
Prompting a Pretrained Transformer Can Be a Universal Approximator

Aleksandar Petrov¹ Philip H.S. Torr¹ Adel Bibi¹

Abstract

Despite the widespread adoption of prompting, prompt tuning and prefix-tuning of transformer models, our theoretical understanding of these fine-tuning methods remains limited. A key question is whether one can arbitrarily modify the behavior of a pretrained model by prompting or prefix-tuning it. Formally, whether prompting and prefix-tuning a pretrained model can universally approximate sequence-to-sequence functions. This paper answers in the affirmative and demonstrates that much smaller pretrained models than previously thought can be universal approximators when prefixed. In fact, prefix-tuning a single attention head is sufficient to approximate any continuous function making the attention mechanism uniquely suited for universal approximation. Moreover, any sequence-to-sequence function can be approximated by prefixing a transformer with depth linear in the sequence length. Beyond these density-type results, we also offer Jackson-type bounds on the length of the prefix needed to approximate a function to a desired precision.

1. Introduction

The scale of modern transformer architectures (Vaswani et al., 2017) is ever-increasing and training competitive models from scratch, even fine-tuning them, is often prohibitively expensive (Lialin et al., 2023). To that end, there has been a proliferation of research aiming at efficient training in general and fine-tuning in particular (Rebuffi et al., 2017; Houlsby et al., 2019; Hu et al., 2021; 2023).

Motivated by the success of few- and zero-shot learning (Wei et al., 2021; Kojima et al., 2022), context-based fine-tuning methods do not change the model parameters. Instead, they modify the way the input is presented. For example, with prompting, one fine-tunes a string of tokens

(a prompt) which is prepended to the user input (Shin et al., 2020; Liu et al., 2023). As optimizing over discrete tokens is difficult, one can optimize the real-valued embeddings instead (soft prompting, prompt tuning, Lester et al. 2021). A generalization to this approach is the optimization over the embeddings of every attention layer (prefix-tuning, Li and Liang 2021). These methods are attractive as they require a few learnable parameters and allow for different prefixes to be used for different samples in the same batch which is not possible with methods that change the model parameters. As every prompt and soft prompt can be expressed as prefix-tuning (Petrov et al., 2024), in this paper, we will focus primarily on prefix-tuning.

While these context-based fine-tuning techniques have seen widespread adoption and are, in some cases, competitive to full fine-tuning (Liu et al., 2022), our understanding of their abilities and restrictions remains limited. How much can the behavior of a model be modified without changing any model parameter? Given a pretrained transformer and an arbitrary target function, how long should the prefix be so that the transformer approximates this function to an arbitrary precision? Differently put, can prefix-tuning of a pretrained transformer be a universal approximator? These are some of the questions we aim to address in this work.

It is well-known that fully-connected neural networks with suitable activation functions can approximate any continuous function (Cybenko, 1989; Hornik et al., 1989; Barron, 1993; Telgarsky, 2015), while Recurrent Neural Networks (RNNs) can approximate dynamical system (Schäfer and Zimmermann, 2006). The attention mechanism (Bahdanau et al., 2015) has also been studied in its own right. Deora et al. (2023) derived convergence and generalization guarantees for gradient-descent training of a single-layer multi-head self-attention model, and Mahdavi et al. (2023) showed that the memorization capacity increases linearly with the number of attention heads. On the other hand, it was shown that attention layers are not expressive enough as they lose rank doubly exponentially with depth if Multi-Layer Perceptrons (MLPs) and residual connections are not present (Dong et al., 2021). However, attention layers, with a hidden size that grows only logarithmically in the sequence lengths, were shown to be good approximators for sparse attention patterns (Likhoshesterov et al., 2021), except for a few tasks that require a linear scaling of the size of the

¹Department of Engineering Science, University of Oxford, UK. Correspondence to: A. Petrov <aleks@robots.ox.ac.uk>.

hidden layers in the sequence length (Sanford et al., 2023).

Considering universal approximation using encoder-only transformers, Yun et al. (2019) showed that transformers are universal approximators of sequence-to-sequence functions by demonstrating that self-attention layers can compute contextual mappings of input sequences. Jiang and Li (2023) demonstrated universality by instead leveraging the Kolmogorov-Albert representation Theorem. Moreover, Alberti et al. (2023) provided universal approximation results for architectures with non-standard attention mechanisms.

Despite this interest in theoretically understanding the approximation properties of the transformer architecture when being trained, much less progress has been made in understanding context-based fine-tuning methods such as prompting, soft prompting, and prefix-tuning. Petrov et al. (2024) have shown that the presence of a prefix cannot change the relative attention over the context and experimentally demonstrate that one cannot learn completely novel tasks with prefix-tuning. In the realm of in-context learning, where input-target pairs are part of the prompt (Brown et al., 2020), Xie et al. (2021) and Yadlowsky et al. (2023) show that the ability to generalize depends on the choice of pre-training tasks. However, these are not universal approximation results. The closest to our objective is the work of Wang et al. (2023). They quantize the input and output spaces allowing them to enumerate all possible sequence-to-sequence functions. All possible functions and inputs can then be hard-coded in a transformer using the constructions by Yun et al. (2019). As this approach relies on memorization, the depth of the model depends on the desired approximation precision ϵ .

In this work, we demonstrate that prefix-tuning can be a universal approximator much more efficiently than previously assumed. In particular:

- i. We show that attention heads are especially suited to model functions over hyperspheres, concretely, prefix-tuning a *single attention head* is sufficient to approximate any smooth continuous function on the hypersphere S^m to any desired precision ϵ ;
- ii. We give a bound on the required prompt length to approximate a smooth target function to a precision ϵ ;
- iii. We demonstrate how this result can be leveraged to approximate general sequence-to-sequence functions with transformers of depth linear in the sequence length and independent of ϵ ;
- iv. We discuss how prefix-tuning may result in element-wise functions which, when combined with cross-element mixing from the pretrained model, may be able to explain the success behind prefix-tuning and prompting and why it works for some tasks and not others.

2. Background Material

2.1. Transformer Architecture

In soft prompting and prefix-tuning, the focus of this work, the sequence fed to a transformer model is split into two parts: a *prefix* sequence $P = (\mathbf{p}_1, \dots, \mathbf{p}_N)$, which is to be learnt or hand-crafted, and an *input* sequence $X = (\mathbf{x}_1, \dots, \mathbf{x}_T)$, where $\mathbf{x}_i, \mathbf{p}_i \in \mathbb{R}^d$. A transformer operating on a sequence consists of alternating attention blocks which operate on the whole sequence and MLPs that operate on individual elements. For a sequence of length $N + T$, an attention head of dimension d is a function $\tilde{u} : \mathbb{R}^{d \times (N+T)} \rightarrow \mathbb{R}^{d \times (N+T)}$. Since we will only be interested in the output at the positions corresponding to the inputs X , we use $u(\cdot; P) : \mathbb{R}^{d \times (N+T)} \rightarrow \mathbb{R}^{d \times T}$ to denote the output of \tilde{u} at the locations corresponding to the input X when prefixed with P . Therefore, the k -th output of u is defined as:

$$[u(X; P)]_k = \frac{\sum_{i=1}^N \exp(\mathbf{x}_k^\top \mathbf{H} \mathbf{p}_i) \mathbf{W}_V \mathbf{p}_i + \sum_{j=1}^T \exp(\mathbf{x}_k^\top \mathbf{H} \mathbf{x}_j) \mathbf{W}_V \mathbf{x}_j}{\sum_{i=1}^N \exp(\mathbf{x}_k^\top \mathbf{H} \mathbf{p}_i) + \sum_{j=1}^T \exp(\mathbf{x}_k^\top \mathbf{H} \mathbf{x}_j)}, \quad (1)$$

where \mathbf{W}_V , the *value* matrix, and \mathbf{H} are in $\mathbb{R}^{d \times d}$. \mathbf{H} is typically split into two lower-rank matrices $\mathbf{H} = \mathbf{W}_Q^\top \mathbf{W}_K$, *query* and *key* matrices. Multiple attention heads can be combined into an attention block but, for simplicity, we will only consider single head attention blocks. A transformer is then constructed by alternating attention heads and MLPs.

We consider *pretrained* transformers but, in the context of this work, these are constructed rather than trained. We refer to the matrices $\mathbf{W}_V, \mathbf{W}_Q, \mathbf{W}_K$ along with the parameters of the MLPs as *pretrained parameters*, and they are fixed throughout and not learnt. The prefix P is the only variable that can be modified to change the behavior of the model.

2.2. Universal Approximation

Let \mathcal{X} and \mathcal{Y} be normed vector spaces. We consider a family of *target functions* which is a subset \mathcal{C} of all mappings $\mathcal{X} \rightarrow \mathcal{Y}$, i.e., $\mathcal{C} \subseteq \mathcal{Y}^{\mathcal{X}}$, with \mathcal{C} is often referred to as a *concept space*. These are the relationships we wish to learn by some simpler candidate functions. Let us denote this set of candidates by $\mathcal{H} \subset \mathcal{Y}^{\mathcal{X}}$, called *hypothesis space*. The problem of approximation is concerned with how well functions in \mathcal{H} approximate functions in \mathcal{C} . There are two main ways to measure how well functions in \mathcal{H} represent functions in \mathcal{C} : density results and approximation rate results (Jiang et al., 2023). Density results show that, given an ϵ , one can find a hypothesis $h \in \mathcal{H}$ approximating any $f \in \mathcal{C}$ with error at most ϵ . Approximation rate results, also called *Jackson-type*, are stronger as they offer a measure of complexity for h to reach a desired precision ϵ . Classically, a Jackson-type

result would provide a minimum width or depth necessary for a neural network to reach a desired precision ϵ . In the context of the present work, the notion of complexity that we care about is the length N of the prefix P . Formally:

Definition 1 (Universal Approximation (Density-Type)). We say that \mathcal{H} is a universal approximator for \mathcal{C} over a compact set $S \subseteq \mathcal{X}$ if for every $f \in \mathcal{C}$ and every $\epsilon > 0$ there exists an $h \in \mathcal{H}$ such that $\sup_{x \in S} \|f(x) - h(x)\| \leq \epsilon$. One typically says that \mathcal{H} is dense in \mathcal{C} .

Lemma 1 (Transitivity). *If \mathcal{A} is dense in \mathcal{B} and \mathcal{B} is dense in \mathcal{C} , then \mathcal{A} is dense in \mathcal{C} .*

Definition 2 (Approximation Rate (Jackson-Type)). Fix a hypothesis space \mathcal{H} . Let $\{\mathcal{H}^N : N \in \mathbb{N}_+\}$ be a collection of subsets of \mathcal{H} such that $\mathcal{H}^N \subset \mathcal{H}^{N+1}$ and $\bigcup_{N \in \mathbb{N}_+} \mathcal{H}^N = \mathcal{H}$. Here, N is a measure of the complexity of the approximation candidates, and \mathcal{H}^N is the subset of hypotheses with complexity at most N . Then, the approximation rate estimate for \mathcal{C} over a compact $S \subseteq \mathcal{X}$ is a bound $Z_{\mathcal{H}}$:

$$N \geq Z_{\mathcal{H}}(f, \epsilon) \implies \inf_{h \in \mathcal{H}^N} \sup_{x \in S} \|f(x) - h(x)\| \leq \epsilon, \forall f \in \mathcal{C}.$$

$Z_{\mathcal{H}}$ gives an upper bound to the minimum hypothesis complexity necessary to reach the target precision ϵ and typically depends on the smoothness of f .

Lemma 2. *A Jackson bound for $\{\mathcal{H}^N \mid N \in \mathbb{N}_+\}$ with finite $Z_{\mathcal{H}}$ for all $f \in \mathcal{C}, \epsilon > 0$ immediately implies that $\bigcup_{N \in \mathbb{N}_+} \mathcal{H}^N = \mathcal{H}$ is dense in \mathcal{C} . Hence, Jackson bounds (Definition 2) are stronger than density results (Definition 1).*

The key hypothesis classes we consider in this work are the set of all prefixed attention heads and the set of prefixed transformers. This is very different from the classical universal approximation setting. The hypothesis classes in the classical universal approximation results consist of all possible parameter values of the model itself (Cybenko, 1989; Yun et al., 2019). When studying universal approximation with prefixing, the model parameters are fixed where prefixes are what can be modified.

Definition 3 (Prefixed Attention Heads Class). This is the class of all attention heads as defined in Equation (1) of dimension d , input/output sequence of length T , prefix of length at most N , and fixed pretrained components $\mathbf{H}, \mathbf{W}_V \in \mathbb{R}^{d \times d}$.

$$\mathcal{H}_{-,d}^{N,T}(\mathbf{H}, \mathbf{W}_V) = \left\{ \begin{array}{l} u : \mathbb{R}^{d \times (N+T)} \rightarrow \mathbb{R}^{d \times T}, \\ [u]_k \text{ as in (1), } \mathbf{p}_i \in \mathbb{R}^d, N' \leq N \end{array} \right\}.$$

For simplicity, we say that $\mathcal{H}_{-,d}^{N,T}$ is dense in \mathcal{C} to imply that there exists a pair $(\mathbf{H}, \mathbf{W}_V)$ such that $\mathcal{H}_{-,d}^{N,T}(\mathbf{H}, \mathbf{W}_V)$ is dense in \mathcal{C} . When considering all possible prefix lengths, we drop the N : $\mathcal{H}_{-,d}^T = \bigcup_{N \in \mathbb{N}} \mathcal{H}_{-,d}^{N,T}$.

Definition 4 (Prefixed Transformers Class). A transformer consists of L layers with each layer l consisting of an attention head with \mathbf{H}^l and \mathbf{W}_V^l followed by an MLP consisting of k_l linear layers, each parameterized as $\mathcal{L}_{k_l}^l(\mathbf{x}) =$

$\mathbf{A}^{l,k} \mathbf{x} + \mathbf{b}^{l,k}$ interspersed with non-linear activation σ . This gives rise to the following hypothesis class when prefixed:

$$\mathcal{H}_{\equiv,d}^{N,T} \left(\left\{ \mathbf{H}^l, \mathbf{W}_V^l, \{(\mathbf{A}^{l,k}, \mathbf{b}^{l,k})\}_{k=1}^{k_l} \right\}_{l=1}^L \right) = \left\{ \begin{array}{l} \mathcal{L}_{k_L}^L \circ \dots \circ \sigma \circ \mathcal{L}_{k_1}^1 \circ h^L \dots \circ h^2 \circ \mathcal{L}_{k_1}^1 \circ \dots \circ \sigma \circ \mathcal{L}_{k_1}^1 \circ h^1 \\ \text{with } h^l \in \mathcal{H}_{-,d}^{N',T}(\mathbf{H}^l, \mathbf{W}_V^l), l = 1, \dots, L, N' \leq N. \end{array} \right\},$$

applying linear layers \mathcal{L} element-wise. Again, we say $\mathcal{H}_{\equiv,d}^{N,T}$ is dense in \mathcal{C} , as a shorthand, to there exists $\{\mathbf{H}^l, \mathbf{W}_V^l, \{(\mathbf{A}^{l,k}, \mathbf{b}^{l,k})\}_{k=1}^{k_l}\}_{l=1}^L$ such that $\mathcal{H}_{\equiv,d}^{N,T}(\{\mathbf{H}^l, \mathbf{W}_V^l, \mathbf{A}^{l,k}, \mathbf{b}^{l,k}\}_{l=1}^L)$ is dense in \mathcal{C} .

In this paper, we consider several different concept classes. For reasons that will become apparent in the following section, we focus on functions whose domain is a hypersphere $S^m = \{\mathbf{y} \in \mathbb{R}^{m+1} \mid \|\mathbf{y}\|_2 = 1\} \subset \mathbb{R}^{m+1}$. We consider both scalar and vector-valued functions on the hypersphere.

Definition 5 (Scalar Functions on the Hypersphere). Define $C(S^m) \subset \mathbb{R}^{S^m}$ to be the space of all continuous functions defined on S^m with bounded norm, i.e.,

$$\|f\|_{\infty} = \sup_{\mathbf{x} \in S^m} |f(\mathbf{x})| < \infty, f \in C(S^m). \quad (2)$$

This is the concept class $\mathcal{C}_{s,m} = C(S^m) \subset \mathbb{R}^{S^m}$.

Definition 6 (Vector-valued Functions on the Hypersphere). The class of vector-valued functions on the hypersphere is:

$$\mathcal{C}_{v,m} = \{f : S^m \rightarrow \mathbb{R}^{m+1} \mid [f]_i \in C(S^m), i = 1, \dots, m+1\}.$$

Transformers are typically used to learn mappings over sequences rather than individual inputs. Hence, we define several sequence-to-sequence concept classes.

