

Metric-Entropy Limits on the Approximation of Nonlinear Dynamical Systems

Yang Pan
yangpan@mins.ee.ethz.ch

Clemens Hutter
huttercl@mins.ee.ethz.ch

Helmut Bölcskei*
hboelcskei@ethz.ch

Abstract

This paper is concerned with fundamental limits on the approximation of nonlinear dynamical systems. Specifically, we show that recurrent neural networks (RNNs) can approximate nonlinear systems—that satisfy a Lipschitz property and forget past inputs fast enough—in metric-entropy-optimal manner. As the sets of sequence-to-sequence mappings realized by the dynamical systems we consider are significantly more massive than function classes generally analyzed in approximation theory, a refined metric-entropy characterization is needed, namely in terms of order, type, and generalized dimension. We compute these quantities for the classes of exponentially- and polynomially Lipschitz fading-memory systems and show that RNNs can achieve them.

This paper is dedicated to Professor Andrew R. Barron on the occasion of his 65th birthday.

1 Introduction

It is well known that neural networks can approximate almost every function arbitrarily well [1–5]. The recently developed Kolmogorov-Donoho rate-distortion theory for deep neural network approximation [6, 7] goes a step further by quantifying how effective such approximations are in terms of the description complexity of the neural networks relative to that of the functions they are to approximate. Specifically, [7] determines the length of the bitstring needed to specify (deep) rectified linear unit (ReLU) networks approximating functions (mapping \mathbb{R}^d to \mathbb{R}) in a given class to within a prescribed error ϵ . It is then shown that for a wide variety of function classes, the length of this bitstring exhibits the same scaling behavior, in ϵ , as the metric entropy of the function class under consideration (see Table 1 below). This means that neural networks are universally Kolmogorov-Donoho optimal for all these function classes.

*H. Bölcskei gratefully acknowledges support by the Lagrange Mathematics and Computing Research Center, Paris, France.

In the present paper, we extend the philosophy of [7] to the approximation of nonlinear sequence-to-sequence mappings through recurrent neural networks (RNNs). Concretely, we consider Lipschitz fading-memory (LFM) systems. In essence, this notion describes systems that gradually forget past inputs, with the speed of memory decay quantified in terms of a Lipschitz property. Such systems find application, inter alia, in finance [8] and material science [9–11]. We first develop tools for quantifying the metric entropy of classes of LFM systems with a given memory decay rate. A general construction of RNNs approximating LFM systems is then shown to yield Kolmogorov-Donoho-optimality for LFM systems of exponentially or polynomially decaying memory.

Related work. Learning of linear dynamical systems has been studied extensively in the literature [12–16]. Notably, [16] provides explicit RNN realizations and approximations for wide classes of linear dynamical systems, including time-varying ones. Going beyond linear systems, learning of nonlinear finite memory systems through RNNs within the finite-state-machine framework was considered in [17]. This program is extended to approximately finite-memory systems in [18] and fading-memory systems in [19–21]. In particular, [21] formalizes the concept of fading-memory systems in control theory, demonstrating that continuous-time fading-memory operators can be approximated by Volterra series. Subsequently, [22] established that discrete-time fading-memory systems can be identified using neural networks. Moreover, [23] demonstrated that echo state networks, a specialized architecture within the RNN family, serve as universal approximators for discrete-time fading-memory systems. None of the studies just reviewed addresses the issue of quantifying the RNN description complexity relative to that of the class of (nonlinear) systems being approximated.

Organization of the paper. The remainder of Section 1 summarizes notation. In Section 2, we introduce our setup and provide a definition of metric-entropy optimality in a very general context, encompassing the approximation of functions as well as dynamical systems. Section 3 develops tools for characterizing the metric entropy of LFM systems. In Section 4, we employ these tools to derive precise scaling results for the metric entropy of exponentially Lipschitz fading-memory (ELFM) and polynomially Lipschitz fading-memory (PLFM) systems. Section 5 presents a construction for the approximation of general LFM systems by RNNs. Finally, in Section 6, we combine the results developed in the previous sections to prove that RNNs can approximate ELFM and PLFM systems in metric-entropy-optimal manner.

Notation. \mathbb{N}_0 and \mathbb{N}_+ denote the set of natural numbers including and excluding 0, respectively. For $N \in \mathbb{N}_0$, $\llbracket N \rrbracket$ stands for the set $\{0, 1, \dots, N\}$, while $\llbracket N \rrbracket^\pm$ refers to $\{-N, \dots, -1, 0, 1, \dots, N\}$. The cardinality of the finite set U is designated by $|U|$. Sequences $x[t] \in \mathbb{R}$ are indexed by $t \in \mathbb{Z}$ or $t \in \mathbb{N}_0$ and we use $\mathbb{R}^{\mathbb{Z}}$ and $\mathbb{R}^{\mathbb{N}_0}$, respectively, to denote the set of such sequences. We refer to the set of all finite-length bitstrings by $\{0, 1\}^*$. The transpose of the matrix A is A^T . For matrices A_1, \dots, A_N , $\text{diag}(A_1, A_2, \dots, A_N)$ denotes the block-diagonal matrix with the A_i on the main diagonal. The $N \times N$ identity matrix is \mathbb{I}_N and 0_N stands for the N -dimensional column vector with all entries equal to 0. For the vector $x \in \mathbb{R}^d$, we let $\|x\|_\infty := \max_{i=1,2,\dots,d} |x_i|$. $\log(\cdot)$ refers to the logarithm to base 2, $\log^{(n)} = \log \circ \dots \circ \log$ is the n -fold iterated logarithm, and $\log^\tau(\cdot) = (\log(\cdot))^\tau$, for $\tau \in \mathbb{R}$. The composition of functions

f_1, f_2 is denoted by $f_2 \circ f_1$ (or $f_1 \circ f_2$). Let $f(\epsilon)$ and $g(\epsilon)$, in both cases for $\epsilon > 0$, be strictly positive for all small enough values of ϵ . We use $f(\epsilon) = o(g(\epsilon))$ to indicate that $\lim_{\epsilon \rightarrow 0} \frac{f(\epsilon)}{g(\epsilon)} = 0$ and we express $\limsup_{\epsilon \rightarrow 0} \frac{f(\epsilon)}{g(\epsilon)} < \infty$ by $f(\epsilon) = O(g(\epsilon))$. Moreover, we write $f(\epsilon) = \Theta(g(\epsilon))$ when both $f(\epsilon) = O(g(\epsilon))$ and $g(\epsilon) = O(f(\epsilon))$. Constants are always understood to be in \mathbb{R} unless explicitly stated otherwise. Finally, we say that a constant is universal if it does not depend on any of the ambient quantities.

2 Problem setup and metric-entropy optimality

2.1 ReLU network approximation

We start by defining ReLU networks.

Definition 2.1 (ReLU network [7]). *Let $L \in \mathbb{N}_+$ and $N_0, N_1, \dots, N_L \in \mathbb{N}_+$. A ReLU (feedforward) neural network Φ is a mapping $\Phi : \mathbb{R}^{N_0} \rightarrow \mathbb{R}^{N_L}$ given by*

$$\Phi = \begin{cases} W_1, & L = 1 \\ W_2 \circ \rho \circ W_1, & L = 2 \\ W_L \circ \rho \circ W_{L-1} \circ \rho \circ \dots \circ \rho \circ W_1, & L \geq 3, \end{cases} \quad (1)$$

where, for $\ell \in \{1, 2, \dots, L\}$, $W_\ell : \mathbb{R}^{N_{\ell-1}} \rightarrow \mathbb{R}^{N_\ell}$, $W_\ell(x) := A_\ell x + b_\ell$, $x \in \mathbb{R}^{N_{\ell-1}}$, are affine transformations with (weight) matrices $A_\ell \in \mathbb{R}^{N_\ell \times N_{\ell-1}}$ and (bias) vectors $b_\ell \in \mathbb{R}^{N_\ell}$, and the ReLU activation function $\rho : \mathbb{R} \rightarrow \mathbb{R}$, $\rho(x) = \max\{0, x\}$ acts component-wise, i.e., $\rho(x_1, \dots, x_N) = (\rho(x_1), \dots, \rho(x_N))$. We denote the set of all ReLU networks with input dimension $N_0 = d$ and output dimension $N_L = d'$ by $\mathcal{N}_{d,d'}$. Moreover, we define the following quantities related to the notion of size of the ReLU network Φ :

- depth $\mathcal{L}(\Phi) := L$,
- the connectivity $\mathcal{M}(\Phi)$ of the network Φ is the total number of non-zero entries in the matrices A_ℓ , $\ell \in \{1, 2, \dots, L\}$, and the vectors b_ℓ , $\ell \in \{1, 2, \dots, L\}$,
- width $\mathcal{W}(\Phi) := \max_{\ell=0, \dots, L} N_\ell$,
- the weight set $\mathcal{K}(\Phi)$ denotes the set of non-zero entries in the matrices A_ℓ , $\ell \in \{1, 2, \dots, L\}$, and the vectors b_ℓ , $\ell \in \{1, 2, \dots, L\}$,
- weight magnitude $\mathcal{B}(\Phi) := \max_{\ell=1, \dots, L} \max \{ \max_{i,j} |(A_\ell)_{ij}|, \|b_\ell\|_\infty \}$.

We next formalize the concept of network weight quantization.

Definition 2.2 (Quantization [7]). *Let $m \in \mathbb{N}_+$ and $\epsilon \in (0, 1/2)$. The network Φ is said to have (m, ϵ) -quantized weights if $\mathcal{K}(\Phi) \subset \left(2^{-m} \lceil \log(\epsilon^{-1}) \rceil \mathbb{Z} \right) \cap [-\epsilon^{-m}, \epsilon^{-m}]$. Moreover, for $a \in \mathbb{R}$,*

we define the (m, ϵ) -quantization mapping rounding a to an integer multiple of $2^{-m} \lceil \log(\epsilon^{-1}) \rceil$ according to

$$\mathcal{Q}_{m, \epsilon}(a) = \left\lceil \frac{a}{2^{-m \lceil \log(\epsilon^{-1}) \rceil}} \right\rceil \cdot 2^{-m \lceil \log(\epsilon^{-1}) \rceil}. \quad (2)$$

Every quantized ReLU network can be represented by a bitstring (see A) specifying the topology of the network along with its quantized weights. The mapping $\mathcal{D}_{\mathcal{N}}$ taking this bitstring back to the quantized ReLU network is referred to as the *canonical neural network decoder*.

Remark. For every ReLU network Φ with (m, ϵ) -quantized weights, there is a bitstring \mathbf{b} of length no more than $C_0 m \log(\epsilon^{-1}) \mathcal{M}(\Phi) \log(\mathcal{M}(\Phi))$ such that $\mathcal{D}_{\mathcal{N}}(\mathbf{b}) = \Phi$, with $C_0 > 0$ a universal constant. This follows by upper-bounding (115) in A.

As we consider the approximation of sequence-to-sequence mappings ($\mathbb{R}^{\mathbb{Z}} \rightarrow \mathbb{R}^{\mathbb{Z}}$), feed-forward networks are not directly applicable because they effect mappings between finite-dimensional spaces, concretely from \mathbb{R}^{N_0} to \mathbb{R}^{N_L} . However, and perhaps surprisingly, simply applying feedforward networks iteratively in a judicious manner turns out to be sufficient for approximating interesting classes of nonlinear sequence-to-sequence mappings in Kolmogorov-Donoho-optimal manner. Concretely, this is effected through recurrent neural networks, defined as follows.

Definition 2.3 (Recurrent neural networks [16, 24]). For $m \in \mathbb{N}_+$, let $\Phi \in \mathcal{N}_{m+1, m+1}$ be a ReLU network of depth $\mathcal{L}(\Phi) \geq 2$. The recurrent neural network (RNN) associated with Φ is the operator $\mathcal{R}_{\Phi} : \mathbb{R}^{\mathbb{N}_0} \rightarrow \mathbb{R}^{\mathbb{N}_0}$ mapping input sequences $(x[t])_{t \in \mathbb{N}_0}$ in \mathbb{R} to output sequences $(y[t])_{t \in \mathbb{N}_0}$ in \mathbb{R} according to

$$\begin{pmatrix} y[t] \\ h[t] \end{pmatrix} = \Phi \left(\begin{pmatrix} x[t] \\ h[t-1] \end{pmatrix} \right), \quad \forall t \in \mathbb{N}_0, \quad (3)$$

where $h[t] \in \mathbb{R}^m$ is the hidden state vector sequence with initial state $h[-1] = 0_m$. We denote the set of all RNNs by \mathcal{N}^R .

Remark. When unfolded in time, an RNN simply amounts to repeated application of Φ .

From Definition 2.3 it is apparent that an RNN \mathcal{R}_{Φ} is fully specified by its associated feedforward network Φ .

Definition 2.4. We define \mathbf{R} as the mapping that takes a ReLU network Φ to its associated RNN \mathcal{R}_{Φ} according to Definition 2.3.

Together with the canonical neural network decoder $\mathcal{D}_{\mathcal{N}}$, we thus obtain a decoder taking a bitstring into an RNN as follows.

Definition 2.5 (Canonical RNN decoder). We define the canonical RNN decoder as $\mathcal{D}_{\mathcal{R}} = \mathbf{R} \circ \mathcal{D}_{\mathcal{N}}$, where \mathbf{R} is as in Definition 2.4 and $\mathcal{D}_{\mathcal{N}}$ is the canonical neural network decoder.

The main point of this paper is to show that the canonical RNN decoder is capable of approximating a wide variety of non-linear sequence-to-sequence mappings ($\mathbb{R}^{\mathbb{Z}} \rightarrow \mathbb{R}^{\mathbb{Z}}$) in metric-entropy-optimal manner. Together with the results in [7], this establishes that ReLU networks optimally approximate a wide range of function classes and sequence-to-sequence mappings.

2.2 Metric-entropy optimality

In this section, we define the notion of metric-entropy optimal approximation. Consider a metric space (\mathcal{X}, ρ) and a compact set $\mathcal{C} \subset \mathcal{X}$. Together, \mathcal{C} and ρ determine an approximation task. Specifically, we wish to approximate elements $f \in \mathcal{C}$ to within a prescribed error $\epsilon > 0$ in the metric ρ by elements $\tilde{f} \in \mathcal{X}$ which can be encoded by finite-length bitstrings $\mathbf{b} \in \{0, 1\}^\ell$. To go from bitstrings to elements of \mathcal{X} , we define decoder mappings as follows.

Definition 2.6. *A decoder $\mathcal{D} : \{0, 1\}^* \rightarrow \mathcal{X}$ is a mapping from bitstrings of arbitrary length to elements of \mathcal{X} .*

We shall frequently want to quantify how well a given decoder \mathcal{D} performs.

Definition 2.7. *Given a metric space (\mathcal{X}, ρ) , a compact set $\mathcal{C} \subset \mathcal{X}$, and a decoder $\mathcal{D} : \{0, 1\}^* \rightarrow \mathcal{X}$, we say that (\mathcal{C}, ρ) is representable by \mathcal{D} , if for every $\epsilon > 0$ and every $f \in \mathcal{C}$, there exist $\ell \in \mathbb{N}_+$ and a bitstring $\mathbf{b} \in \{0, 1\}^\ell$ such that*

$$\rho(\mathcal{D}(\mathbf{b}), f) \leq \epsilon.$$

Furthermore, we set

$$L(\epsilon; \mathcal{D}, \mathcal{C}, \rho) := \min \{ \ell' \in \mathbb{N}_+ \mid \forall f \in \mathcal{C}, \exists \ell \leq \ell', \exists \mathbf{b} \in \{0, 1\}^\ell \text{ s.t. } \rho(\mathcal{D}(\mathbf{b}), f) \leq \epsilon \}.$$

Remark. *This setting allows us to fix a decoder \mathcal{D} (e.g., the canonical neural network decoder) and then study how well \mathcal{D} performs on different (\mathcal{C}, ρ) . That is, \mathcal{D} does not depend on \mathcal{C}, ρ, f , or ϵ .*

The quantity $L(\epsilon; \mathcal{D}, \mathcal{C}, \rho)$ measures how bit-efficient the decoder \mathcal{D} is in representing \mathcal{C} with respect to ρ . It is now natural to ask what the minimum required number of bits, independently of \mathcal{D} , is for representing \mathcal{C} with respect to ρ . The concept of metric entropy [25, 26] allows to answer this question.

Definition 2.8. *Let (\mathcal{X}, ρ) be a metric space and $\mathcal{C} \subset \mathcal{X}$ compact. The set $\{x_1, x_2, \dots, x_N\} \subset \mathcal{C}$ (respectively $\{x_1, x_2, \dots, x_N\} \subset \mathcal{X}$) is an ϵ -covering (respectively ϵ -net) for (\mathcal{C}, ρ) if, for each $x \in \mathcal{C}$, there exists an $i \in \{1, 2, \dots, N\}$ so that $\rho(x, x_i) \leq \epsilon$. The ϵ -covering number $N(\epsilon; \mathcal{C}, \rho)$ (respectively the exterior ϵ -covering number $N^{\text{ext}}(\epsilon; \mathcal{C}, \rho)$) is the cardinality of a smallest ϵ -covering (respectively smallest ϵ -net) for (\mathcal{C}, ρ) .*

In general, it is hard to obtain precise expressions for covering numbers. One therefore typically resorts to characterizations of their asymptotic behavior as $\epsilon \rightarrow 0$. In [7], where sets of functions are considered, this is done through the concept of optimal exponents. Here, however, we are concerned with sets of systems, which are much more massive and hence require a refined framework for quantifying the asymptotic behavior of their covering numbers. Thus, inspired by [27, Section II.C], we use the following notions.

Definition 2.9 (Order, type, and generalized dimension). *Consider a metric space (\mathcal{X}, ρ) and a compact set $\mathcal{C} \subset \mathcal{X}$. Then, (\mathcal{C}, ρ) is said to be of order $\kappa \in \mathbb{N}_+$ and type $\lambda \in \mathbb{N}_+$ if the quantity*

$$\mathfrak{d} := \limsup_{\epsilon \rightarrow 0} \frac{\log^{(\kappa+1)} N^{\text{ext}}(\epsilon; \mathcal{C}, \rho)}{\log^\lambda(\epsilon^{-1})} \quad (4)$$

is strictly positive and finite. In this case, we call \mathfrak{d} the generalized dimension.

Order, type, and generalized dimension provide measures for the “description complexity” of (\mathcal{C}, ρ) with the order κ the coarsest one. For a given order, the type λ constitutes a finer measure, and for fixed order and type, the generalized dimension \mathfrak{d} is the finest measure [27].

Whenever the optimal exponent according to [7, Definition IV.1] is well-defined (i.e., strictly positive and finite), the underlying set has order and type equal to one and generalized dimension equal to the inverse of the optimal exponent (Lemma B.1). Based on this insight, we obtain Table 1, which lists the generalized dimension for the sets considered in [7, Table I].

Returning to the previous discussion, we are now able to characterize the minimum number of bits required by any decoder to represent (\mathcal{C}, ρ) .

Lemma 2.10. *Consider the metric space (\mathcal{X}, ρ) , the compact set $\mathcal{C} \subset \mathcal{X}$ of order κ , type λ , and generalized dimension \mathfrak{d} , and assume that (\mathcal{C}, ρ) is representable by a decoder \mathcal{D} . Then, it holds that*

$$\limsup_{\epsilon \rightarrow 0} \frac{\log^{(\kappa)} L(\epsilon; \mathcal{D}, \mathcal{C}, \rho)}{\log^\lambda(\epsilon^{-1})} \geq \mathfrak{d}. \quad (5)$$

Proof. See C.1. □

It is natural to say that a decoder \mathcal{D} is optimal if it satisfies (5) with equality.

Definition 2.11. *Consider the metric space (\mathcal{X}, ρ) and the compact set $\mathcal{C} \subset \mathcal{X}$ of order κ and type λ with generalized dimension \mathfrak{d} . We say that (\mathcal{C}, ρ) is optimally representable by the decoder \mathcal{D} , if (\mathcal{C}, ρ) is representable by \mathcal{D} and*

$$\limsup_{\epsilon \rightarrow 0} \frac{\log^{(\kappa)} L(\epsilon; \mathcal{D}, \mathcal{C}, \rho)}{\log^\lambda(\epsilon^{-1})} = \mathfrak{d}. \quad (6)$$

We now recall a remarkable universal optimality property of ReLU networks, namely all the function classes listed in Table 1 are optimally representable, in the sense of Definition 2.11, by the canonical neural network decoder. This is a simple reformulation of the results in [7]; we provide the details of this reformulation in B. In the present paper, we establish that RNNs (Definition 2.3), with inner ReLU networks, extend this universality to the approximation of a wide range of nonlinear dynamical systems.

2.3 Lipschitz fading-memory systems

We proceed to characterize the class of dynamical systems considered in this paper and start by defining their domain.

	Metric	\mathcal{C}		κ	λ	\mathfrak{d}
$\{\mathbb{R} \rightarrow \mathbb{R}\}$	$L^2([0, 1])$	L^2 -Sobolev	$\mathcal{U}(W_2^m([0, 1]))$	1	1	$1/m$
$\{\mathbb{R} \rightarrow \mathbb{R}\}$	$L^2([0, 1])$	Hölder	$\mathcal{U}(C^\alpha([0, 1]))$	1	1	$1/\alpha$
$\{\mathbb{R} \rightarrow \mathbb{R}\}$	$L^2([0, 1])$	Bump Algebra	$\mathcal{U}(B_{1,1}^1([0, 1]))$	1	1	1
$\{\mathbb{R} \rightarrow \mathbb{R}\}$	$L^2([0, 1])$	Bounded Variation	$\mathcal{U}(BV([0, 1]))$	1	1	1
$\{\mathbb{R}^d \rightarrow \mathbb{R}\}$	$L^2(\Omega)$	L^p -Sobolev	$\mathcal{U}(W_p^m(\Omega))$	1	1	$\frac{d}{m}$
$\{\mathbb{R}^d \rightarrow \mathbb{R}\}$	$L^2(\Omega)$	Besov	$\mathcal{U}(B_{p,q}^m(\Omega))$	1	1	$\frac{d}{m}$
$\{\mathbb{R}^d \rightarrow \mathbb{R}\}$	$L^2(\Omega)$	Modulation	$\mathcal{U}(M_{p,p}^s(\mathbb{R}^d))$	1	1	$(\frac{1}{p} - \frac{1}{2} + \frac{2s}{d})^{-1}$
$\{\mathbb{R}^d \rightarrow \mathbb{R}\}$	$L^2(\Omega)$	Cartoon functions	$\mathcal{E}^\beta([-\frac{1}{2}, \frac{1}{2}]^d)$	1	1	$\frac{2(d-1)}{\beta}$

Table 1: Generalized dimension for the sets considered in [7]. Here, $\mathcal{U}(X) = \{f \in X : \|f\|_X \leq 1\}$ denotes the unit ball in the space X and $\Omega \subseteq \mathbb{R}^d$ is a Lipschitz domain.

