} := -\frac{d_f}{\delta}.$$

Finally, the third (output) layer $w \in \mathbb{R}^{2n}$, with dimensions indexed by $(h \in \{1, 2\}, i \in [n])$, multiplies the indicators of each x_i by the desired $f(x_i)$:

$$w_{(1,i)} := f(x_i), \quad w_{(2,i)} := -f(x_i).$$

For any $x_0 \in \mathbb{R}^{d_f}$, let $B_\delta(x_0)$ be the set of x such that $\|x - x_0\|_\infty \leq \delta$. By this construction, for each $x \in B_\delta(x_i)$, we have $f(x) = x_i$, as required. \square

Note that we use 3-layer ReLU networks for function approximation in order to minimize the introduction of unnecessary notation. Some minor remarks:

- It would be routine to replace this construction with any architecture which can represent an arbitrary function *approximately* (Hornik et al., 1989; Cybenko, 1989); this includes the 2-layer feedforward networks (and nonlinear activations other than the ReLU) which are typically used by Transformers in practice.
- It is possible to embed this construction in $f_{\text{tf+mlp}}$ with a 2-layer ReLU network, by using (W_C, W_1, W_2) and introducing a nonlinearity after W_C , without changing the results.
- When $d_f = 1$, W_2 is unnecessary (one can represent f directly using the bump function basis).

B.4. Beyond Boolean domains

These representational results are stated with a Boolean domain $\{0, 1\}^T$ for clarity and simplicity of presentation; this is not essential or fundamental to these constructions. Generalizations to the following input domains are straightforward:

- Discrete categorical tokens $\{1, \dots, M\}^T$. Use M orthogonal token embeddings, instead of only 2. The “sparse function approximation” construction uses sM^s instead of $s2^s$ parameters. The smaller (poly(s)-parameter) representation of symmetric Boolean functions has no clear analogue.
- Continuous inputs on a bounded domain (e.g. $[0, 1]^T$). The (fixed or trainable) positional embeddings can still select the s sparse indices. Replace the discrete ReLU network construction with any continuous universal function approximation construction. The “bag-of-vectors” formulation has no clear analogue.

The capacity results from Section 4 hold for any input domain, as long as the embeddings are bounded in the $\|\cdot\|_{2,\infty}$ norm. Note that Transformers are predominantly used with discrete inputs.

B.5. Proofs

Throughout the constructions in each case, we will refer to standard coordinate bases in several spaces:

- $E_0, E_1 \in \mathbb{R}^d$ denote the embeddings of the 0, 1 tokens $E_{0,\cdot}, E_{1,\cdot}$.
- $e_1^{(k)}, \dots, e_k^{(k)}$ denotes the standard basis in \mathbb{R}^k .
- $e_i^{(d)}$ denotes the standard basis in \mathbb{R}^d .
- $e_1^{(T)}, \dots, e_T^{(T)}, e_{[\text{CLS}]}$ denotes the standard basis in \mathbb{R}^{T+1} with the special $[\text{CLS}]$ index.
- Recall that the v_i form a Δ -approximate orthonormal basis for \mathbb{R}^d , $v_{[\text{CLS}]}, e_0, e_1$ are exactly orthogonal to each of them as well as each other, and d is chosen such that these conditions can be met.

Let $n(i)$ be a unique bijection between \mathcal{I} and $[s]$. Let $v_{\mathcal{I}} := \sum_{i \in \mathcal{I}} v_i$.

Approximate vector equality. We will use $u \approx_\varepsilon v$ to denote that two vectors $u, v \in \mathbb{R}^{d_u}$ satisfy $\|u - v\|_\infty \leq \varepsilon$.

Proof of Lemma B.2. We construct attention heads such that the softmax mixture always selects the indices in \mathcal{I} .