Definition 7 (General Sequence-to-sequence Functions). Given a fixed sequence length $T \in \mathbb{N}_{>0}$, we define the sequence-to-sequence function class as:

$$\mathcal{C}_{T,m} = \{f : (S^m)^T \rightarrow (\mathbb{R}^{m+1})^T \mid f \text{ continuous and bounded}\}.$$

We will also consider the subset of element-wise functions:

Definition 8 (Element-wise functions). Element-wise functions operate over sequences of inputs but apply the exact same function independently to all inputs:

$$\mathcal{C}_{\parallel,T,m} = \left\{ f \in \mathcal{C}_{T,m} \mid \begin{array}{l} \text{there exists } g \in \mathcal{C}_{v,m}, \text{ such that} \\ f(\mathbf{x}_1, \dots, \mathbf{x}_T) = (g(\mathbf{x}_1), \dots, g(\mathbf{x}_T)) \\ \text{for all } (\mathbf{x}_1, \dots, \mathbf{x}_T) \in (S^m)^T \end{array} \right\}.$$

3. Universal Approximation with a Single Attention Head

In this section, we will restrict ourselves to the setting when the input sequence is of length $T=1$, i.e., $X=(\mathbf{x})$. General sequence-to-sequence functions will be discussed in

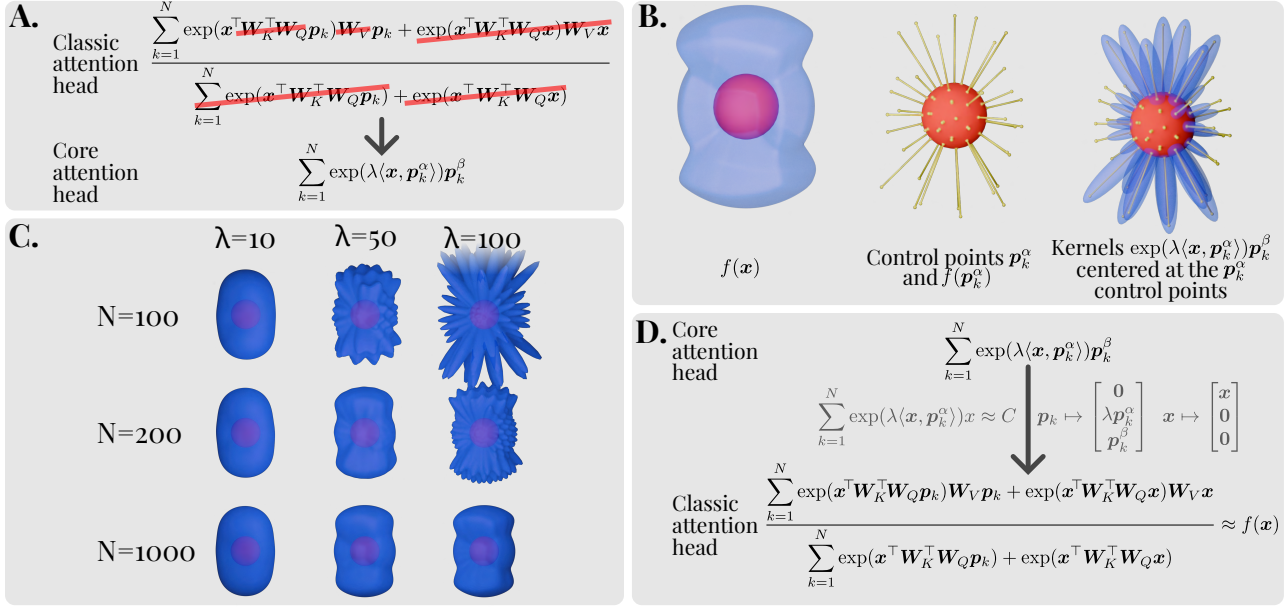


Figure 1. Approximating functions on the hypersphere with a single attention head. **A.** We simplify the classical attention head into a *core attention head*. **B.** The $\exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) \mathbf{p}_k^\beta$ terms act like kernels when \mathbf{x} is restricted to a hypersphere. We can approximate a function f by placing N control points $\mathbf{p}_1^\alpha, \dots, \mathbf{p}_N^\alpha$ and centering a kernel at each of them. **C.** Increasing λ results in less smoothing, while increasing N results in more control points and hence better approximation. With large enough λ and N , we can approximate f to any desired accuracy. **D.** With the normalization term in classical attention close to a constant, and giving \mathbf{x} , \mathbf{p}_k^α and \mathbf{p}_k^β orthogonal subspaces, core attention can be represented as classical attention. Hence, a classical attention head can also approximate f with arbitrary precision.

Section 4. We will show that a single attention head can approximate any continuous function on the hypersphere, or that $\mathcal{H}_{-,m+1}^1$ is dense in $\mathcal{C}_{s,m}$. To do this, we first simplify the classical attention head in Equation (1), resulting in what we call a *core attention head*. Then, we show that each of the terms in the core attention act as a kernel, meaning that it can approximate any function in $\mathcal{C}_{s,m}$. Finally, we show that any core attention head can be approximated by a classical attention head, hence, $\mathcal{H}_{-,m+1}^1$ is indeed dense in $\mathcal{C}_{s,m}$. The complete pipeline is illustrated in Figure 1.

To illuminate the approximation abilities of the attention head mechanism we relax it a bit. That is, we allow for different values of the prefix positions when computing the attention (the \exp terms in Equation (1)) and when computing the value (the right multiplication with \mathbf{W}_V). We will also drop the terms depending only on \mathbf{x} , set $\mathbf{H} = \lambda \mathbf{I}_d$, $\lambda > 0$, and $\mathbf{W}_V = \mathbf{I}_d$. We refer to this relaxed version as a *split attention head* with its corresponding hypothesis class:

$$h_{\parallel}(\mathbf{x}) = \frac{\sum_{k=1}^{N'} \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) \mathbf{p}_k^\beta}{\sum_{k=1}^{N'} \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle)}. \quad (3)$$

Definition 9 (Split Attention Head Class).

$$\mathcal{H}_{\parallel,d}^N = \left\{ h_{\parallel} \text{ as in (3), } \mathbf{p}_k^\alpha, \mathbf{p}_k^\beta \in \mathbb{R}^d, N' \leq N, \lambda > 0 \right\}.$$

We will later show that a split head can be represented by

a classical attention head. For now, let us simplify a bit further: we drop the denominator, resulting in our *core attention head*:

$$h_{\otimes}(\mathbf{x}) = \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) \mathbf{p}_k^\beta, \quad (4)$$

which gives rise to the hypothesis class:

$$\mathcal{H}_{\otimes,d}^N = \left\{ \mathbf{x} \mapsto \sum_{k=1}^{N'} \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) \mathbf{p}_k^\beta, \text{ where } \mathbf{p}_k^\alpha, \mathbf{p}_k^\beta \in \mathbb{R}^d, N' \leq N, \lambda > 0 \right\}.$$

We also have their scalar-valued counterparts:

$$h_{\circ}(\mathbf{x}) = \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) p_k^\beta, \quad (5)$$

$$\mathcal{H}_{\circ,d}^N = \left\{ \mathbf{x} \mapsto \sum_{k=1}^{N'} \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) p_k^\beta, \text{ where } p_k^\alpha \in \mathbb{R}^d, p_k^\beta \in \mathbb{R}, N' \leq N, \lambda > 0 \right\}.$$

As the dot product is a notion of similarity, one can interpret h_{\otimes} in Equation (4) and h_{\circ} in Equation (5) as interpolators. The \mathbf{p}_i^α vectors act as control points, while the \mathbf{p}_i^β vectors designate the output value at the location of the corresponding control point. The dot product with the input \mathbf{x} controls how much each control point should contribute to the final result, with control points closer to \mathbf{x} (larger dot product) contributing more.

Unfortunately, it is not generally true that higher dot product means smaller distance, hence the above interpretation fails

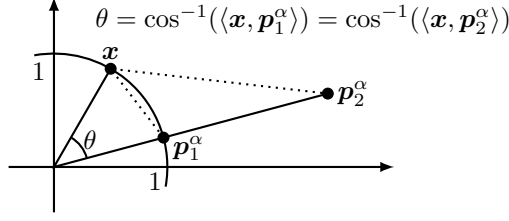


Figure 2. The dot product is a measure of closeness over the hypersphere. We want large dot product for points with lower distances. That is not the case for general $\mathbf{p}_1^\alpha, \mathbf{p}_2^\alpha \in \mathbb{R}^{m+1}$: above we show larger dot product for points which are further away, i.e., $\langle \mathbf{x}, \mathbf{p}_1^\alpha \rangle < \langle \mathbf{x}, \mathbf{p}_2^\alpha \rangle$ despite $\|\mathbf{x} - \mathbf{p}_1^\alpha\|_2 < \|\mathbf{x} - \mathbf{p}_2^\alpha\|_2$. However, if we restrict $\mathbf{x}, \mathbf{p}_i^\alpha$, and \mathbf{p}_j^α to the hypersphere S^m , then the dot product measures the cosine between \mathbf{x} and \mathbf{p}_i which is truly a measure of closeness: $\langle \mathbf{x}, \mathbf{p}_i^\alpha \rangle < \langle \mathbf{x}, \mathbf{p}_j^\alpha \rangle \iff \|\mathbf{x} - \mathbf{p}_i^\alpha\|_2 > \|\mathbf{x} - \mathbf{p}_j^\alpha\|_2$.

in \mathbb{R}^{m+1} . To see this, consider two control points $\mathbf{p}_1^\alpha, \mathbf{p}_2^\alpha \in \mathbb{R}^{m+1}$ such that $\mathbf{p}_2^\alpha = t\mathbf{p}_1^\alpha$, with $t > 1$. Then for $\mathbf{x} = \mathbf{p}_1^\alpha$ we would have $\langle \mathbf{x}, \mathbf{p}_1^\alpha \rangle = \|\mathbf{p}_1^\alpha\|_2^2 < \langle \mathbf{x}, \mathbf{p}_2^\alpha \rangle = t\|\mathbf{p}_1^\alpha\|_2^2$; the dot product is smaller for \mathbf{p}_1^α , the control point that is closer to \mathbf{x} , than for the much further away \mathbf{p}_2^α (see Figure 2). Therefore, the further away control point has a larger contribution than the closer point, which is at odds with the interpolation behaviour we desire. In general, the contribution of control points with larger norms will “dominate” the one of points with smaller norms. This has been observed for the attention mechanism in general by Demeter et al. (2020).

Fortunately, the domination of larger norm control points \mathbf{p}_i^α is not an issue if all control points have the same norm. In particular, if \mathbf{x} and \mathbf{p}_i^α lie on the unit hypersphere $S^m = \{\mathbf{y} \in \mathbb{R}^{m+1} \mid \|\mathbf{y}\|_2 = 1\}$ then $\langle \mathbf{x}, \mathbf{p}_i^\alpha \rangle = \cos(\angle(\mathbf{x}, \mathbf{p}_i^\alpha))$ and it has the desired property that the closer \mathbf{x} is to \mathbf{p}_i^α , the higher their dot product. By doing this, we restrict h_\otimes to be a function from the hypersphere S^m to \mathbb{R}^{m+1} . While this might seem artificial, modern transformer architectures do operate over hyperspheres as LayerNorm projects activations to S^m (Brody et al., 2023).

The central result of this section is that the functions in the form of Equation (5) can approximate any continuous function defined on the hypersphere, i.e., $\mathcal{H}_{\odot, m+1} = \bigcup_{N=1}^{\infty} \mathcal{H}_{\odot, m+1}^N$ is dense in $\mathcal{C}_{s, m}$ (Definition 5) and $\mathcal{H}_{\otimes, m+1} = \bigcup_{N=1}^{\infty} \mathcal{H}_{\otimes, m+1}^N$ is dense in $\mathcal{C}_{v, m}$ (Definition 6). Furthermore, we offer a Jackson-type approximation rate result which gives us a bound on the necessary prefix length N to achieve a desired approximation quality.

Theorem 1 (Jackson-type Bound for Universal Approximation on the Hypersphere). *Let $f \in C(S^m)$ be a continuous function on S^m , $m \geq 8$ with modulus of continuity*

$$\omega(f; t) = \sup\{|f(\mathbf{x}) - f(\mathbf{y})| \mid \mathbf{x}, \mathbf{y} \in S^m, \cos^{-1}(\langle \mathbf{x}, \mathbf{y} \rangle) \leq t\} \leq Lt,$$

for some $L > 0$. Then, for any $\epsilon > 0$, there exist

$\mathbf{p}_1^\alpha, \dots, \mathbf{p}_N^\alpha \in S^m$ and $p_1^\beta, \dots, p_N^\beta \in \mathbb{R}$ such that

$$\sup_{\mathbf{x} \in S^m} \left| f(\mathbf{x}) - \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) p_k^\beta \right| \leq \epsilon,$$

where $\lambda = \Lambda(\epsilon/2)$ with

$$\Lambda(\sigma) = \frac{(8LC_R + m\sigma + \sigma) \left(1 - \frac{\sigma^2}{8LC_H C_R + 2\sigma C_H}\right)^{\frac{\sigma}{4LC_R + \sigma}}}{\sigma \left(1 - \frac{\sigma^2}{8LC_H C_R + 2\sigma C_H}\right)^{\frac{2\sigma}{4LC_R + \sigma}}} = \mathcal{O}\left(\frac{L^3 C_H}{\sigma^4}\right),$$

and any $N \geq N(\lambda, \epsilon)$ with

$$N(\lambda, \epsilon) = \Phi(m) \left(\frac{3\pi(L + \lambda \|f\|_\infty)^{c_{m+1}(\lambda)} \exp(\lambda)}{\epsilon} \right)^{2(m+1)} = \mathcal{O}(\epsilon^{-10-14m-4m^2}),$$

with C_H being a constant depending on the smoothness of f (formally defined in the proof), C_R being a constant not depending on f or ϵ , $\Phi(m) = \mathcal{O}(m \log m)$ being a function that depends only on the dimension m and c_{m+1} being a normalization function.

Corollary 1. $\mathcal{H}_{\odot, m+1}$ is dense in $\mathcal{C}_{s, m}$

Proof. Theorem 1 holds for all $\epsilon > 0$ and Lemma 2. \square

Theorem 1 is a Jackson-type result as Equation (7) gives the number N of control points needed to approximate f with accuracy ϵ . This corresponds to the length of the prefix sequence. Moreover, the smoother the target f is, i.e., the smaller L, C_H , the shorter the prefix length N . Thus, our construction uses only as much prefix positions as necessary.

The proof of Theorem 1 follows closely (Ng and Kwong, 2022). While they only provide a density result, we offer a Jackson-type bound which is non-trivial and may be of an independent interest. The idea behind the proof is as following. We first approximate f with its convolution with a kernel having the form of the terms in Equation (5):

$$(f * K_\lambda^{\text{vMF}})(\mathbf{x}) = \int_{S^m} c_{m+1}(\lambda) \exp(\lambda \langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y}) d\omega_m(\mathbf{y}). \quad (8)$$

The larger the λ is, the closer $f * K_\lambda^{\text{vMF}}$ is to f and hence the smaller the approximation error (Menegatto, 1997). $\Lambda(\epsilon/2)$ gives the smallest value for λ such that this error is $\epsilon/2$. Equation (8) can then be approximated with sums: we partition S^m into N sets V_1, \dots, V_N small enough that f does not vary too much within each set. Each control point \mathbf{p}_k^α is placed in its corresponding V_k . Then, $\exp(\lambda \langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y})$ can be approximated with $\exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) f(\mathbf{p}_k^\alpha)$ when \mathbf{y} is in the k -th set V_k . Hence, Equation (8) can be approximated with $\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) C f(\mathbf{p}_k^\alpha)$ for some suitable constant C . By increasing N we can reduce the error of approximating the convolution with the sum. Equation (7) gives us the minimum N needed so that this error is $\epsilon/2$. Hence, we have error of at most $\epsilon/2$ from approximating f with the convolution and $\epsilon/2$ from approximating the convolution with the sum, resulting in our overall error being bounded

by ϵ . The full proof is in Appendix B and is illustrated in Figure 4. The theorem can be extended to vector-valued functions in $\mathcal{C}_{v,m}$ with a multiplicative factor $1/\sqrt{m+1}$:

Corollary 2. *Let $f : S^m \rightarrow \mathbb{R}^{m+1}$, $m \geq 8$ be such that each component f_i satisfies the conditions in Theorem 1. Define $\|f\|_\infty = \max_{1 \leq i \leq m+1} \|f_i\|_\infty$. Then, for any $\epsilon > 0$, there exist $\mathbf{p}_1^\alpha, \dots, \mathbf{p}_N^\alpha \in S^m$ and $\mathbf{p}_1^\beta, \dots, \mathbf{p}_N^\beta \in \mathbb{R}^{m+1}$ such that*

$$\sup_{\mathbf{x} \in S^m} \left\| f(\mathbf{x}) - \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) \mathbf{p}_k^\beta \right\|_2 \leq \epsilon,$$

with $\lambda = \Lambda(\epsilon/2\sqrt{m+1})$ for any $N \geq N(\lambda, \epsilon/\sqrt{m+1})$. That is, $\mathcal{H}_{\otimes, m+1}$ is dense in $\mathcal{C}_{v,m}$ with respect to the $\|\cdot\|_2$ norm.

Thanks to Theorem 1 and Corollary 2, we know that functions in $\mathcal{C}_{v,m}$ can be approximated by core attention (Equation (4)). We only have to demonstrate that a core attention head can be represented as a classical attention head (Equation (1)). We do this by reversing the simplifications we made when constructing the core attention head.