Definition 2.12. For fixed $D > 0$, we denote the set of admissible input signals by $\mathcal{S} := [-D, D]^{\mathbb{Z}}$, that is, for every $x[\cdot] \in \mathcal{S}$, it holds that $|x[t]| \leq D, \forall t \in \mathbb{Z}$.

The quantity $D > 0$ is taken to be fixed throughout the paper and the dependence of \mathcal{S} on D is not explicitly indicated.

First, the systems $G : \mathcal{S} \rightarrow \mathbb{R}^{\mathbb{Z}}$ under consideration are causal.

Definition 2.13 (Causality). A system $G : \mathcal{S} \rightarrow \mathbb{R}^{\mathbb{Z}}$ is causal, if for each $T \in \mathbb{Z}$, for every pair $x, x' \in \mathcal{S}$ with $x[t] = x'[t], \forall t \in \mathbb{Z}$ with $t \leq T$, it holds that $(Gx)[T] = (Gx')[T]$.

Second, we require time-invariance.

Definition 2.14 (Time-invariance). A system $G : \mathcal{S} \rightarrow \mathbb{R}^{\mathbb{Z}}$ is time-invariant, if for every $\tau \in \mathbb{Z}$, it holds that

$$\mathbf{T}_\tau Gx = G\mathbf{T}_\tau x, \quad \forall x \in \mathcal{S},$$

with the shift operator $\mathbf{T}_\tau : \mathbb{R}^{\mathbb{Z}} \rightarrow \mathbb{R}^{\mathbb{Z}}$ defined as $(\mathbf{T}_\tau x)[t] = x[t - \tau]$.

Next, we follow Volterra, who suggested that [19, p. 188] “a first extremely natural postulate is to suppose that the influence of the [input] a long time before the given moment gradually fades out.” This property was termed “fading memory” in [21], and here we introduce a more quantitative version thereof, namely the concept of “Lipschitz fading memory”. The concept is inspired by examples in [8, 28–30], which will be discussed in more detail later.

Definition 2.15 (Lipschitz fading-memory). We say that $(w[t])_{t \in \mathbb{N}_0}$ is a weight sequence if it is non-increasing and satisfies $w[t] \in (0, 1], \forall t \in \mathbb{N}_0$, and $\lim_{t \rightarrow \infty} w[t] = 0$. A system $G : \mathcal{S} \rightarrow \mathbb{R}^{\mathbb{Z}}$ has Lipschitz fading-memory (LFM) with respect to the weight sequence w if

$$|(Gx)[t] - (Gy)[t]| \leq \sup_{\tau \geq 0} w[\tau] |x[t - \tau] - y[t - \tau]|, \quad \forall t \in \mathbb{Z}, \forall x, y \in \mathcal{S}.$$

The class of LFM systems considered in the remainder of the paper can now formally be defined as follows.

Definition 2.16 (Lipschitz fading-memory systems). *Given a weight sequence $w[\cdot]$, we define*

$$\mathcal{G}(w) = \{G : \mathcal{S} \rightarrow \mathbb{R}^{\mathbb{Z}} \mid G \text{ is causal, time-invariant, has Lipschitz fading-memory w.r.t. } w, \text{ and satisfies } (G0)[t] = 0, \forall t \in \mathbb{Z}\}. \quad (7)$$

As we will want to approximate LFM systems $G \in \mathcal{G}(w)$ by RNNs, we need a metric that quantifies approximation quality. This metric should take into account that the RNNs we consider start running at time $t = 0$ and will, moreover, be of worst-case nature.

Definition 2.17. *Let $\mathcal{S}^+ := \{s \in \mathcal{S} \mid s[t] = 0, \forall t < 0\}$. For $G, G' \in \{\mathbb{R}^{\mathbb{Z}} \rightarrow \mathbb{R}^{\mathbb{Z}}\}$, we define the metric*

$$\rho_*(G, G') = \sup_{x \in \mathcal{S}^+} \sup_{t \in \mathbb{N}_0} |(Gx)[t] - (G'x)[t]|.$$

We hasten to add that the restriction to one-sided input signals in Definition 2.17 and to taking the supremum over $t \in \mathbb{N}_0$ in the output signals does not impact the hardness of the approximation task, as shown by the next result.

Lemma 2.18. *Let $(w[t])_{t \in \mathbb{N}_0}$ be a weight sequence. For $G, G' \in \mathcal{G}(w)$, we have*

$$\rho_*(G, G') = \sup_{x \in \mathcal{S}} \sup_{t \in \mathbb{Z}} |(Gx)[t] - (G'x)[t]|.$$

Proof. See C.2. □

We are now ready to formally state the main goal of this paper, which is to prove that $(\mathcal{G}(w), \rho_*)$ is optimally representable by the canonical RNN decoder in Definition 2.5. In fact, we will be seeking a quantitative version of this statement comparing the description complexity of the class $\mathcal{G}(w)$ to that of the RNNs approximating it.

3 Metric entropy of LFM systems

In this section, we study the (ϵ -)scaling behavior of $N^{\text{ext}}(\epsilon; \mathcal{G}(w), \rho_*)$ for general weight sequences w . This will be effected by deriving an upper bound on $N^{\text{ext}}(\epsilon; \mathcal{G}(w), \rho_*)$ through the construction of a covering and a matching (in terms of scaling behavior) lower bound by identifying an explicit packing. We first define the concept of packings.

Definition 3.1. *Let (\mathcal{X}, ρ) be a metric space and $\mathcal{C} \subset \mathcal{X}$ compact. An ϵ -packing for (\mathcal{C}, ρ) is a set $\{x_1, x_2, \dots, x_N\} \subset \mathcal{C}$ such that $\rho(x_i, x_j) > \epsilon$, for all distinct i, j . The ϵ -packing number $M(\epsilon; \mathcal{C}, \rho)$ is the cardinality of a largest ϵ -packing for (\mathcal{C}, ρ) .*

We shall frequently make use of the following two results relating the packing, covering, and exterior covering numbers.

Lemma 3.2 ([25], Theorem IV). *Let (\mathcal{X}, ρ) be a metric space and $\mathcal{C} \subset \mathcal{X}$ compact. For all $\epsilon > 0$, we have*

$$M(2\epsilon; \mathcal{C}, \rho) \leq N^{\text{ext}}(\epsilon; \mathcal{C}, \rho) \leq N(\epsilon; \mathcal{C}, \rho) \leq M(\epsilon; \mathcal{C}, \rho). \quad (8)$$

Lemma 3.3 ([25], p. 93). *Let (\mathcal{X}, ρ_X) and (\mathcal{Y}, ρ_Y) be metric spaces and consider the compact sets $\mathcal{C}_X \subset \mathcal{X}$ and $\mathcal{C}_Y \subset \mathcal{Y}$. Assume that there exists an isometric isomorphism $f : \mathcal{C}_X \rightarrow \mathcal{C}_Y$, i.e., f is bijective and for every pair $a, b \in \mathcal{C}_X$, one has $\rho_Y(f(a), f(b)) = \rho_X(a, b)$. Then,*

$$N(\epsilon; \mathcal{C}_X, \rho_X) = N(\epsilon; \mathcal{C}_Y, \rho_Y) \quad \text{and} \quad M(\epsilon; \mathcal{C}_X, \rho_X) = M(\epsilon; \mathcal{C}_Y, \rho_Y). \quad (9)$$

Lemma 3.3 will allow us to work with a simplified metric space $(\mathcal{G}_0(w), \rho_0)$ instead of the original one $(\mathcal{G}(w), \rho_*)$. Concretely, we exploit the properties of LFM systems to effect this reduction as follows. First, as LFM systems are causal, their output at time t depends on the history of inputs up to and including time t only. Second, time-invariance implies that the mapping taking the history of the input signal to the current output at time t does not change with t and we can therefore restrict ourselves to $t = 0$ w.l.o.g. Thus, the mapping realized by an LFM system is completely characterized by the response to signals in the set

$$\mathcal{S}^- := \{s \in \mathcal{S} \mid \forall \ell \in \mathbb{N}_+ : s[\ell] = 0\}. \quad (10)$$

Using this insight, the simplified metric space can be defined as follows.

Definition 3.4. *Let $(w[t])_{t \in \mathbb{N}_0}$ be a weight sequence. We define the metric space $(\mathcal{G}_0(w), \rho_0)$ according to*

$$\mathcal{G}_0(w) = \{g : \mathcal{S}^- \rightarrow \mathbb{R} \mid |g(x) - g(x')| \leq \|x - x'\|_w, \forall x, x' \in \mathcal{S}^-, g(0) = 0\}, \quad (11)$$

$$\text{where} \quad \|x - x'\|_w := \sup_{t \in \mathbb{N}_0} w[t] |(x[-t] - x'[-t])|. \quad (12)$$

The metric ρ_0 is given by

$$\rho_0(g, g') = \sup_{x \in \mathcal{S}^-} |g(x) - g'(x)|, \quad \text{for } g, g' \in \{\mathcal{S}^- \rightarrow \mathbb{R}\}. \quad (13)$$

Next, we define the projection operator $\mathcal{P} : \mathcal{S} \rightarrow \mathcal{S}^-$ as

$$(\mathcal{P}x)[t] = x[t] \cdot \mathbb{1}_{\{t \leq 0\}}$$

and formalize the isometric isomorphism between functionals $g \in \mathcal{G}_0(w)$ and systems $G \in \mathcal{G}(w)$ as follows.

Lemma 3.5. *Let $w[\cdot]$ be a weight sequence. The map*

$$\mathcal{I} : \mathcal{G}_0(w) \rightarrow \mathcal{G}(w) \quad (14)$$

$$g \rightarrow G := (x \rightarrow \{g(\mathcal{P}\mathbf{T}_{-t}x)\}_{t \in \mathbb{Z}}) \quad (15)$$

is an isometric isomorphism between $(\mathcal{G}_0(w), \rho_0)$ and $(\mathcal{G}(w), \rho_*)$. Furthermore, $N(\epsilon; \mathcal{G}_0(w), \rho_0) = N(\epsilon; \mathcal{G}(w), \rho_*)$ and $M(\epsilon; \mathcal{G}_0(w), \rho_0) = M(\epsilon; \mathcal{G}(w), \rho_*)$, for all $\epsilon > 0$.

Proof. See C.3. □

In the remainder of this section, we first lower-bound $M(\epsilon; \mathcal{G}_0(w), \rho_0)$, then upper-bound $N(\epsilon; \mathcal{G}_0(w), \rho_0)$, and finally use Lemmata 3.3 and 3.5 to translate these bounds into bounds on $N^{\text{ext}}(\epsilon; \mathcal{G}(w), \rho_*)$. The lower bound is established as follows.

Lemma 3.6. *Let $w[\cdot]$ be a weight sequence. The ϵ -packing number of $(\mathcal{G}_0(w), \rho_0)$ satisfies*

$$\log M(\epsilon; \mathcal{G}_0(w), \rho_0) \geq \left(\prod_{\ell=0}^T \left\lceil \frac{2Dw[\ell]}{\epsilon} \right\rceil \right) - 1,$$

with $T := \max\{T' \in \mathbb{N}_0 \mid w[T'] > \frac{\epsilon}{2D}\}$.

Proof. The proof is taken from [25] and is detailed, for completeness, in C.4. □

To upper-bound $N(\epsilon; \mathcal{G}_0(w), \rho_0)$, we construct an ϵ -net for $(\mathcal{G}_0(w), \rho_0)$. This construction is again inspired by [25] but we need to modify it to ensure that the elements of the ϵ -net can efficiently be realized by ReLU networks. To be specific, we employ piecewise linear mappings to approximate LFM systems instead of piecewise constant mappings as considered in [25]; this requires significant adjustments to the proof in [25, Section 7.2].

We start by introducing the “spike” function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$ considered in [31], and defined as

$$\phi(z) = \max\{1 + \min\{z_1, \dots, z_d, 0\} - \max\{z_1, \dots, z_d, 0\}, 0\}. \quad (16)$$

An illustration of spike functions for $d = 1$ and $d = 2$ is provided in Figure 1. The idea underlying the construction of such “spike” functions can be traced back to [32], where the convex set


$$\mathcal{T} := \{z \in \mathbb{R}^d \mid \max\{z_1, \dots, z_d, 0\} - \min\{z_1, \dots, z_d, 0\} \leq 1\} \quad (17)$$

is considered and shown to be the union of $(d + 1)!$ simplices, each having 0 as a vertex, given by


$$\{z \in \mathcal{T} \mid z_{\sigma(0)} \leq z_{\sigma(1)} \leq \dots \leq z_{\sigma(d)}\}, \quad (18)$$

where σ is a permutation of the integers $0, 1, \dots, d$ and $z_0 := 0$. In [32] this result is employed to approximate continuous functions mapping \mathbb{R}^d to \mathbb{R} by functions that are piecewise linear on the simplices (18).

We remark that the spike function (16) is a composition of affine functions and min/max functions, which, as shown in Section 5, renders it uniquely suitable for realization through ReLU networks. To be specific, in Lemmata C.2 and 5.1, we provide concrete constructions of spike functions using ReLU networks. To the best of our knowledge, these constructions are novel.



(a) $d = 1$



(b) $d = 2$

Figure 1: Spike functions in dimensions 1 and 2.

We proceed to show how a *partition of unity* (p.o.u.) can be realized as a weighted linear combination of shifted spike functions, a property that will be of key importance in the RNN constructions described in Section 5.

Lemma 3.7. *Consider the spike function*

$$\phi(z) = \max\{1 + \min\{z_1, \dots, z_d, 0\} - \max\{z_1, \dots, z_d, 0\}, 0\}, \quad (19)$$

the lattice

$$\mathfrak{M} = \llbracket M_1 \rrbracket^\pm \times \dots \times \llbracket M_d \rrbracket^\pm \subset \mathbb{R}^d, \text{ where } M_\ell \in \mathbb{N}_+, \text{ for } \ell \in \{1, 2, \dots, d\}, \quad (20)$$

and the set

$$\Xi := \{\phi(\cdot - n)\}_{n \in \mathfrak{M}}. \quad (21)$$

Then, Ξ forms a p.o.u. on $\prod_{\ell=1}^d [-M_\ell, M_\ell]$, i.e.,

(i) $0 \leq \phi(z - n) \leq 1$, for $z \in \mathbb{R}^d$ and $n \in \mathfrak{M}$;

(ii) $\phi(\cdot - n)$ is compactly supported, specifically $\text{supp}(\phi(\cdot - n)) \subset n + [-1, 1]^d$;

(iii) it holds that

$$\sum_{n \in \mathfrak{M}} \phi(z - n) = 1, \quad \text{for } z \in \prod_{\ell=1}^d [-M_\ell, M_\ell].$$

Proof. To prove (i), we note that $\phi(z) \geq 0$ by definition and

$$1 + \min\{z_1, \dots, z_d, 0\} - \max\{z_1, \dots, z_d, 0\} \leq 1.$$

For (ii), it suffices to show that

$$\phi(z) = 0, \quad \text{for } z \in \mathbb{R}^d \setminus [-1, 1]^d.$$

To this end, we pick $z \in \mathbb{R}^d \setminus [-1, 1]^d$ arbitrarily and fix an arbitrary $\ell \in \{1, 2, \dots, d\}$ such that $|z_\ell| > 1$. Assume that $z_\ell > 1$. (The case $z_\ell < -1$ follows similarly.) Then,

$$1 + \min\{z_1, \dots, z_d, 0\} - \max\{z_1, \dots, z_d, 0\} \leq 1 + 0 - z_\ell < 0,$$

which by (19) implies $\phi(z) = 0$.

We proceed to prove (iii). Since $\text{supp}(\phi(\cdot - n)) \subset n + [-1, 1]^d$, we have

$$\sum_{n \in \mathfrak{N}} \phi(z - n) = \sum_{n \in \prod_{\ell=1}^d \{[z_\ell], \lceil z_\ell \rceil\}} \phi(z - n), \quad \text{for } z \in \prod_{\ell=1}^d [-M_\ell, M_\ell].$$

Defining $\bar{z} \in \mathbb{R}^d$ according to $\bar{z}_\ell = [z_\ell]$, for $\ell = 1, 2, \dots, d$, and noting that $\phi(z - n) = \phi((z - \bar{z}) - (n - \bar{z}))$, it suffices to show that

$$\sum_{n \in \{0, 1\}^d} \phi(z - n) = 1, \quad \text{for } z \in [0, 1]^d \text{ and } n \in \{0, 1\}^d.$$

As $\min\{x_1, \dots, x_d\}$ and $\max\{x_1, \dots, x_d\}$ are permutation-invariant, so is ϕ by (19). In what follows, we can therefore assume, w.l.o.g., that $z_1 \geq z_2 \geq \dots \geq z_d$.

Now, with e_k the k -th unit vector in \mathbb{R}^d , let

$$A := \left\{ 0, e_1, e_1 + e_2, \dots, \sum_{i=1}^k e_i, \dots, \sum_{i=1}^d e_i \right\} \subset \{0, 1\}^d.$$

We claim that

$$\phi(z - n) = 0, \quad \text{for } z \in [0, 1]^d \text{ and } n \in \{0, 1\}^d \setminus A. \quad (22)$$

This can be verified as follows. First, thanks to

$$n \in \{0, 1\}^d \setminus A \Leftrightarrow \exists i, j \in \{1, 2, \dots, d\}, i < j, \text{ s.t. } n_i = 0, n_j = 1,$$

we get for $z \in [0, 1]^d$ and $n \in \{0, 1\}^d \setminus A$,

$$\begin{aligned} \min\{z_1 - n_1, z_2 - n_2, \dots, z_d - n_d, 0\} &\leq z_j - n_j = z_j - 1, \\ \max\{z_1 - n_1, z_2 - n_2, \dots, z_d - n_d, 0\} &\geq z_i - n_i = z_i, \\ \Rightarrow 1 + \min\{z_1 - n_1, z_2 - n_2, \dots, z_d - n_d, 0\} - \max\{z_1 - n_1, z_2 - n_2, \dots, z_d - n_d, 0\} &\leq 0, \end{aligned}$$

and hence $\phi(z - n) = 0$. It thus suffices to show that

$$\sum_{n \in A} \phi(z - n) = 1, \quad \text{for } z \in [0, 1]^d. \quad (23)$$

Now, a direct calculation yields, for $z \in [0, 1]^d$,

$$\phi(z) = 1 - z_1, \quad (24)$$

$$\phi\left(z - \sum_{i=1}^k e_i\right) = z_k - z_{k+1}, \quad \text{for } k = 1, \dots, d-1, \quad (25)$$

$$\phi\left(z - \sum_{i=1}^d e_i\right) = z_d. \quad (26)$$

Summing (24)-(26) results in (23). \square

Next, we establish a traversal property of the lattice \mathfrak{M} .

Definition 3.8 (Regular path). *For every $d \in \mathbb{N}_+$, $M_\ell \in \mathbb{N}_+$, $\ell \in \{1, \dots, d\}$, and corresponding lattice $\mathfrak{M} = \llbracket M_1 \rrbracket^\pm \times \dots \times \llbracket M_d \rrbracket^\pm$, we call a path of lattice points $n_1 \leftrightarrow n_2 \leftrightarrow \dots \leftrightarrow n_{|\mathfrak{M}|}$ regular for \mathfrak{M} if*

(i) *the path visits each lattice point in \mathfrak{M} exactly once,*

(ii) *n_{i+1} and n_i , for each $i = 1, \dots, |\mathfrak{M}| - 1$, differ in exactly one position, specifically by $+1$ or -1 .*

Lemma 3.9. *For every $d \in \mathbb{N}_+$, $M_\ell \in \mathbb{N}_+$, $\ell \in \{1, \dots, d\}$, and corresponding lattice $\mathfrak{M} = \llbracket M_1 \rrbracket^\pm \times \dots \times \llbracket M_d \rrbracket^\pm$, there exists a regular path $n_1 \leftrightarrow n_2 \leftrightarrow \dots \leftrightarrow n_{|\mathfrak{M}|}$ for $|\mathfrak{M}|$.*

Proof. We write \mathfrak{M}_d for the d -dimensional lattice $\llbracket M_1 \rrbracket^\pm \times \dots \times \llbracket M_d \rrbracket^\pm$ to emphasize the dependence on the dimension d and prove the statement by induction over d . The base case $d = 1$ follows by simply considering the path $-M_1 \leftrightarrow -M_1 + 1 \leftrightarrow \dots \leftrightarrow 0 \leftrightarrow \dots \leftrightarrow M_1 - 1 \leftrightarrow M_1$ for lattice \mathfrak{M}_1 . Assume now that the statement holds for $d = k$, i.e., there exists a path $n_1 \leftrightarrow \dots \leftrightarrow n_{|\mathfrak{M}_k|}$ that is regular for lattice \mathfrak{M}_k . Then, for $d = k + 1$, the path $(n_1, -M_{k+1}) \leftrightarrow \dots \leftrightarrow (n_{|\mathfrak{M}_k|}, -M_{k+1}) \leftrightarrow (n_{|\mathfrak{M}_k|}, -M_{k+1} + 1) \leftrightarrow \dots \leftrightarrow (n_1, -M_{k+1} + 1) \leftrightarrow (n_1, -M_{k+1} + 2) \leftrightarrow \dots \leftrightarrow (n_{|\mathfrak{M}_k|}, -M_{k+1} + 2) \leftrightarrow \dots$ is regular for lattice \mathfrak{M}_{k+1} . \square

We are now ready to describe our construction of the ϵ -net for $(\mathcal{G}_0(w), \rho_0)$. In fact, we shall specify not only the elements of the ϵ -net, but also the mapping taking a given functional $g \in \mathcal{G}_0(w)$ to an ϵ -close element of the net. Counting the number of ball centers needed to ensure that every $g \in \mathcal{G}_0(w)$ is in the ϵ -vicinity of some ball center, then yields the cardinality of the ϵ -net. The mapping proceeds by constructing a functional \tilde{g} which is approximately faithful with respect to g on amplitude-discretized and time-truncated input signals. The formal result is as follows.