Single head, deterministic P . We seek to approximate the 1-bounded, $(2/s)$ -strictly monotone function

$$\frac{1}{s} \sum_{i \in \mathcal{I}} \chi_i,$$

where $\chi_i = +1$ if $b_i = 0$ and -1 if $b_i = 1$. Set

$$W_Q := Rv_{[\text{CLS}]}e_1^{(k)\top}, \quad W_K := v_{\mathcal{I}}e_1^{(k)\top}, \quad W_V := (E_0 - E_1)e_1^{(k)\top}, \quad W_C := e_1^{(k)}e_1^{(d)\top},$$

where R will be chosen later. Then, by approximate orthogonality,

$$v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top = v_{[\text{CLS}]}^\top W_Q W_K^\top (P + E_b)^\top = v_{[\text{CLS}]}^\top W_Q W_K^\top P^\top \approx_{Rs\Delta} R \sum_{i \in \mathcal{I}} e_i^{(T)\top}.$$

By Lemma B.7,

$$\left\| \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) - \frac{1}{s} \sum_{i \in \mathcal{I}} e_i^{(T)\top} \right\|_1 \leq \frac{2T}{\exp(R - 2Rs\Delta)}.$$

Finally, we have

$$XW_V W_C = E_b W_V W_C = \left(\sum_{i \in [T]} \chi_i e_i^{(T)} \right) e_1^{(d)\top},$$

so that by Hölder's inequality,

$$f_{\text{tf-head}}(X) \circ \Pi_1 = \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) XW_V W_C e_1^{(d)} \approx_{\frac{2T}{\exp(R-2Rs\Delta)}} \frac{1}{s} \sum_{i \in \mathcal{I}} \chi_i.$$

To get $(\gamma/4)$ -uniform approximation, we choose

$$R = \frac{\log\left(\frac{8T}{\gamma}\right)}{1 - s\Delta}.$$

Multiple heads, deterministic P . For $h = 1, \dots, s$, and the same R as above:

$$W_Q^{[h]} := Rv_{[\text{CLS}]}e_1^{(k)\top}, \quad W_K^{[h]} := v_{n-1(h)}e_1^{(k)\top}, \quad W_V^{[h]} := (E_0 - E_1)e_2^{(k)\top}, \quad W_C^{[h]} := e_1^{(k)}e_h^{(d)\top}.$$

This is the same construction as above, but each head only selects one of the coordinates in \mathcal{I} . Thus, by the same analysis,

$$f_{\text{tf-head}}(X) \circ \Pi_s \approx_{\frac{2T}{\exp(R-2Rs\Delta)}} \sum_{i \in \mathcal{I}} \chi_i e_{n(i)}^{(d)}.$$

This function is clearly 1-bounded and 2-injective. □

Proof of Lemma B.3. The constructions closely follow Lemma B.2, but are simpler.

Single head, trainable P . For each $i \in \mathcal{I}$, set the trainable positional embeddings to be

$$P_{i,:} := \begin{cases} v_1 & i \in \mathcal{I} \\ 0 & \text{otherwise} \end{cases}.$$

Set

$$W_Q := Rv_{[\text{CLS}]}e_1^{(k)\top}, \quad W_K := v_1e_1^{(k)\top}, \quad W_V := (E_0 - E_1)e_1^{(k)\top}, \quad W_C := e_1^{(k)}e_1^{(d)\top}.$$

Now, we have (with equality)

$$v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top = R \sum_{i \in \mathcal{I}} e_i^{(T)\top},$$

so that Lemma B.7 gives

$$\left\| \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) - \frac{1}{s} \sum_{i \in \mathcal{I}} e_i^{(T)\top} \right\|_1 \leq \frac{2T}{\exp(R)}.$$

Like before, we have

$$f_{\text{tf-head}}(X) \circ \Pi_1 = \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) X W_V W_C e_1^{(d)} \approx \frac{2T}{\exp(R)} \frac{1}{s} \sum_{i \in \mathcal{I}} \chi_i.$$