Let's start by bringing the normalization term back, resulting in $\mathcal{H}_{\parallel, d}^N$, the split attention head hypothesis (Definition 9). Intuitively, $\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle)$ is almost constant when the \mathbf{p}_k^α are uniformly distributed over the sphere as the distribution of distances from \mathbf{x} to \mathbf{p}_k^α will be similar, regardless of where \mathbf{x} lies. We can bound how far $\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle)$ is from being a constant and adjust the approximation error to account for it. Appendix C has the full proof.

Theorem 2. *Let $f : S^m \rightarrow \mathbb{R}^{m+1}$, $m \geq 8$ be such that each component f_i satisfies the conditions in Theorem 1. Then, for any $0 < \epsilon < 2\|f\|_\infty$, there exist $\mathbf{p}_1^\alpha, \dots, \mathbf{p}_N^\alpha \in S^m$ such that*

$$\sup_{\mathbf{x} \in S^m} \left\| f(\mathbf{x}) - \frac{\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle) \mathbf{p}_k^\beta}{\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle)} \right\|_2 \leq \epsilon,$$

with

$$\lambda = \Lambda\left(\frac{2\epsilon L}{2L + \|f\|_\infty}\right)$$

$$\mathbf{p}_k^\beta = f(\mathbf{p}_k^\alpha), \quad \forall k = 1, \dots, N,$$

for any $N \geq N(\lambda, \epsilon/\sqrt{m+1})$. That is, $\mathcal{H}_{\parallel, m+1}$ is dense in $\mathcal{C}_{v,m}$ with respect to the $\|\cdot\|_2$ norm.

An interesting observation is that adding the normalization term has not affected the asymptotic behavior of λ and hence also of the prefix length N . Furthermore, notice how the value \mathbf{p}_i^β at the control point \mathbf{p}_i^α is simply $f(\mathbf{p}_i^\alpha)$, the target function evaluated at this control point.

We ultimately care about the ability of the classical attention head (Definition 3) to approximate functions in $\mathcal{C}_{v,m}$ by prefixing. Hence, we need to bring back the terms depending only on the input \mathbf{x} , combine \mathbf{p}_k^α and \mathbf{p}_k^β parts into a single

prefix \mathbf{p}_k and bring back the \mathbf{H} and \mathbf{W}_V matrices. One can do this by considering an attention head with a hidden dimension $3(m+1)$ allowing us to place \mathbf{x} , \mathbf{p}_k^α and \mathbf{p}_k^β in different subspaces of the embedding space. To do this, define a pair of embedding and projection operations:

$$\Pi : S^m \rightarrow \mathbb{R}^{3(m+1)} \quad \Pi^{-1} : \mathbb{R}^{3(m+1)} \rightarrow \mathbb{R}^{m+1}$$

$$\mathbf{x} \mapsto \begin{bmatrix} \mathbf{I}_{m+1} \\ \mathbf{0}_{m+1} \\ \mathbf{0}_{m+1} \end{bmatrix} \mathbf{x} \quad \mathbf{x} \mapsto \begin{bmatrix} \mathbf{I}_{m+1} \\ \mathbf{0}_{m+1} \\ \mathbf{0}_{m+1} \end{bmatrix}^\top \mathbf{x}.$$

Lemma 3. $\Pi^{-1} \circ \mathcal{H}_{-, 3(m+1)}^1 \circ \Pi$ is dense in $\mathcal{H}_{\parallel, m+1}$, with the composition applied to each function in the class.

Proof. We can prove something stronger. For all $f \in \mathcal{H}_{\parallel, m+1}$ there exists a $g \in \mathcal{H}_{-, 3(m+1)}^1$ such that $f = \Pi^{-1} \circ g \circ \Pi$. If $f \in \mathcal{H}_{\parallel, m+1}$, then

$$f(\mathbf{x}) = \frac{\sum_{k=1}^N \mathbf{p}_k^\beta \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle)}{\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{p}_k^\alpha \rangle)}, \quad \forall \mathbf{x} \in S^m$$

for some $N, \lambda, \mathbf{p}_i^\alpha, \mathbf{p}_i^\beta$. Define:

$$\mathbf{p}_k = \begin{bmatrix} \mathbf{0} \\ \lambda \mathbf{p}_k^\alpha \\ \mathbf{p}_k^\beta \end{bmatrix} \in \mathbb{R}^{3(m+1)},$$

$$\mathbf{H} = \begin{bmatrix} M\mathbf{I} & \mathbf{I} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \mathbf{W}_V = \begin{bmatrix} \mathbf{0} & \mathbf{0} & \mathbf{I} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{bmatrix} \in \mathbb{R}^{3(m+1) \times 3(m+1)},$$

With M a negative constant tending to $-\infty$. Then:

$$g(\mathbf{x}) = \frac{\sum_{i=1}^N \exp(\mathbf{x}^\top \mathbf{H} \mathbf{p}_i) \mathbf{W}_V \mathbf{p}_i + \exp(\mathbf{x}^\top \mathbf{H} \mathbf{x}) \mathbf{W}_V \mathbf{x}}{\sum_{i=1}^N \exp(\mathbf{x}^\top \mathbf{H} \mathbf{p}_i) + \exp(\mathbf{x}^\top \mathbf{H} \mathbf{x})},$$

is in $\mathcal{H}_{-, 3(m+1)}^1$ and $f = \Pi^{-1} \circ g \circ \Pi$. As this holds for all $f \in \mathcal{H}_{\parallel, m+1}$, it follows that $\mathcal{H}_{\parallel, m+1} \subset \Pi^{-1} \circ \mathcal{H}_{-, 3(m+1)}^1 \circ \Pi$. Hence, $\Pi^{-1} \circ \mathcal{H}_{-, 3(m+1)}^1 \circ \Pi$ is dense in $\mathcal{H}_{\parallel, m+1}$. \square

Lemma 3 shows that every split attention head can be exactly represented as 3 times bigger classical attention head. Note that our choice for \mathbf{H} and \mathbf{W}_V is not unique. Equivalent constructions are available by multiplying each component by an invertible matrix, effectively changing the basis. Finally, the embedding and projection operations can be represented as MLPs and hence can be embedded in a transformer architecture. Now, we can provide the final result of this section, namely that the standard attention head of a transformer can approximate any vector-valued function on the hypersphere:

Theorem 3. Let $f : S^m \rightarrow \mathbb{R}^{m+1}$, $m \geq 8$ be such that each component f_i satisfies the conditions in Theorem 1. Then, for any $0 < \epsilon \leq 2\|f\|_\infty$, there exists an attention head $h \in \mathcal{H}_{-,3(m+1)}^{N,1}$ such that

$$\sup_{\mathbf{x} \in S^m} \|f(\mathbf{x}) - (\Pi^{-1} \circ h \circ \Pi)(\mathbf{x})\|_2 \leq \epsilon, \quad (9)$$

for any $N \geq N(\lambda, \epsilon/\sqrt{m+1})$. That is, $\Pi^{-1} \circ \mathcal{H}_{-,3(m+1)}^1 \circ \Pi$ is dense in $\mathcal{C}_{v,m}$ with respect to the $\|\cdot\|_2$ norm.

Proof. The density result follows directly from Theorem 2 and Lemma 3 and transitivity (Lemma 1). The Jackson bound is the same as in Theorem 2 as transforming the split attention head to a classical attention head is exact and does not contribute further error. \square

Therefore, we have shown that a single attention head with a hidden dimension $3(m+1)$ can approximate any continuous function $f : C(S^m) \rightarrow \mathbb{R}^{m+1}$ to an arbitrary accuracy. This is for *fixed* pretrained components, that is, \mathbf{H} and \mathbf{W}_V are as given in the proof of Lemma 3 and depend neither on the input \mathbf{x} nor on the target function f . Therefore, the behavior of the attention head is fully controlled by the prefix. This is a Jackson-type result, with the length N of the prefix given in Theorem 2. To the best of our knowledge, Theorem 3 is the first bound on the necessary prefix length to achieve a desired accuracy of function approximation using an attention head. Most critically, Theorem 3 demonstrates that attention heads are more expressive than commonly thought. A *single* attention head with a very simple structure can be a universal approximator.

4. Universal Approximation of Sequence-to-Sequence Functions

The previous section showed how we can approximate any continuous $f : S^m \rightarrow \mathbb{R}^{m+1}$ with a single attention head. Still, one typically uses the transformer architecture for operations over sequences rather than over single inputs (the case with $T \geq 1$). We will now show how we can leverage Theorem 3 to model general sequence-to-sequence functions. First, we show the simpler case of functions that apply the exact same mapping to all inputs. We then show how to model general sequence-to-sequence functions using a variant of the Kolmogorov–Arnold theorem.

Element-wise functions Theorem 3 can be extended to element-wise functions where the exact same function is applied to each element in the input sequence, i.e., the concept class $\mathcal{C}_{\parallel,T,m}$ from Definition 8. If $f \in \mathcal{C}_{\parallel,T,m}$, then there exists a $g \in \mathcal{C}_{v,m}$ such that $f(\mathbf{x}_1, \dots, \mathbf{x}_T) = (g(\mathbf{x}_1), \dots, g(\mathbf{x}_T))$. By Theorem 3, there exists a prefix $\mathbf{p}_1, \dots, \mathbf{p}_N$ that approximates g . As the construction in Lemma 3 prevents interactions between two dif-

ferent inputs \mathbf{x}_i and \mathbf{x}_j , an attention head $h^T \in \mathcal{H}_{-,3(m+1)}^{N,T}$ for a T -long input (Equation (1)) with the exact same prefix $\mathbf{p}_1, \dots, \mathbf{p}_N$ approximates f :

Corollary 3. $\Pi^{-1} \circ \mathcal{H}_{-,3(m+1)}^T \circ \Pi$ is dense in $\mathcal{C}_{\parallel,T,m}$ with respect to the $\|\cdot\|_2$ norm applied element-wise. That is, for every $\epsilon > 0$, there exists $h^T \in \mathcal{H}_{-,3(m+1)}^{N,T}$ such that:

$$\sup_{\{\mathbf{x}_i\} \in (S^m)^T} \max_{1 \leq k \leq T} \| [f(\{\mathbf{x}_i\}) - (\Pi^{-1} \circ h^T \circ \Pi)(\{\mathbf{x}_i\})]_k \|_2 \leq \epsilon,$$

with Π and Π^{-1} applied element-wise, $[\cdot]_k$ selecting the k -th element, and approximate rate bound on N as in Theorem 3.

General sequence-to-sequence functions Ultimately, we are interested in modeling arbitrary functions from sequences of inputs $(\mathbf{x}_1, \dots, \mathbf{x}_T)$ to sequences of outputs $(\mathbf{y}_1, \dots, \mathbf{y}_T)$, that is, the $\mathcal{C}_{T,m}$. We will use a version of the Kolmogorov–Arnold representation Theorem. The Theorem is typically defined on functions over the unit hypercube $[0, 1]^m$. As there exists a homeomorphism between $[0, 1]^m$ and a subset of S^m (Lemma D.1), for simplicity, we will ignore this technical detail. Our construction requires only $T + 2$ attention layers, each with a single head.

The original Kolmogorov–Arnold representation theorem (Kolmogorov, 1957) identifies every continuous function $f : [0, 1]^d \rightarrow \mathbb{R}$ with univariate functions $g_q, \psi_{p,q}$ such that:

$$f(x_1, \dots, x_d) = \sum_{q=0}^{2d} g_q \left(\sum_{p=1}^d \psi_{p,q}(x_p) \right).$$

In other words, multivariate functions can be represented as sums and compositions of univariate functions. As transformers are good at summing and attention heads are good at approximating functions, they can approximate functions of this form. However, g_q and $\psi_{p,q}$ are generally not well-behaved (Girosi and Poggio, 1989), so we will use the construction by Schmidt–Hieber (2021) instead.

Lemma 4 (Theorem 2 in (Schmidt–Hieber, 2021)). *For a fixed d , there exists a monotone functions $\psi : [0, 1] \rightarrow C$ (the Cantor set) such that for any function $f : [0, 1]^d \rightarrow \mathbb{R}$, we can find a function $g : C \rightarrow \mathbb{R}$ such that*

$$i. f(x_1, \dots, x_d) = g \left(3 \sum_{p=1}^d 3^{-p} \psi(x_p) \right), \quad (10)$$

- ii. if f is continuous, then g is also continuous,
- iii. if $|f(\mathbf{x}) - f(\mathbf{y})| \leq Q\|\mathbf{x} - \mathbf{y}\|_\infty$, for all $\mathbf{x}, \mathbf{y} \in [0, 1]^d$ and some Q , then $|g(x) - g(y)| \leq 2Q, \forall x, y \in C$.

In comparison with the original Kolmogorov–Arnold theorem, we need a single inner function ψ which does not depend on the target function f and only one outer function g . Furthermore, both ψ and g are Lipschitz. Hence, we can approximate them with our results from Section 3.

We need to modify Lemma 4 a bit to make it fit the sequence-to-sequence setting. First, flatten a sequence of T $(m+1)$ -dimensional vectors into a single vector in $[0, 1]^{(m+1)T}$. Second, define $\Psi_d : [0, 1]^d \rightarrow \mathbb{R}^d$ to be the element-wise application of ψ : $\Psi_d(\{x_i\}_{i=1}^d) = (\psi(x_i))_{i=1}^d$. We can also define $G_i : C \rightarrow \mathbb{R}^{m+1}$, $i = 1, \dots, T$ and extend Equation (10) for our setting:

$$\begin{aligned} f(\mathbf{x}_1, \dots, \mathbf{x}_T) &= (G_1(R), \dots, G_T(R)), \text{ with} \\ R &= 3 \sum_{i=1}^T 3^{-(i-1)(m+1)} \sum_{p=1}^{m+1} 3^p \psi(\mathbf{x}_{i,p}) \quad (11) \\ &= 3 \sum_{i=1}^T 3^{-(i-1)(m+1)} \begin{bmatrix} 3^{-1} \\ \vdots \\ 3^{-(m+1)} \end{bmatrix}^\top \Psi_{m+1}(\mathbf{x}_i). \end{aligned}$$

Equation (11) can now be represented with a transformer with $T+2$ attention layers. Ψ_{m+1} is applied element-wise, hence, all $\Psi_{m+1}(\mathbf{x}_i)$ can be computed in parallel with a single attention head (Corollary 3). The dot product with the $[3^{-1} \dots 3^{-(m+1)}]$ vector can be computed using a single MLP. The product with the $3^{-(i-1)(m+1)}$ scalar is a bit more challenging as it depends on the position in the sequence. However, if we concatenate position encodings to the input, another MLP can use them to compute this factor and the multiplication. The outer sum over the T inputs and the multiplication by 3 can be achieved with a single attention head. Hence, using only 2 attention layers, we have compressed the whole sequence in a single scalar R .¹

The only thing left is to apply G_1, \dots, G_T to R to compute each of the T outputs. As each one of these is Lipschitz, we can approximate each with a single attention head using Theorem 3. Each G_i is different and would need its own set of prefixes, requiring T attention heads arranged in T attention layers. Using the positional encodings, each layer can compute the output for its corresponding position and pass the input unmodified for the other positions. The overall prefix size would be the longest of the prefixes necessary to approximate $\Psi_{m+1}, G_1, \dots, G_T$.

Hence, we have constructed an architecture that can approximate any sequence-to-sequence function $f \in \mathcal{C}_{T,m}$ with only $T+2$ attention layers. Thus, $\mathcal{H}_{\equiv, d}^T$ is dense in $\mathcal{C}_{T,m}$.

5. Discussion and Conclusions

Comparison with prior work Just like us, Wang et al. (2023) show that prefix-tuning can be a universal approximator. Their approach relies on discretizing the input space and the set of sequence-to-sequence functions to a given

¹Yun et al. (2019) use a similar approach but use discretization to enumerate all possible sequences and require $\mathcal{O}(\epsilon^{-m})$ attention layers. In our continuous setting, R is computed with 2 layers.

precision depending on ϵ , resulting in a finite number of pairs of functions and inputs, each having a unique corresponding output. Then, using the results of Yun et al. (2019), they construct a *meta-transformer* which maps each of the function-input pairs to their corresponding output. This approach has several limitations: i) the model has exponential depth $\mathcal{O}(T\epsilon^{-m})$; ii) reducing the approximation error ϵ requires increasing the model depth; iii) the prefix length is fixed, hence a constant function and a highly non-smooth function would have equal prefix lengths, and iv) it effectively has memorized all possible functions and inputs, explaining the exponential size of their constructions. In contrast, we show that memorization is not needed: attention heads are naturally suited for universal approximation. Section 4 showed that $T+2$ layers are enough, we require shorter prefixes for more smooth functions and reducing the approximation error ϵ can be done by increasing the prefix length, without modifying the pretrained model.

Petrov et al. (2024) have shown that prefix-tuning cannot change the relative attention patterns over the input tokens and hence cannot learn tasks with new attention patterns. This appears to be a limitation but Von Oswald et al. (2023) and Akyürek et al. (2022) proved that there exist attention heads that can learn any linear model, samples of which are given as a prefix. In this work, we showed the existence of a “universal” attention head (\mathbf{H} and \mathbf{W}_V in Lemma 3) that can be used to emulate any new function defined as a prefix.