Lemma 3.10. For every $\epsilon > 0$ and $s \geq 1$, let

$$\begin{aligned} T &:= \max \left\{ \ell \in \mathbb{N}_0 \mid w[\ell] > \frac{\epsilon}{D} \frac{s}{s+1} \right\}, \\ \delta_\ell &:= \frac{s}{s+1} \frac{\epsilon}{w[\ell]}, \quad \forall \ell \in \llbracket T \rrbracket, \\ N_\ell &:= \left\lceil \frac{D}{\delta_\ell} \right\rceil, \quad \forall \ell \in \llbracket T \rrbracket, \end{aligned}$$

and define the mapping

$$\begin{aligned} k &: S_- \rightarrow \mathbb{R}^{T+1} \\ k_\ell(x) &= \frac{x[-\ell]}{\delta_\ell}, \quad \forall \ell \in \llbracket T \rrbracket. \end{aligned}$$

Furthermore, consider the lattice

$$\mathfrak{N} := \llbracket N_0 \rrbracket^\pm \times \cdots \times \llbracket N_T \rrbracket^\pm \subset \mathbb{R}^{T+1}.$$

Then, there exists a set $\mathcal{U}_0 \subset \mathbb{R}^{|\mathfrak{N}|}$ with $|\mathcal{U}_0| = (2 \lceil \frac{s+1}{2} \rceil + 1)^{|\mathfrak{N}|-1}$ such that

$$\mathcal{U} := \left\{ \tilde{g} : S_- \rightarrow \mathbb{R} \mid \tilde{g}(x) = \sum_{n \in \mathfrak{N}} \hat{g}_n \phi(k(x) - n), \{ \hat{g}_n \}_{n \in \mathfrak{N}} \in \mathcal{U}_0 \right\}$$

is an ϵ -net for $(\mathcal{G}_0(w), \rho_0)$. Furthermore, it holds that

$$|\mathcal{U}| = \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{|\mathfrak{N}|-1} = \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{\left[\prod_{\ell=0}^T (2 \lceil \frac{Dw[\ell]}{\epsilon} \frac{s+1}{s} \rceil + 1) \right] - 1}.$$

Proof. The proof proceeds in three steps.

Step 1: We pick amplitude-discretized and time-truncated signals from the set \mathcal{S}^- on the lattice \mathfrak{N} and approximately interpolate the functional $g \in \mathcal{G}_0(w)$ on these signals, using the elements of the p.o.u. $\Xi = \{ \phi(\cdot - n) \}_{n \in \mathfrak{N}}$ as interpolation basis functions, to get a new functional \hat{g} .

Fix $g \in \mathcal{G}_0(w)$ arbitrarily. For each $n \in \mathfrak{N}$, define $\hat{x}_n \in \mathcal{S}^-$ according to

$$\hat{x}_n[-\ell] = \begin{cases} \delta_\ell n_\ell, & \text{if } \ell \in \llbracket T \rrbracket \\ 0, & \text{if } \ell \in \mathbb{Z} \setminus \llbracket T \rrbracket \end{cases} \quad (27)$$

and let $\hat{g} : S_- \rightarrow \mathbb{R}$ be given by

$$\hat{g}(x) = \sum_{n \in \mathfrak{N}} g(\hat{x}_n) \phi(k(x) - n). \quad (28)$$

We first prove that

$$|\hat{g}(x) - g(x)| \leq \frac{s\epsilon}{s+1}, \quad \text{for } x \in \mathcal{S}^-. \quad (29)$$

By Lemma 3.7, $\{\phi(\cdot - n)\}_{n \in \mathfrak{N}}$ constitutes a p.o.u. on $\prod_{\ell=0}^T[-N_\ell, N_\ell]$, and hence

$$\phi(k(x) - n) = 0, \quad \text{for all } x \in \mathcal{S}^- \text{ s.t. } \|k(x) - n\|_\infty > 1. \quad (30)$$

Furthermore, for $x \in S_-$ and $n \in \mathfrak{N}$ such that $\|k(x) - n\|_\infty \leq 1$, we have

$$\left| \frac{x[-\ell]}{\delta_\ell} - n_\ell \right| \leq 1, \quad \text{for } \ell \in \llbracket T \rrbracket,$$

which gives

$$\begin{aligned} w[\ell] |x[-\ell] - \hat{x}_n[-\ell]| &\leq w[\ell] \delta_\ell = \frac{s\epsilon}{s+1}, \quad \text{for } \ell \in \llbracket T \rrbracket, \\ w[\ell] |x[-\ell] - \hat{x}_n[-\ell]| &= w[\ell] |x[-\ell]| \\ &\leq \frac{\epsilon}{D} \frac{s}{s+1} D = \frac{s\epsilon}{s+1}, \quad \text{for } \ell > T, \end{aligned}$$

and therefore

$$\|x - \hat{x}_n\|_w \leq \frac{s\epsilon}{s+1}. \quad (31)$$

Hence, we can bound the approximation error $|g(x) - \hat{g}(x)|$, for $x \in \mathcal{S}^-$, according to

$$|g(x) - \hat{g}(x)| \stackrel{(28), \text{p.o.u.}}{=} \left| \sum_{n \in \mathfrak{N}} (g(x) - g(\hat{x}_n)) \phi(k(x) - n) \right| \quad (32)$$

$$\stackrel{(30)}{=} \left| \sum_{\|k(x) - n\|_\infty \leq 1} (g(x) - g(\hat{x}_n)) \phi(k(x) - n) \right| \quad (33)$$

$$\stackrel{(\phi \geq 0)}{\leq} \sum_{\|k(x) - n\|_\infty \leq 1} |g(x) - g(\hat{x}_n)| \phi(k(x) - n) \quad (34)$$

$$\stackrel{\text{Lipschitz}}{\leq} \sum_{\|k(x) - n\|_\infty \leq 1} \|x - \hat{x}_n\|_w \phi(k(x) - n) \quad (35)$$

$$\stackrel{(31)}{\leq} \frac{s\epsilon}{s+1} \sum_{\|k(x) - n\|_\infty \leq 1} \phi(k(x) - n) \quad (36)$$

$$\stackrel{(30)}{=} \frac{s\epsilon}{s+1} \sum_{n \in \mathfrak{N}} \phi(k(x) - n) \quad (37)$$

$$\stackrel{\text{p.o.u.}}{=} \frac{s\epsilon}{s+1}. \quad (38)$$

Step 2: We modify the interpolation weights $g(\hat{x}_n)$ in \hat{g} to weights \hat{g}_n , for every $n \in \mathfrak{N}$, along a fixed path of \mathfrak{N} , to get a new functional \tilde{g} that approximates g to within an error of at

most ϵ .

Specifically, we construct, by induction, a functional \tilde{g} such that

$$|\tilde{g}(x) - \hat{g}(x)| \leq \Delta := \frac{\epsilon}{s+1}. \quad (39)$$

To this end, we let

$$\tilde{g}(x) := \sum_{n \in \mathfrak{N}} \hat{g}_n \phi(k(x) - n) \quad (40)$$

and then find, for every $n \in \mathfrak{N}$, a value \hat{g}_n so that

$$|g(\hat{x}_n) - \hat{g}_n| \leq \Delta, \quad (41)$$

which, in turn, yields

$$\begin{aligned} |\hat{g}(x) - \tilde{g}(x)| &= \left| \sum_{n \in \mathfrak{N}} (g(\hat{x}_n) - \hat{g}_n) \phi(k(x) - n) \right| \\ &\leq \sum_{n \in \mathfrak{N}} |g(\hat{x}_n) - \hat{g}_n| \phi(k(x) - n) \\ &\leq \Delta \sum_{n \in \mathfrak{N}} \phi(k(x) - n) \\ &\stackrel{\text{p.o.u.}}{=} \Delta. \end{aligned}$$

Therefore,

$$|\tilde{g}(x) - g(x)| \leq |\tilde{g}(x) - \hat{g}(x)| + |\hat{g}(x) - g(x)| \leq \Delta + \frac{s\epsilon}{s+1} = \epsilon. \quad (42)$$

It remains to specify how the values \hat{g}_n , $n \in \mathfrak{N}$, can be obtained such that (41) holds. This will be done by performing the mapping from $g(\hat{x}_n)$ to \hat{g}_n along a judiciously chosen path $n_1 \leftrightarrow n_2 \leftrightarrow \dots \leftrightarrow n_{|\mathfrak{N}|}$ that is regular for \mathfrak{N} (Lemma 3.9). To this end, we distinguish two cases.

- *Case 1: The path $n_1 \leftrightarrow n_2 \leftrightarrow \dots \leftrightarrow n_{|\mathfrak{N}|}$ starts or ends at $0 \in \mathfrak{N}$.*

In particular, we assume, w.l.o.g., that $n_1 = 0$. Next, we find the values \hat{g}_{n_k} satisfying (41) inductively over the index $k = 1, 2, \dots, |\mathfrak{N}|$. The base case $k = 1$ is immediate as we can simply set $\hat{g}_{n_1} = 0$ and thereby obtain

$$|g(\hat{x}_{n_1}) - \hat{g}_{n_1}| \stackrel{(11)}{=} |g(0) - 0| = 0 \leq \Delta.$$

Now, assume that (41) holds for some n_k , $k \in \{1, 2, \dots, |\mathfrak{N}|\}$, on the path $n_1 \rightarrow \dots \rightarrow n_{|\mathfrak{N}|}$. Then, by Lemma 3.9, n_{k+1} differs in exactly one position from n_k , namely by $+1$ or -1 ,

which, thanks to (12) and (27), yields $\|\widehat{x}_{n_{k+1}} - \widehat{x}_{n_k}\|_w \leq \frac{\epsilon s}{s+1}$. Upon noting that

$$\begin{aligned} |g(\widehat{x}_{n_{k+1}}) - \widehat{g}_{n_k}| &= \left| \underbrace{g(\widehat{x}_{n_{k+1}}) - g(\widehat{x}_{n_k})}_{\text{Lipschitz}} + \underbrace{g(\widehat{x}_{n_k}) - \widehat{g}_{n_k}}_{(41)} \right| \\ &\leq \|\widehat{x}_{n_{k+1}} - \widehat{x}_{n_k}\|_w + \Delta \\ &\leq \frac{\epsilon s}{s+1} + \Delta = (s+1)\Delta, \end{aligned}$$

we can conclude that there exists an

$$\begin{aligned} m &\in \left\{ -2 \left\lceil \frac{s+1}{2} \right\rceil, -2 \left\lceil \frac{s+1}{2} \right\rceil + 2, \dots, 2 \left\lceil \frac{s+1}{2} \right\rceil \right\}, \\ \text{s.t. } \widehat{g}_{n_{k+1}} &= \widehat{g}_{n_k} + m\Delta \quad \text{and} \quad |\widehat{g}_{n_{k+1}} - g(\widehat{x}_{n_{k+1}})| \leq \Delta. \end{aligned} \quad (43)$$

This completes the induction.

- *Case 2: The path $n_1 \leftrightarrow n_2 \leftrightarrow \dots \leftrightarrow n_{|\mathfrak{N}|}$ does not start or end at $0 \in \mathfrak{N}$.*

With $n_i = 0$, for some $i \in \{2, 3, \dots, |\mathfrak{N}| - 1\}$, and following the spirit of *Case 1*, we split the path $n_1 \leftrightarrow n_2 \leftrightarrow \dots \leftrightarrow n_{|\mathfrak{N}|}$ into $\text{path}_{\leftarrow} := n_i \rightarrow n_{i-1} \rightarrow \dots \rightarrow n_1$ and $\text{path}_{\rightarrow} := n_i \rightarrow n_{i+1} \rightarrow \dots \rightarrow n_{|\mathfrak{N}|}$. The idea is to prove (41) by performing induction across path_{\leftarrow} and $\text{path}_{\rightarrow}$ separately. This can be done following the same procedure as in *Case 1*.

Step 3: Repeat *Steps 1* and *2* for all $g \in \mathcal{G}_0(w)$ and collect all the resulting $\{\widehat{g}_n\}_{n \in \mathfrak{N}}$ in the set \mathcal{U}_0 . For *Case 1* in *Step 2*, we need to store one value for \widehat{g}_{n_1} and, according to (43), $(2 \lceil \frac{s+1}{2} \rceil + 1)$ increments or decrements from \widehat{g}_{n_i} to $\widehat{g}_{n_{i+1}}$, for $i = 0, 1, \dots, |\mathfrak{N}| - 1$. This yields a total of

$$\left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{|\mathfrak{N}|-1}$$

values. For *Case 2*, noting that path_{\leftarrow} and $\text{path}_{\rightarrow}$ are of length $i - 1$ and $|\mathfrak{N}| - i$, respectively, and applying the same argument as above, there is again a total of

$$\left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{i-1} \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{|\mathfrak{N}|-i} = \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{|\mathfrak{N}|-1}$$

increments or decrements. The set \mathcal{U}_0 is hence of cardinality

$$|\mathcal{U}_0| = \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{|\mathfrak{N}|-1} = \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{\left[\prod_{\ell=0}^T (2 \lceil \frac{Dw[\ell] s+1}{\epsilon} \rceil + 1) \right] - 1}. \quad (44)$$

Finally, setting

$$\mathcal{U} = \left\{ \tilde{g} : S_- \rightarrow \mathbb{R} \mid \tilde{g}(x) = \sum_{n \in \mathfrak{N}} \widehat{g}_n \phi(k(x) - n), \{ \widehat{g}_n \}_{n \in \mathfrak{N}} \in \mathcal{U}_0 \right\},$$

concludes the proof. \square

Based on the ϵ -net constructed in Lemma 3.10, we can now upper-bound the metric entropy of $(\mathcal{G}_0(w), \rho_0)$ as follows.

Corollary 3.11. *For every $s \geq 1$, the covering number of $(\mathcal{G}_0(w), \rho_0)$ satisfies*

$$\log N(\epsilon; \mathcal{G}_0(w), \rho_0) \leq \log \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right) \prod_{\ell=0}^T \left(2 \left\lceil \frac{2Dw[\ell]}{\epsilon} \frac{s+1}{s} \right\rceil + 1 \right),$$

where $T := \max \{ \ell \in \mathbb{N}_0 \mid w[\ell] > \frac{\epsilon}{2D} \frac{s}{s+1} \}$.

Proof. For every $\epsilon > 0$, Lemma 3.10 delivers an $\frac{\epsilon}{2}$ -net for $(\mathcal{G}_0(w), \rho_0)$ with

$$\left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{\left[\prod_{\ell=0}^T (2 \left\lceil \frac{2Dw[\ell]}{\epsilon} \frac{s+1}{s} \right\rceil + 1) \right] - 1}$$

elements. By Definition 2.8, it hence follows that

$$N^{\text{ext}}(\epsilon/2; \mathcal{G}_0(w), \rho_0) \leq \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{\left[\prod_{\ell=0}^T (2 \left\lceil \frac{2Dw[\ell]}{\epsilon} \frac{s+1}{s} \right\rceil + 1) \right]}.$$

Application of Lemma 3.2 then yields an upper bound on the covering number according to

$$N(\epsilon; \mathcal{G}_0(w), \rho_0) \leq N^{\text{ext}}(\epsilon/2; \mathcal{G}_0(w), \rho_0) \leq \left(2 \left\lceil \frac{s+1}{2} \right\rceil + 1 \right)^{\left[\prod_{\ell=0}^T (2 \left\lceil \frac{2Dw[\ell]}{\epsilon} \frac{s+1}{s} \right\rceil + 1) \right]},$$

which finishes the proof. \square

We conclude the developments in this section by deriving a lower bound and an upper bound on the exterior covering number of $(\mathcal{G}(w), \rho_*)$.

Theorem 3.12. *The exterior covering number of $(\mathcal{G}(w), \rho_*)$ satisfies*

$$\left(\prod_{\ell=0}^{T'} \left\lceil \frac{Dw[\ell]}{\epsilon} \right\rceil \right) - 1 \leq \log N^{\text{ext}}(\epsilon; \mathcal{G}(w), \rho_*) \leq \log (3) \prod_{\ell=0}^{T''} \left(2 \left\lceil \frac{4Dw[\ell]}{\epsilon} \right\rceil + 1 \right), \quad (45)$$

where $T' := \max \{ \ell \in \mathbb{N}_0 \mid w[\ell] > \frac{\epsilon}{D} \}$ and $T'' := \max \{ \ell \in \mathbb{N}_0 \mid w[\ell] > \frac{\epsilon}{4D} \}$.

Proof. The result follows by noting that

$$\begin{aligned}
\left(\prod_{\ell=0}^{T'} \left\lceil \frac{Dw[\ell]}{\epsilon} \right\rceil \right) - 1 &\stackrel{\text{Lemma 3.6}}{\leq} \log M(2\epsilon; \mathcal{G}_0(w), \rho_0) \\
&\stackrel{\text{Lemma 3.5}}{=} \log M(2\epsilon; \mathcal{G}(w), \rho_*) \\
&\stackrel{\text{Lemma 3.2}}{\leq} \log N^{\text{ext}}(\epsilon; \mathcal{G}(w), \rho_*) \leq \log N(\epsilon; \mathcal{G}(w), \rho_*) \\
&\stackrel{\text{Lemma 3.5}}{=} \log N(\epsilon; \mathcal{G}_0(w), \rho_0) \\
&\stackrel{\text{Corollary 3.11}}{\leq} \log(3) \prod_{\ell=0}^{T''} \left(2 \left\lceil \frac{4Dw[\ell]}{\epsilon} \right\rceil + 1 \right).
\end{aligned} \tag{46}$$

□

4 Metric entropy of ELFM and PLFM systems

We now discuss two specific classes of LFM systems, namely ELFM and PLFM systems. Specifically, we characterize the description complexity of these two classes by computing their type, order, and generalized dimension. The corresponding results will then serve as a reference for the RNN approximations in Section 5. Specifically, we will establish, in Section 6, that RNNs can approximate ELFM and PLFM systems in metric-entropy-optimal manner.

4.1 Metric entropy of ELFM systems

The concept of ELFM systems is inspired, inter alia, by applications in finance, such as those discussed in [8], where asset pricing decisions are influenced by past observations. Instead of relying solely on finite-length observations, the model integrates infinite past observations by endowing them with exponentially decaying memory. Similar settings are also considered in random walk models [29, 30]. These examples, when appropriately adapted to the setup in the present paper, fit into the setting of LFM systems according to Definition 2.16, with respect to exponentially decaying weight sequences. We formally define exponentially decaying memory and the corresponding ELFM systems as follows.

Definition 4.1. For $a \in (0, 1]$ and $b > 0$, let

$$w_{a,b}^{(e)}[t] := ae^{-bt}, \quad \text{for all } t \in \mathbb{N}_0.$$

An LFM system with respect to $\{w_{a,b}^{(e)}[t]\}_{t \in \mathbb{N}_0}$ is said to be ELFM. We write $\mathcal{G}(w_{a,b}^{(e)})$ for the class of systems that are ELFM with respect to the weight sequence $w_{a,b}^{(e)}[t]$.

The remainder of this section is devoted to computing the order, type, and generalized dimension of $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$. To this end, we first establish an auxiliary result.

Lemma 4.2. *Let $a \in (0, 1]$, $b, c, d > 0$ and consider the weight sequence $w_{a,b}^{(e)}$ as per Definition 4.1. Let*

$$T := \max \left\{ t \in \mathbb{N}_0 \mid w_{a,b}^{(e)}[t] > \frac{\epsilon}{d} \right\}.$$

Then,

$$\log \left(\prod_{\ell=0}^T \frac{cw_{a,b}^{(e)}[\ell]}{\epsilon} \right) = \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})). \quad (47)$$

Proof. See C.5. □

We are now ready to state the main result of this section quantifying the massiveness of the class of ELFM systems.

Lemma 4.3. *Let $a \in (0, 1]$ and $b > 0$. The class of ELFM systems $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ is of order 1 and type 2, with generalized dimension*

$$\mathfrak{d} = \frac{1}{2b \log(e)}.$$

Proof. Consider $\epsilon \in (0, \epsilon_0)$ with $\epsilon_0 = \frac{Dw_{a,b}^{(e)}[0]}{2} = \frac{aD}{2}$. By Theorem 3.12, the exterior covering number of $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ satisfies

$$\left(\prod_{\ell=0}^{T'} \left\lceil \frac{Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right\rceil \right) - 1 \leq \log N^{\text{ext}}(\epsilon; \mathcal{G}(w_{a,b}^{(e)}), \rho_*) \leq \log(3) \prod_{\ell=0}^{T''} \left(2 \left\lceil \frac{4Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right\rceil + 1 \right), \quad (48)$$

where

$$T' := \max \left\{ \ell \in \mathbb{N}_0 \mid w_{a,b}^{(e)}[\ell] > \frac{\epsilon}{D} \right\} \text{ and } T'' := \max \left\{ \ell \in \mathbb{N}_0 \mid w_{a,b}^{(e)}[\ell] > \frac{\epsilon}{4D} \right\}.$$

We can further lower-bound the left-most term in (48) according to

$$\begin{aligned} \left(\prod_{\ell=0}^{T'} \left\lceil \frac{Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right\rceil \right) - 1 &\geq \left(\prod_{\ell=0}^{T'} \frac{Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right) - 1 \\ &\geq \frac{1}{2} \prod_{\ell=0}^{T'} \frac{Dw_{a,b}^{(e)}[\ell]}{\epsilon}, \end{aligned} \quad (49)$$

where (49) follows from

$$\frac{1}{2} \prod_{\ell=0}^{T'} \frac{Dw_{a,b}^{(e)}[\ell]}{\epsilon} \geq 1,$$

which, in turn, is a consequence of

$$\begin{aligned}\frac{Dw_{a,b}^{(e)}[0]}{\epsilon} &\geq \frac{Dw_{a,b}^{(e)}[0]}{\epsilon_0} = 2, \\ \frac{Dw_{a,b}^{(e)}[\ell]}{\epsilon} &\geq \frac{Dw_{a,b}^{(e)}[T']}{\epsilon} > 1, \quad \text{for } \ell \in \llbracket T' \rrbracket \setminus \{0\}.\end{aligned}$$

Similarly, we can further upper-bound the right-most term in (48) according to

$$\begin{aligned}\log(3) \prod_{\ell=0}^{T''} \left(2 \left\lceil \frac{4Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right\rceil + 1 \right) &\leq \log(3) \prod_{\ell=0}^{T''} \left(\frac{8Dw_{a,b}^{(e)}[\ell]}{\epsilon} + 3 \right) \\ &\leq \log(3) \prod_{\ell=0}^{T''} \frac{20Dw_{a,b}^{(e)}[\ell]}{\epsilon},\end{aligned}\tag{50}$$

where (50) follows from

$$\frac{4Dw_{a,b}^{(e)}[\ell]}{\epsilon} \geq \frac{4Dw_{a,b}^{(e)}[T'']}{\epsilon} > 1, \quad \text{for } \ell \in \llbracket T'' \rrbracket.$$

Combining (48)–(50), taking logarithms one more time, and dividing by $\log^2(\epsilon^{-1})$, yields

$$\begin{aligned}\frac{\log \left(\prod_{\ell=0}^{T'} \frac{Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right) - 1}{\log^2(\epsilon^{-1})} &\leq \frac{\log^{(2)} N^{\text{ext}}(\epsilon; \mathcal{G}(w_{a,b}^{(e)}), \rho_*)}{\log^2(\epsilon^{-1})} \\ &\leq \frac{\log \left(\prod_{\ell=0}^{T''} \frac{20Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right) + \log^{(2)}(3)}{\log^2(\epsilon^{-1})}.\end{aligned}\tag{51}$$

Taking the limit $\epsilon \rightarrow 0$ and applying Lemma 4.2 to the lower and the upper bound in (51), we obtain

$$\lim_{\epsilon \rightarrow 0} \frac{\log^{(2)} N^{\text{ext}}(\epsilon; \mathcal{G}(w_{a,b}^{(e)}), \rho_*)}{\log^2(\epsilon^{-1})} = \frac{1}{2b \log(e)},$$

which implies that $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ is of order 1 and type 2, with generalized dimension

$$\mathfrak{d} = \frac{1}{2b \log(e)}.$$

This concludes the proof. \square

The generalized dimension being inversely proportional to b reflects that faster memory decay rates, i.e., larger b , result in system classes that are less massive. Additionally, we note that the description complexity of ELFM systems is primarily determined by the order being equal to 1 and the type equal to 2. Compared to the function classes in Table 1, which are all

of order 1 and type 1, this shows that the class of ELFM systems is significantly more massive than unit balls in function spaces.