To get $(\gamma/4)$ -uniform approximation, we choose

$$R = \log \left(\frac{8T}{\gamma} \right).$$

Multiple heads, trainable P . For each $i \in \mathcal{I}$, set the trainable positional embeddings to be

$$P_{i,:} := \begin{cases} e_{n(i)}^{(d)} & i \in \mathcal{I} \\ 0 & \text{otherwise} \end{cases}.$$

For $h = 1, \dots, s$, and the same R as above:

$$W_Q^{[h]} := R v_{[\text{CLS}]} e_1^{(k)\top}, \quad W_K^{[h]} := e_h^{(d)} e_1^{(k)\top}, \quad W_V^{[h]} := (E_0 - E_1) e_1^{(k)\top}, \quad W_C^{[h]} := e_1^{(k)} e_h^{(d)\top}.$$

This is the same construction as above, but each head only selects one of the coordinates in \mathcal{I} . Thus, by the same analysis,

$$f_{\text{tf-head}}(X) \circ \Pi_s \approx \frac{2T}{\exp(R)} \sum_{i \in \mathcal{I}} \chi_i e_{n(i)}^{(d)}.$$

□

Proof of Lemma B.4. This input mapping does not use position embeddings, and does not need multiple heads to implement arbitrary (non-symmetric) functions. The constructed monotone and injective functions are slightly different, but the proof strategy is very similar. The key difference is that the softmax mixture is uniform only on the positions $i \in \mathcal{I}$ where $b_i = 1$.

Bag of vectors, scalar output. The function we will approximate is defined as follows:

$$g_{\mathcal{I}}(r) := \frac{r - s}{r + 1}, \quad \text{where } r = \sum_{i \in \mathcal{I}} b_i, \quad s = |\mathcal{I}|.$$

Note that this function is $(1/s)$ -strictly monotone, and has absolute value bounded by s . Set

$$W_Q := R v_{[\text{CLS}]} e_1^{(k)\top}, \quad W_K := (v_{\mathcal{I}} + v_{[\text{CLS}]}) e_1^{(k)\top},$$

$$W_V := \sum_{i \in \mathcal{I}} v_i e_{n(i)}^{(k)\top} - v_{[\text{CLS}]} \left(\sum_{i \in \mathcal{I}} e_{n(i)}^{(k)} \right)^\top, \quad W_C := \sum_{i \in \mathcal{I}} e_{n(i)}^{(k)} e_1^{(d)\top},$$

where R will be chosen later. Then, by approximate orthogonality,

$$v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top \approx_{Rs\Delta} R \left(v_{[\text{CLS}]} + \sum_{i \in \mathcal{I}} b_i e_i^{(T)\top} \right),$$

so that by Lemma B.7,

$$\left\| \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) - \frac{1}{r+1} \left(e_{[\text{CLS}]}^{(T)\top} + \sum_{i \in \mathcal{I}} b_i e_i^{(T)\top} \right) \right\|_1 \leq \frac{2T}{\exp(R - 2Rs\Delta)}.$$

Finally, we have

$$XW_V W_C e_1^{(d)} = -s e_{[\text{CLS}]}^{(T)} + \sum_{i \in [T]} b_i \cdot v_i^\top v_{\mathcal{I}} \cdot e_i^{(T)} \approx_{s\Delta} -s e_{[\text{CLS}]}^{(T)} + \sum_{i \in \mathcal{I}} b_i e_i^{(T)},$$

so that

$$\begin{aligned} & |f_{\text{tf-head}}(X) \circ \Pi_1 - g_{\mathcal{I}}(r)| \leq \\ & \left\| \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) - \frac{1}{r+1} \left(e_{[\text{CLS}]}^{(T)\top} + \sum_{i \in \mathcal{I}} b_i e_i^{(T)\top} \right) \right\|_1 \left(\|XW_V W_C e_1^{(d)}\|_\infty + s\Delta \right) \\ & \quad + \left\| \text{softmax}(v_{[\text{CLS}]}^\top W_Q W_K^\top X^\top) \right\|_1 (s\Delta) \\ & \leq \frac{2Ts(1+\Delta)}{\exp(R - 2Rs\Delta)} + s\Delta. \end{aligned}$$