Prefixes have been observed to have larger norms than token embeddings (Bailey et al., 2023). Our results provide an explanation to that. While the control points \mathbf{p}_k^α are in S^m and hence have norm 1, in Lemma 3 we fold λ into them. Recall that the less smooth f is, the higher the concentration parameter λ has to be in order to reduce the influence of one control point on the locations far from it. Hence, the less smooth f is, the larger the norm of the prefixes.

Connection to prompting and safety implications

While this work focused on prefix-tuning, the results can extend to prompting. Observe that prefix tuning (where we have a distinct prefix) can be reduced to soft prompting (where only the first layer is prefixed) by using an appropriate attention mechanism and position embeddings. Hence, if a function $f \in \mathcal{C}_{T,m}$ requires N prefixes to be approximated to precision ϵ with prefix-tuning, it would require $\mathcal{O}(TN)$ soft tokens to be approximated with soft prompting. Finally, observe that a soft token can be encoded with a sequence of hard tokens, the number of hard tokens per soft token depends on the required precision and the vocabulary size V . Hence, f could be approximated with $\mathcal{O}(\log_V(\epsilon^{-1})mTN)$ hard tokens. Therefore, our universal approximation results may translate to prompting. This raises concerns as to whether it is at all possible to prevent a transformer model from exhibiting undesirable behaviors

(Zou et al., 2023; Wolf et al., 2023; Chao et al., 2023). Furthermore, this means that transformer-based agents might have the technical possibility to collude in undetectable and uninterpretable manner (de Witt et al., 2023). Still, our results require specific form of the attention and value matrices and, hence, it is not clear whether these risk translate to real-world models.

Prefix-Tuning and Prompting a Pretrained Transformer might be Less efficient than Training it Typically, with neural networks one expects that the number of trainable parameters would grow as $\mathcal{O}(\epsilon^{-m})$ (Schmidt-Hieber, 2021). Indeed that is the case for universal approximation with a transformer when one learns the key, query and value matrices and the MLP parameters as shown by Yun et al. (2019). However, as Equation (7) shows, our construction results in the trainable parameters (prefix length in our case) growing as $\mathcal{O}(\epsilon^{-10-14m-4m^2})$. That m^2 term indicates worse asymptotic efficiency of prefix-tuning and prompting compared to training a transformer. However, our approach may not be tight. Thus, it remains an open question if a tighter Jackson bound exists or if prefix-tuning and prompting inherently require more trainable parameters to reach the same approximation accuracy as training a transformer.

Prefix-tuning and prompting may work by combining prefix-based element-wise maps with pretrained cross-element mixing The construction for general sequence-to-sequence functions in Section 4 is highly unlikely to occur in transformers pretrained on real data as it requires very specific parameter values. While the element-wise setting (Corollary 3) is more plausible, it cannot approximate general sequence-to-sequence functions. Hence, neither result explains why prefix-tuning works in practice. To this end, we hypothesise that prompting and prefix-tuning, can modify how single tokens are processed (akin to fine-tuning only MLPs), while the cross-token information mixing happens with pretrained attention patterns. Therefore, prompting and prefix-tuning can easily learn novel tasks as long as no new attention patterns are required. Our findings suggest a method for guaranteeing that a pretrained model possesses the capability to act as a token-wise universal approximator. This can be achieved by ensuring each layer of the model includes at least one attention head conforming to the structure in Lemma 3.

Limitations. We assume a highly specific pretrained model which is unlikely to occur in practice when pretraining with real-world data. Hence, the question of, given a real-world pretrained transformer, which is the class of functions it can approximate with prefix-tuning is still open. This is an inverse (Bernstein-type, Jiang et al. 2023) bound and is considerably more difficult to derive.

Impact Statement

This paper presents theoretical understanding about how the approximation abilities of the transformer architecture. Our results show that, under some conditions, prompting and prefix-tuning can arbitrarily modify the behavior of a model. This may have implications on how we design safety and security measures for transformer-based systems. However, whether these theoretical risks could manifest in realistic pretrained models remains an open problem.

Acknowledgements

We would like to thank Tom Lamb for spotting several mistakes and helping us rectify them. This work is supported by a UKRI grant Turing AI Fellowship (EP/W002981/1) and the EPSRC Centre for Doctoral Training in Autonomous Intelligent Machines and Systems (EP/S024050/1). AB has received funding from the Amazon Research Awards. We also thank the Royal Academy of Engineering and FiveAI.

References

- Ekin Akyürek, Dale Schuurmans, Jacob Andreas, Tengyu Ma, and Denny Zhou. 2022. **What learning algorithm is in-context learning? Investigations with linear models.** In *International Conference on Learning Representations*.
- Silas Alberti, Niclas Dern, Laura Thesing, and Gitta Kutyniok. 2023. **Sumformer: Universal approximation for efficient transformers.** In *Proceedings of 2nd Annual Workshop on Topology, Algebra, and Geometry in Machine Learning (TAG-ML)*.
- Donald E Amos. 1974. **Computation of modified Bessel functions and their ratios.** *Mathematics of Computation*, 28(125):239–251.
- Kendall Atkinson and Weimin Han. 2012. *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*.
- Yogesh J Bagul and Satish K Panchal. 2018. **Certain inequalities of Kober and Lazarević type.** *Research Group in Mathematical Inequalities and Applications Research Report Collection*, 21(8).
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. **Neural machine translation by jointly learning to align and translate.** In *International Conference on Learning Representations*.
- Luke Bailey, Gustaf Ahdritz, Anat Kleiman, Siddharth Swaroop, Finale Doshi-Velez, and Weiwei Pan. 2023. **Soft prompting might be a bug, not a feature.** In *Workshop on Challenges in Deployable Generative AI at International Conference on Machine Learning*.

- Alex Barnett. 2021. [Lower bounds on the modified Bessel function of the first kind](#). Mathematics Stack Exchange.
- Andrew R Barron. 1993. [Universal approximation bounds for superpositions of a sigmoidal function](#). *IEEE Transactions on Information Theory*, 39(3):930–945.
- Shaked Brody, Uri Alon, and Eran Yahav. 2023. [On the expressivity role of LayerNorm in transformers’ attention](#). *arXiv preprint arXiv:2305.02582*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. [Language models are few-shot learners](#). *Advances in Neural Information Processing Systems*.
- Patrick Chao, Alexander Robey, Edgar Dobriban, Hamed Hassani, George J. Pappas, and Eric Wong. 2023. [Jail-breaking black box large language models in twenty queries](#). *arXiv preprint arXiv:2310.08419*.
- George Cybenko. 1989. [Approximation by superpositions of a sigmoidal function](#). *Mathematics of control, signals and systems*, 2(4):303–314.
- Feng Dai and Yuan Xu. 2013. *Approximation Theory and Harmonic Analysis on Spheres and Balls*.
- Christian Schroeder de Witt, Samuel Sokota, J. Zico Kolter, Jakob Foerster, and Martin Strohmeier. 2023. [Perfectly secure steganography using minimum entropy coupling](#). In *International Conference on Learning Representations*.
- David Demeter, Gregory Kimmel, and Doug Downey. 2020. [Stolen probability: A structural weakness of neural language models](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*.
- Puneesh Deora, Rouzbeh Ghaderi, Hossein Taheri, and Christos Thrampoulidis. 2023. [On the optimization and generalization of multi-head attention](#). *arXiv preprint arXiv:2310.12680*.
- Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. 2021. [Attention is not all you need: Pure attention loses rank doubly exponentially with depth](#). In *International Conference on Machine Learning*.
- Ricardo Estrada. 2014. [On radial functions and distributions and their Fourier transforms](#). *Journal of Fourier Analysis and Applications*, 20(2):301–320.
- Uriel Feige and Gideon Schechtman. 2002. [On the optimality of the random hyperplane rounding technique for MAX CUT](#). *Random Structures & Algorithms*, 20(3):403–440.
- Paul Funk. 1915. Beiträge zur Theorie der Kugelfunktionen. *Mathematische Annalen*, 77:136–152.
- Federico Girosi and Tomaso Poggio. 1989. [Representation properties of networks: Kolmogorov’s theorem is irrelevant](#). *Neural Computation*, 1(4):465–469.
- E Hecke. 1917. Über orthogonal-invariante Integralgleichungen. *Mathematische Annalen*, 78:398–404.
- Kurt Hornik, Maxwell Stinchcombe, and Halbert White. 1989. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-efficient transfer learning for NLP](#). In *International Conference on Machine Learning*.
- Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2021. [LoRA: Low-rank adaptation of large language models](#). In *International Conference on Learning Representations*.
- Zhiqiang Hu, Yihuai Lan, Lei Wang, Wanyu Xu, Ee-Peng Lim, Roy Ka-Wei Lee, Lidong Bing, and Soujanya Poria. 2023. [LLM-Adapters: An adapter family for parameter-efficient fine-tuning of large language models](#). *arXiv preprint arXiv:2304.01933*.
- Haotian Jiang and Qianxiao Li. 2023. [Approximation theory of transformer networks for sequence modeling](#). *arXiv preprint arXiv:2305.18475*.
- Haotian Jiang, Qianxiao Li, Zhong Li, and Shida Wang. 2023. [A brief survey on the approximation theory for sequence modelling](#). *arXiv preprint arXiv:2302.13752*.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). *Advances in Neural Information Processing Systems*.
- Andrei Nikolaevich Kolmogorov. 1957. On the representation of continuous functions of many variables by superposition of continuous functions of one variable and addition. In *Doklady Akademii Nauk*, volume 114, pages 953–956. Russian Academy of Sciences.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*.
- Shengqiao Li. 2010. [Concise formulas for the area and volume of a hyperspherical cap](#). *Asian Journal of Mathematics & Statistics*, 4(1):66–70.

- Xiang Lisa Li and Percy Liang. 2021. **Prefix-Tuning: Optimizing continuous prompts for generation**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*.
- Vladislav Lialin, Vijeta Deshpande, and Anna Rumshisky. 2023. **Scaling down to scale up: A guide to parameter-efficient fine-tuning**. *arXiv preprint arXiv:2303.15647*.
- Valerii Likhoshesterov, Krzysztof Choromanski, and Adrian Weller. 2021. **On the expressive power of self-attention matrices**. *arXiv preprint arXiv:2106.03764*.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2023. **Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing**. *ACM Computing Surveys*.
- Xiao Liu, Kaixuan Ji, Yicheng Fu, Weng Tam, Zhengxiao Du, Zhilin Yang, and Jie Tang. 2022. **P-Tuning: Prompt tuning can be comparable to fine-tuning across scales and tasks**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*.
- Sadegh Mahdavi, Renjie Liao, and Christos Thrampoulidis. 2023. **Memorization capacity of multi-head attention in transformers**. *arXiv preprint arXiv:2306.02010*.
- Valdir Antônio Menegatto. 1997. **Approximation by spherical convolution**. *Numerical Functional Analysis and Optimization*, 18(9-10):995–1012.
- Tin Lok James Ng and Kwok-Kun Kwong. 2022. **Universal approximation on the hypersphere**. *Communications in Statistics – Theory and Methods*, 51(24):8694–8704.
- Aleksandar Petrov, Philip HS Torr, and Adel Bibi. 2024. **When do prompting and prefix-tuning work? A theory of capabilities and limitations**. In *International Conference on Learning Representations*.
- David L Ragozin. 1971. **Constructive polynomial approximation on spheres and projective spaces**. *Transactions of the American Mathematical Society*, 162:157–170.
- Sylvestre-Alvise Rebuffi, Hakan Bilen, and Andrea Vedaldi. 2017. **Learning multiple visual domains with residual adapters**. In *Advances in Neural Information Processing Systems*.
- C. A. Rogers. 1963. **Covering a sphere with spheres**. *Mathematika*, 10(2):157–164.
- Clayton Sanford, Daniel Hsu, and Matus Telgarsky. 2023. **Representational strengths and limitations of transformers**. *arXiv preprint arXiv:2306.02896*.
- Anton Maximilian Schäfer and Hans Georg Zimmermann. 2006. **Recurrent neural networks are universal approximators**. In *Artificial Neural Networks–ICANN 2006: 16th International Conference*.
- Johannes Schmidt-Hieber. 2021. **The Kolmogorov–Arnold representation theorem revisited**. *Neural Networks*, 137:119–126.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. **AutoPrompt: Eliciting knowledge from language models with automatically generated prompts**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.
- Matus Telgarsky. 2015. **Representation benefits of deep feedforward networks**. *arXiv preprint arXiv:1509.08101*.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. **Attention is all you need**. In *Advances in Neural Information Processing Systems*.
- Johannes Von Oswald, Eyvind Niklasson, Ettore Randazzo, Joao Sacramento, Alexander Mordvintsev, Andrey Zhmoginov, and Max Vladymyrov. 2023. **Transformers learn in-context by gradient descent**. In *International Conference on Machine Learning*.
- Yihan Wang, Jatin Chauhan, Wei Wang, and Cho-Jui Hsieh. 2023. **Universality and limitations of prompt tuning**. In *Advances in Neural Information Processing Systems*.
- Jason Wei, Maarten Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M Dai, and Quoc V Le. 2021. **Finetuned language models are zero-shot learners**. In *International Conference on Learning Representations*.
- Yotam Wolf, Noam Wies, Oshri Avnery, Yoav Levine, and Amnon Shashua. 2023. **Fundamental limitations of alignment in large language models**. *arXiv preprint arXiv:2304.11082*.
- Sang Michael Xie, Aditi Raghunathan, Percy Liang, and Tengyu Ma. 2021. **An explanation of in-context learning as implicit Bayesian inference**. In *International Conference on Learning Representations*.
- Steve Yadlowsky, Lyric Doshi, and Nilesh Tripuraneni. 2023. **Pretraining data mixtures enable narrow model selection capabilities in transformer models**. *arXiv preprint arXiv:2311.00871*.

Chulhee Yun, Srinadh Bhojanapalli, Ankit Singh Rawat, Sashank Reddi, and Sanjiv Kumar. 2019. *Are transformers universal approximators of sequence-to-sequence functions?* In *International Conference on Learning Representations*.

Andy Zou, Zifan Wang, J Zico Kolter, and Matt Fredrikson. 2023. *Universal and transferable adversarial attacks on aligned language models.* *arXiv preprint arXiv:2307.15043*.

A. Background on Analysis on the Sphere

As mentioned in the main text, the investigation of the properties of attention heads naturally leads to analysing functions over the hypersphere. To this end, our results require some basic facts about the analysis on the hypersphere. We will review them in this appendix. For a comprehensive reference, we recommend (Atkinson and Han, 2012) and (Dai and Xu, 2013).

Define $\mathbb{P}_k(\mathbb{R}^{m+1})$ to be the space of polynomials of degree at most k . The restriction of a polynomial $p \in \mathbb{P}_k(\mathbb{R}^{m+1})$ to the unit hypersphere $S^m = \{\mathbf{x} \in \mathbb{R}^{m+1} \mid \|\mathbf{x}\|_2 = 1\}$ is called a *spherical polynomial*. We can thus define the space of spherical polynomials:

$$\mathbb{P}_k(S^m) = \{p|_{S^m} \text{ for } p \in \mathbb{P}_k(\mathbb{R}^{m+1})\}.$$

Define by $\mathbb{H}_k(\mathbb{R}^{m+1})$ the space of polynomials of degree k that are homogeneous:

$$\mathbb{H}_k(\mathbb{R}^{m+1}) = \text{span} \left\{ (x_1, \dots, x_{m+1}) \mapsto x_1^{\alpha_1} \times \dots \times x_{m+1}^{\alpha_{m+1}} \mid \sum_{i=1}^{m+1} \alpha_i = k \right\}.$$

Its restriction to the sphere $\mathbb{H}_k(S^m)$ is defined analogously to $\mathbb{P}_k(S^m)$. Finally, we can define the space $\mathbb{Y}(\mathbb{R}^{m+1})$ of harmonic homogeneous polynomials:

$$\mathbb{Y}_k(\mathbb{R}^{m+1}) = \left\{ p \in \mathbb{H}_k(\mathbb{R}^{m+1}) \mid \frac{\partial^2}{\partial x^2} p(\mathbf{x}) = 0, \forall \mathbf{x} \in \mathbb{R}^{m+1} \right\}.$$

$\mathbb{Y}_k(S^m)$ which is the restriction of $\mathbb{Y}_k(\mathbb{R}^{m+1})$ to S^m is the set of *spherical harmonics* of degree k . Spherical harmonics are the higher-dimensional extension of Fourier series.

Notably, even though

$$\mathbb{Y}_k(S^m) \subset \mathbb{H}_k(S^m) \subset \mathbb{P}_k(S^m),$$

the restriction of any polynomial on S^m is a sum of spherical harmonics:

$$\mathbb{P}_k(S^m) = \mathbb{Y}_0(S^m) \oplus \dots \oplus \mathbb{Y}_k(S^m),$$

with \oplus being the direct sum (Atkinson and Han, 2012, Corollary 2.19).