4.2 Metric entropy of PLFM systems

Next, we consider systems with polynomially decaying memory, a concept used, e.g., in the context of PDEs [33,34]. Specifically, the references [33,34] are concerned with Volterra integro-differential equations [28] of polynomially decaying memory kernels. These examples, when suitably adjusted to the framework in the present paper, align with dynamical systems that are LFM with respect to polynomially decaying weight sequences. We formalize the concept of polynomially decaying memory and PLFM systems as follows.

Definition 4.4. For $q \in (0, 1]$ and $p > 0$, let

$$w_{p,q}^{(p)}[t] := \frac{q}{(1+t)^p}, \quad \text{for all } t \in \mathbb{N}_0.$$

An LFM system with respect to $\{w_{p,q}^{(p)}[t]\}_{t \in \mathbb{N}_0}$ is said to be PLFM. We write $\mathcal{G}(w_{p,q}^{(p)})$ for the class of systems that are PLFM with respect to the weight sequence $w_{p,q}^{(p)}[t]$.

As in the previous section, we first need an auxiliary result.

Lemma 4.5. Let $q \in (0, 1]$, $p, c, d > 0$ and consider $w_{p,q}^{(p)}$ as per Definition 4.4. Let

$$T := \max \left\{ t \in \mathbb{N}_0 \mid w_{p,q}^{(p)}[t] > \frac{\epsilon}{d} \right\}.$$

Then,

$$\log \left(\prod_{\ell=0}^T \frac{c w_{p,q}^{(p)}[\ell]}{\epsilon} \right) = \Theta(\epsilon^{-1/p}). \quad (52)$$

Proof. See C.6. □

We obtain the order, type, and generalized dimension of PLFM systems as follows.

Lemma 4.6. Let $q \in (0, 1]$ and $p > 0$. The class of PLFM systems $(\mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ is of order 2 and type 1, with generalized dimension

$$\mathfrak{d} = \frac{1}{p}.$$

Proof. See C.7. □

Compared to the class of ELFM systems, which exhibits order 1, the class of PLFM systems is more massive as it has order 2. This is intuitively meaningful as polynomial decay is significantly slower than exponential decay, rendering PLFM systems to depend more strongly on past inputs. Additionally, the generalized dimension exhibiting inverse proportionality to p , reflects that faster (polynomial) decay, i.e., larger p , leads to smaller description complexity.

5 Approximating LFM systems through RNNs

We now proceed to realize the elements in the ϵ -net for $(\mathcal{G}_0(w), \rho_0)$ constructed in Lemma 3.10 by ReLU networks. Based on the connection between $(\mathcal{G}_0(w), \rho_0)$ and $(\mathcal{G}(w), \rho_*)$ identified in Section 3, this will then allow us to construct RNNs approximating general LFM systems.

5.1 Approximating $(\mathcal{G}_0(w), \rho_0)$ with ReLU networks

The building block in Lemma 3.10 for approximating functionals in $\mathcal{G}_0(w)$ is the p.o.u. $\Xi = \{\phi(\cdot - n)\}_{n \in \mathfrak{M}}$, which is why we first focus on constructing ReLU networks realizing Ξ . As the elements of Ξ are shifted versions of the spike function ϕ and shifts, by virtue of being affine transformations, are trivially realized by single-layer ReLU networks, it suffices to find a ReLU network realization of ϕ . Note that this argument made use of the fact that compositions of ReLU networks are again ReLU networks (see Lemma C.3).

Lemma 5.1. *For $d \in \mathbb{N}_+$, consider the spike function $\phi : \mathbb{R}^d \rightarrow \mathbb{R}$,*

$$\phi(z) = \max\{1 + \min\{z_1, \dots, z_d, 0\} - \max\{z_1, \dots, z_d, 0\}, 0\}. \quad (53)$$

There exists a ReLU network $\Phi \in \mathcal{N}_{d,1}$ with $\mathcal{L}(\Phi) = \lceil \log(d+1) \rceil + 4$, $\mathcal{M}(\Phi) \leq 60d - 28$, $\mathcal{W}(\Phi) \leq 6d$, $\mathcal{K}(\Phi) = \{1, -1\}$, and $\mathcal{B}(\Phi) = 1$, such that

$$\Phi(z) = \phi(z), \quad \text{for all } z \in \mathbb{R}^d. \quad (54)$$

Proof. Let Φ_d^{max} be the ReLU network realization of $\max\{z_1, z_2, \dots, z_d\}$ according to Lemma C.2. Then, using $\min\{z_1, z_2, \dots, z_d\} = -\max\{-z_1, -z_2, \dots, -z_d\}$, we obtain

$$\begin{aligned} \phi(z) &= \rho(1 - \rho(\Phi_d^{max}(-z)) - \rho(\Phi_d^{max}(z))) \\ &= \underbrace{((W_3 \circ \rho \circ W_2 \circ \rho \circ P(\Phi_d^{max}, \Phi_d^{max}))}_{=: \Phi_2} \circ \underbrace{W_1}_{=: \Phi_1})(z), \end{aligned} \quad (55)$$

where $W_1(z) = \begin{pmatrix} \mathbb{I}_d & -\mathbb{I}_d \end{pmatrix}^T z$, $W_2(z) = \begin{pmatrix} -1 & -1 \end{pmatrix} z + 1$, $W_3(z) = z$, and $P(\Phi_d^{max}, \Phi_d^{max})$ is the parallelization of two Φ_d^{max} -networks according to Lemma C.4 such that $P(\Phi_d^{max}, \Phi_d^{max})(z) = (\Phi_d^{max}(z), \Phi_d^{max}(z))^T$. Now, by Lemma C.3, there exists a ReLU network Φ realizing $\Phi_2 \circ \Phi_1$ in

(55), with

$$\begin{aligned}
\mathcal{L}(\Phi) &\stackrel{\text{Lemma C.3}}{=} \mathcal{L}(\Phi_2) + \mathcal{L}(\Phi_1) \\
&\stackrel{\text{Lemma C.2, C.4}}{=} \lceil \log(d+1) \rceil + 4, \\
\mathcal{M}(\Phi) &\stackrel{\text{Lemma C.3}}{\leq} 2\mathcal{M}(\Phi_1) + 2\mathcal{M}(\Phi_2) \\
&\stackrel{\text{Lemma C.4}}{=} 2\mathcal{M}(W_3) + 2\mathcal{M}(W_2) + 2\mathcal{M}(W_1) + 4\mathcal{M}(\Phi_d^{max}) \\
&\stackrel{\text{Lemma C.2}}{\leq} 60d - 28, \\
\mathcal{W}(\Phi) &\stackrel{\text{Lemma C.3}}{\leq} \max\{4d, \max\{\mathcal{W}(\Phi_2), \mathcal{W}(\Phi_1)\}\} \\
&\stackrel{\text{Lemma C.4}}{\leq} \max\{4d, 2\mathcal{W}(\Phi_d^{max}), \mathcal{W}(W_3), \mathcal{W}(W_2), \mathcal{W}(W_1)\} \\
&\stackrel{\text{Lemma C.2}}{\leq} 6d, \\
\mathcal{K}(\Phi) &\stackrel{\text{Lemma C.3}}{\subset} \mathcal{K}(\Phi_2) \cup (-\mathcal{K}(\Phi_2)) \cup \mathcal{K}(\Phi_1) \cup (-\mathcal{K}(\Phi_1)) \\
&\stackrel{\text{Lemma C.3, C.4}}{=} \left(\bigcup_{i=1}^3 \mathcal{K}(W_i) \cup (-\mathcal{K}(W_i)) \right) \cup \mathcal{K}(\Phi_d^{max}) \cup (-\mathcal{K}(\Phi_d^{max})) \\
&\stackrel{\text{Lemma C.2}}{=} \{1, -1\}, \\
\mathcal{B}(\Phi) &= \max_{b \in \mathcal{K}(\Phi)} |b| = 1.
\end{aligned}$$

This concludes the proof. \square

We now show how the elements of the ϵ -net constructed in Lemma 3.10 can be realized by ReLU networks. In particular, our construction will be found, in Section 6, to be encodable by bitstrings of length scaling (in ϵ) in a metric-entropy optimal manner.

Lemma 5.2. *For every $s \geq 1$, there exists $\epsilon_0 > 0$, such that for $\epsilon \in (0, \epsilon_0)$, with*

$$T := \max \left\{ \ell \in \mathbb{N}_0 \mid w[\ell] > \frac{\epsilon}{D} \frac{s}{s+1} \right\},$$

for every $g \in \mathcal{G}_0(w)$, there exists a $\Phi \in \mathcal{N}_{T+1,1}$ with $(2, \epsilon)$ -quantized weights (Definition 2.2) satisfying

$$\left| \Phi \left(\{x[-\ell]\}_{\ell=0}^T \right) - g(x) \right| \leq \epsilon, \quad \text{for all } x \in \mathcal{S}^-.$$

Moreover,

$$\mathcal{L}(\Phi) = \lceil \log(T+2) \rceil + 6 \quad \text{and} \quad \mathcal{M}(\Phi) \leq 244(T+1) \prod_{\ell=0}^T \left(\frac{2Dw[\ell]s+1}{\epsilon} \frac{s+1}{s} + 4 \right). \quad (56)$$

Proof. Fix $g \in \mathcal{G}_0(w)$. To construct a $(2, \epsilon)$ -quantized ReLU network approximating g , we

follow the spirit of the proof of Lemma 3.10 and first consider

$$\hat{g}(x) = \sum_{n \in \mathfrak{N}} g(\hat{x}_n) \phi \left(\left\{ \frac{x[-\ell]}{\delta_\ell} - n_\ell \right\}_{\ell=0}^T \right), \quad (57)$$

with

$$\begin{aligned} \delta_\ell &:= \frac{s}{s+1} \frac{\epsilon}{w[\ell]}, \quad \forall \ell \in \llbracket T \rrbracket, \\ N_\ell &:= \left\lceil \frac{D}{\delta_\ell} \right\rceil, \quad \forall \ell \in \llbracket T \rrbracket, \\ \mathfrak{N} &:= \llbracket N_0 \rrbracket^\pm \times \cdots \times \llbracket N_T \rrbracket^\pm \subset \mathbb{R}^{T+1}, \\ \hat{x}_n[-\ell] &:= \begin{cases} \delta_\ell n_\ell, & \text{if } \ell \in \llbracket T \rrbracket, \\ 0, & \text{else.} \end{cases} \end{aligned} \quad (58)$$

It was shown in (29) that

$$|\hat{g}(x) - g(x)| \leq \frac{s\epsilon}{s+1}, \quad \text{for } x \in \mathcal{S}^-. \quad (59)$$

We next quantize the parameters δ_ℓ^{-1} and the lattice \mathfrak{N} in (58) according to

$$\begin{aligned} \tilde{\delta}_\ell^{-1} &:= \mathcal{Q}_{2,\epsilon}(\delta_\ell^{-1}), \\ \tilde{N}_\ell &= \left\lceil \frac{D}{\tilde{\delta}_\ell} \right\rceil, \quad \ell \in \llbracket T \rrbracket, \\ \tilde{\mathfrak{N}} &= \llbracket \tilde{N}_0 \rrbracket^\pm \times \cdots \times \llbracket \tilde{N}_T \rrbracket^\pm, \end{aligned} \quad (60)$$

and adjust the grid points \hat{x}_n to

$$\tilde{x}_n[-\ell] = \begin{cases} \tilde{\delta}_\ell n_\ell, & \text{if } \ell \in \llbracket T \rrbracket, \\ 0, & \text{else.} \end{cases} \quad (61)$$

Furthermore, we quantize $g(\tilde{x}_n)$ according to

$$\tilde{g}_n := \mathcal{Q}_{2,\epsilon}(g(\tilde{x}_n)) \quad (62)$$

and consider the function

$$\tilde{g}(x) = \sum_{n \in \tilde{\mathfrak{N}}} \tilde{g}_n \phi \left(\left\{ \frac{x[-\ell]}{\tilde{\delta}_\ell} - n_\ell \right\}_{\ell=0}^T \right) =: f \left(\{x[-\ell]\}_{\ell=0}^T \right). \quad (63)$$

For ease of notation, we define $\tilde{k} : \mathcal{S}_- \rightarrow \mathbb{R}^{T+1}$ as $\tilde{k}_\ell(x) = \tilde{\delta}_\ell^{-1} x[-\ell]$, for $\ell \in \llbracket T \rrbracket$. Next, we show that

$$|\tilde{g}(x) - g(x)| \leq \epsilon. \quad (64)$$

This follows from

$$|\tilde{g}(x) - g(x)| = \left| \sum_{\|\tilde{k}(x) - n\|_\infty \leq 1} (g(x) - \tilde{g}_n) \phi(\tilde{k}(x) - n) \right| \quad (65)$$

$$\leq \sum_{\|\tilde{k}(x) - n\|_\infty \leq 1} (|g(x) - g(\tilde{x}_n)| + |g(\tilde{x}_n) - \tilde{g}_n|) \phi(\tilde{k}(x) - n) \quad (66)$$

$$\leq \sum_{\|\tilde{k}(x) - n\|_\infty \leq 1} \left(\|x - \tilde{x}_n\|_w + \frac{\epsilon}{s+1} \right) \phi(\tilde{k}(x) - n) \quad (67)$$

$$\leq \epsilon \sum_{\|\tilde{k}(x) - n\|_\infty \leq 1} \phi(\tilde{k}(x) - n) \quad (68)$$

$$\leq \epsilon. \quad (69)$$

Here, (65) and (69) are by the p.o.u. property of ϕ and (67) is a consequence of the Lipschitz property of g according to (11) and

$$\begin{aligned} |g(\tilde{x}_n) - \tilde{g}_n| &= |g(\tilde{x}_n) - \mathcal{Q}_{2,\epsilon}(g(\tilde{x}_n))| \\ &\leq 2^{-2 \lceil \log(\epsilon^{-1}) \rceil} \stackrel{\epsilon \in (0, \epsilon_0), (189)}{\leq} \frac{\epsilon}{s+1}. \end{aligned} \quad (70)$$

Further, (68) follows from the fact that for $\|\tilde{k}(x) - n\|_\infty \leq 1$,

$$\left| \frac{x^{[-\ell]}}{\tilde{\delta}_\ell} - n_\ell \right| \leq 1, \quad \text{for } \ell \in \llbracket T \rrbracket,$$

and hence

$$\|x - \tilde{x}_n\|_w \leq \max_{\ell \in \llbracket T \rrbracket} \tilde{\delta}_\ell w[\ell] \leq \frac{s\epsilon}{s+1}, \quad (71)$$

where the second inequality in (71) is by

$$\tilde{\delta}_\ell \stackrel{(60)}{=} (\mathcal{Q}_{2,\epsilon}(\delta_\ell^{-1}))^{-1} \leq \delta_\ell = \frac{s}{s+1} \frac{\epsilon}{w[\ell]}.$$

Based on (63), we can rewrite (64) according to

$$\left| f \left(\{x^{[-\ell]}\}_{\ell=0}^T \right) - g(x) \right| \leq \epsilon, \quad \text{for } x \in S_-. \quad (72)$$

It hence suffices to construct a ReLU network Φ that realizes f , which then, thanks to (72), approximates g to within an error of at most ϵ . To this end, assume that $n^1, n^2, \dots, n^{|\tilde{\mathfrak{N}}|}$ is an arbitrary, but fixed, enumeration of the elements of $\tilde{\mathfrak{N}}$. Set $\tilde{W}^\Sigma(x) = \Lambda x$ and $\tilde{W}_{n^i}(x) = \hat{B}x + \hat{b}_{n^i}$,

with

$$\begin{aligned}\Lambda &= \left(\tilde{g}_{n^1} \quad \tilde{g}_{n^2} \quad \dots \quad \tilde{g}_{n^{|\tilde{\mathfrak{N}}|}} \right), \\ \widehat{B} &= \text{diag} \left(\tilde{\delta}_0^{-1}, \tilde{\delta}_1^{-1}, \dots, \tilde{\delta}_T^{-1} \right), \\ \hat{b}_{n^i} &= \left(-n_0^i \quad -n_1^i \quad \dots \quad -n_T^i \right)^T.\end{aligned}\tag{73}$$

Moreover, let Ψ be a ReLU network realizing the spike function ϕ according to Lemma 5.1 and define the following ReLU networks

$$\begin{aligned}\Phi_2 &= \widetilde{W}^\Sigma, \\ \Phi_{1,2}^{n^i} &= \Psi, \quad \text{for } i = 1, \dots, |\tilde{\mathfrak{N}}|, \\ \Phi_{1,1}^{n^i} &= \widehat{W}_{n^i}, \quad \text{for } i = 1, \dots, |\tilde{\mathfrak{N}}|.\end{aligned}\tag{74}$$

Next, we use Lemma C.3 to compose $\Phi_{1,2}^{n^i}$ and $\Phi_{1,1}^{n^i}$ in order to realize the shifted versions of the spike function according to

$$\left(\Phi_{1,2}^{n^i} \circ \Phi_{1,1}^{n^i} \right) \left(\{x[-\ell]\}_{\ell=0}^T \right) = \phi \left(\left\{ \frac{x[-\ell]}{\tilde{\delta}_\ell} - n_\ell^i \right\}_{\ell=0}^T \right), \quad \text{for } i = 1, 2, \dots, |\tilde{\mathfrak{N}}|.$$

Then, we apply Lemma C.4 to construct a ReLU network as the parallelization of the compositions $\Phi_{1,2}^{n^i} \circ \Phi_{1,1}^{n^i}$, for $i = 1, 2, \dots, |\tilde{\mathfrak{N}}|$,

$$\Phi_1 := P \left(\left(\Phi_{1,2}^{n^1} \circ \Phi_{1,1}^{n^1} \right), \left(\Phi_{1,2}^{n^2} \circ \Phi_{1,1}^{n^2} \right), \dots, \left(\Phi_{1,2}^{n^{|\tilde{\mathfrak{N}}|}} \circ \Phi_{1,1}^{n^{|\tilde{\mathfrak{N}}|}} \right) \right).\tag{75}$$

Finally, we use Lemma C.3 again to get a ReLU network that composes Φ_2 and Φ_1 according to

$$\Phi = \Phi_2 \circ \Phi_1 = \Phi_2 \circ P \left(\left(\Phi_{1,2}^{n^1} \circ \Phi_{1,1}^{n^1} \right), \left(\Phi_{1,2}^{n^2} \circ \Phi_{1,1}^{n^2} \right), \dots, \left(\Phi_{1,2}^{n^{|\tilde{\mathfrak{N}}|}} \circ \Phi_{1,1}^{n^{|\tilde{\mathfrak{N}}|}} \right) \right),\tag{76}$$

thereby realizing the linear combination

$$\sum_{n \in \tilde{\mathfrak{N}}} \tilde{g}_n \phi \left(\left\{ \frac{x[-\ell]}{\tilde{\delta}_\ell} - n_\ell \right\}_{\ell=0}^T \right) \stackrel{(63)}{=} f \left(\{x[-\ell]\}_{\ell=0}^T \right).$$

To conclude the proof, we verify that Φ , indeed, has $(2, \epsilon)$ -quantized weights, compute $\mathcal{L}(\Phi)$, and derive an upper bound on $\mathcal{M}(\Phi)$. We defer the corresponding details to C.9. \square

5.2 Approximating $(\mathcal{G}(w), \rho_*)$ with RNNs

Having constructed ReLU networks that realize elements of $\mathcal{G}_0(w)$ according to Lemma 5.2, we are now ready to describe the realization of systems in $\mathcal{G}(w)$ through RNNs. This will be done by employing the connection between $\mathcal{G}(w)$ and $\mathcal{G}_0(w)$, as established in Lemma 3.5. Specifically, we construct RNNs that suitably remember past inputs and produce approximations of the desired output.

Theorem 5.3. For every $s \geq 1$, there exists $\epsilon_0 > 0$, such that for $\epsilon \in (0, \epsilon_0)$, with

$$T := \max \left\{ \ell \in \mathbb{N}_0 \mid w[\ell] > \frac{\epsilon}{D} \frac{s}{s+1} \right\},$$

for every $G \in \mathcal{G}(w)$, there is an RNN \mathcal{R}_Ψ associated with a ReLU network $\Psi \in \mathcal{N}_{T+1, T+1}$, satisfying

$$\rho_*(\mathcal{R}_\Psi, G) \leq \epsilon.$$

Moreover, \mathcal{R}_Ψ has $(2, \epsilon)$ -quantized weights and there exists a universal constant $C > 0$ such that

$$\mathcal{M}(\Psi) \leq C(T+1)^2 \prod_{\ell=0}^T \left(\frac{2Dw[\ell]}{\epsilon} \frac{s+1}{s} + 4 \right). \quad (77)$$

Proof. Fix $s \geq 1$ and $G \in \mathcal{G}(w)$ arbitrarily. We proceed in two steps.

Step 1: We construct a ReLU network $\Phi : \mathbb{R}^{T+1} \rightarrow \mathbb{R}$ such that

$$\sup_{x \in \mathcal{S}^+} \sup_{t \in \mathbb{N}_0} |\Phi(\{x[t-\ell]\}_{\ell=0}^T) - G(x)[t]| \leq \epsilon. \quad (78)$$

To this end, we first note that by Lemma 3.5, one can find a $g \in \mathcal{G}_0(w)$ so that

$$g(\mathcal{P}\mathbf{T}_{-t}x) = G(x)[t], \quad \text{for all } x \in \mathcal{S} \text{ and } t \in \mathbb{Z}. \quad (79)$$

Furthermore, by Lemma 5.2, there exists $\epsilon_0 > 0$, such that for every $\epsilon \in (0, \epsilon_0)$, there is a ReLU network Φ satisfying

$$\left| \Phi(\{z[-\ell]\}_{\ell=0}^T) - g(z) \right| \leq \epsilon, \quad \text{for all } z \in \mathcal{S}^-. \quad (80)$$

Next, fix an input $x \in \mathcal{S}^+$ and a time index $t \in \mathbb{N}_0$ and let

$$z' := \mathcal{P}\mathbf{T}_{-t}\{x\}. \quad (81)$$

Note that

$$\begin{aligned} z'[\ell] &= 0, \quad \text{for } \ell \in \mathbb{N}_+, \text{ and hence } z' \in \mathcal{S}^-, \\ z'[-\ell] &= x[t-\ell], \quad \text{for } \ell \in \mathbb{N}_0. \end{aligned} \quad (82)$$

Inserting z' from (81) into (80) and using (82) and (79), it follows that

$$\left| \Phi(\{x[t-\ell]\}_{\ell=0}^T) - G(x)[t] \right| \leq \epsilon.$$

As $x \in \mathcal{S}^+$ and $t \in \mathbb{N}_0$ were arbitrary, this proves (78).