To get $(\gamma/4)$ -uniform approximation, we choose

$$R = \frac{\log\left(\frac{8Ts(1+\Delta)}{\gamma - s\Delta}\right)}{1 - s\Delta}.$$

Bag of vectors, s -dimensional output. We use the same construction as above, except

$$W_C := \sum_{i \in \mathcal{I}} e_{n(i)}^{(k)} e_{n(i)}^{(d)\top}.$$

This will allow us to approximate the function

$$g'_{\mathcal{I}}(b) = \frac{1}{r+1} \sum_{i \in \mathcal{I}} (b_i - 1) e_{n(i)}^{(d)},$$

which is $(1/s)$ -injective and has absolute value is bounded by 1. Then, for each $i \in \mathcal{I}$, we have

$$XW_V W_C e_i^{(d)} = -e_{[\text{CLS}]}^{(T)} + v_i^\top v_{\mathcal{I}} \cdot e_i^{(T)} \approx_{s\Delta} -e_{[\text{CLS}]}^{(T)} + b_i e_i^{(T)}.$$

Repeating the above analysis for each coordinate, we have

$$f_{\text{tf-head}}(X) \circ \Pi_s \approx_\varepsilon g'_{\mathcal{I}}(r),$$

where a slightly tighter bound

$$\varepsilon = \frac{2T(1+s\Delta)}{\exp(R - 2Rs\Delta)} + s\Delta$$

comes from the fact that $\|XW_V W_C e_i^{(d)}\|_\infty$ is now bounded by 1 instead of s . The previous choice of R suffices for $(\gamma/4)$ -uniform approximation. \square

Proof of Lemma B.5. This follows by instantiating Lemma B.8 with $\delta = \gamma/8, d_f = 1, n = s + 1$. Notice that a $(\gamma/4)$ -uniform approximation of a γ -strictly monotone function satisfies the conditions needed for Lemma B.8. \square

Proof of Lemma B.6. This follows by instantiating Lemma B.8 with $\delta = \gamma/8, d_f = s, n = 2^s$. Notice that a $(\gamma/4)$ -uniform approximation of a γ -injective function satisfies the conditions needed for Lemma B.8. \square

C. Details for experiments

C.1. Empirical scaling laws (Figure 6)

In this section, we provide details for the sample complexity scaling law experiments, which are the main empirical verification of the $\log T$ dependence of the sample complexity arising from the analysis.

Data. Synthetic supervised learning tasks corresponding to learning a 3-sparse conjunction were generated by the protocol described in the main paper, parameterized by sample size m and context T : in each trial, one of the $\binom{T}{3}$ subsets of indices was selected uniformly at random⁹, under i.i.d. Bernoulli inputs $x_i \sim \text{Bern}(p)$. Here, $p = (1/2)^{1/3}$ was chosen so that the classes are balanced $\Pr[y = 0] = \Pr[y = 1]$. m samples were drawn this distribution to form a training set (rejecting training sets which were compatible with multiple hypotheses), and 10^4 samples were drawn from the same distribution as a holdout validation set. The grid of problem instances was constructed as follows: $T \in \{100, 150, 200, \dots, 750, 800\} \cup \{900, 1000, 1100\}$, $m \in \{50, 60, 70, \dots, 200\}$.

Training. A 1-layer Transformer network (with a scalar output at a special trainable token $x_{[\text{CLS}]}$ at an extra index $[\text{CLS}]$) was trained with Adam (Kingma & Ba, 2014) and full-batch gradients of the cross entropy loss for binary classification. 40 independent trials were performed (re-randomizing the dataset generation); each trial was restarted (with fresh random initialization and dropout masks) 5 times before being labeled as a failure. Cross-validation was performed on a holdout sample of size 10^4 every 10 iterations. At the end of 1000 training iterations, the trial was counted as a success if the maximum validation accuracy throughout training was greater than 0.99. (In 100% of runs, the training loss was driven to $< 10^{-4}$, with 100% training accuracy, within 1000 iterations.)