We define $C(S^m)$ to be the space of all continuous functions defined on S^m with the uniform norm

$$\|f\|_\infty = \sup_{\mathbf{x} \in S^m} |f(\mathbf{x})|, \quad f \in C(S^m). \quad (12)$$

Similarly, $\mathcal{L}_p(S^m), 1 \leq p < \infty$ is the space of all functions defined on S^m which are integrable with respect to the standard surface measure dw_m . The norm in this space is:

$$\|f\|_p = \left(\frac{1}{w_m} \int_{S^m} |f(\mathbf{x})|^p dw_m(\mathbf{x}) \right)^{1/p}, \quad f \in \mathcal{L}_p(S^m), \quad (13)$$

with the surface area being

$$w_m = \int_{S^m} dw_m = \frac{2\pi^{(m+1)/2}}{\Gamma((m+1)/2)}. \quad (14)$$

We will use V_m to denote any of these two spaces and $\|\cdot\|_m$ the corresponding norm.

A key property of spherical harmonics is that sums of spherical harmonics can uniformly approximate the functions in $C(S^m)$. In other words, the span of $\bigcup_{k=0}^{\infty} \mathbb{Y}_k(S^m)$ is dense in $C(S^m)$ with respect to the uniform norm $\|\cdot\|_\infty$. Hence, any $f \in C(S^m)$ can be expressed as a series of spherical harmonics:

$$f(\mathbf{x}) = \sum_{k=0}^{\infty} Y_k^m(\mathbf{x}), \quad \text{with } Y_k^m \in \mathbb{Y}_k(S^m), \quad \forall k.$$

We will also make a heavy use of the concept of spherical convolutions. Define the space of kernels $\mathcal{L}^{1,m}$ to consist of all measurable functions K on $[-1, 1]$ with norm

$$\|K\|_{1,m} = \frac{w_{m-1}}{w_m} \int_{-1}^1 |K(t)|(1-t^2)^{(m-2)/2} dt < \infty.$$

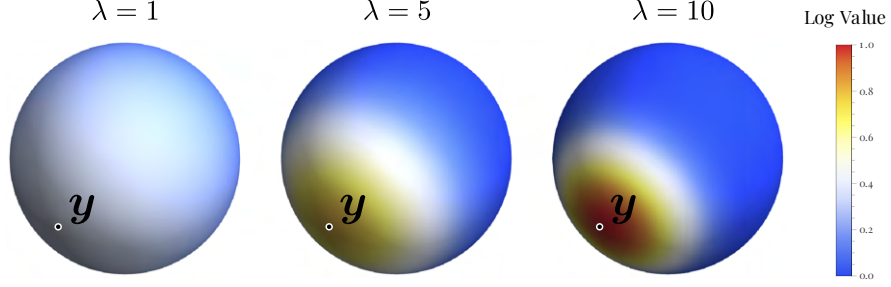


Figure 3. Plots of the von Mises-Fisher kernel $K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle)$ for $\lambda = 1, 5, 10$ and fixed \mathbf{y} in three dimensions ($m = 2$). The larger λ is, the more concentrated the kernel is around \mathbf{y} .

Definition A.1 (Spherical convolution). The spherical convolution $K * f$ of a kernel K in $\mathcal{L}^{1,m}$ with a function $f \in V_m$ is defined by:

$$(K * f)(\mathbf{x}) = \frac{1}{w_m} \int_{S^m} K(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y}) dw_m(\mathbf{y}), \quad \mathbf{x} \in S^m.$$

Spherical convolutions map functions $f \in V_m$ to functions in V_m . Furthermore, the spherical harmonics are eigenfunctions of the function generated by a kernel in $\mathcal{L}^{1,m}$:

Lemma A.1 (Funk and Hecke’s formula (Funk, 1915; Hecke, 1917; Estrada, 2014)).

$$K * Y_k^m = a_k^m(K) Y_k^m, \quad \text{when } K \in \mathcal{L}^{1,m}, Y_k^m \in \mathbb{Y}_k(S^m), k = 0, 1, \dots,$$

where $a_k^m(K)$ are the coefficients in the series expansion in terms of Gegenbauer polynomials associated with the kernel K :

$$a_k^m(K) = \frac{w_{m-1}}{w_m} \int_{-1}^1 K(t) \frac{Q_k^{(m-1)/2}(t)}{Q_k^{(m-1)/2}(1)} (1-t^2)^{(m-2)/2} dt, \quad k = 0, 1, \dots \quad (15)$$

Here, $Q_k^{(m-1)/2}$ is the Gegenbauer polynomial of degree k .

Note also that with a change of variables we have:

$$\int_{S^m} K(\langle \mathbf{x}, \mathbf{y} \rangle) dw_m(\mathbf{y}) = w_{m-1} \int_{-1}^1 K(t) (1-t^2)^{(m-2)/2} dt. \quad (16)$$

Ideally, we would like a kernel that acts as an identity for the convolution operation. In this case, we would have $\|K * f - f\|_\infty = 0, f \in \mathcal{C}(S^m)$ which would be rather convenient. However, there is no such kernel in the spherical setting (Menegatto, 1997). The next best thing is to construct a sequence of kernels $\{K_n\} \in \mathcal{L}^{1,m}$ such that $\|K_n * f - f\|_m \rightarrow 0$ as $n \rightarrow \infty$ for all $f \in V_m$. This sequence of kernels is called an *approximate identity*. The specific such sequence of kernels we will use is based on the von Mises-Fisher distribution as this gives us the $\exp(\langle \mathbf{x}, \mathbf{y} \rangle)$ form that we also observe in the transformer attention mechanism.

Definition A.2 (von Mises-Fisher kernels, (Ng and Kwong, 2022)). We define the sequence of von Mises-Fisher kernels as:

$$K_\lambda^{\text{vMF}}(t) = c_{m+1}(\lambda) \exp(\lambda t), \quad t \in [-1, 1],$$

where

$$c_{m+1}(\lambda) = \frac{w_m \lambda^{\frac{m+1}{2}-1}}{(2\pi)^{\frac{m+1}{2}} I_{\frac{m+1}{2}-1}(\lambda)},$$

with I_v being the modified Bessel function at order v .

Note that a von Mises-Fisher kernel can also be expressed in terms of points on S^m . In particular, for a fixed $\mathbf{y} \in S^m$ we have $K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle), \mathbf{x} \in S^m$. The parameter λ is a “peakiness” parameter: the large λ is, the closer $K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle)$ approximates

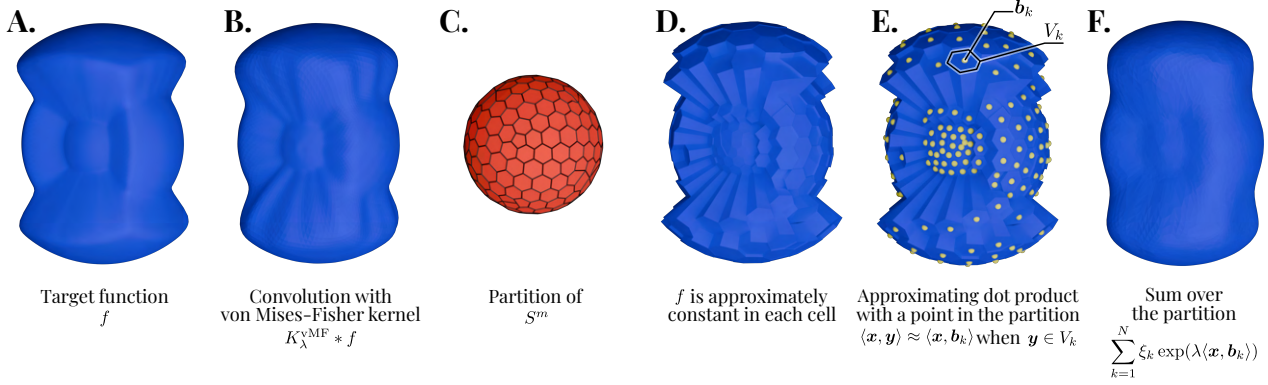


Figure 4. Intuition behind the proof of our Jackson-type bound for universal approximation on the hypersphere. **A.** We want to approximate a function f over the hypersphere S^m . This illustration is in three-dimensional space, so $m = 2$. **B.** In order to get the $\exp(\lambda \langle \cdot, \mathbf{y} \rangle)$ form that we want, we convolve f with the $K_\lambda^{\text{vMF}}(t) = c_{m+1}(\lambda) \exp(\lambda t)$ kernel. **C.** We partition S^m into N cells V_1, \dots, V_N . **D.** Our choice of N is such that f does not vary too much in each cell and hence can be approximated by a function that is constant in each V_k . **E.** As each cell is small, the dot product of \mathbf{x} with any point in the cell V_k can be approximated by the dot product of \mathbf{x} with a fixed point $\mathbf{b}_k \in V_k$. **F.** This allows us to approximate the integral in the convolution K_λ^{vMF} with a finite sum.

the delta function centered at \mathbf{y} , as can be seen in Figure 3. It is easy to check that $\|K_\lambda^{\text{vMF}}\|_{1,m} = 1$, $\forall \lambda > 1, m > 1$ and hence the sequence is in $\mathcal{L}^{1,m}$, meaning they are valid kernels. Ng and Kwong (2022, Lemma 4.2) show that $\{K_\lambda^{\text{vMF}}\}$ is indeed an approximate identity, i.e., $\|K_\lambda^{\text{vMF}} * f - f\|_m \rightarrow 0$ as $\lambda \rightarrow \infty$ for all $f \in V_m$.² As we want a Jackson-type result however, we will need to upper bound the error $\|K_\lambda^{\text{vMF}} * f - f\|_m$ as a function of λ , that is a non-asymptotic result on the quality of the approximation by spherical convolutions with K_λ^{vMF} . We do that in Lemma B.5.

B. A Jackson-type Bound for Universal Approximation on the Unit Hypersphere

The overarching goal in this section is to provide a Jackson-type (Definition 2) bound for approximating functions $f : S^m \rightarrow \mathbb{R}^{m+1}$ on the hypersphere $S^m = \{\mathbf{x} \in \mathbb{R}^{m+1} \mid \|\mathbf{x}\|_2 = 1\}$ by functions of the form

$$h(\mathbf{x}) = \sum_{k=1}^N \xi_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \quad (17)$$

To this end, we will leverage results from approximation on the hypersphere using spherical convolutions by Menegatto (1997) and recent results on the universal approximation on the hypersphere by Ng and Kwong (2022). While these two works inspire the general proof strategy, they only offer uniform convergence (i.e., density-type results, Definition 1). Instead, we offer a non-asymptotic analysis and develop the first approximation rate results on the sphere for functions of the form of Equation (17), i.e., Jackson-type results (Definition 2).

The high-level idea of the proof is to split the goal into approximating f with the convolution $f * K_\lambda^{\text{vMF}}$ and approximating the convolution $f * K_\lambda^{\text{vMF}}$ with a sum of terms that have the $\exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)$ structure resembling the kernel K_λ^{vMF} (Definition A.2):

$$\sup_{\mathbf{x} \in S^m} \left\| f(\mathbf{x}) - \sum_{k=1}^N \xi_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right\| \leq \underbrace{\|f - f * K_\lambda^{\text{vMF}}\|_\infty}_{\text{Equation (19) / Lemma B.5}} + \underbrace{\left\| f * K_\lambda^{\text{vMF}} - \sum_{k=1}^N \xi_k \exp(\lambda \langle \cdot, \mathbf{b}_k \rangle) \right\|_\infty}_{\text{Lemma B.7}}. \quad (18)$$

This is also illustrated in Figure 4.

Let's focus on the first term in Equation (18). It can be further decomposed into three terms by introducing $W_q \in \mathbb{P}_q(S^m)$,

²The w_m term in the normalization constant $c_{m+1}(\lambda)$ is not in (Ng and Kwong, 2022). However, without it K_λ^{vMF} are not an approximate identity.

the best approximation of f with a spherical polynomial of degree q :

$$\|K_\lambda^{\text{vMF}} * f - f\|_\infty \leq \underbrace{\|K_\lambda^{\text{vMF}} * f - K_\lambda^{\text{vMF}} * W_q\|_\infty}_{\text{Lemma B.2}} + \underbrace{\|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty}_{\text{Lemma B.4}} + \underbrace{\|W_q - f\|_\infty}_{\text{Lemma B.1}}. \quad (19)$$

There are a number of Jackson-type results for how well finite sums of spherical polynomials approximate functions $f \in V_m$ (the last term in Equation (19)). In particular, they are interested in bounding

$$\min_{W_q \in \mathbb{P}_q(S^m)} \|f - W_q\|_p, \quad 1 \leq p \leq \infty. \quad (20)$$

We will use a simple bound by [Ragozin \(1971\)](#):

Lemma B.1 (Ragozin bound). *For $f \in C(S^m)$ and $q \in \mathbb{N}_{>0}$ it holds that:*

$$\min_{W_q \in \mathbb{P}_q(S^m)} \|f - W_q\|_\infty \leq C_R \omega\left(f; \frac{1}{q}\right), \quad (21)$$

for some constant C_R that does not depend on f or q and ω being the first modulus of continuity of f defined as:

$$\omega(f; t) = \sup\{|f(\mathbf{x}) - f(\mathbf{y})| \mid \mathbf{x}, \mathbf{y} \in S^m, \cos^{-1}(\mathbf{x}^\top \mathbf{y}) \leq t\}.$$

We recommend [Atkinson and Han \(2012, Chapter 4\)](#) and [Dai and Xu \(2013, Chapter 4\)](#) for an overview of the various bounds proposed for Equation (20) depending on the continuity properties of f and its derivatives. In particular, the above bound could be improved with a term $1/n^k$ if f has k continuous derivatives ([Ragozin, 1971](#)).

We can upper-bound the first term in Equation (19) by recalling that the norm of the kernel K_λ^{vMF} is 1:

Lemma B.2.

$$\|K_\lambda^{\text{vMF}} * f - K_\lambda^{\text{vMF}} * W_q\|_m \leq \|f - W_q\|_m.$$

Hence:

$$\|K_\lambda^{\text{vMF}} * f - K_\lambda^{\text{vMF}} * W_q\|_\infty \leq \|f - W_q\|_\infty \leq C_R \omega\left(f; \frac{1}{q}\right).$$

Proof. Convolution is linear so $\|K_\lambda^{\text{vMF}} * f - K_\lambda^{\text{vMF}} * W_q\|_m = \|K_\lambda^{\text{vMF}} * (f - W_q)\|_m$. Using the Hölder inequality ([Dai and Xu, 2013, Theorem 2.1.2](#)) we get $\|K_\lambda^{\text{vMF}} * f - K_\lambda^{\text{vMF}} * W_q\|_m \leq \|K_\lambda^{\text{vMF}}\|_{1,m} \|f - W_q\|_m$. As $\|K_\lambda^{\text{vMF}}\|_{1,m} = 1$ for all $\lambda > 0, m > 1$, we obtain the inequality in the lemma. For the uniform norm, we also use the Ragozin bound from [Lemma B.1](#). \square

Only the second term in Equation (19) is left. However, before we tackle it, we will need a helper lemma that bounds the eigenvalues of the von Mises-Fisher kernel (Equation (15)):

Lemma B.3 (Bounds on the eigenvalues $a_k^m(K_\lambda^{\text{vMF}})$). *The eigenvalues a_k^m , as defined in Equation (15), for the sequence of von Mises-Fisher kernels (Definition A.2) are bounded from below and above as:*

$$0 < \left(\frac{\lambda}{\left(\frac{m-1}{2} + k\right) + \sqrt{\lambda^2 + \left(\frac{m-1}{2} + k\right)^2}} \right)^k \leq a_k^m(K_\lambda^{\text{vMF}}) \leq 1$$

Proof. We have

$$\begin{aligned} a_k^m(K_\lambda^{\text{vMF}}) &= \frac{w_{m-1}}{w_m} \int_{-1}^1 K_\lambda^{\text{vMF}}(t) \frac{Q_k^{(m-1)/2}(t)}{Q_k^{(m-1)/2}(1)} (1-t^2)^{(m-2)/2} dt \\ &= \frac{w_{m-1}}{w_m} \int_{-1}^1 c_{m+1}(\lambda) \exp(\lambda t) \frac{Q_k^{(m-1)/2}(t)}{Q_k^{(m-1)/2}(1)} (1-t^2)^{(m-2)/2} dt \end{aligned}$$

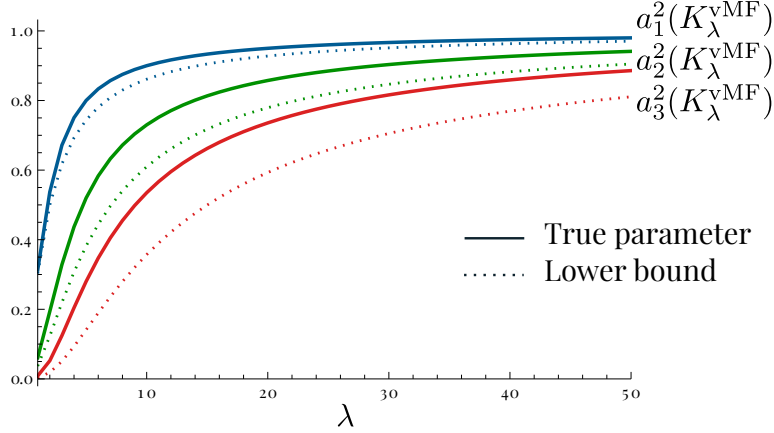


Figure 5. The coefficients a_k^m for the von Mises-Fisher kernels K_λ^{vMF} for $m = 2$ and $k \in \{1, 2, 3\}$ as well as the lower bound from Lemma B.3.

$$\stackrel{\star}{=} \frac{I_{\frac{m-1}{2}+k}(\lambda)}{I_{\frac{m-1}{2}}(\lambda)},$$

with \star solved using Mathematica. From here we can see that $a_0^m(K_\lambda^{\text{vMF}}) = 1$ for all $m > 1, \lambda > 1$. Furthermore, for $v > 1$ and $\lambda > 0$ the modified Bessel function of the first kind $I_v(\lambda)$ is monotonically decreasing as v increases. Therefore, $a_k^m(K_\lambda^{\text{vMF}}) \leq a_0^m(K_\lambda^{\text{vMF}}) = 1$, which gives us the upper bound in the lemma.