Step 2: We construct an RNN \mathcal{R}_Ψ realizing the mapping $x \rightarrow \Phi(\{x[t-\ell]\}_{\ell=0}^T)_{t \in \mathbb{N}_0}$.

Fix $x \in \mathcal{S}^+$ arbitrarily. Recall the RNN Definition 2.3. The basic idea is to identify an RNN

\mathcal{R}_Ψ which, for every time step $t \in \mathbb{N}_0$,

- delivers the output

$$y[t] = \Phi \left(\{x[t - \ell]\}_{\ell=0}^T \right) \quad (83)$$

- and memorizes the T past inputs $x[t], x[t - 1], \dots, x[t - T + 1]$ in the hidden state vector $h[t]$, i.e.,

$$h_\ell[t] = x[t - \ell + 1], \quad \text{for } \ell \in \{1, 2, \dots, T\}. \quad (84)$$

To memorize the T past inputs, we note that the one-layer neural network

$$\Phi_T(z) = \begin{pmatrix} \mathbb{I}_T & 0_T \end{pmatrix} z, \quad \text{for } z \in \mathbb{R}^{T+1}, \quad (85)$$

satisfies

$$\begin{aligned} \Phi_T \left(\{x[t - \ell]\}_{\ell=0}^T \right) &= \begin{pmatrix} \mathbb{I}_T & 0_T \end{pmatrix} (x[t], x[t - 1], \dots, x[t - T])^T \\ &= (x[t], x[t - 1], \dots, x[t - T + 1])^T \in \mathbb{R}^T. \end{aligned} \quad (86)$$

Now, we apply Lemma C.5 to augment Φ_T to depth $\mathcal{L}(\Phi)$ without changing its input-output relation. This results in the ReLU network Φ_T^* . Then, we apply Lemma C.4 to parallelize Φ and Φ_T^* leading to the desired ReLU network

$$\Psi = P(\Phi, \Phi_T^*). \quad (87)$$

By Definition 2.3, the corresponding RNN \mathcal{R}_Ψ effects the input-output mapping according to

$$\begin{pmatrix} y[t] \\ h[t] \end{pmatrix} = \Psi \left(\begin{pmatrix} x[t] \\ h[t - 1] \end{pmatrix} \right), \quad \text{for all } t \in \mathbb{N}_0.$$

With these choices, (83) and (84) can now be proven by induction over $t \in \mathbb{N}_0$. The base case is immediate as $x[t] = 0$, for $t < 0$, owing to $x \in \mathcal{S}^+$ and, by Definition 2.3, $h[-1] = 0_T$. To establish the induction step, we assume that (83) and (84) hold for $t - 1$ with $t \in \mathbb{N}_+$, i.e.,

$$\begin{aligned} y[t - 1] &= \Phi \left(\{x[t - 1 - \ell]\}_{\ell=0}^T \right), \\ h_\ell[t - 1] &= x[t - \ell], \quad \text{for } \ell \in \{1, 2, \dots, T\}. \end{aligned}$$

Now, for time step t , we note that

$$\begin{aligned}
\Psi \left(\begin{pmatrix} x[t] \\ h[t-1] \end{pmatrix} \right) &= P(\Phi, \Phi_T^*) \left(\begin{pmatrix} x[t] \\ h[t-1] \end{pmatrix} \right) \\
&= \begin{pmatrix} \Phi(\{x[t-\ell]\}_{\ell=0}^T) \\ \Phi_T^*(\{x[t-\ell]\}_{\ell=0}^T) \end{pmatrix} \\
&= \begin{pmatrix} \Phi(\{x[t-\ell]\}_{\ell=0}^T) \\ (x[t], x[t-1], \dots, x[t-T+1])^T \end{pmatrix} \\
&= \begin{pmatrix} \Phi(\{x[t-\ell]\}_{\ell=0}^T) \\ h[t] \end{pmatrix}.
\end{aligned}$$

As x was arbitrary, this completes the induction and thereby Step 2. To conclude, we combine the results in Steps 1 and 2 according to

$$\rho_*(\mathcal{R}_\Psi, G) = \sup_{x \in \mathcal{S}^+} \sup_{t \in \mathbb{N}_0} |y[t] - G(x)[t]| \quad (88)$$

$$= \sup_{x \in \mathcal{S}^+} \sup_{t \in \mathbb{N}_0} |\Phi(\{x[t-\ell]\}_{\ell=0}^T) - G(x)[t]| \quad (89)$$

$$\leq \epsilon, \quad (90)$$

where (89) follows from (83) and (90) is by (78). Furthermore, we have

$$\begin{aligned}
\mathcal{K}(\Psi) &\stackrel{(87), \text{Lemma C.4}}{=} \mathcal{K}(\Phi) \cup (\mathcal{K}(\Phi_T^*)) \\
&\stackrel{\text{Lemma C.5}}{\subset} \mathcal{K}(\Phi) \cup \mathcal{K}(\Phi_T) \cup (-\mathcal{K}(\Phi_T)) \cup \{1, -1\} \\
&\stackrel{(85)}{=} \mathcal{K}(\Phi) \cup \{1, -1\} \\
&\stackrel{\text{Lemma 5.2}}{\subset} 2^{-2 \lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z} \cap [-\epsilon^{-2}, \epsilon^{-2}].
\end{aligned}$$

Thus, Ψ has $(2, \epsilon)$ -quantized weights. Finally, we get an upper bound on $\mathcal{M}(\Psi)$ according to

$$\begin{aligned}
\mathcal{M}(\Psi) &\stackrel{(87), \text{Lemma C.4}}{=} \mathcal{M}(\Phi) + \mathcal{M}(\Phi_T^*) \\
&\stackrel{(85), \text{Lemma C.5}}{\leq} \mathcal{M}(\Phi) + \mathcal{M}(\Phi_T) + T\mathcal{W}(\Phi_T) + 2T(\mathcal{L}(\Phi) - \mathcal{L}(\Phi_T)) \\
&\stackrel{(85), \text{Lemma 5.2}}{\leq} C(T+1)^2 \prod_{\ell=0}^T \left(\frac{2Dw[\ell]}{\epsilon} \frac{s+1}{s} + 4 \right),
\end{aligned}$$

with the universal constant $C > 0$ chosen sufficiently large. \square

6 Metric-entropy-optimal approximation of ELFM and PLFM systems

So far we have characterized the description complexity of ELFM and PLFM systems based on order, type, and generalized dimension (Section 4) and we constructed RNNs approximating general LFM systems (Section 5). We are now ready to state the main results of the paper, namely that the RNNs we constructed are optimal for ELFM and PLFM system approximation in terms of description complexity.

To this end, we first compute the number of bits needed by the canonical RNN decoder in Definition 2.5 to obtain the RNN constructed in Theorem 5.3, specifically its topology and quantized weights. The following result holds for general LFM systems and will later be particularized to ELFM and PLFM systems.

Corollary 6.1. *The class of LFM systems $(\mathcal{G}(w), \rho_*)$ is representable by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$ with*

$$L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w), \rho_*) \leq C_1 M \log(M) \log(\epsilon^{-1}),$$

where

$$M := (T + 1)^2 \prod_{\ell=0}^T \left(\frac{12Dw[\ell]}{\epsilon} \right),$$

$$T := \max \left\{ \ell \in \mathbb{N}_0 \mid w[\ell] > \frac{\epsilon}{2D} \right\},$$

and $C_1 > 0$ is a universal constant.

Proof. Applying Theorem 5.3 and setting $s = 1$, it follows that there exists an $\epsilon_0 > 0$, such that for every $\epsilon \in (0, \epsilon_0)$ and every $G \in \mathcal{G}(w)$, we can find an RNN \mathcal{R}_{Ψ} , associated with a ReLU network $\Psi \in \mathcal{N}_{T+1, T+1}$, satisfying

$$\rho_*(\mathcal{R}_{\Psi}, G) \leq \epsilon. \quad (91)$$

Moreover, \mathcal{R}_{Ψ} has $(2, \epsilon)$ -quantized weights and the number of non-zero weights in Ψ can be upper-bounded according to

$$\mathcal{M}(\Psi) \leq C(T + 1)^2 \prod_{\ell=0}^T \left(\frac{4Dw[\ell]}{\epsilon} + 4 \right). \quad (92)$$

By the definition of the canonical neural network decoder, Remark 2.1, and Definition 2.5, there exists a bitstring $b \in \{0, 1\}^L$ with

$$L \leq 2C_0 \mathcal{M}(\Psi) \log(\mathcal{M}(\Psi)) \log(\epsilon^{-1}), \quad (93)$$

such that

$$\mathcal{D}_{\mathcal{R}}(b) = \mathcal{R}_{\Psi}. \quad (94)$$

Combining (91), (93), (94), and Definition 2.7, upon noting that G was chosen arbitrarily, it follows that $(\mathcal{G}(w), \rho_*)$ is representable by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$, with the minimum number of required bits satisfying

$$L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w), \rho_*) \leq 2C_0 \mathcal{M}(\Psi) \log(\mathcal{M}(\Psi)) \log(\epsilon^{-1}) \quad (95)$$

$$\begin{aligned} &\leq 2C_0 C(T+1)^2 \left[\prod_{\ell=0}^T \left(\frac{12Dw[\ell]}{\epsilon} \right) \right] \log(\epsilon^{-1}) \\ &\quad \cdot \log \left(C(T+1)^2 \left(\prod_{\ell=0}^T \left(\frac{12Dw[\ell]}{\epsilon} \right) \right) \right) \end{aligned} \quad (96)$$

$$\begin{aligned} &\leq 2C_0 CC'(T+1)^2 \left[\prod_{\ell=0}^T \left(\frac{12Dw[\ell]}{\epsilon} \right) \right] \log(\epsilon^{-1}) \\ &\quad \cdot \log \left((T+1)^2 \prod_{\ell=0}^T \left(\frac{12Dw[\ell]}{\epsilon} \right) \right) \end{aligned} \quad (97)$$

$$= C_1 M \log(M) \log(\epsilon^{-1}), \quad (98)$$

where (95) is a consequence of (93) and Definition 2.7, (96) follows from (92) and the fact that $2Dw[\ell]/\epsilon > 1$, for $\ell \in \llbracket T \rrbracket$, (97) holds upon choosing the universal constant C' sufficiently large, namely s.t.

$$\log \left(C(T+1)^2 \left(\prod_{\ell=0}^T \left(\frac{12Dw[\ell]}{\epsilon} \right) \right) \right) \leq C' \log \left((T+1)^2 \left(\prod_{\ell=0}^T \left(\frac{12Dw[\ell]}{\epsilon} \right) \right) \right),$$

and (98) follows by setting $C_1 = 2C_0 CC'$. \square

6.1 RNNs optimally approximate ELFM systems

We now particularize the result in Corollary 6.1 to ELFM systems, which allows us to determine the growth rate, with respect to ϵ , of the minimum number of bits $L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ needed to represent $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$. A comparison with the description complexity of $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ established in Lemma 4.3 then leads to the conclusion that RNNs are capable of approximating ELFM systems in metric-entropy-optimal manner.

Theorem 6.2. *Let $a \in (0, 1]$ and $b > 0$. The class of ELFM systems (Definition 4.1) $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ is optimally representable by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$ (Definition 2.5).*

Proof. By Corollary 6.1, the minimum number of bits $L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ needed to represent $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$ satisfies

$$L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w_{a,b}^{(e)}), \rho_*) \leq C_1 M \log(M) \log(\epsilon^{-1}), \quad (99)$$

where

$$M := (T + 1)^2 \prod_{\ell=0}^T \left(\frac{12Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right),$$

$$T := \max \left\{ \ell \in \mathbb{N}_0 \mid w_{a,b}^{(e)}[\ell] > \frac{\epsilon}{2D} \right\}.$$

By Lemma 4.3, $(\mathcal{G}(w_{a,b}^{(e)}), \rho_*)$ is of order 1 and type 2, with generalized dimension

$$\mathfrak{d} = \frac{1}{2b \log(e)}.$$

Thus, by Lemma 2.10 and Definition 2.11, it suffices to prove that

$$\lim_{\epsilon \rightarrow 0} \frac{\log(C_1 M \log(M) \log(\epsilon^{-1}))}{\log^2(\epsilon^{-1})} = \frac{1}{2b \log(e)}. \quad (100)$$

To this end, we first note that

$$\log(M) = 2 \log(T + 1) + \log \left(\prod_{\ell=0}^T \left(\frac{12Dw_{a,b}^{(e)}[\ell]}{\epsilon} \right) \right) \quad (101)$$

$$= O(\log^{(2)}(\epsilon^{-1})) + \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})) \quad (102)$$

$$= \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})), \quad (103)$$

where (102) follows from Lemma 4.2 and $T + 1 = O(\log(\epsilon^{-1}))$ thanks to (154). Now, we rewrite the numerator in (100) according to

$$\log(C_1 M \log(M) \log(\epsilon^{-1})) = \log(M) + \log^{(2)}(M) + o(\log^2(\epsilon^{-1})) \quad (104)$$

$$= \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1}))$$

$$+ \log \left(\frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})) \right) \quad (105)$$

$$= \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})). \quad (106)$$

Dividing (106) by $\log^2(\epsilon^{-1})$ and taking $\epsilon \rightarrow 0$, concludes the proof. \square

6.2 RNNs optimally approximate PLFM systems

We next particularize the result in Corollary 6.1 to PLFM systems and will see that the minimum number of bits $L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ required to represent $(\mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$ grows significantly faster (with respect to the prescribed error ϵ) than for

ELFM systems. This can be attributed to the fact that the memory of PLFM systems decays much more slowly. Nonetheless, as shown next, RNNs are capable of approximating PLFM systems in metric-entropy-optimal manner.

Theorem 6.3. *Let $q \in (0, 1]$ and $p > 0$. The class of PLFM systems $(\mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ (Definition 4.4) is optimally representable by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$ (Definition 2.5).*

Proof. By Corollary 6.1, the minimum number of bits $L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ needed to represent $(\mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ by the canonical RNN decoder $\mathcal{D}_{\mathcal{R}}$ satisfies

$$L(\epsilon; \mathcal{D}_{\mathcal{R}}, \mathcal{G}(w_{p,q}^{(p)}), \rho_*) \leq C_1 M \log(M) \log(\epsilon^{-1}), \quad (107)$$

where

$$M := (T + 1)^2 \prod_{\ell=0}^T \left(\frac{12Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right),$$

$$T := \max \left\{ \ell \in \mathbb{N}_0 \mid w_{p,q}^{(p)}[\ell] > \frac{\epsilon}{2D} \right\}.$$

By Lemma 4.6, $(\mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ is of order 2 and type 1, with generalized dimension

$$\mathfrak{d} = \frac{1}{p}.$$

Thus, by Lemma 2.10 and Definition 2.11, it suffices to prove that

$$\lim_{\epsilon \rightarrow 0} \frac{\log^{(2)}(C_1 M \log(M) \log(\epsilon^{-1}))}{\log(\epsilon^{-1})} = \frac{1}{p}. \quad (108)$$

To this end, we first note that

$$\log(M) = 2 \log(T + 1) + \log \left(\prod_{\ell=0}^T \left(\frac{12Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right) \right) \quad (109)$$

$$= O(\log(\epsilon^{-1})) + \Theta(\epsilon^{-1/p}) \quad (110)$$

$$= \Theta(\epsilon^{-1/p}), \quad (111)$$

where (110) follows from Lemma 4.5 and (160). Now, we rewrite the numerator in (108) according to

$$\log^{(2)}(C_1 M \log(M) \log(\epsilon^{-1})) = \log \left(\log(M) + \log^{(2)}(M) + o(\log(\epsilon^{-1})) \right) \quad (112)$$

$$= \log \left(\Theta(\epsilon^{-1/p}) + \log(\Theta(\epsilon^{-1/p})) + o(\log(\epsilon^{-1})) \right) \quad (113)$$

$$= \log(\Theta(\epsilon^{-1/p})), \quad (114)$$

where (114) follows from

$$\Theta(\epsilon^{-1/p}) + \log(\Theta(\epsilon^{-1/p})) + o(\log(\epsilon^{-1})) = \Theta(\epsilon^{-1/p}).$$

Finally, dividing (114) by $\log(\epsilon^{-1})$ and taking $\epsilon \rightarrow 0$, concludes the proof. \square

7 Conclusion

Returning to Table 1, we note that it can be complemented by our results for ELFM and PLFM systems as summarized in Table 2 below. As both classes of systems in Table 2 are optimally representable by the canonical RNN decoder (Definition 2.5), we can conclude that, remarkably, the metric-entropy-optimal universality of ReLU networks extends from function classes to nonlinear dynamical systems, in the latter case simply by embedding the ReLU network in a recurrence.

Metric	\mathcal{C}	κ	λ	\mathfrak{d}
$\{\mathbb{R}^{\mathbb{Z}} \rightarrow \mathbb{R}^{\mathbb{Z}}\} \quad \rho_*$ (Def. 2.17)	ELFM systems (Section 4.1) $\mathcal{G}(w_{a,b}^{(e)})$	1	2	$\frac{1}{2b \log(e)}$
$\{\mathbb{R}^{\mathbb{Z}} \rightarrow \mathbb{R}^{\mathbb{Z}}\} \quad \rho_*$ (Def. 2.17)	PLFM systems (Section 4.2) $\mathcal{G}(w_{p,q}^{(p)})$	2	1	$\frac{1}{p}$

Table 2: Scaling behavior of the covering numbers (Definition 2.9) for classes of nonlinear dynamical systems.

Finally, we remark that many of the results in this paper apply to general weight sequences $w[\cdot]$, specifically the bounds in Section 3 on the exterior covering number of $(\mathcal{G}(w), \rho_*)$ as well as the RNN constructions in Section 5.

A Representing ReLU networks by bitstrings

Definition A.1. Let Φ be a ReLU network with (m, ϵ) -quantized weights. Denote the number of non-zero weights by $M := \mathcal{M}(\Phi)$ and the depth by $L := \mathcal{L}(\Phi)$. We organize the bitstring representation of Φ in 6 segments as follows.

[Segment 1] The bitstring starts with M 1's followed by a single 0.

[Segment 2] L is specified in binary representation. As $L \leq M$, it suffices to allocate $\lceil \log(M) \rceil$ bits.

[Segment 3] $N_0, \dots, N_L \leq M$ are specified in binary representation using a total of $(M + 1)\lceil \log(M) \rceil$ bits.

[Segment 4] The topology of the network, i.e., the locations of the non-zero entries in the A_ℓ and b_ℓ , $\ell \in \{1, \dots, L\}$, is encoded as follows. We denote the bitstring corresponding to the binary representation of an integer $i \in \{1, \dots, M\}$ by $b(i) \in \{0, 1\}^{\lceil \log(M) \rceil}$. For $\ell \in \{1, \dots, L\}$, $i \in \{1, \dots, N_\ell\}$, $j \in \{1, \dots, N_{\ell-1}\}$, a non-zero entry $(A_\ell)_{ij}$ is indicated by $[b(\ell), b(i), b(j)]$ and a non-zero entry $(b_\ell)_i$ by $[b(\ell), b(i), b(i)]$. Thus, encoding the topology of the network requires a total of $3\lceil \log(M) \rceil M$ bits.

[Segment 5] The quantity $m\lceil \log(\epsilon^{-1}) \rceil$ is represented by a bitstring of that many 1's followed by a single 0.

[Segment 6] The value of each non-zero weight and bias is represented by a bitstring of length $B_\epsilon = 2m\lceil \log(\epsilon^{-1}) \rceil + 1$.

The overall length of the bitstring is now given by

$$\underbrace{M + 1}_{\text{Segment 1}} + \underbrace{\lceil \log(M) \rceil}_{\text{Segment 2}} + \underbrace{(M + 1)\lceil \log(M) \rceil}_{\text{Segment 3}} + \underbrace{3\lceil \log(M) \rceil M}_{\text{Segment 4}} + \underbrace{m\lceil \log(\epsilon^{-1}) \rceil + 1}_{\text{Segment 5}} + \underbrace{MB_\epsilon}_{\text{Segment 6}}. \quad (115)$$

The ReLU network Φ can be recovered by sequentially reading out M, L , the N_ℓ , the topology, the quantity $m\lceil \log(\epsilon^{-1}) \rceil$, and the quantized weights from the overall bitstring. It is not difficult to verify that the bitstring is crafted such that this yields unique decodability.

B Comparison with [7]

As mentioned in Section 2.2, compared to [7], we use refined notions of massiveness of sets, as the system classes we consider are significantly more massive than the function classes dealt with in [7]. We next detail how the results from [7] fit into our framework.

In [7] the scaling behavior of covering numbers is quantified in terms of the optimal exponent γ^* [7, Definition IV.1]. The following result relates γ^* to the generalized dimension (Definition 2.9) employed here.

Lemma B.1. *Let $d \in \mathbb{N}_+$, $\Omega \subseteq \mathbb{R}^d$, and let $\mathcal{C} \subset L^2(\Omega)$ be compact and such that the optimal exponent $\gamma^*(\mathcal{C})$ according to [7, Definition IV.1] is finite and non-zero. Then, \mathcal{C} is, with respect to the metric $\rho(f, g) := \|f - g\|_{L^2(\Omega)}$, of order $\kappa = 1$, type $\lambda = 1$, and generalized dimension*

$$\mathfrak{d} = \frac{1}{\gamma^*}.$$

Proof. For order $\kappa = 1$ and type $\lambda = 1$, the generalized dimension is given by

$$\mathfrak{d} = \limsup_{\epsilon \rightarrow 0} \frac{\log^{(2)} N^{\text{ext}}(\epsilon; \mathcal{C}, \rho)}{\log(\epsilon^{-1})}. \quad (116)$$

We note that by [35, Remark 5.10, Equation (5.5)] and Lemma 3.2, it holds that

$$\gamma^* = \sup \{ \gamma > 0 : \log N^{\text{ext}}(\epsilon; \mathcal{C}, \rho) \in O(\epsilon^{-1/\gamma}), \epsilon \rightarrow 0 \}. \quad (117)$$

We first establish that $\mathfrak{d} \leq \frac{1}{\gamma^*}$. To this end, fix $\Delta > 0$ arbitrarily and observe that $\log N^{\text{ext}}(\epsilon; \mathcal{C}, \rho) \in O(\epsilon^{-1/(\gamma^* - \Delta)})$. Hence, there exist $\epsilon_0, C > 0$ such that

$$\log N^{\text{ext}}(\epsilon; \mathcal{C}, \rho) \leq C\epsilon^{-1/(\gamma^* - \Delta)}, \quad \forall \epsilon \in (0, \epsilon_0),$$

and thus

$$\mathfrak{d} = \limsup_{\epsilon \rightarrow 0} \frac{\log^{(2)} N^{\text{ext}}(\epsilon; \mathcal{C}, \rho)}{\log(\epsilon^{-1})} \leq \limsup_{\epsilon \rightarrow 0} \left(\frac{1}{\gamma^* - \Delta} + \frac{\log(C)}{\log(\epsilon^{-1})} \right) = \frac{1}{\gamma^* - \Delta}. \quad (118)$$

As $\Delta > 0$ was arbitrary, we have established that $\mathfrak{d} \leq \frac{1}{\gamma^*}$.