Architecture. Like (Chen et al., 2021a), our experimental setup is based on a popular PyTorch implementation (<https://github.com/karpathy/minGPT>), with some optimizations for faster 1-layer training and inference. This implementation includes widely-used architectural details (GeLU activations; dropout) which were not discussed in the theoretical analysis; refer to the referenced repository for details. All hyperparameter settings left undiscussed are taken from the default settings in this codebase.

Hyperparameters. A fixed architecture was used ($d = 64, k = 4, 16$ parallel heads), with trainable positional embeddings initialized with Gaussian entries $\mathcal{N}(0, \sigma^2)$, $\sigma = 0.02$, 3 input token embeddings (corresponding to 0, 1, $[\text{CLS}]$), and 2 output embeddings (corresponding to the possible labels 0, 1). For regularization mechanisms, typical choices were used: 0.1 for {attention, embedding, output} dropout; 10^{-4} weight decay. The Adam optimizer was instantiated with typical parameters $\eta = 10^{-3}, \beta_1 = 0.9, \beta_2 = 0.999$.

Infrastructure and computational costs. All experiments were performed on an internal cluster, with NVIDIA Tesla P100, NVIDIA Tesla P40, and NVIDIA RTX A6000 GPUs. Each training run took at most 10 minutes (with most computation time spent on cross-validation), for a total of ~ 150 GPU-hours.

C.2. Other figures

Example training curves. Figure 3 (left) shows the best training curves (in terms of highest validation accuracy) out of 5 restarts for all 40 replicates, in the settings $T = 300, m = 200$ and $T = 300, m = 50$. As the sample size m decreases, the trained models overfit with higher probability; when they overfit, they differ significantly in validation accuracy.

Attention weights. In Figure 3, the normalized attention weights at the $[\text{CLS}]$ position are shown for a Transformer model trained for 500 iterations with $T = 50, m = 300$ (achieving 100% validation accuracy on 10^4 holdout samples), for 100 validation samples (which each induce different attention weights, shown in the scatter plot). One key difference (for simplicity of visualization) is that this network only has one attention head, with embedding dimensions $d = k = 64$.

Parity. These experiments use the same architecture as the main AND experiments, except the. In the loss curves presented in Figure 4, gradient-based training is done on streaming online losses (so that there is no training/validation split),

⁹Note that due to the permutation-symmetry of the Transformer architecture (as long as the position embeddings are initialized with a permutation-symmetric distribution), it is equivalent to select $\mathcal{I} = [s]$. Also, by symmetry of the initial token embeddings, AND and OR are interchangeable in these experiments and results.

with batch size 2048.

D. Additional related work

In this section, we discuss some additional related work.

Domains beyond language. We provide a few references on the successful application of Transformers to domains beyond natural language processing; this frontier is continually and rapidly evolving.

- With minimal changes compared to the analogous natural language settings, Transformer sequence models have been applied to theorem proving (Polu & Sutskever, 2020) and program synthesis (Chen et al., 2021b).
- Beyond sequence modeling: In self-supervised learning of image representations, the Vision Transformer (ViT) (Dosovitskiy et al., 2020) has sometimes outperformed older convolution-based architectures, particularly when pre-trained on massive datasets. Further architectural variants have been proposed for vision and other continuous modalities (Tolstikhin et al., 2021; Lee-Thorp et al., 2021; Jaegle et al., 2021b;a; d’Ascoli et al., 2021).
- Natural sciences: A state-of-the-art pipeline for protein structure prediction (Jumper et al., 2021) features a self-attention-based component.
- Reinforcement learning: Transformer architectures have shown promise for planning (Janner et al., 2021) and offline RL (Chen et al., 2021a).