For the lower bound, we will use the following bound on the ratio of modified Bessel functions by Amos (1974, Eq. 9):

$$\frac{I_{v+1}(x)}{I_v(x)} \geq \frac{x}{(v+1) + \sqrt{x^2 + (v+1)^2}}, \quad v \geq 0, x \geq 0.$$

As mentioned above, $0 \leq \frac{I_{v+1}(x)}{I_v(x)} \leq 1$. Furthermore, these ratios are decreasing as v increases, i.e., $\frac{I_{v+2}(x)}{I_{v+1}(x)} \leq \frac{I_{v+1}(x)}{I_v(x)}$ for all $v \geq 0$ and $x \geq 0$ (Amos, 1974, Eq. 10). Combining these facts gives us:

$$\frac{I_{v+k}(x)}{I_v(x)} \geq \left(\frac{I_{v+k}(x)}{I_{v+k-1}(x)} \right)^k = \left(\frac{x}{(v+k) + \sqrt{x^2 + (v+k)^2}} \right)^k. \quad (22)$$

We can now give the lower bound for $a_k^m(K_n^{\text{vMF}})$ using Equation (22):

$$a_k^m(K_\lambda^{\text{vMF}}) = \frac{I_{\frac{m-1}{2}+k}(\lambda)}{I_{\frac{m-1}{2}}(\lambda)} \geq \left(\frac{\lambda}{\left(\frac{m-1}{2} + k\right) + \sqrt{\lambda^2 + \left(\frac{m-1}{2} + k\right)^2}} \right)^k.$$

The lower bound for $m = 2$ is plotted in Figure 5. □

We can now provide a bound for the second term in Equation (19):

Lemma B.4. Take an $f \in C(S^m)$. Furthermore, assume that there exists a constant $C_H \geq 0$ that upper-bounds the norms of the spherical harmonics of any best polynomial approximation W_q of f :

$$\text{for all } q \geq 1, W_q = \sum_{k=0}^q Y_k^m, Y_k^m \in \mathbb{Y}(S^m), \|Y_k^m\|_\infty \leq C_H, \forall k = 0, \dots, q, \text{ where } W_q = \arg \min_{h \in \mathbb{P}_q(S^m)} \|f - h\|_\infty.$$

Then

$$\|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty \leq C_H q \left(1 - \left(\frac{\lambda}{\left(\frac{m-1}{2} + q\right) + \sqrt{\lambda^2 + \left(\frac{m-1}{2} + q\right)^2}} \right)^q \right).$$

Proof. Using that W_q is a spherical polynomial of degree q and hence can be expressed as a sum of spherical harmonics $W_q = \sum_{k=0}^q Y_k^m$, we get:

$$\begin{aligned}
 \|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty &= \left\| K_\lambda^{\text{vMF}} * \sum_{k=0}^q Y_k^m(x) - \sum_{k=0}^q Y_k^m(x) \right\|_\infty \\
 &= \left\| \sum_{k=0}^q (K_\lambda^{\text{vMF}} * Y_k^m(x) - Y_k^m(x)) \right\|_\infty \\
 &\leq \sum_{k=0}^q \| (K_\lambda^{\text{vMF}} * Y_k^m(x) - Y_k^m(x)) \|_\infty && \text{(Triangle inequality)} \\
 &= \sum_{k=0}^q \| (a_k^m(K_\lambda^{\text{vMF}}) Y_k^m(x) - Y_k^m(x)) \|_\infty && \text{(from Lemma A.1)} \\
 &= \sum_{k=0}^q \| (a_k^m(K_\lambda^{\text{vMF}}) - 1) Y_k^m(x) \|_\infty \\
 &= \sum_{k=0}^q |a_k^m(K_\lambda^{\text{vMF}}) - 1| \|Y_k^m(x)\|_\infty \\
 &\leq \sum_{k=0}^q |a_k^m(K_\lambda^{\text{vMF}}) - 1| C_H \\
 &\leq C_H \sum_{k=0}^q \left(1 - \left(\frac{\lambda}{\left(\frac{m-1}{2} + k \right) + \sqrt{\lambda^2 + \left(\frac{m-1}{2} + k \right)^2}} \right)^k \right) && \text{(from Lemma B.3)} \\
 &\leq C_H \sum_{k=0}^q \left(1 - \underbrace{\left(\frac{\lambda}{\left(\frac{m-1}{2} + q \right) + \sqrt{\lambda^2 + \left(\frac{m-1}{2} + q \right)^2}} \right)^k}_B \right) && \text{(as } k \leq q) \\
 &= C_H \left(q + 1 - \sum_{k=0}^q B^k \right) \\
 &\leq C_H (q + 1 - (1 + qB^q)) && \text{(using that } 0 < B < 1) \\
 &\leq C_H q (1 - B^q).
 \end{aligned}$$

□

We can finally combine Lemmas B.1, B.2 and B.4 in order to provide an upper bound to Equation (19):

Lemma B.5 (Bound on $\|K_\lambda^{\text{vMF}} * f - f\|_\infty$). *Take an $f \in V_m$ with modulus of continuity $\omega(f; t) \leq Lt$. As in Lemma B.4, assume that there exists a constant $C_H \geq 0$ that upper-bounds the norms of the spherical harmonics of any best polynomial approximation W_q of f :*

$$\text{for all } q \geq 1, W_q = \sum_{k=0}^q Y_k^m, Y_k^m \in \mathbb{Y}(S^m), \|Y_k^m\|_\infty \leq C_H, \forall k = 0, \dots, q, \text{ where } W_q = \arg \min_{h \in \mathbb{P}_q(S^m)} \|f - h\|_\infty.$$

If $\lambda \geq \Lambda(\epsilon)$, where

$$\Lambda(\epsilon) = \frac{(8LC_R + m\epsilon + \epsilon) \left(1 - \frac{\epsilon^2}{8LC_H C_R + 2\epsilon C_H} \right)^{\frac{4LC_R + \epsilon}{4LC_R + \epsilon}}}{\epsilon \left(1 - \left(1 - \frac{\epsilon^2}{8LC_H C_R + 2\epsilon C_H} \right)^{\frac{2\epsilon}{4LC_R + \epsilon}} \right)} = \mathcal{O} \left(\frac{L^3 C_H C_R^3}{\epsilon^4} \right). \quad (23)$$

then $\|f - K_\lambda^{\text{vMF}} * f\|_\infty \leq \epsilon$.

Proof. As we want to upper-bound Equation (19) with ϵ , we will split our ϵ budget over the three terms.

For the first and the third terms, using the Ragozin bound from Lemma B.1 we have:

$$\|f - W_q\|_\infty \leq C_R \omega\left(f; \frac{1}{q}\right) \leq \frac{C_R L}{q}, \quad q \geq 1. \quad (24)$$

We want to select an integer q large enough so that $\|f - W_q\|_\infty \leq \epsilon/4$. That is $q = \lceil \frac{4C_R L}{\epsilon} \rceil$. This will be how we bound the first and last terms in Equation (19).

Let's focus on the second term. Pick W_q to be the best approximation from the Ragozin bound. From Lemma B.4 we have that

$$\|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty \leq C_H q (1 - B^q),$$

where

$$B = \frac{\lambda}{\left(\frac{m-1}{2} + q\right) + \sqrt{\lambda^2 + \left(\frac{m-1}{2} + q\right)^2}}.$$

The error budget we want to allocate for the $\|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty$ term is $\epsilon/2$. Hence:

$$B \geq \left(1 - \frac{\epsilon}{2C_H q}\right)^{1/q} = D \implies \|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty \leq \frac{\epsilon}{2}. \quad (25)$$

We just need to find the minimum value for λ such that Equation (25) holds. We have:

$$B = \frac{\lambda}{\underbrace{\left(\frac{m-1}{2} + q\right)}_E + \sqrt{\lambda^2 + \left(\frac{m-1}{2} + q\right)^2}} = \frac{\lambda}{E + \sqrt{\lambda^2 + E^2}}. \quad (26)$$

Then, combining Equations (25) and (26) we get:

$$\begin{aligned} \frac{\lambda}{E + \sqrt{\lambda^2 + E^2}} &\geq D \\ \lambda &\geq \frac{2DE}{1 - D^2}. \end{aligned}$$

Finally, replacing D and E with the expressions in Equations (25) and (26), upper-bounding $q = \lceil \frac{4C_R L}{\epsilon} \rceil$ as $q \geq \frac{4C_R L}{\epsilon} + 1$ and simplifying the expression we get our final bound for λ . If $\lambda \geq \Lambda(\epsilon)$, with

$$\Lambda(\epsilon) = \frac{(8LC_R + m\epsilon + \epsilon) \left(1 - \frac{\epsilon^2}{8LC_H C_R + 2\epsilon C_H}\right)^{\frac{\epsilon}{4LC_R + \epsilon}}}{\epsilon \left(1 - \left(1 - \frac{\epsilon^2}{8LC_H C_R + 2\epsilon C_H}\right)^{\frac{2\epsilon}{4LC_R + \epsilon}}\right)},$$

then $\|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty \leq \epsilon/2$.

Hence, for any $\lambda \geq \Lambda(\epsilon)$ we have:

$$\begin{aligned} \|K_\lambda^{\text{vMF}} * f - f\|_\infty &\leq \|K_\lambda^{\text{vMF}} * f - K_\lambda^{\text{vMF}} * W_q\|_\infty + \|K_\lambda^{\text{vMF}} * W_q - W_q\|_\infty + \|W_q - f\|_\infty \\ &\leq \|f - W_q\|_\infty + \frac{\epsilon}{2} + \|W_q - f\|_\infty \\ &\leq \frac{\epsilon}{4} + \frac{\epsilon}{2} + \frac{\epsilon}{4} \\ &= \epsilon. \end{aligned}$$

This concludes our bound on Equation (19).

Finally, to give the asymptotic behavior of $\Lambda(\epsilon)$ as $\epsilon \rightarrow 0$ we observe that the Taylor series expansion of Λ around $\epsilon = 0$ is:

$$\Lambda(\epsilon) = \frac{128L^3 C_H C_R^3}{\epsilon^4} + \frac{16L^2(m-1)C_H C_R^2}{\epsilon^3} + \mathcal{O}\left(\frac{1}{\epsilon^2}\right),$$

hence:

$$\Lambda(\epsilon) = \mathcal{O}\left(\frac{L^3 C_H C_R^3}{\epsilon^4}\right).$$

□

Lemma B.5 is our bound on Equation (19) which is also the first term of Equation (18). Recall that bounding Equation (18) is our ultimate goal. Hence, we are halfway done with our proof. Let's focus now on the second term in Equation (18), that is, how well we can approximate the convolution of f with the von Mises-Fisher kernel using a finite sum:

$$\left\| f * K_\lambda^{\text{vMF}} - \sum_{k=1}^N \xi_k \exp(\lambda \langle \cdot, \mathbf{b}_k \rangle) \right\|_\infty.$$

The basic idea behind bounding this term is that we can partition the hypersphere S^m into N sets $(\{V_1, \dots, V_N\})$, each small enough so that, for a fixed $\mathbf{x} \in S^m$, the $K(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y})$ term in the convolution

$$(K * f)(\mathbf{x}) = \frac{1}{w_m} \int_{S^m} K(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y}) dw_m(\mathbf{y})$$

is almost the same for all values $\mathbf{y} \in V_k$ in that element of the partition. Hence we can approximate the integral over the partition with estimate over a single point \mathbf{b}_k :

$$\int_{V_k} K(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y}) dw_m(\mathbf{y}) \approx |V_k| K(\langle \mathbf{x}, \mathbf{b}_k \rangle) f(\mathbf{b}_k).$$

The rest of this section will make this formal.

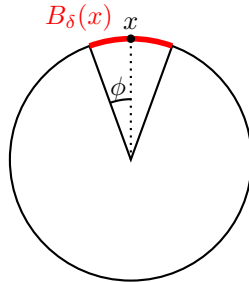
First, in order to construct our partition $\{V_1, \dots, V_N\}$ of S^m we will first construct a cover of S^m . Then, our partition will be such that each element V_k is a subset of the corresponding element of the cover of S^m . In this way, we can control the maximum size of the elements of the cover.

Lemma B.6. Consider a cover $\{B_\delta(\mathbf{b}_1), \dots, B_\delta(\mathbf{b}_{N_\delta^m})\}$ of S^m by N_δ^m hyperspherical caps $B_\delta(\mathbf{x}) = \{\mathbf{y} \in S^m \mid \langle \mathbf{x}, \mathbf{y} \rangle \geq 1 - \delta\}$ for $0 < \delta < 1$, centred at $\mathbf{b}_1, \dots, \mathbf{b}_{N_\delta^m} \in S^m$. By cover we mean that $\bigcup_{i=1}^{N_\delta^m} B_\delta(\mathbf{b}_i) = S^m$, with N_δ^m being the smallest number of hyperspherical caps to cover S^m (its covering number). Then, for $m \geq 8$ we have:

$$\frac{2}{I_{(\delta(2-\delta))}(\frac{m}{2}, \frac{1}{2})} \leq N_\delta^m < \frac{\Phi(m)}{(\delta(2-\delta))^{\frac{m+1}{2}}} < \frac{\Phi(m)}{\delta^{m+1}},$$

with I being the regularized incomplete beta function and $\Phi(m) = \mathcal{O}(m \log m)$ being a function that depends only on the dimension m .

Proof. Define $\phi = \cos^{-1}(1 - \delta)$:



Naturally, the area of the caps needs to be at least as much as the area of the hypersphere for the set of caps to be a cover. This gives us our lower bound. The area of a cap with colatitude angle ϕ as above is (Li, 2010):

$$w_m^\phi = \frac{1}{2} w_m I_{\sin^2 \phi} \left(\frac{m}{2}, \frac{1}{2} \right).$$

As $\sin \phi = \sin(\cos^{-1}(1 - \delta)) = \sqrt{1 - (1 - \delta)^2} = \sqrt{\delta(2 - \delta)}$, we have our lower bound:

$$N_\delta^m \geq \frac{w_m}{w_m^\phi} = \frac{2}{I_{\sin^2 \phi}(\frac{m}{2}, \frac{1}{2})} = \frac{2}{I_{(\delta(2-\delta))}(\frac{m}{2}, \frac{1}{2})}$$

For the upper bound, we can use the observation that if a unit *ball* is covered with *balls* of radius r , then the unit *sphere* is also covered with *caps* of radius r . From (Rogers, 1963, intermediate result from the proof of theorem 3) we have that for $m \geq 8$ and $1/r \geq m + 1$, a unit ball can be covered by less than

$$e((m + 1) \log(m + 1) + (m + 1) \log \log(m + 1) + 5(m + 1)) \frac{1}{r^{m+1}} = \Phi(m) \frac{1}{r^{m+1}}$$

balls of radius r . For high m , this is a pretty good approximation since most of the volume of the hypersphere lies near its surface. Our caps B_δ can fit inside balls of radius $r = \sin \phi$. Hence, we have the upper bound:

$$N_\delta^m < \frac{\Phi(m)}{\sin^{m+1} \phi} = \frac{\Phi(m)}{(\delta(2 - \delta))^{\frac{m+1}{2}}}.$$

□

Now we can use a partition resulting from this covering in order to bound the error between the integral and its Riemannian sum approximation:

Lemma B.7 (Approximation via Riemann sums). *Let $g(x, y) : S^m \times S^m \rightarrow \mathbb{R}$ with $m \geq 8$ be a continuous function with modulus of continuity for both arguments $\omega(g(\cdot; \mathbf{y}), t) \leq Lt, \forall \mathbf{y} \in S^m$ and $\omega(g(\mathbf{x}, \cdot); t) \leq Lt, \forall \mathbf{x} \in S^m$. Take any $0 < \delta < 1$. Then, there exists a partition $\{V_1, \dots, V_{N_\delta^m}\}$ of S^m into $N_\delta^m = \lceil \Phi(m)/\delta^{m+1} \rceil$ subsets, as well as $\mathbf{b}_1, \dots, \mathbf{b}_{N_\delta^m} \in S^m$ such that:*

$$\max_{\mathbf{x} \in S^m} \frac{1}{w_m} \left| \int_{S^m} g(\mathbf{x}, \mathbf{y}) dw_m(\mathbf{y}) - \sum_{k=1}^{N_\delta^m} g(\mathbf{x}, \mathbf{b}_k) w_m(V_k) \right| \leq 3L \cos^{-1}(1 - \delta).$$

Here, $\Phi(m) = \mathcal{O}(m \log m)$ is a function that depends only on the dimension m .