Next, we show that $\mathfrak{d} \geq \frac{1}{\gamma^*}$. Again, fix $\Delta > 0$ arbitrarily. By (116), there exists an $\epsilon_0 > 0$ such that for all $\epsilon \in (0, \epsilon_0)$,

$$\frac{\log^{(2)} N^{\text{ext}}(\epsilon; \mathcal{C}, \rho)}{\log(\epsilon^{-1})} \leq \mathfrak{d} + \Delta. \quad (119)$$

This implies

$$\log N^{\text{ext}}(\epsilon; \mathcal{C}, \rho) \leq \epsilon^{-(\mathfrak{d} + \Delta)}, \quad \forall \epsilon \in (0, \epsilon_0), \quad (120)$$

and thus $\log N^{\text{ext}}(\epsilon; \mathcal{C}, \rho) \in O(\epsilon^{-1/(\mathfrak{d} + \Delta)^{-1}})$. Hence, $\gamma^* \geq (\mathfrak{d} + \Delta)^{-1}$. As Δ was arbitrary, this finalizes the proof. \square

Table 1 in the present paper now follows from [7, Table I] by application of Lemma B.1. Furthermore, [7] shows through the transference principle [7, Section VII] that a wide range of function classes, including those in Table 1, are *optimally representable by neural networks* [7, Definition VI.5]. The following Lemma hence allows us to conclude that every row in Table 1 is optimally representable (according to Definition 2.11) by the canonical neural network decoder.

Lemma B.2. *Let $d \in \mathbb{N}_+$, $\Omega \subseteq \mathbb{R}^d$, and let $\mathcal{C} \subset L^2(\Omega)$ be compact. If the function class $\mathcal{C} \subset L^2(\Omega)$ is optimally representable by neural networks according to [7, Definition VI.5],*

then (\mathcal{C}, ρ) is optimally representable by the canonical neural network decoder with respect to the metric $\rho(f, g) := \|f - g\|_{L^2(\Omega)}$.

Proof. First, we note that by Lemma B.1, \mathcal{C} is of order $\kappa = 1$, type $\lambda = 1$, and has generalized dimension $\mathfrak{d} = \frac{1}{\gamma^*(\mathcal{C})}$. By [7, Definition VI.5], we have $\gamma^*(\mathcal{C}) = \gamma_{\mathcal{N}}^{*,\text{eff}}(\mathcal{C})$, with $\gamma_{\mathcal{N}}^{*,\text{eff}}(\mathcal{C})$ as per [7, Definition VI.3]. Next, fix $\Delta > 0$ arbitrarily. Now, following the proof of [7, Theorem VI.4, p. 2602] with $\gamma_{\mathcal{N}}^{*,\text{eff}} - \Delta$ in place of γ , we can conclude the existence of a polynomial π^* , a constant C , and a mapping $\Psi : (0, \frac{1}{2}) \times \mathcal{C} \rightarrow \mathcal{N}_{d,1}$ with the following properties. For every $f \in \mathcal{C}$ and every $\epsilon \in (0, \frac{1}{2})$, the network $\tilde{\Phi}_{\epsilon,f} = \Psi(\epsilon, f)$ has $(\lceil \pi^*(\log(\epsilon^{-1})) \rceil, \epsilon)$ -quantized weights and satisfies

$$\|f - \tilde{\Phi}_{\epsilon,f}\|_{L^2(\Omega)} \leq \epsilon \quad \text{and} \quad \mathcal{M}(\tilde{\Phi}_{\epsilon,f}) \leq C\epsilon^{-1/(\gamma_{\mathcal{N}}^{*,\text{eff}} - \Delta)} =: M_\epsilon.$$

Thus, $\tilde{\Phi}_{\epsilon,f}$ can, according to Remark 2.1, be reconstructed uniquely by the canonical neural network decoder $\mathcal{D}_{\mathcal{N}}$ from a bitstring of length no more than

$$C_0 \lceil \pi^*(\log(\epsilon^{-1})) \rceil \log(\epsilon^{-1}) M_\epsilon \log(M_\epsilon).$$

Therefore, \mathcal{C} is representable by the canonical neural network decoder $\mathcal{D}_{\mathcal{N}}$ with $L(\epsilon; \mathcal{D}_{\mathcal{N}}, \mathcal{C}, \rho) \leq C_0 M_\epsilon \log(M_\epsilon) \log^{q_0}(\epsilon^{-1})$, where q_0 is a constant depending on π^* only, and hence

$$\log L(\epsilon; \mathcal{D}_{\mathcal{N}}, \mathcal{C}, \rho) \leq \frac{1}{\gamma_{\mathcal{N}}^{*,\text{eff}} - \Delta} \log(\epsilon^{-1}) + o(\log(\epsilon^{-1})). \quad (121)$$

As $\Delta > 0$ was arbitrary, we thus get

$$\limsup_{\epsilon \rightarrow 0} \frac{L(\epsilon; \mathcal{D}_{\mathcal{N}}, \mathcal{C}, \rho)}{\log(\epsilon^{-1})} \leq \frac{1}{\gamma_{\mathcal{N}}^{*,\text{eff}}}, \quad (122)$$

which, together with Lemma 2.10 and the fact that $\mathfrak{d} = \frac{1}{\gamma^*} = \frac{1}{\gamma_{\mathcal{N}}^{*,\text{eff}}}$ implies that \mathcal{C} is optimally representable by the canonical neural network decoder according to Definition 2.11. \square

C Proofs

C.1 Proof of Lemma 2.10

To simplify notation, we let $L(\epsilon) := L(\epsilon; \mathcal{D}, \mathcal{C}, \rho)$. We first establish that

$$L(\epsilon) \geq \log(N^{\text{ext}}(\epsilon; \mathcal{C}, \rho)) - 1, \quad \forall \epsilon > 0. \quad (123)$$

By way of contradiction, assume that

$$2^{L(\epsilon)+1} < N^{\text{ext}}(\epsilon; \mathcal{C}, \rho), \quad \text{for some } \epsilon > 0.$$

It then follows from Definition 2.7 that for this ϵ , for every $f \in \mathcal{C}$, there is an integer $\ell \leq L(\epsilon)$ and a bitstring $\mathbf{b}_f \in \{0, 1\}^\ell$, such that $\rho(\mathcal{D}(\mathbf{b}_f), f) \leq \epsilon$. This directly implies that the set

$$\mathcal{U} := \left\{ \mathcal{D}(\mathbf{b}) \mid \mathbf{b} \in \bigcup_{\ell=0}^{L(\epsilon)} \{0, 1\}^\ell \right\},$$

is an ϵ -net for \mathcal{C} . Furthermore,

$$|\mathcal{U}| \leq \sum_{\ell=0}^{L(\epsilon)} 2^\ell = \frac{2^{L(\epsilon)+1} - 1}{2 - 1} \leq 2^{L(\epsilon)+1} < N^{\text{ext}}(\epsilon; \mathcal{C}, \rho).$$

Hence, \mathcal{U} is an ϵ -net of cardinality strictly smaller than $N^{\text{ext}}(\epsilon; \mathcal{C}, \rho)$, which stands in contradiction to Definition 2.8 and so (123) must hold. This, in turn, implies that

$$\begin{aligned} \limsup_{\epsilon \rightarrow 0} \frac{\log^{(\kappa)} L(\epsilon)}{\log^\lambda(\epsilon^{-1})} &\geq \limsup_{\epsilon \rightarrow 0} \frac{\log^{(\kappa)} (\log(N^{\text{ext}}(\epsilon; \mathcal{C}, \rho_*) - 1))}{\log^\lambda(\epsilon^{-1})} \\ &= \limsup_{\epsilon \rightarrow 0} \frac{\log^{(\kappa+1)} N^{\text{ext}}(\epsilon; \mathcal{C}, \rho_*)}{\log^\lambda(\epsilon^{-1})} = \mathfrak{d}, \end{aligned}$$

as desired. □

C.2 Proof of Lemma 2.18

Fix $G, G' \in \mathcal{G}(w)$. First, note that

$$\begin{aligned} \rho_*(G, G') &= \sup_{x \in \mathcal{S}^+} \sup_{t \in \mathbb{N}_0} |(Gx)[t] - (G'x)[t]| \\ &\leq \sup_{x \in \mathcal{S}} \sup_{t \in \mathbb{Z}} |(Gx)[t] - (G'x)[t]|. \end{aligned} \tag{124}$$

Next, arbitrarily fix $\Delta > 0$. By definition of the supremum, it follows that there exist $x_0 \in \mathcal{S}$ and $e \in \mathbb{Z}$, such that

$$\sup_{x \in \mathcal{S}} \sup_{t \in \mathbb{Z}} |(Gx)[t] - (G'x)[t]| - \Delta/2 \leq |(Gx_0)[e] - (G'x_0)[e]|. \tag{125}$$

Since $\lim_{t \rightarrow \infty} w[t] = 0$, there exists $T > 0$ so that $w[t] \leq \Delta/(4D)$, for all $t > T$. Next, define

$$x_1[t] = x_0[t - (T - e)] \in \mathcal{S} \quad \text{and} \quad y[t] = x_1[t] \cdot \mathbb{1}_{\{t \geq 0\}} \in \mathcal{S}^+.$$

We then get

$$\sup_{\tau \geq 0} w[\tau] |(x_1[T - \tau] - y[T - \tau])| = \sup_{\tau > T} w[\tau] |x_1[T - \tau]| \leq \sup_{\tau > T} w[\tau] D \leq \Delta/4, \tag{126}$$

where we used that $|x_1[\cdot]| \leq D$ by Definition 2.12. We furthermore obtain

$$\begin{aligned} |(Gx_1)[T] - (G'x_1)[T]| &\leq |(Gx_1)[T] - (Gy)[T]| + |(G'x_1)[T] - (G'y)[T]| \\ &\quad + |(Gy)[T] - (G'y)[T]| \end{aligned} \tag{127}$$

$$\leq \Delta/4 + \Delta/4 + |(Gy)[T] - (G'y)[T]| \tag{128}$$

$$= \Delta/2 + |(Gy)[T] - (G'y)[T]|, \tag{129}$$

where in (127) we used the triangle inequality and in (128) we invoked the Lipschitz fading-memory property of G and G' in combination with (126). Next, note that by time-invariance of G and G' , it holds that

$$\begin{aligned} (Gx_1)[T] &= (G\mathbf{T}_{(T-e)}x_0)[T] = (\mathbf{T}_{(T-e)}Gx_0)[T] = (Gx_0)[e], \\ (G'x_1)[T] &= (G'\mathbf{T}_{(T-e)}x_0)[T] = (\mathbf{T}_{(T-e)}G'x_0)[T] = (G'x_0)[e]. \end{aligned} \tag{130}$$

Combining all these results yields

$$\begin{aligned} \sup_{x \in \mathcal{S}} \sup_{t \in \mathbb{Z}} |(Gx)[t] - (G'x)[t]| - \Delta/2 &\stackrel{(125)}{\leq} |(Gx_0)[e] - (G'x_0)[e]| \\ &\stackrel{(130)}{=} |(Gx_1)[T] - (G'x_1)[T]| \\ &\stackrel{(129)}{\leq} \Delta/2 + |(Gy)[T] - (G'y)[T]| \\ &\stackrel{y \in \mathcal{S}^+}{\leq} \Delta/2 + \sup_{x \in \mathcal{S}^+} \sup_{t \in \mathbb{N}_0} |(Gx)[t] - (G'x)[t]| \\ &\stackrel{\text{Def. 2.17}}{=} \Delta/2 + \rho_*(G, G'). \end{aligned} \tag{131}$$

As $\Delta > 0$ was arbitrary, we have thus established that

$$\sup_{x \in \mathcal{S}} \sup_{t \in \mathbb{Z}} |(Gx)[t] - (G'x)[t]| \leq \rho_*(G, G'),$$

which, together with (124), completes the proof. \square

C.3 Proof of Lemma 3.5

Recall the projection operator $\mathcal{P} : \mathcal{S} \rightarrow \mathcal{S}^-$ defined according to

$$(\mathcal{P}x)[t] = x[t] \cdot \mathbb{1}_{\{t \leq 0\}}.$$

First, we need to show that $\mathcal{I} : \mathcal{G}_0(w) \rightarrow \mathcal{G}(w)$ given by

$$g \rightarrow \mathcal{I}(g), \quad \text{with } ((\mathcal{I}(g))(x))[t] = g(\mathcal{P}\mathbf{T}_{-t}x),$$

is a well-defined mapping from $\mathcal{G}_0(w)$ to $\mathcal{G}(w)$, i.e., we need to verify that for every $g \in \mathcal{G}_0(w)$, indeed $\mathcal{I}(g) \in \mathcal{G}(w)$. This will be done by establishing that $\mathcal{I}(g)$ satisfies the conditions of Definition 2.16. We first verify that $\mathcal{I}(g)$ is causal. Note that for every $T \in \mathbb{Z}$, for every $x, x' \in \mathcal{S}$ with $x[t] = x'[t], \forall t \leq T$, it holds that $\mathcal{P}\mathbf{T}_{-T}x = \mathcal{P}\mathbf{T}_{-T}x'$, and hence

$$((\mathcal{I}(g))(x))[T] = g(\mathcal{P}\mathbf{T}_{-T}x) = g(\mathcal{P}\mathbf{T}_{-T}x') = ((\mathcal{I}(g))(x'))[T].$$

Thus, $\mathcal{I}(g)$ is, indeed, causal according to Definition 2.13. Next, we verify time-invariance as follows:

$$\begin{aligned} (\mathbf{T}_\tau(\mathcal{I}(g)(x)))[t] &= ((\mathcal{I}(g))(x))[t - \tau] \\ &= g(\mathcal{P}\mathbf{T}_{\tau-t}x) \\ &= g(\mathcal{P}\mathbf{T}_{-t}\mathbf{T}_\tau x) \\ &= ((\mathcal{I}(g))(\mathbf{T}_\tau x))[t]. \end{aligned}$$

The Lipschitz fading-memory property of $\mathcal{I}(g)$ according to Definition 2.15 is established by noting that

$$|((\mathcal{I}(g))(x))[t] - ((\mathcal{I}(g))(x'))[t]| \tag{132}$$

$$\begin{aligned} &= |g(\mathcal{P}\mathbf{T}_{-t}x) - g(\mathcal{P}\mathbf{T}_{-t}x')| \\ &\leq \sup_{\tau \geq 0} w[\tau] |((\mathcal{P}\mathbf{T}_{-t}x)[- \tau] - (\mathcal{P}\mathbf{T}_{-t}x')[- \tau])| \\ &= \sup_{\tau \geq 0} w[\tau] |x[t - \tau] - x'[t - \tau]|, \quad \forall t \in \mathbb{Z}, \forall x, x' \in \mathcal{S}, \end{aligned} \tag{133}$$

where in (133) we used (11). Finally, we observe that, by (11),

$$((\mathcal{I}(g))(0))[t] = g(\mathcal{P}\mathbf{T}_{-t}0) = g(0) = 0, \quad \forall t \in \mathbb{Z}.$$

In summary, we have thus shown that $\mathcal{I}(g) \in \mathcal{G}(w)$ and hence \mathcal{I} is, indeed, well-defined. Furthermore, we have

$$\rho_*(\mathcal{I}(g), \mathcal{I}(g')) = \sup_{x \in \mathcal{S}} \sup_{t \in \mathbb{Z}} |((\mathcal{I}(g))(x))[t] - ((\mathcal{I}(g'))(x))[t]| \tag{134}$$

$$\begin{aligned} &= \sup_{x \in \mathcal{S}} \sup_{t \in \mathbb{Z}} |g(\mathcal{P}\mathbf{T}_{-t}x) - g'(\mathcal{P}\mathbf{T}_{-t}x)| \\ &= \sup_{y \in \mathcal{S}} |g(\mathcal{P}y) - g'(\mathcal{P}y)| \\ &= \sup_{z \in \mathcal{S}^-} |g(z) - g'(z)| \\ &= \rho_0(g, g'), \quad \forall g \in \mathcal{G}_0(w), \forall g' \in \mathcal{G}_0(w), \end{aligned} \tag{135}$$

where in (134) we invoked Lemma 2.18 and in (135) we used that \mathcal{S} is closed under time shifts. This establishes that \mathcal{I} is isometric and consequently injective.

Next, we prove that \mathcal{I} is surjective. To this end, we fix $G \in \mathcal{G}(w)$ arbitrarily, consider

$$g(s) := (Gs)[0], \quad \forall s \in \mathcal{S}_-, \quad (136)$$

and show that $g \in \mathcal{G}_0(w)$ as well as $\mathcal{I}(g) = G$. First, we establish that $g \in \mathcal{G}_0(w)$. The Lipschitz property of g can be verified according to

$$\begin{aligned} |g(x) - g(x')| &= |(Gx)[0] - (Gx')[0]| \\ &\leq \sup_{\tau \geq 0} w[\tau] |x[-\tau] - x'[-\tau]| \end{aligned} \quad (137)$$

$$= \|x - x'\|_w, \quad \forall x, x' \in \mathcal{S}^-, \quad (138)$$

where in (137) we used the fact that $G \in \mathcal{G}(w)$ has Lipschitz fading memory (Definition 2.15) and (138) is by (12). Furthermore, we have $g(0) = (G0)[0] = 0$ and hence $g \in \mathcal{G}_0(w)$.

It remains to show that $\mathcal{I}(g) = G$. To this end, we fix $x \in \mathcal{S}$ and $t \in \mathbb{Z}$, both arbitrarily, and prove that

$$((\mathcal{I}(g))(x))[t] = (Gx)[t].$$

First, note that $(\mathcal{P}\mathbf{T}_{-t}x)[t'] = (\mathbf{T}_{-t}x)[t']$, for all $t' \leq 0$. By causality of G , we conclude that $(G\mathcal{P}\mathbf{T}_{-t}x)[0] = (G\mathbf{T}_{-t}x)[0]$, which yields

$$\begin{aligned} ((\mathcal{I}(g))(x))[t] &= g(\mathcal{P}\mathbf{T}_{-t}x) \\ &= (G\mathcal{P}\mathbf{T}_{-t}x)[0] \end{aligned} \quad (139)$$

$$\begin{aligned} &= (G\mathbf{T}_{-t}x)[0] = (\mathbf{T}_{-t}Gx)[0] \\ &= (Gx)[t], \end{aligned} \quad (140)$$

where in (139) we used (136), and in (140) we invoked the causality and the time-invariance of G . As $x \in \mathcal{S}$ and $t \in \mathbb{Z}$ were arbitrary, this proves that $\mathcal{I}(g) = G$. Furthermore, since $G \in \mathcal{G}(w)$ was arbitrary, we have established that \mathcal{I} is surjective. Thus, \mathcal{I} is, indeed, an isometric isomorphism between $(\mathcal{G}(w), \rho_*)$ and $(\mathcal{G}_0(w), \rho_0)$. Application of Lemma 3.3 then yields $N(\epsilon; \mathcal{G}_0(w), \rho_0) = N(\epsilon; \mathcal{G}(w), \rho_*)$ and $M(\epsilon; \mathcal{G}_0(w), \rho_0) = M(\epsilon; \mathcal{G}(w), \rho_*)$. \square

C.4 Proof of Lemma 3.6

The proof relies on the following auxiliary result.

Lemma C.1. *Let $w[\cdot]$ be a weight sequence. The packing number of \mathcal{S}^- w.r.t. $\|\cdot\|_w$ satisfies*

$$M(\epsilon; \mathcal{S}^-, \|\cdot\|_w) \geq \prod_{\ell=0}^T \left\lceil \frac{2Dw[\ell]}{\epsilon} \right\rceil,$$

with $T := \max\{T' \in \mathbb{N}_0 \mid w[T'] > \frac{\epsilon}{2D}\}$.