Expressive power of Transformers. Several works establish results on the representational power of self-attention architectures in regimes where the statistical guarantees are necessarily weak or vacuous (i.e. there are too many functions in the class). (Dehghani et al., 2018; Yun et al., 2019; Bhattamishra et al., 2020a;b) establish universal function approximation and Turing-completeness, which have been known for previous architectures (Siegelmann & Sontag, 1995). Our work is a significantly finer-grained analysis, in which we establish a hierarchy of function classes (indexed by sparsity s) representable by these architectures, with tight (in terms of T) statistical guarantees. (Hron et al., 2020; Yang, 2020) analyze properties of the kernels induced by Transformers at the infinite-width limit. Quoting the discussion following Theorem 4.1 of (Wei et al., 2021), who show statistically meaningful approximations of TM using Transformers: “*The correct norm-based Rademacher complexity bound to use for Transformers is unclear.*” The connection between Transformers and circuits also appears in (Elhage et al., 2021; Olsson et al., 2022), with a different technical approach (interpreting the classes of computations performed by attention weights and heads). (Likhoshesterov et al., 2021) analyze the sparsity patterns representable by a self-attention head, with results superficially similar to ours: when the embedding dimension is at least logarithmic in the context length, all sparse matrices can be approximately realized by an attention head. However, their analysis is not about the capacity of the function class: it quantifies over the input X , and holds the parameters (W_Q, W_K, \dots) to be constant (rather than vice versa). This finding serves as an interesting complement to our result: even though the attention mixture weights can take on exponentially many sparsity patterns for distinct inputs, the generalization error scales as $\log(T)$.

Interpreting attention mixtures. A line of empirical work (“BERTology”) has made progress on understanding and interpreting state-of-the-art Transformer language models by examining the activations of their attention mechanisms (Clark et al., 2019; Tenney et al., 2019; Rogers et al., 2020). In some cases, these works have found instances in which Transformers seem to have learned features that are reminiscent of (sparse) hand-crafted features used in natural language processing, without explicit supervision. Our analysis formalizes the intuition that self-attention heads can represent sparse interactions within the context in a statistically meaningful way.

Other theoretical work on self-attention. Kerg et al. (2020) analyze self-attention and its benefits for learning long-term dependencies by establishing gradient norms bounds and showing how attention helps address the problem of gradient vanishing in recurrent networks. In contrast to our results that analyze the statistical and representational properties of attention-based architectures, this work focuses on the computational aspects of gradient-based methods on recurrent networks with self-attention.

Other attention-based architectures. Our analysis is amenable to computationally efficient variants of the Transformer which use parsimonious (e.g. low-rank) approximations of the softmax kernel, like the Performer (Choromanski et al., 2020). Building upon the success of modern attention-based architectures, a large body of work (e.g. Goyal et al. (2020; 2021), and the works cited within) has sought to design architectures which induce model sparsity and modularity. Our analysis is relevant to any architecture that uses a softmax (or similar) bottleneck for statistical capacity, and could inform design principles for norm-based capacity control of these architectures.

Synthetic experiments with Transformers. Power et al. (2021) train small Transformer networks on synthetic algebraic tasks, and discover an abrupt phase transition from overfitting to correct generalization similar to ours. (Tay et al., 2020) propose some synthetic tasks for benchmarking the ability of Transformer variants to capture long-range dependences. (Chen et al., 2021a) present a synthetic demonstration of extrapolation (inferring a maximum-reward path from random walk observations) when using Transformers for offline reinforcement learning. (Lu et al., 2021) probe the transfer learning capabilities of pretrained Transformers, and consider some simple Boolean tasks. Our experimental protocol of learning sparse Boolean functions provides a simple and fundamental setting for elucidating computational and statistical properties of sequence modeling architectures.