Proof. This proof is a non-asymptotic version of the proof of Lemma 4.3 from Ng and Kwong (2022). First, we can use Lemma B.6 to construct a covering $\{B_\delta(\mathbf{b}_1), \dots, B_\delta(\mathbf{b}_{N_\delta^m})\}$ of S^m . If we have a covering of S^m it is trivial to construct a partition of it $\{V_1, \dots, V_{N_\delta^m}\}, \bigcup_{k=1}^{N_\delta^m} V_k = S^m, V_i \cap V_j = \emptyset, i \neq j$ such that $V_k \subseteq B_\delta(\mathbf{b}_k), \forall k$. This partition can also be selected to be such that all elements of it have the same measure $w_m(V_1) = w_m(V_i), \forall i$ (Feige and Schechtman, 2002, Lemma 21). While this is not necessary for this proof, we will use this equal measure partition in Lemma C.2 and Theorem C.1.

We can then use the triangle inequality to split the term we want to bound in three separate terms:

$$\begin{aligned} \left| \int_{S^m} g(\mathbf{x}, \mathbf{y}) dw_m(\mathbf{y}) - \sum_{k=1}^{N_\delta^m} g(\mathbf{x}, \mathbf{b}_k) w_m(V_k) \right| &\leq \left| \int_{S^m} g(\mathbf{x}, \mathbf{y}) dw_m(\mathbf{y}) - \int_{S^m} g(\mathbf{b}_*, \mathbf{y}) dw_m(\mathbf{y}) \right| \\ &+ \left| \int_{S^m} g(\mathbf{b}_*, \mathbf{y}) dw_m(\mathbf{y}) - \sum_{k=1}^{N_\delta^m} g(\mathbf{b}_*, \mathbf{b}_k) w_m(V_k) \right| \\ &+ \left| \sum_{k=1}^{N_\delta^m} g(\mathbf{b}_*, \mathbf{b}_k) w_m(V_k) - \sum_{k=1}^{N_\delta^m} g(\mathbf{x}, \mathbf{b}_k) w_m(V_k) \right|, \end{aligned} \quad (27)$$

where \mathbf{b}_* is the center of one of the caps whose corresponding partition contains \mathbf{x} , i.e., $\mathbf{b}_* = \mathbf{b}_i \iff \mathbf{x} \in V_i$. Due to $\{V_1, \dots, V_{N_\delta^m}\}$ being a partition, \mathbf{b}_* is well-defined as \mathbf{x} is in exactly one of the elements of the partition.

Observe also that the modulus of continuity gives us a Lipschitz-like bound, i.e., if $\langle \mathbf{x}, \mathbf{y} \rangle \geq 1 - \delta$ for $\mathbf{x}, \mathbf{y} \in S^m$ and $\omega(f; t) \leq Lt$, then

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq \omega(f; \cos^{-1}(\langle \mathbf{x}, \mathbf{y} \rangle)) \leq \omega(f; \cos^{-1}(1 - \delta)) \leq L \cos^{-1}(1 - \delta). \quad (28)$$

Let's start with the first term in Equation (27). Using the fact that we selected \mathbf{b}_* to be such that $\langle \mathbf{b}_*, \mathbf{x} \rangle \geq 1 - \delta$ and Equation (28), we have:

$$\begin{aligned} \left| \int_{S^m} g(\mathbf{x}, \mathbf{y}) dw_m(\mathbf{y}) - \int_{S^m} g(\mathbf{b}_*, \mathbf{y}) dw_m(\mathbf{y}) \right| &= \left| \int_{S^m} (g(\mathbf{x}, \mathbf{y}) - g(\mathbf{b}_*, \mathbf{y})) dw_m(\mathbf{y}) \right| \\ &\leq \int_{S^m} |g(\mathbf{x}, \mathbf{y}) - g(\mathbf{b}_*, \mathbf{y})| dw_m(\mathbf{y}) \\ &\leq \int_{S^m} L \cos^{-1}(1 - \delta) dw_m(\mathbf{y}) \\ &= L \cos^{-1}(1 - \delta) \int_{S^m} dw_m(\mathbf{y}) \\ &= L \cos^{-1}(1 - \delta) w_m. \end{aligned}$$

We can similarly upper-bound the second term of Equation (27) using also the fact that $\{V_k\}$ is a partition of S^m :

$$\begin{aligned} \left| \int_{S^m} g(\mathbf{b}_*, \mathbf{y}) dw_m(\mathbf{y}) - \sum_{k=1}^{N_\delta^m} g(\mathbf{b}_*, \mathbf{b}_k) w_m(V_k) \right| &= \left| \sum_{k=1}^{N_\delta^m} \int_{V_k} g(\mathbf{b}_*, \mathbf{y}) dm_w(\mathbf{y}) - \sum_{k=1}^{N_\delta^m} g(\mathbf{b}_*, \mathbf{b}_k) w_m(V_k) \right| \\ &= \left| \sum_{k=1}^{N_\delta^m} \int_{V_k} (g(\mathbf{b}_*, \mathbf{y}) - g(\mathbf{b}_*, \mathbf{b}_k)) dm_w(\mathbf{y}) \right| \\ &\leq \sum_{k=1}^{N_\delta^m} \int_{V_k} |g(\mathbf{b}_*, \mathbf{y}) - g(\mathbf{b}_*, \mathbf{b}_k)| dm_w(\mathbf{y}) \\ &\leq \sum_{k=1}^{N_\delta^m} \int_{V_k} L \cos^{-1}(1 - \delta) dm_w(\mathbf{y}) \\ &= L \cos^{-1}(1 - \delta) \sum_{k=1}^{N_\delta^m} \int_{V_k} dm_w(\mathbf{y}) \\ &= L \cos^{-1}(1 - \delta) w_m. \end{aligned}$$

And analogously, for the third term we get:

$$\begin{aligned} \left| \sum_{k=1}^{N_\delta^m} g(\mathbf{b}_*, \mathbf{b}_k) w_m(V_k) - \sum_{k=1}^{N_\delta^m} g(\mathbf{x}, \mathbf{b}_k) w_m(V_k) \right| &= \left| \sum_{k=1}^{N_\delta^m} (g(\mathbf{b}_*, \mathbf{b}_k) - g(\mathbf{x}, \mathbf{b}_k)) w_m(V_k) \right| \\ &\leq \sum_{k=1}^{N_\delta^m} |g(\mathbf{b}_*, \mathbf{b}_k) - g(\mathbf{x}, \mathbf{b}_k)| w_m(V_k) \\ &\leq L \cos^{-1}(1 - \delta) \sum_{k=1}^{N_\delta^m} w_m(V_k) \\ &= L \cos^{-1}(1 - \delta) w_m. \end{aligned}$$

Finally, observing that the above bounds do not depend on the choice of $\mathbf{x} \in S^m$ and combining the three results we obtain our desired bound. \square

By observing that we can set $g(\mathbf{x}, \mathbf{y}) = K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y})$, it becomes clear how Lemma B.7 can be used to bound the second term in Equation (18). For that we will also need to know what is the modulus of continuity of the von Mises-Fisher kernels K_λ^{vMF} .

Lemma B.8 (Modulus of continuity of K_λ^{vMF}). *The von Mises-Fisher kernels K_λ^{vMF} have modulus of continuity $\omega(K_\lambda^{\text{vMF}}; t) \leq c_{m+1}(\lambda) \exp(\lambda)$.*

Proof. Recall that $K_\lambda^{\text{vMF}}(u)$ is defined on $u \in [-1, 1]$. $K_\lambda^{\text{vMF}}(u)$ and its derivative are both monotonically increasing in u . Hence:

$$\omega(K_\lambda^{\text{vMF}}; t) = \sup \{ |K_\lambda^{\text{vMF}}(\langle \mathbf{z}, \mathbf{x} \rangle) - K_\lambda^{\text{vMF}}(\langle \mathbf{z}, \mathbf{y} \rangle) \mid \mathbf{x}, \mathbf{y} \in S^m, \cos^{-1}(\langle \mathbf{x}, \mathbf{y} \rangle) \leq t \}.$$

Using the mean value theorem we know there exists a $\mathbf{d} \in S^m$ such that

$$\begin{aligned} |K_\lambda^{\text{vMF}}(\langle \mathbf{z}, \mathbf{x} \rangle) - K_\lambda^{\text{vMF}}(\langle \mathbf{z}, \mathbf{y} \rangle)| &= |(K_\lambda^{\text{vMF}})'(\langle \mathbf{z}, \mathbf{d} \rangle) (\langle \mathbf{z}, \mathbf{x} \rangle - \langle \mathbf{z}, \mathbf{y} \rangle)| \\ &= (K_\lambda^{\text{vMF}})'(\langle \mathbf{z}, \mathbf{d} \rangle) |\langle \mathbf{z}, \mathbf{x} \rangle - \langle \mathbf{z}, \mathbf{y} \rangle| \\ &= \lambda c_{m+1}(\lambda) \exp(\lambda \langle \mathbf{z}, \mathbf{d} \rangle) |\langle \mathbf{z}, \mathbf{x} \rangle - \langle \mathbf{z}, \mathbf{y} \rangle| \\ &\leq \lambda c_{m+1}(\lambda) \exp(\lambda) |\langle \mathbf{z}, \mathbf{x} \rangle - \langle \mathbf{z}, \mathbf{y} \rangle| \\ &= \lambda c_{m+1}(\lambda) \exp(\lambda) \|\mathbf{z}\|_2 \|\mathbf{x} - \mathbf{y}\|_2 \cos(\text{angle}(\mathbf{z}, \mathbf{x} - \mathbf{y})) \\ &\leq \lambda c_{m+1}(\lambda) \exp(\lambda) \|\mathbf{x} - \mathbf{y}\|_2. \end{aligned}$$

Using the law of cosines and that the angle between \mathbf{x} and \mathbf{y} is less than t :

$$\begin{aligned} \|\mathbf{x} - \mathbf{y}\|_2 &= \sqrt{\|\mathbf{x}\|_2^2 + \|\mathbf{y}\|_2^2 - 2\|\mathbf{x}\|_2\|\mathbf{y}\|_2 \cos(\text{angle}(\mathbf{x}, \mathbf{y}))} \\ &= \sqrt{2 - 2 \cos(\text{angle}(\mathbf{x}, \mathbf{y}))} \\ &\leq \sqrt{2 - 2 \cos t} \\ &\leq t. \end{aligned}$$

Hence:

$$\omega(K_\lambda^{\text{vMF}}; t) \leq \lambda c_{m+1}(\lambda) \exp(\lambda) t.$$

□

Our final result, a bound on Equation (18), combines Lemma B.5 and Lemma B.7, each bounding one of the two terms in Equation (18).

Theorem B.1 (Jackson-type bound for universal approximation on the hypersphere, Theorem 1 in the main text). *Let $f \in C(S^m)$ be a continuous function on S^m with modulus of continuity $\omega(f; t) \leq Lt$ for some $L \in \mathbb{R}_{>0}$ and $m \geq 8$. Assume that there exists a constant $C_H \geq 0$ that upper-bounds the norms of the spherical harmonics of any best polynomial approximation W_q of f :*

$$\text{for all } q \geq 1, W_q = \sum_{k=0}^q Y_k^m, Y_k^m \in \mathbb{Y}(S^m), \|Y_k^m\|_\infty \leq C_H, \forall k = 0, \dots, q, \text{ where } W_q = \arg \min_{h \in \mathbb{P}_q(S^m)} \|f - h\|_\infty.$$

Then, for any $\epsilon > 0$, there exist $\xi_1, \dots, \xi_N \in \mathbb{R}$ and $\mathbf{b}_1, \dots, \mathbf{b}_N \in S^m$ such that

$$\sup_{\mathbf{x} \in S^m} \left| f(\mathbf{x}) - \sum_{k=1}^N \xi_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right| \leq \epsilon,$$

where $\lambda = \Lambda(\epsilon/2)$ (Equation (23)) and for any N such that

$$N \geq N(\lambda, \epsilon) = \Phi(m) \left(\frac{3\pi(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)}{\epsilon} \right)^{2(m+1)} = \mathcal{O}(\epsilon^{-10-14m-4m^2}). \quad (29)$$

Proof. Recall the decomposition in Equation (18). We will split our error budget ϵ in half. Hence, we first select λ such that approximating f with its convolution with K_λ^{vMF} results in an error at most $\lambda/2$. Then, using this λ , we find how finely we need to partition S^m in order to be able to approximate the convolution with a sum up to an error $\epsilon/2$.

Let's select how "peaky" we need the kernel K_λ^{vMF} to be, that is, how big should λ be. From Lemma B.5 we have that if $\lambda = \Lambda(\epsilon/2)$, then we would have $\|f - f * K_\lambda^{\text{vMF}}\|_\infty \leq \epsilon/2$.

Now, for the second term in Equation (18), consider Lemma B.7 with $g(\mathbf{x}, \mathbf{y}) = K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y})$. From Lemma B.8 we have that the modulus of continuity of K_λ^{vMF} is $\omega(K_\lambda^{\text{vMF}}; t) \leq t \lambda c_{m+1}(\lambda) \exp(\lambda)$. Hence, we have modulus of continuity for $g(\mathbf{x}, \mathbf{y})$ being bounded as:

$$\begin{aligned} \omega(g; t) &\leq \|K_\lambda^{\text{vMF}}\|_\infty \omega(f; t) + \|f\|_\infty \omega(K_\lambda^{\text{vMF}}; t) \\ &\leq K_\lambda^{\text{vMF}}(1) Lt + \|f\|_\infty \lambda c_{m+1}(\lambda) \exp(\lambda) t \\ &= c_{m+1}(\lambda) \exp(\lambda) Lt + \|f\|_\infty \lambda c_{m+1}(\lambda) \exp(\lambda) t \\ &= c_{m+1}(\lambda) \exp(\lambda) (L + \lambda \|f\|_\infty) t. \end{aligned}$$

Take

$$\delta = \left(\frac{2\epsilon}{6\pi(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)} \right)^2.$$

Then, by Lemma B.7, there exists a partition $\{V_1, \dots, V_N\}$ of S^m and $\mathbf{b}_1, \dots, \mathbf{b}_N \in S^m$ for N as in the lemma such that:

$$\max_{\mathbf{x} \in S^m} \left| \frac{1}{w_m} \int_{S^m} K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y}) dw_m(\mathbf{y}) - \frac{1}{w_m} \sum_{k=1}^N K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{b}_k \rangle) f(\mathbf{b}_k) w_m(V_k) \right| \leq 3(L + \lambda \|f\|_\infty) c_{m+1} \exp(\lambda) \cos^{-1}(1 - \delta).$$

As $(2x/\pi)^2 < 1 - \cos(x)$ (Bagul and Panchal, 2018, Theorem 1), we have:

$$\delta = \left(\frac{2\epsilon}{6\pi(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)} \right)^2 < 1 - \cos\left(\frac{\epsilon}{6(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)} \right). \quad (30)$$

Hence:

$$\begin{aligned} &\max_{\mathbf{x} \in S^m} \left| \frac{1}{w_m} \int_{S^m} K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle) f(\mathbf{y}) dw_m(\mathbf{y}) - \frac{1}{w_m} \sum_{k=1}^N K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{b}_k \rangle) f(\mathbf{b}_k) w_m(V_k) \right| \\ &\leq 3(L + \lambda \|f\|_\infty) c_{m+1} \exp(\lambda) \cos^{-1}(1 - \delta) \\ &< 3(L + \lambda \|f\|_\infty) c_{m+1} \exp(\lambda) \cos^{-1}\left(\cos\left(\frac{\epsilon}{6(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)} \right) \right) \\ &= \epsilon/2. \end{aligned}$$

Combining the two results we have:

$$\begin{aligned} \left\| f - \frac{1}{w_m} \sum_{k=1}^N K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{b}_k \rangle) f(\mathbf{y}) w_m(V_k) \right\|_\infty &\leq \|f - f * K_\lambda^{\text{vMF}}\|_\infty + \left\| f * K_\lambda^{\text{vMF}} - \frac{1}{w_m} \sum_{k=1}^N K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{b}_k \rangle) f(\mathbf{y}) w_m(V_k) \right\|_\infty \\ &\leq \epsilon/2 + \epsilon/2 \\ &= \epsilon. \end{aligned}$$

Now, the only thing left is to show that this expression can be expressed in the form of Equation (17).