Proof. For $t \in \llbracket T \rrbracket$, we let $N_t := \left\lceil \frac{2Dw[t]}{\epsilon} \right\rceil \in \left[\frac{2Dw[t]}{\epsilon}, \frac{2Dw[t]}{\epsilon} + 1 \right)$ and $\delta_t = \frac{2D}{N_t - 1} > \frac{\epsilon}{w[t]}$. Next, we

define the set

$$\mathcal{U} = \{x_{i_0, \dots, i_T} \mid i_t \in \llbracket N_t - 1 \rrbracket, \text{ for } t \in \llbracket T \rrbracket\}, \text{ where} \quad (141)$$

$$x_{i_0, \dots, i_T} := \begin{cases} -D + i_t \delta_t, & -T \leq t \leq 0 \\ 0, & \text{else} \end{cases}, \quad (142)$$

and show that \mathcal{U} constitutes an ϵ -packing for $(\mathcal{S}^-, \|\cdot\|_w)$. First, we establish that $\mathcal{U} \subset \mathcal{S}^-$ by verifying that, for all $x_{i_0, \dots, i_T} \in \mathcal{U}$, $x_{i_0, \dots, i_T}[t] \in [-D, D]$, for all $t \in \mathbb{Z}$, and $x_{i_0, \dots, i_T}[t] = 0$, for all $t \in \mathbb{N}_+$. Indeed, for $t \in \{-T, \dots, 0\}$, we have

$$-D \leq x_{i_0, \dots, i_T}[t] = -D + i_t \delta_t \leq -D + (N_t - 1)\delta_t = -D + 2D = D,$$

and for $t \notin \{-T, \dots, 0\}$,

$$x_{i_0, \dots, i_T}[t] = 0.$$

Next, we show that for distinct $x_{i_0, \dots, i_T}, x_{j_0, \dots, j_T} \in \mathcal{U}$ (i.e., there is at least one $\ell \in \llbracket T \rrbracket$ such that $i_\ell \neq j_\ell$), it holds that $\|x_{i_0, \dots, i_T} - x_{j_0, \dots, j_T}\|_w > \epsilon$. Indeed, for any such ℓ , we have

$$\begin{aligned} \|x_{i_0, \dots, i_T} - x_{j_0, \dots, j_T}\|_w &= \sup_{t \in \mathbb{N}_0} w[t] |(x_{i_0, \dots, i_T}[-t] - x_{j_0, \dots, j_T}[-t])| \\ &\geq w[\ell] |(x_{i_0, \dots, i_T}[-\ell] - x_{j_0, \dots, j_T}[-\ell])| \\ &= w[\ell] |(-D + i_\ell \delta_\ell + D - j_\ell \delta_\ell)| \\ &= (i_\ell - j_\ell) \delta_\ell w[\ell] > \epsilon, \end{aligned} \quad (143)$$

where (143) follows from $\delta_\ell > \frac{\epsilon}{w[\ell]}$. This establishes that \mathcal{U} , indeed, constitutes an ϵ -packing for $(\mathcal{S}^-, \|\cdot\|_w)$, and we therefore have

$$M(\epsilon; \mathcal{S}^-, \|\cdot\|_w) \geq |\mathcal{U}| = \prod_{\ell=0}^T N_\ell = \prod_{\ell=0}^T \left\lceil \frac{2Dw[\ell]}{\epsilon} \right\rceil.$$

□

Proof of Lemma 3.6. Let $M := \prod_{\ell=0}^T \left\lceil \frac{2Dw[\ell]}{\epsilon} \right\rceil$ and take $\{x_1, \dots, x_M\}$ to be an ϵ -packing for $(\mathcal{S}^-, \|\cdot\|_w)$ according to Lemma C.1. Hence, with $\delta := \min_{\ell \neq j} \|x_\ell - x_j\|_w > \epsilon$, there exists an $\epsilon' \in (\epsilon, \delta)$. Next, define balls of radius $\epsilon'/2$ centered at the packing points $\{x_1, \dots, x_M\}$ according to $\mathcal{S}_\ell = \{x[\cdot] \in \mathcal{S}^- \mid \|x - x_\ell\|_w \leq \epsilon'/2\}$, for $\ell \in \{1, \dots, M\}$. We now show that these balls are non-overlapping. Indeed, assuming that, for $j \neq \ell$, there exists an $x \in \mathcal{S}^-$ with $\|x - x_j\|_w < \epsilon'/2$ and $\|x - x_\ell\|_w < \epsilon'/2$, leads to the contradiction

$$\epsilon' < \delta \leq \|x_\ell - x_j\|_w = \|x_\ell - x + x - x_j\|_w \leq \|x_\ell - x\|_w + \|x - x_j\|_w \leq \epsilon'/2 + \epsilon'/2.$$

As the balls \mathcal{S}_ℓ , $\ell \in \{1, \dots, M\}$, are non-overlapping, the signal $x[\cdot] = 0$ is contained in at most one ball \mathcal{S}_ℓ which we take to be \mathcal{S}_M w.l.o.g. Now, we define the set \mathcal{U} , with elements indexed

by bitstring subscripts, according to

$$\mathcal{U} = \{g_{\alpha_1, \dots, \alpha_{M-1}}(\cdot) \mid \alpha_\ell \in \{0, 1\}, \text{ for } \ell \in \{1, \dots, M-1\}\}, \text{ where}$$

$$g_{\alpha_1, \dots, \alpha_{M-1}}(x) := \begin{cases} (2\alpha_\ell - 1)(\epsilon'/2 - \|x - x_\ell\|_w), & \text{for } x \in \mathcal{S}_\ell, \ell \in \{1, \dots, M-1\} \\ 0, & \text{else,} \end{cases}$$

and show that \mathcal{U} constitutes an ϵ -packing for $(\mathcal{G}_0(w), \rho_0)$. First, we establish that $\mathcal{U} \subset \mathcal{G}_0(w)$ by verifying that every $g_{\alpha_1, \dots, \alpha_{M-1}}(\cdot) \in \mathcal{U}$ satisfies the conditions in (11). Indeed, $g_{\alpha_1, \dots, \alpha_{M-1}}(0) = 0$ because the zero signal is either in no ball or in \mathcal{S}_M . Next, we show that $|g_{\alpha_1, \dots, \alpha_{M-1}}(x) - g_{\alpha_1, \dots, \alpha_{M-1}}(x')| \leq \|x - x'\|_w, \forall x, x' \in \mathcal{S}^-$. This will be done by distinguishing cases. First, assume that x and x' are contained in the same ball \mathcal{S}_ℓ , for some $\ell \in \{1, \dots, M-1\}$. Then, we have

$$\begin{aligned} |g_{\alpha_1, \dots, \alpha_{M-1}}(x) - g_{\alpha_1, \dots, \alpha_{M-1}}(x')| &= |2\alpha_\ell - 1| \cdot \left| \|x' - x_\ell\|_w - \|x - x_\ell\|_w \right| \\ &\leq \|(x' - x_\ell) - (x - x_\ell)\|_w = \|x' - x\|_w, \end{aligned} \quad (144)$$

where we used the reverse triangle inequality. Next, assume that $x \in \mathcal{S}_\ell$ and $x' \in \mathcal{S}_j$, with $\ell, j \in \{1, \dots, M-1\}$ and $\ell \neq j$. We let $z(\mu) := x + \mu(x' - x)$, $\mu \in [0, 1]$. Since $z(0) = x \in \mathcal{S}_\ell$ and $z(1) = x' \notin \mathcal{S}_\ell$ (because the balls are non-overlapping), there must be a $\mu_1 \in (0, 1)$ such that $\|z(\mu_1) - x_\ell\|_w = \epsilon'/2$. As $z(\mu_1) \notin \mathcal{S}_j$ and $z(1) = x' \in \mathcal{S}_j$, there must similarly be a $\mu_2 \in (\mu_1, 1)$ so that $\|z(\mu_2) - x_j\|_w = \epsilon'/2$. For notational simplicity, we let $z_1 := z(\mu_1)$, $z_2 := z(\mu_2)$, and $g := g_{\alpha_1, \dots, \alpha_{M-1}}$. Next, we bound

$$\begin{aligned} |g(x) - g(x')| &\leq |g(x) - g(z_1)| + |g(z_1) - g(z_2)| + |g(z_2) - g(x')| \\ &\leq |g(x) - g(z_1)| + 0 + |g(z_2) - g(x')| \end{aligned} \quad (145)$$

$$\leq \|x - z_1\|_w + \|z_2 - x'\|_w \quad (146)$$

$$= \|\mu_1(x - x')\|_w + \|(1 - \mu_2)(x - x')\|_w \quad (147)$$

$$= (1 + \mu_1 - \mu_2)\|x - x'\|_w < \|x - x'\|_w, \quad (148)$$

where in (145) we used that $g(z_1) = g(z_2) = 0$, in (146) we applied (144) upon noting that x, z_1 and z_2, x' are contained in the same ball each, in (147) we inserted the definition of z_1 and z_2 , and in (148) we used $\mu_2 > \mu_1$. Finally, assume that $x \in \mathcal{S}_\ell$, for some $\ell \in \{1, \dots, M-1\}$, and $x' \notin \mathcal{S}_j$, for all $j \in \{1, \dots, M-1\}$. We let $z(\mu) := x + \mu(x' - x)$, $\mu \in [0, 1]$. Since $z(0) = x \in \mathcal{S}_\ell$ and $z(1) = x' \notin \mathcal{S}_\ell$, there must be a $\mu_1 \in (0, 1)$ such that $\|z(\mu_1) - x_\ell\|_w = \epsilon'/2$. Again, we set $z_1 = z(\mu_1)$ and bound

$$\begin{aligned} |g(x) - g(x')| &\leq |g(x) - g(z_1)| + |g(z_1) - g(x')| \\ &\leq |g(x) - g(z_1)| + 0 \end{aligned} \quad (149)$$

$$\leq \|x - z_1\|_w = \|\mu_1(x - x')\|_w \quad (150)$$

$$= \mu_1\|x - x'\|_w < \|x - x'\|_w, \quad (151)$$

where in (149) we used that $g(z_1) = g(x') = 0$, in (150) we applied (144) upon noting that x, z_1 are contained in the same ball, and in (151) we used $\mu_1 < 1$. We have thus established that $|g_{\alpha_1, \dots, \alpha_{M-1}}(x) - g_{\alpha_1, \dots, \alpha_{M-1}}(x')| \leq \|x - x'\|_w$, $\forall x, x' \in \mathcal{S}^-$, and hence $\mathcal{U} \subset \mathcal{G}_0(w)$.

Next, we show that for distinct $g_{\alpha_1, \dots, \alpha_{M-1}}, g_{\beta_1, \dots, \beta_{M-1}} \in \mathcal{U}$ (i.e., there is at least one $\ell \in \{1, \dots, M-1\}$ such that $\alpha_\ell \neq \beta_\ell$), it holds that $\rho_0(g_{\alpha_1, \dots, \alpha_{M-1}}, g_{\beta_1, \dots, \beta_{M-1}}) > \epsilon$. Indeed, for any such ℓ , we have

$$\rho_0(g_{\alpha_1, \dots, \alpha_{M-1}}, g_{\beta_1, \dots, \beta_{M-1}}) = \sup_{\tilde{x} \in \mathcal{S}^-} |g_{\alpha_1, \dots, \alpha_{M-1}}(\tilde{x}) - g_{\beta_1, \dots, \beta_{M-1}}(\tilde{x})| \quad (152)$$

$$\begin{aligned} &\geq |g_{\alpha_1, \dots, \alpha_{M-1}}(x_\ell) - g_{\beta_1, \dots, \beta_{M-1}}(x_\ell)| \\ &= |2(\alpha_\ell - \beta_\ell)\epsilon'/2| = \epsilon' > \epsilon, \end{aligned} \quad (153)$$

where in (152) we used (13) and in (153) we inserted the particular choice $\tilde{x} = x_\ell$ to lower-bound the sup. This establishes that \mathcal{U} , indeed, constitutes an ϵ -packing for $(\mathcal{G}_0(w), \rho_0)$ and we therefore have

$$\log M(\epsilon; \mathcal{G}_0(w), \rho_0) \geq \log |\mathcal{U}| = M - 1 = \left(\prod_{\ell=0}^T \left\lceil \frac{2Dw[\ell]}{\epsilon} \right\rceil \right) - 1.$$

□

C.5 Proof of Lemma 4.2

Note that $w_{a,b}^{(e)}[t] > \epsilon/d$ gives $t < \frac{\log(\frac{ad}{\epsilon})}{b \log(e)}$, and thereby

$$T = \max \left\{ t \in \mathbb{N}_0 \mid t < \frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right\} = \left\lceil \frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right\rceil - 1. \quad (154)$$

Next, we have

$$\begin{aligned} \log \left(\prod_{\ell=0}^T \frac{cw_{a,b}^{(e)}[\ell]}{\epsilon} \right) &= \log \left(\prod_{\ell=0}^T \frac{ace^{-b\ell}}{\epsilon} \right) \\ &= (T+1) \log(\epsilon^{-1}) - b \log(e) \frac{T(T+1)}{2} + (T+1) \log(ac) \\ &= \left\lceil \frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right\rceil \log(\epsilon^{-1}) - b \log(e) \frac{\left\lceil \frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right\rceil \left(\left\lceil \frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right\rceil - 1 \right)}{2} \\ &\quad + \left\lceil \frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right\rceil \log(ac). \end{aligned} \quad (155)$$

Assuming that ϵ is sufficiently small to guarantee that $\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} - 1 \geq 0$, we can upper-bound (155) according to

$$\log \left(\prod_{\ell=0}^T \frac{cw_{a,b}^{(e)}[\ell]}{\epsilon} \right) \tag{156}$$

$$\begin{aligned} &\leq \left(\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} + 1 \right) \log(\epsilon^{-1}) - b \log(e) \frac{\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \left(\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} - 1 \right)}{2} + \left(\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} + 1 \right) \log(ac) \\ &= \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + \left(\frac{\log(ad)}{b \log(e)} + 1 - \frac{1}{2b \log(e)} (2 \log(ad) - b \log(e)) \right. \\ &\quad \left. + \frac{\log(ac)}{b \log(e)} \right) \log(\epsilon^{-1}) - \frac{\log(ad) (\log(ad) - b \log(e))}{2b \log(e)} + \log(ac) \left(\frac{\log(ad)}{b \log(e)} + 1 \right) \\ &= \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})). \end{aligned} \tag{157}$$

In the same spirit, we can lower-bound (155) as

$$\log \left(\prod_{\ell=0}^T \frac{cw_{a,b}^{(e)}[\ell]}{\epsilon} \right) \tag{158}$$

$$\begin{aligned} &\geq \left(\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right) \log(\epsilon^{-1}) - b \log(e) \frac{\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \left(\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} + 1 \right)}{2} + \left(\frac{\log(\frac{ad}{\epsilon})}{b \log(e)} \right) \log(ac) \\ &= \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + \left(\frac{\log(ad)}{b \log(e)} - \frac{1}{2b \log(e)} (2 \log(ad) + b \log(e)) \right. \\ &\quad \left. + \frac{\log(ac)}{b \log(e)} \right) \log(\epsilon^{-1}) - \frac{\log(ad) (\log(ad) + b \log(e))}{2b \log(e)} + \log(ac) \left(\frac{\log(ad)}{b \log(e)} \right) \\ &= \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})). \end{aligned} \tag{159}$$

Finally, combining (157) and (159) yields

$$\log \left(\prod_{\ell=0}^T \frac{cw_{a,b}^{(e)}[\ell]}{\epsilon} \right) = \frac{1}{2b \log(e)} \log^2(\epsilon^{-1}) + o(\log^2(\epsilon^{-1})),$$

as desired. \square

C.6 Proof of Lemma 4.5

We start by noting that

$$T = \max \left\{ t \in \mathbb{N}_0 \mid w_{p,q}^{(p)} > \frac{\epsilon}{d} \right\} = \left\lceil \left(\frac{dq}{\epsilon} \right)^{1/p} \right\rceil - 2. \tag{160}$$

Next,

$$\begin{aligned}
\log \left(\prod_{\ell=0}^T \frac{c w_{p,q}^{(p)}[\ell]}{\epsilon} \right) &= (T+1) \log(c) + (T+1) \log(\epsilon^{-1}) + \sum_{\ell=0}^T \log(w_{p,q}^{(p)}[\ell]) \\
&= (T+1) \log(cq) + (T+1) \log(\epsilon^{-1}) - p \log((T+1)!) \\
&= (T+1) \log(cq) + (T+1) \log(\epsilon^{-1}) - p((T+1) \log(T+1) \\
&\quad - (T+1) \log(e)) + O(\log(T+1)) \tag{161}
\end{aligned}$$

$$\begin{aligned}
&= (T+1) \log(cqe^p) + (T+1) \left(\log(\epsilon^{-1}) - p \log \left(\left\lceil \left(\frac{dq}{\epsilon} \right)^{1/p} \right\rceil - 1 \right) \right) \\
&\quad + O(\log(\epsilon^{-1})), \tag{162}
\end{aligned}$$

where (161) follows from the logarithm form of Stirling's approximation, namely

$$\log(n!) = n \log(n) - n \log(e) + O(\log(n)). \tag{163}$$

Now, assuming that ϵ is sufficiently small to guarantee that $1 \leq \frac{1}{2}(dq/\epsilon)^{1/p}$, applying $\log((dq/\epsilon)^{1/p} - 1) \geq \log(\frac{1}{2}(dq/\epsilon)^{1/p}) \geq 0$, we can upper-bound (162) as

$$\begin{aligned}
&(T+1) \log(cqe^p) + (T+1) \left(\log(\epsilon^{-1}) - p \log \left(\left\lceil \left(\frac{dq}{\epsilon} \right)^{1/p} \right\rceil - 1 \right) \right) + O(\log(\epsilon^{-1})) \\
&\leq (T+1) \log(cqe^p) + (T+1) \left(\log(\epsilon^{-1}) - p \log \left(\frac{1}{2} \left(\frac{dq}{\epsilon} \right)^{1/p} \right) \right) + O(\log(\epsilon^{-1})) \\
&= (T+1)(\log(ce^p/d) + p) + O(\log(\epsilon^{-1})) \\
&= \left(\left\lceil \left(\frac{dq}{\epsilon} \right)^{1/p} \right\rceil - 1 \right) (\log(ce^p/d) + p) + O(\log(\epsilon^{-1})) \\
&\leq \left(\frac{dq}{\epsilon} \right)^{1/p} (\log(ce^p/d) + p) + O(\log(\epsilon^{-1})). \tag{164}
\end{aligned}$$

Furthermore, we can lower-bound (162) according to

$$\begin{aligned}
& (T+1) \log(cqe^p) + (T+1) \left(\log(\epsilon^{-1}) - p \log \left(\left(\frac{dq}{\epsilon} \right)^{1/p} \right) \right) + O(\log(\epsilon^{-1})) \\
&= (T+1) \log(ce^p/d) + O(\log(\epsilon^{-1})) \\
&= \left(\left\lceil \left(\frac{dq}{\epsilon} \right)^{1/p} \right\rceil - 1 \right) \log(ce^p/d) + O(\log(\epsilon^{-1})) \\
&\geq \left(\left(\frac{dq}{\epsilon} \right)^{1/p} - 1 \right) \log(ce^p/d) + O(\log(\epsilon^{-1})) \\
&= \left(\frac{dq}{\epsilon} \right)^{1/p} \log(ce^p/d) + O(\log(\epsilon^{-1})). \tag{165}
\end{aligned}$$

Combining (162), (164), and (165) yields the desired result

$$\log \left(\prod_{\ell=0}^T \frac{cw_{p,q}^{(p)}[\ell]}{\epsilon} \right) = \Theta(\epsilon^{-1/p}).$$

□

C.7 Proof of Lemma 4.6

Consider $\epsilon \in (0, \epsilon_0)$ with $\epsilon_0 = \frac{Dw_{p,q}^{(p)}[0]}{2} = \frac{aD}{2}$. By Theorem 3.12, the exterior covering number of $(\mathcal{G}(w_{p,q}^{(p)}), \rho)$ satisfies

$$\left(\prod_{\ell=0}^{T'} \left\lceil \frac{Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right\rceil \right) - 1 \leq \log N^{\text{ext}}(\epsilon; \mathcal{G}(w_{p,q}^{(p)}), \rho_*) \leq \log(3) \prod_{\ell=0}^{T''} \left(2 \left\lceil \frac{4Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right\rceil + 1 \right), \tag{166}$$

where

$$T' := \max \left\{ \ell \in \mathbb{N}_0 \mid w_{p,q}^{(p)}[\ell] > \frac{\epsilon}{D} \right\} \text{ and } T'' := \max \left\{ \ell \in \mathbb{N}_0 \mid w_{p,q}^{(p)}[\ell] > \frac{\epsilon}{4D} \right\}.$$

We can further lower-bound the left-most term in (166) according to

$$\begin{aligned}
\left(\prod_{\ell=0}^{T'} \left\lceil \frac{Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right\rceil \right) - 1 &\geq \left(\prod_{\ell=0}^{T'} \frac{Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right) - 1 \\
&\geq \frac{1}{2} \prod_{\ell=0}^{T'} \frac{Dw_{p,q}^{(p)}[\ell]}{\epsilon}, \tag{167}
\end{aligned}$$

where (167) follows from

$$\frac{1}{2} \prod_{\ell=0}^{T'} \frac{Dw_{p,q}^{(p)}[\ell]}{\epsilon} \geq 1,$$

which, in turn, is a consequence of

$$\begin{aligned} \frac{Dw_{p,q}^{(p)}[0]}{\epsilon} &\geq \frac{Dw_{p,q}^{(p)}[0]}{\epsilon_0} = 2, \\ \frac{Dw_{p,q}^{(p)}[\ell]}{\epsilon} &\geq \frac{Dw_{p,q}^{(p)}[T']}{\epsilon} > 1, \quad \text{for } \ell \in \llbracket T' \rrbracket \setminus \{0\}. \end{aligned}$$

Similarly, we can further upper-bound the right-most term in (166) as

$$\begin{aligned} \log(3) \left(\prod_{\ell=0}^{T''} \left(2 \left\lceil \frac{4Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right\rceil + 1 \right) \right) &\leq \log(3) \left(\prod_{\ell=0}^{T''} \left(\frac{8Dw_{p,q}^{(p)}[\ell]}{\epsilon} + 3 \right) \right) \\ &\leq \log(3) \left(\prod_{\ell=0}^{T''} \frac{20Dw_{p,q}^{(p)}[\ell]}{\epsilon} \right), \end{aligned} \quad (168)$$

where (168) follows from

$$\frac{4Dw_{p,q}^{(p)}[\ell]}{\epsilon} \geq \frac{4Dw_{p,q}^{(p)}[T'']}{\epsilon} > 1, \quad \text{for } \ell \in \llbracket T'' \rrbracket.$$

Combining (166)–(168), taking logarithms two more times, dividing the results by $\log(\epsilon^{-1})$, and applying Lemma 4.5, we obtain

$$\frac{\log(\Theta(\epsilon^{-1/p}))}{\log(\epsilon^{-1})} \leq \frac{\log^{(3)} N^{\text{ext}}(\epsilon; \mathcal{G}(w_{p,q}^{(p)}), \rho_*)}{\log(\epsilon^{-1})} \leq \frac{\log(\Theta(\epsilon^{-1/p}))}{\log(\epsilon^{-1})}. \quad (169)$$

Taking the limit $\epsilon \rightarrow 0$, then yields

$$\lim_{\epsilon \rightarrow 0} \frac{\log^{(3)} N^{\text{ext}}(\epsilon; \mathcal{G}(w_{p,q}^{(p)}), \rho_*)}{\log(\epsilon^{-1})} = \frac{1}{p},$$

which implies that $(\mathcal{G}(w_{p,q}^{(p)}), \rho_*)$ is of order 2 and type 1, with generalized dimension

$$\mathfrak{d} = \frac{1}{p}.$$

□

C.8 Auxiliary results on ReLU networks

Lemma C.2. For $d \in \mathbb{N}_+$ with $d \geq 2$, consider the following functions,

$$\begin{aligned} f_d^{\min}(z) &:= \min\{z_1, z_2, \dots, z_d\}, & \text{for } z \in \mathbb{R}^d, \\ f_d^{\max}(z) &:= \max\{z_1, z_2, \dots, z_d\}, & \text{for } z \in \mathbb{R}^d. \end{aligned}$$

Then, there exist ReLU networks $\Phi_d^{\min} \in \mathcal{N}_{d,1}$ and $\Phi_d^{\max} \in \mathcal{N}_{d,1}$, with $\mathcal{L}(\Phi_d^{\min}) = \mathcal{L}(\Phi_d^{\max}) = \lceil \log(d+1) \rceil + 1$, $\mathcal{M}(\Phi_d^{\min}) = \mathcal{M}(\Phi_d^{\max}) \leq 14d - 9$, $\mathcal{W}(\Phi_d^{\min}) = \mathcal{W}(\Phi_d^{\max}) \leq 3d$, $\mathcal{K}(\Phi_d^{\min}) = \mathcal{K}(\Phi_d^{\max}) = \{1, -1\}$, $\mathcal{B}(\Phi_d^{\min}) = \mathcal{B}(\Phi_d^{\max}) = 1$, such that

$$\begin{aligned} \Phi_d^{\min}(z) &= f_d^{\min}(z), & \text{for all } z \in \mathbb{R}^d, \\ \Phi_d^{\max}(z) &= f_d^{\max}(z), & \text{for all } z \in \mathbb{R}^d. \end{aligned}$$