$$\begin{aligned} \frac{1}{w_m} \sum_{k=1}^{N_\delta^m} K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{b}_k \rangle) f(\mathbf{b}_k) w_m(V_k) &= \sum_{k=1}^{N_\delta^m} \frac{1}{w_m} c_{m+1}(\lambda) \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) f(\mathbf{b}_k) w_m(V_k) \\ &= \sum_{k=1}^{N_\delta^m} \xi_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle), \end{aligned}$$

with

$$\xi_k = c_{m+1}(\lambda) f(\mathbf{b}_k) \frac{w_m(V_k)}{w_m}. \quad (31)$$

If we have chosen a partition of equal measure this further simplifies to

$$\xi_k = \frac{c_{m+1}(\lambda)}{N} f(\mathbf{b}_k).$$

Hence, for this choice of Λ , N and \mathbf{b}_k and ξ_k constructed as above, we indeed have

$$\sup_{\mathbf{x} \in S^m} \left| f(\mathbf{x}) - \sum_{k=1}^N \xi_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right| \leq \epsilon.$$

Finally, let's study the asymptotic growth of N as $\epsilon \rightarrow 0$. We have:

$$N(\lambda, \epsilon) = \Phi(m) \left(\frac{3\pi(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)}{\epsilon} \right)^{2(m+1)}.$$

$\Phi(m)$ is constant in ϵ so we can ignore it. Expanding c_{m+1} and dropping the terms that do not depend on ϵ gives us:

$$\mathcal{O} \left(\frac{(L + \lambda \|f\|_\infty) \lambda^{\frac{m+1}{2} - 1} \exp(\lambda)}{\epsilon I_{\frac{m+1}{2} - 1}(\lambda)} \right)^{2(m+1)}. \quad (32)$$

The asymptotics of the modified Bessel function of the first kind are difficult to analyse. However, as we care about an upper bound, we can simplify the expression by lower-bounding $I_\nu(\lambda)$ using Equation (22) and that $I_0(\lambda) \geq C \exp(\lambda)/\sqrt{\lambda}$ for $\lambda > 1/2$ (Barnett, 2021):

$$I_\nu(\lambda) \geq C \left(\frac{\sqrt{\nu^2 + \lambda^2} - \nu}{\lambda} \right)^{\nu+1} \frac{\exp(\sqrt{\nu^2 + \lambda^2})}{\sqrt{\lambda}},$$

for some constant C . Plugging this in Equation (32), replacing λ with its asymptotic growth ϵ^{-4} and taking the Taylor series expansion at for $\epsilon \rightarrow 0$ gives us:

$$\mathcal{O}(\epsilon^{-10-14m-4m^2}).$$

□

We can easily extend Theorem B.1 to vector-valued functions:

Corollary B.1 (Corollary 2 in the main text). *Let $f : S^m \rightarrow \mathbb{R}^{m+1}$, $m \geq 8$ be such that each component f_i is in $C(S^m)$, $i = 1, \dots, m+1$ and satisfies the conditions in Theorem B.1. Furthermore, define $\|f\|_\infty = \max_{1 \leq i \leq m+1} \|f_i\|_\infty$. Then, for any $\epsilon > 0$, there exist $\xi_1, \dots, \xi_N \in \mathbb{R}^{m+1}$ and $\mathbf{b}_1, \dots, \mathbf{b}_N \in S^m$ such that*

$$\sup_{\mathbf{x} \in S^m} \left\| f(\mathbf{x}) - \sum_{k=1}^N \xi_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right\|_2 \leq \epsilon,$$

with $\lambda = \Lambda(\epsilon/2\sqrt{m+1})$ for any $N \geq N(\lambda, \epsilon/\sqrt{m+1})$.

Proof. The proof is the same as for Theorem B.1. As the concentration parameter λ of the kernels K_λ^{vMF} depends only on the smoothness properties of the individual components and these are assumed to be the same, the same kernel choice can be used for all components f_i . Furthermore, the choice of partition is independent of the function to be approximated and depends only on the concentration parameter of the kernel. Hence, we can also use the same partition for all components f_i . We only need to take into account that:

$$\|x - y\|_2 = \sqrt{\sum_{i=1}^{m+1} |x_i - y_i|^2} \leq \sqrt{(m+1)\epsilon^2} = \sqrt{m+1}\epsilon, \quad \forall x, y \in \mathbb{R}^{m+1}, |x_i - y_i| \leq \epsilon, i = 1, \dots, m+1.$$

which results in the factor of $\sqrt{m+1}$.

□

C. A Jackson-type Bound for Approximation with a Split Attention Head

Lemma C.1. Let $a, b : \mathbb{R}^d \rightarrow \mathbb{R}$, $c, d : \mathbb{R} \rightarrow \mathbb{R}$, $c(x), d(x) \neq 0, \forall x \in \mathbb{R}$ and $\epsilon_1, \epsilon_2 \geq 0$ be such that:

$$\begin{aligned} \sup_{\mathbf{y} \in \mathbb{R}^d} \|a(\mathbf{y}) - b(\mathbf{y})\|_2 &\leq \epsilon_1 \\ \sup_{x \in \mathbb{R}} |c(x) - d(x)| &= |c - d|_\infty \leq \epsilon_2. \end{aligned}$$

Then for all $x \in \mathbb{R}$ and $\mathbf{y} \in \mathbb{R}^d$:

$$\left\| \frac{a(\mathbf{y})}{c(x)} - \frac{b(\mathbf{y})}{d(x)} \right\|_2 \leq \frac{\epsilon_1 |c|_\infty + \epsilon_2 \sup_{\mathbf{y} \in \mathbb{R}^d} \|a(\mathbf{y})\|_2}{|c(x) d(x)|}.$$

Proof. For a fixed $x \in \mathbb{R}$ and $\mathbf{y} \in \mathbb{R}^d$, using the triangle inequality gives us

$$\begin{aligned} \left\| \frac{a(\mathbf{y})}{c(x)} - \frac{b(\mathbf{y})}{d(x)} \right\|_2 &= \left\| \frac{a(\mathbf{y}) d(x) - b(\mathbf{y}) c(x)}{c(x) d(x)} \right\|_2 \\ &= \frac{\|a(\mathbf{y}) d(x) - b(\mathbf{y}) c(x)\|_2}{|c(x) d(x)|} \\ &= \frac{\|a(\mathbf{y}) d(x) - a(\mathbf{y}) c(x) + a(\mathbf{y}) c(x) - b(\mathbf{y}) c(x)\|_2}{|c(x) d(x)|} \\ &\leq \frac{\|a(\mathbf{y}) (d(x) - c(x))\|_2 + \|c(x) (a(\mathbf{y}) - b(\mathbf{y}))\|_2}{|c(x) d(x)|} \\ &\leq \frac{\epsilon_2 \|a(\mathbf{y})\|_2 + \epsilon_1 |c(x)|}{|c(x) d(x)|} \\ &\leq \frac{\epsilon_1 |c|_\infty + \epsilon_2 \sup_{\mathbf{y} \in \mathbb{R}^d} \|a(\mathbf{y})\|_2}{|c(x) d(x)|}. \end{aligned}$$

And as this holds for all $x \in \mathbb{R}$ and $\mathbf{y} \in \mathbb{R}^d$, the inequality in the lemma follows. \square

Lemma C.2. Let $f : S^m \rightarrow \mathbb{R}^{m+1}$, $m \geq 8$ satisfy the requirements in Corollary B.1. Then, given an $\epsilon > 0$ and taking λ , N , and $\mathbf{b}_1, \dots, \mathbf{b}_N \in S^m$ as prescribed by the Corollary, we have that $\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)$ is close to being a constant:

$$\sup_{\mathbf{x} \in S^m} \left| 1 - \frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right| \leq \frac{\epsilon}{2\sqrt{m+1}(L + \lambda \|f\|_\infty)}. \quad (33)$$

Proof. We can use Lemma B.7 by taking $g(\mathbf{x}, \mathbf{y}) = K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle)$. From Lemma B.8 we have that the modulus of continuity of K_λ^{vMF} is $\omega(K_\lambda^{\text{vMF}}; t) \leq t c_{m+1}(\lambda) \exp(\lambda)$. Observe that using Equation (16) we have

$$\int_{S^m} g(\mathbf{x}, \mathbf{y}) dw_m(\mathbf{y}) = \int_{S^m} K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{y} \rangle) dw_m(\mathbf{y}) = w_{m-1} \int_{-1}^1 K_\lambda^{\text{vMF}}(t) (1-t^2)^{(m-2)/2} dt = w_m.$$

The value for δ has to be selected as in Corollary B.1 (Equation (30)):

$$\delta = \left(\frac{2\epsilon}{6\pi\sqrt{m+1}(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)} \right)^2 < 1 - \cos \left(\frac{\epsilon}{6\sqrt{m+1}(L + \lambda \|f\|_\infty) c_{m+1}(\lambda) \exp(\lambda)} \right).$$

Now, using the same partition from Lemma B.7, and recalling that we constructed it such that each element of the partition has the same measure $w_m(V_1) = w_m(V_i), \forall i$, we have:

$$\max_{\mathbf{x} \in S^m} \left| \frac{1}{w_m} \int_{S^m} g(\mathbf{x}, \mathbf{y}) dw_m(\mathbf{y}) - \frac{1}{w_m} \sum_{k=1}^N K_\lambda^{\text{vMF}}(\langle \mathbf{x}, \mathbf{b}_k \rangle) w_m(V_k) \right|$$

$$\begin{aligned}
 &= \max_{\mathbf{x} \in S^m} \left| 1 - \frac{1}{w_m} \sum_{k=1}^N c_{m+1}(\lambda) \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) w_m(V_k) \right| \\
 &= \max_{\mathbf{x} \in S^m} \left| 1 - \frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right| \\
 &\leq 3c_{m+1} \exp(\lambda) \cos^{-1}(1 - \delta) \\
 &< 3c_{m+1} \exp(\lambda) \cos^{-1} \left(\cos \left(\frac{\epsilon}{6\sqrt{m+1}(L + \lambda\|f\|_\infty)c_{m+1}(\lambda) \exp(\lambda)} \right) \right) \\
 &= \frac{3c_{m+1} \exp(\lambda) \epsilon}{6\sqrt{m+1}(L + \lambda\|f\|_\infty)c_{m+1}(\lambda) \exp(\lambda)} \\
 &= \frac{\epsilon}{2\sqrt{m+1}(L + \lambda\|f\|_\infty)}.
 \end{aligned}$$

□

Theorem C.1. Let $f : S^m \rightarrow \mathbb{R}^{m+1}$, $m \geq 8$ satisfies the conditions in Corollary B.1. Define $\|f\|_\infty = \max_{1 \leq i \leq m+1} \|f_i\|_\infty$. For any $\epsilon > 0$, there exist $\mathbf{b}_1, \dots, \mathbf{b}_N \in S^m$ such that f can be uniformly approximated to an error at most ϵ :

$$\sup_{\mathbf{x} \in S^m} \left\| f(\mathbf{x}) - \frac{\sum_{k=1}^N \boldsymbol{\xi}_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)}{\sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)} \right\|_2 \leq \epsilon,$$

with:

$$\begin{aligned}
 \lambda &= \Lambda \left(\frac{2\epsilon L}{2L + \|f\|_\infty} \right) \text{ with } \Lambda \text{ from Equation (23),} \\
 N &\geq \Phi(m) \left(\frac{3\pi(L + \lambda\|f\|_\infty)\sqrt{m+1} c_{m+1}(\lambda) \exp(\lambda)}{\epsilon} \right)^{2(m+1)}, \\
 \boldsymbol{\xi}_k &= f(\mathbf{b}_k), \forall k = 1, \dots, N.
 \end{aligned}$$

Proof. From Corollary B.1 we know that $\sum_{k=1}^N \boldsymbol{\xi}_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)$ approximates $f(\mathbf{x})$ and from Lemma C.2 we know that $\frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)$ approximates 1. Using Lemma C.1 we can combine the two results to bound how well

$$\frac{\sum_{k=1}^N \boldsymbol{\xi}_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)}{\frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)}$$

approximates $f(\mathbf{x})/1 = f(\mathbf{x})$. The fact that $\frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)$ is not identically 1 means that we will need to increase the precision of approximating the numerator by reducing ϵ in order to account for the additional error coming from the denominator. In particular, we have

$$\begin{aligned}
 \sup_{\mathbf{x} \in S^m} \left\| f(\mathbf{x}) - \sum_{k=1}^N \boldsymbol{\xi}_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right\|_2 &\leq \epsilon', \\
 \sup_{\mathbf{x} \in S^m} \left| 1 - \frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right| &\leq \frac{\epsilon'}{2\sqrt{m+1}(L + \lambda\|f\|_\infty)}, \\
 \sup_{\mathbf{x} \in S^m} \|f(\mathbf{x})\|_2 &\leq \sqrt{m+1}\|f\|_\infty, \\
 \sup_{\mathbf{x} \in S^m} \left| \frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle) \right| &\leq 1 + \frac{\epsilon'}{2\sqrt{m+1}(L + \lambda\|f\|_\infty)}.
 \end{aligned}$$

Hence, applying Lemma C.1 gives us:

$$\left\| f(\mathbf{x}) - \frac{\sum_{k=1}^N \boldsymbol{\xi}_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)}{\frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)} \right\|_2 = \left\| \frac{f(\mathbf{x})}{1} - \frac{\sum_{k=1}^N \boldsymbol{\xi}_k \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)}{\frac{c_{m+1}(\lambda)}{N} \sum_{k=1}^N \exp(\lambda \langle \mathbf{x}, \mathbf{b}_k \rangle)} \right\|_2$$

$$\begin{aligned}
 &\leq \frac{\epsilon' + \frac{\epsilon'}{2\sqrt{m+1}(L+\lambda\|f\|_\infty)}\sqrt{m+1}\|f\|_\infty}{1 + \frac{\epsilon'}{2\sqrt{m+1}(L+\lambda\|f\|_\infty)}} \\
 &= \frac{\epsilon'\sqrt{m+1}(\|f\|_\infty + 2L + 2\lambda\|f\|_\infty)}{2\sqrt{m+1}(\lambda\|f\|_\infty + L) + \epsilon'} \\
 &\leq \frac{\epsilon'}{2} \left(\frac{\|f\|_\infty}{\lambda\|f\|_\infty + L} + 2\frac{\lambda\|f\|_\infty + L}{\lambda\|f\|_\infty + L} \right) \\
 &\leq \frac{\epsilon'}{2} \left(2 + \frac{\|f\|_\infty}{L} \right),
 \end{aligned}$$

where we used that $\epsilon', \|f\|_\infty L > 0$, which is the case in realistic scenarios. Therefore, if we want this error to be upper bounded by ϵ , we need to select

$$\epsilon' \leq \frac{2\epsilon L}{2L + \|f\|_\infty}.$$

From Corollary B.1 (Equation (23)) that can be achieved by selecting

$$\lambda = \Lambda \left(\frac{2\epsilon L}{2L + \|f\|_\infty} \right)$$

and

$$N \geq \Phi(m) \left(\frac{3\pi(L + \lambda\|f\|_\infty)\sqrt{m+1} c_{m+1}(\lambda) \exp(\lambda)}{\epsilon} \right)^{2(m+1)}.$$

Finally, observe that the $c_{m+1}(\lambda)/N$ factor can be folded in the ξ_k terms (Equation (31)):

$$\xi_k = \frac{N}{c_{m+1}(\lambda)} f(\mathbf{b}_k) c_{m+1}(\lambda) \frac{w_m(V_k)}{w_m} = f(\mathbf{b}_k),$$

with ξ_k nicely reducing to be the evaluation of f at the corresponding control point \mathbf{b}_k . \square

D. Additional Results

Lemma D.1. Define the stereographic projection and its inverse:

$$\begin{aligned}
 \Sigma_m : \bar{S}^m &\rightarrow \mathbb{R}^m \\
 (x_1, \dots, x_{m+1}) &\mapsto \left(\frac{x_1}{1 - x_{m+1}}, \dots, \frac{x_m}{1 - x_{m+1}} \right) \\
 \\
 \Sigma_m^{-1} : \mathbb{R}^m &\rightarrow S^m \\
 (y_1, \dots, y_m) &\mapsto \left(\frac{2y_1}{\sum_{i=1}^m y_i^2 + 1}, \dots, \frac{2y_m}{\sum_{i=1}^m y_i^2 + 1}, \frac{\sum_{i=1}^m y_i^2 - 1}{\sum_{i=1}^m y_i^2 + 1} \right)
 \end{aligned}$$

with \bar{S}^m the part of S^m that gets mapped to $[0, 1]^m$, i.e., $\bar{S}^m = \Sigma_m^{-1}([0, 1]^m)$. Σ_m and Σ_m^{-1} are continuous and inverses of each other and there exist L_m^Σ and $L_m^{\Sigma^{-1}}$ such that $\omega(\Sigma_m; t) \leq L_m^\Sigma t$ and $\omega(\Sigma_m^{-1}; t) \leq L_m^{\Sigma^{-1}} t$. Furthermore, $\Sigma_m \circ \mathcal{H}_{H,3(m+1)}^1 \circ \Sigma_m^{-1}$ is dense in the set of continuous functions $[0, 1]^m \rightarrow \mathbb{R}^m$.