Proof. We only need to show the result for the max function as $\min\{z_1, z_2, \dots, z_d\} = -\max\{-z_1, -z_2, \dots, -z_d\}$. First, we realize $\max\{x_1, x_2\}$ according to

$$\max\{x_1, x_2\} = x_1 + \rho(x_2 - x_1) \tag{170}$$

by the ReLU network

$$\widehat{\Phi} = \widehat{W}_2 \circ \rho \circ \widehat{W}_1, \tag{171}$$

with

$$\begin{aligned} \widehat{W}_1(x) &= \begin{pmatrix} 1 & 0 \\ -1 & 0 \\ -1 & 1 \end{pmatrix} x =: A_1 x, \\ \widehat{W}_2(x) &= \begin{pmatrix} 1 & -1 & 1 \end{pmatrix} x =: A_2 x. \end{aligned} \tag{172}$$

Now, for arbitrary $d \in \mathbb{N}_+$ with $d \geq 2$, choose $\ell \in \mathbb{N}_+$ such that $2^\ell \leq d < 2^{\ell+1}$. Then, we double the first $k = 2^{\ell+1} - d$ elements and retain the remaining $d - k$ elements as follows

$$\max\{x_1, \dots, x_d\} = \max\{x_1, x_1, \dots, x_k, x_k, x_{k+1}, x_{k+2}, \dots, x_d\}. \tag{173}$$

The primary reason for doubling elements is that by extending the set $\{x_1, \dots, x_d\}$ to a set whose cardinality is a power of 2, we can utilize (170) together with a divide-and-conquer approach to determine the maximum value of the set $\{x_1, \dots, x_d\}$. This then leads to a ReLU network realization of depth scaling logarithmically in d . Now, let $W_{\ell+1}(x) := A_2 x$, $W_j(x) := B_j x$, for

$j \in \{\ell, \ell - 1, \dots, 1\}$, and $W_0(x) := A_0x$, with

$$\begin{aligned}
B_o &= A_1 \text{diag}(A_2, A_2) \\
&= \begin{pmatrix} 1 & -1 & 1 & 0 & 0 & 0 \\ -1 & 1 & -1 & 0 & 0 & 0 \\ -1 & 1 & -1 & 1 & -1 & 1 \end{pmatrix}, \\
B_j &= \text{diag}(\underbrace{B_o, \dots, B_o}_{2^{\ell-j}}), \\
A_0 &= \text{diag}(\underbrace{A_1, \dots, A_1}_{2^\ell}) \text{diag}(\underbrace{\mathbb{1}_2, \dots, \mathbb{1}_2}_k, \mathbb{I}_{d-k}) \\
&= \text{diag}(\underbrace{A_1 \mathbb{1}_2, \dots, A_1 \mathbb{1}_2}_k, \underbrace{A_1, \dots, A_1}_{2^\ell - k}), \\
A_1 \mathbb{1}_2 &= (1, -1, 0)^T.
\end{aligned} \tag{174}$$

Combining (170)-(174), we can show that

$$\Phi_d^{\max} := W_{\ell+1} \circ \rho \circ W_\ell \circ \rho \circ \dots \circ \rho \circ W_1 \circ \rho \circ W_0 = f_d^{\max} \tag{175}$$

as follows:

- (i) We double the first $k = 2^{\ell+1} - d$ elements and retain the remaining $d - k$ elements of $\{x_1, x_2, \dots, x_d\}$ according to

$$\begin{aligned}
&\begin{pmatrix} x_1 & x_1 & \dots & x_k & x_k & x_{k+1} & x_{k+2} & \dots & x_d \end{pmatrix}^T \\
&= \text{diag}(\underbrace{\mathbb{1}_2, \dots, \mathbb{1}_2}_k, \mathbb{I}_{d-k}) \begin{pmatrix} x_1 & x_2 & \dots & x_d \end{pmatrix}^T.
\end{aligned}$$

For simplicity, we denote the resulting set $\{x_1, x_1, \dots, x_k, x_k, x_{k+1}, x_{k+2}, \dots, x_d\}$ as $\{y_1, y_2, \dots, y_{2^{\ell+1}}\}$.

- (ii) Application of (171) to the 2^ℓ pairs in $\{y_1, y_2, \dots, y_{2^{\ell+1}}\}$, results in

$$\{\max\{y_1, y_2\}, \max\{y_3, y_4\}, \dots, \max\{y_{2^{\ell+1}-1}, y_{2^{\ell+1}}\}\}$$

and hence reduces the number of elements from $2^{\ell+1}$ to 2^ℓ . This reduction is applied iteratively until we get $\max\{x_1, x_2, \dots, x_d\}$. The compositions $A_1 \circ \text{diag}(A_2, A_2)$, formed in the iterative application of (171), constitute the main diagonal elements of the matrices B_k , $k \in \{\ell, \ell - 1, \dots, 1\}$. We illustrate this part of the procedure using the simplest

possible example, namely for $\ell = 1$ and hence $\max\{y_1, y_2, y_3, y_4\}$. Specifically, we have

$$\begin{aligned}
& \max\{y_1, y_2, y_3, y_4\} = \max\{\max\{y_1, y_2\}, \max\{y_3, y_4\}\} \\
& = \max\left\{\text{diag}(A_2, A_2) \rho\left(\text{diag}(A_1, A_1) \left((y_1, y_2), (y_3, y_4)\right)^T\right)\right\} \\
& = A_2 \rho\left(A_1 \text{diag}(A_2, A_2) \rho\left(\text{diag}(A_1, A_1) \left((y_1, y_2), (y_3, y_4)\right)^T\right)\right) \\
& = (W_2 \circ \rho \circ W_1 \circ \rho \circ W_0)(y).
\end{aligned}$$

The proof is concluded by noting that

$$\begin{aligned}
\mathcal{L}(\Phi_d^{\max}) &\stackrel{(175)}{=} \ell + 1 = \lceil \log(d+1) \rceil + 1, \\
\mathcal{M}(\Phi_d^{\max}) &\stackrel{(175)}{=} \sum_{j=0}^{\ell+1} \mathcal{M}(W_j) = \mathcal{M}(A_2) + \sum_{j=1}^{\ell} \mathcal{M}(B_j) + \mathcal{M}(A_0) \\
&\stackrel{(174), (172)}{=} 3 + 2k + 4(2^\ell - k) + 12 \sum_{j=1}^{\ell} 2^{\ell-j} \leq 14d - 9, \\
\mathcal{W}(\Phi_d^{\max}) &\stackrel{(175)}{=} \max_{j=0,1,\dots,\ell+1} \mathcal{W}(W_j) \stackrel{(174), (172)}{=} 3 \cdot 2^\ell \leq 3d, \\
\mathcal{K}(\Phi_d^{\max}) &\stackrel{(175)}{\subset} \bigcup_{j=0}^{\ell+1} \mathcal{K}(W_j) \stackrel{(174), (172)}{=} \{1, -1\}, \\
\mathcal{B}(\Phi_d^{\max}) &= \max_{b \in \mathcal{K}(\Phi_d^{\max})} |b| = 1.
\end{aligned}$$

□

Lemma C.3 (Composition of ReLU networks [7]). *For $i = 1, \dots, n$, let $d_i \in \mathbb{N}_+$ and $\Phi_i \in \mathcal{N}_{d_i, d_{i+1}}$. Then, there exists a network $\Psi \in \mathcal{N}_{d_1, d_{n+1}}$ with*

$$\Psi(x) = (\Phi_n \circ \Phi_{n-1} \circ \dots \circ \Phi_1)(x), \quad \text{for all } x \in \mathbb{R}^{d_1}, \quad (176)$$

satisfying

$$\begin{aligned}
\mathcal{L}(\Psi) &= \sum_{i=1}^n \mathcal{L}(\Phi_i), \\
\mathcal{M}(\Psi) &\leq 2 \sum_{i=1}^n \mathcal{M}(\Phi_i), \\
\mathcal{W}(\Psi) &\leq \max\left\{\max_{i=1,\dots,n-1} \{2d_i\}, \max_{i=1,\dots,n} \{\mathcal{W}(\Phi_i)\}\right\}, \\
\mathcal{K}(\Psi) &\subset \bigcup_{i=1}^n (\mathcal{K}(\Phi_i) \cup (-\mathcal{K}(\Phi_i))), \\
\mathcal{B}(\Psi) &= \max_{i=1,\dots,n} \mathcal{B}(\Phi_i).
\end{aligned} \quad (177)$$

Proof. Follows along the same lines as the proof of [7, Lemma II.3].

□

Lemma C.4 (Parallelization of ReLU networks of the same depth [7]). For $i = 1, \dots, n$, let $d_i, d'_i \in \mathbb{N}_+$ and $\Phi_i \in \mathcal{N}_{d_i, d'_i}$ with $\mathcal{L}(\Phi_i) = L$. Then, there exists a ReLU network $P(\Phi_1, \Phi_2, \dots, \Phi_n) \in \mathcal{N}_{\sum_{i=1}^n d_i, \sum_{i=1}^n d'_i}$ with

$$P(\Phi_1, \Phi_2, \dots, \Phi_n)(x) = (\Phi_1(x), \Phi_2(x), \dots, \Phi_n(x))^T, \quad \text{for all } x \in \mathbb{R}^{\sum_{i=1}^n d_i}, \quad (178)$$

satisfying

$$\begin{aligned} \mathcal{L}(P(\Phi_1, \Phi_2, \dots, \Phi_n)) &= L, \\ \mathcal{M}(P(\Phi_1, \Phi_2, \dots, \Phi_n)) &= \sum_{i=1}^n \mathcal{M}(\Phi_i), \\ \mathcal{W}(P(\Phi_1, \Phi_2, \dots, \Phi_n)) &\leq \sum_{i=1}^n \mathcal{W}(\Phi_i), \\ \mathcal{K}(P(\Phi_1, \Phi_2, \dots, \Phi_n)) &= \bigcup_{i=1}^n \mathcal{K}(\Phi_i), \\ \mathcal{B}(P(\Phi_1, \Phi_2, \dots, \Phi_n)) &= \max_{i=1, \dots, n} \mathcal{B}(\Phi_i). \end{aligned} \quad (179)$$

Proof. Follows along similar lines as the proof of [7, Lemma II.5]. \square

Lemma C.5 (Augmenting network depth [7]). Let $d_1, d_2, K \in \mathbb{N}_+$, and $\Phi \in \mathcal{N}_{d_1, d_2}$ with $\mathcal{L}(\Phi) < K$. Then, there exists a network $\Psi \in \mathcal{N}_{d_1, d_2}$ with

$$\begin{aligned} \mathcal{L}(\Psi) &= K, \\ \mathcal{M}(\Psi) &\leq \mathcal{M}(\Phi) + d_2 \mathcal{W}(\Phi) + 2d_2(K - \mathcal{L}(\Phi)), \\ \mathcal{W}(\Psi) &= \max\{2d_2, \mathcal{W}(\Phi)\}, \\ \mathcal{K}(\Psi) &\subset (\mathcal{K}(\Phi) \cup (-\mathcal{K}(\Phi)) \cup \{1, -1\}), \\ \mathcal{B}(\Psi) &= \max\{1, \mathcal{B}(\Phi)\}, \end{aligned} \quad (180)$$

satisfying $\Psi(x) = \Phi(x)$, for all $x \in \mathbb{R}^{d_1}$.

Proof. Follows along the same lines as the proof of [7, Lemma II.4]. \square

C.9 Remainder of the proof of Lemma 5.2

We verify that $\Phi = \Phi_2 \circ \Phi_1$ defined in (76) has $(2, \epsilon)$ -quantized weights.

- The weights in $\Phi_2 = \widetilde{W}^\Sigma$ defined in (73):

$$\begin{aligned} \tilde{g}_n &\stackrel{(62)}{=} \mathcal{Q}_{2, \epsilon}(g(\tilde{x}_n)) \stackrel{\text{Definition 2.2}}{=} \left\lceil g(\tilde{x}_n) / 2^{-2 \lceil \log(\epsilon^{-1}) \rceil} \right\rceil \cdot 2^{-2 \lceil \log(\epsilon^{-1}) \rceil} \\ &\in 2^{-2 \lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z}, \\ |\tilde{g}_n| &\leq |g(\tilde{x}_n)| + 2^{-2 \lceil \log(\epsilon^{-1}) \rceil} \stackrel{(11)}{\leq} D + 2^{-2 \lceil \log(\epsilon^{-1}) \rceil} \stackrel{(189)}{\leq} D + \epsilon^2 \\ &\stackrel{(189)}{\leq} D + 1 \leq \epsilon^{-2}, \end{aligned}$$

which yields

$$\mathcal{K}(\Phi_2) \subset 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z} \cap [-\epsilon^{-2}, \epsilon^{-2}]. \quad (181)$$

- The weights in Φ_1 defined in (75):

$$\mathcal{K}(\Phi_1) = \bigcup_{n \in \tilde{\mathfrak{N}}} (\mathcal{K}(\Phi_{1,2}^n) \cup \mathcal{K}(\Phi_{1,1}^n)) = \mathcal{K}(\Psi) \cup \left(\bigcup_{n \in \tilde{\mathfrak{N}}} \mathcal{K}(\widehat{W}_n) \right). \quad (182)$$

Recalling that Ψ realizes the spike function and applying Lemma 5.1, we obtain

$$\mathcal{K}(\Psi) = \{1, -1\} \subset 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z} \cap [-\epsilon^{-2}, \epsilon^{-2}]. \quad (183)$$

Based on (73), we get

$$\bigcup_{n \in \tilde{\mathfrak{N}}} \mathcal{K}(\widehat{W}_n) = \bigcup_{\ell=0}^T \left(\bigcup_{n \in \tilde{\mathfrak{N}}} \{\tilde{\delta}_\ell^{-1}, n_\ell\} \right). \quad (184)$$

Moreover, for all $\ell = 0, \dots, T$ and $n \in \tilde{\mathfrak{N}}$, we have

$$\begin{aligned} \tilde{\delta}_\ell^{-1} &\stackrel{(60)}{=} \mathcal{Q}_{2,\epsilon}(\delta_\ell^{-1}) \stackrel{\text{Definition 2.2}}{=} \left\lfloor \frac{s+1}{s} \frac{w[\ell]}{\epsilon} / 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \right\rfloor \cdot 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \\ &\in 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z}, \\ |\tilde{\delta}_\ell^{-1}| &\leq \frac{s+1}{s} \frac{w[\ell]}{\epsilon} + 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \stackrel{\text{Definition 2.15}}{\leq} \frac{s+1}{s} \epsilon^{-1} + \epsilon^2 \stackrel{(189)}{\leq} \epsilon^{-2}. \end{aligned}$$

Hence,

$$\tilde{\delta}_\ell^{-1} \in 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z} \cap [-\epsilon^{-2}, \epsilon^{-2}], \quad (185)$$

and

$$\begin{aligned} n_\ell &\in \mathbb{Z} \subset 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z}, \\ |n_\ell| &\leq \tilde{N}_\ell \stackrel{(58)}{=} \left\lceil \frac{D}{\tilde{\delta}_\ell} \right\rceil \leq D |\tilde{\delta}_\ell^{-1}| + 1 \\ &\leq D \left(\frac{s+1}{s} \epsilon^{-1} + \epsilon^2 \right) + 1 \\ &\stackrel{(189)}{\leq} \left(D \frac{s+1}{s} + D + 1 \right) \epsilon^{-1} \stackrel{(189)}{\leq} \epsilon^{-2}, \end{aligned}$$

which yields

$$n_\ell \in 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z} \cap [-\epsilon^{-2}, \epsilon^{-2}]. \quad (186)$$

Based on (184), (185), and (186), we have

$$\bigcup_{n \in \tilde{\mathfrak{N}}} \mathcal{K}(\widehat{W}_n) \subset 2^{-2\lceil \log(\epsilon^{-1}) \rceil} \mathbb{Z} \cap [-\epsilon^{-2}, \epsilon^{-2}]. \quad (187)$$

Combining (182), (183), (184), and (187) yields

$$\mathcal{K}(\Phi_1) \subset 2^{-2} \lceil \log(\epsilon^{-1}) \rceil \mathbb{Z} \cap [-\epsilon^{-2}, \epsilon^{-2}]. \quad (188)$$

Using (181) and (188), Lemma C.3 shows that $\Phi = \Phi_2 \circ \Phi_1$, indeed, has $(2, \epsilon)$ -quantized weights.

Finally, we compute $\mathcal{L}(\Phi)$ and derive an upper bound on $\mathcal{M}(\Phi)$ according to

$$\begin{aligned} \mathcal{L}(\Phi) &\stackrel{(76), \text{Lemma C.3}}{=} \mathcal{L}(\Phi_2) + \mathcal{L}(\Phi_1) \\ &\stackrel{(75), (74), \text{Lemma C.4}}{=} \mathcal{L}(\Psi) + 2 \\ &\stackrel{\text{Lemma 5.1}}{=} \lceil \log(T+2) \rceil + 6, \\ \mathcal{M}(\Phi) &\stackrel{\text{Lemma C.3}}{\leq} 2(\mathcal{M}(\Phi_2) + \mathcal{M}(\Phi_1)) \\ &\stackrel{\text{Lemma C.4}}{=} 2 \left(\mathcal{M}(\Phi_2) + \sum_{i=1}^{|\tilde{\mathfrak{N}}|} \left(2\mathcal{M}(\Psi) + 2\mathcal{M}(\widehat{W}_{n^i}) \right) \right) \\ &\stackrel{(73)-(76), \text{Lemma 5.1}}{\leq} 2|\tilde{\mathfrak{N}}| (1 + 120(T+1) - 56 + 2(T+1)) \\ &\leq 244(T+1)|\tilde{\mathfrak{N}}| \\ &\stackrel{(58)}{\leq} 244(T+1) \prod_{\ell=0}^T (2D\tilde{\delta}_\ell^{-1} + 3) \\ &\stackrel{(58), (60)}{\leq} 244(T+1) \prod_{\ell=0}^T \left(2D \left(\frac{s+1}{s} \frac{w[\ell]}{\epsilon} + \epsilon^2 \right) + 3 \right) \\ &\stackrel{(189)}{\leq} 244(T+1) \prod_{\ell=0}^T \left(\frac{2Dw[\ell]}{\epsilon} \frac{s+1}{s} + 4 \right). \end{aligned}$$

The proof is concluded by setting

$$\epsilon_0 = \min \left\{ 1, \frac{1}{2}, \frac{1}{s+1}, \sqrt{\frac{1}{D+1}}, \left(D \frac{s+1}{s} + D+1 \right)^{-1} \right\}. \quad (189)$$

□

Acknowledgments

The authors would like to thank Erwin Riegler for useful suggestions regarding the proof of Lemma C.2 and Thomas Allard for help with Table 1.

References

- [1] G. Cybenko, “Approximation by superpositions of a sigmoidal function,” *Mathematics of Control, Signals, and Systems*, vol. 2, no. 4, pp. 303–314, 1989.
- [2] K.-I. Funahashi, “On the approximate realization of continuous mappings by neural networks,” *Neural Networks*, vol. 2, no. 3, pp. 183–192, 1989.
- [3] K. Hornik, M. Stinchcombe, and H. White, “Multilayer feedforward networks are universal approximators,” *Neural Networks*, vol. 2, no. 5, pp. 359–366, 1989.
- [4] A. R. Barron, “Universal approximation bounds for superpositions of a sigmoidal function,” *IEEE Transactions on Information Theory*, vol. 39, no. 3, p. 930–945, May 1993.
- [5] ———, “Approximation and estimation bounds for artificial neural networks,” *Machine Learning*, vol. 14, no. 1, p. 115–133, Jan. 1994.
- [6] H. Bölcskei, P. Grohs, G. Kutyniok, and P. Petersen, “Optimal approximation with sparsely connected deep neural networks,” *SIAM Journal on Mathematics of Data Science*, vol. 1, no. 1, pp. 8–45, Jan. 2019.
- [7] D. Elbrächter, D. Perekrestenko, P. Grohs, and H. Bölcskei, “Deep neural network approximation theory,” *IEEE Transactions on Information Theory*, vol. 67, no. 5, pp. 2581–2623, May 2021.
- [8] S. Nagel and Z. Xu, “Asset pricing with fading memory,” *The Review of Financial Studies*, vol. 35, no. 5, pp. 2190–2245, 2022.
- [9] E. H. Dill, “Simple materials with fading memory,” *Academic Press, New York*, 1975.
- [10] C. C. Wang, “The principle of fading memory,” *Archive for Rational Mechanics and Analysis*, vol. 18, pp. 343–366, 1965.
- [11] W. A. Day, *The Thermodynamics of Simple Materials with Fading Memory*. Springer Science & Business Media, 2013, vol. 22.
- [12] A. Bakshi, A. Liu, A. Moitra, and M. Yau, “A new approach to learning linear dynamical systems,” in *Proceedings of the 55th Annual ACM Symposium on Theory of Computing*, 2023, pp. 335–348.
- [13] B. Boots, G. J. Gordon, and S. Siddiqi, “A constraint generation approach to learning stable linear dynamical systems,” *Advances in Neural Information Processing Systems*, vol. 20, 2007.
- [14] E. Hazan, K. Singh, and C. Zhang, “Learning linear dynamical systems via spectral filtering,” *Advances in Neural Information Processing Systems*, vol. 30, 2017.

- [15] A. Wagenmaker and K. Jamieson, “Active learning for identification of linear dynamical systems,” in *Conference on Learning Theory*. PMLR, 2020, pp. 3487–3582.
- [16] C. Hutter, R. Gül, and H. Bölcskei, “Metric entropy limits on recurrent neural network learning of linear dynamical systems,” *Applied and Computational Harmonic Analysis*, vol. 59, pp. 198–223, 2022.
- [17] C. L. Giles, B. G. Horne, and T. Lin, “Learning a class of large finite state machines with a recurrent neural network,” *Neural Networks*, vol. 8, no. 9, pp. 1359–1365, 1995.
- [18] I. W. Sandberg, “Approximately finite memory and the circle criterion,” *IEEE Transactions on Circuits and Systems I: Fundamental Theory and Applications*, vol. 41, no. 7, pp. 473–476, 1994.
- [19] V. Volterra, *Theory of Functionals and of Integral and Integro-Differential Equations*, ser. Dover Books on Mathematics. New York: Dover Publications, 1959.
- [20] N. Wiener, *Nonlinear problems in random theory*. Cambridge, Massachusetts, USA: The MIT Press, 1966.
- [21] S. Boyd and L. Chua, “Fading memory and the problem of approximating nonlinear operators with Volterra series,” *IEEE Transactions on Circuits and Systems*, vol. 32, no. 11, pp. 1150–1161, 1985.
- [22] M. B. Matthews, “Approximating nonlinear fading-memory operators using neural network models,” *Circuits, Systems and Signal Processing*, vol. 12, no. 2, pp. 279–307, 1993.
- [23] L. Grigoryeva and J.-P. Ortega, “Echo state networks are universal,” *Neural Networks*, vol. 108, pp. 495–508, 2018.
- [24] R. Pascanu, C. Gulcehre, K. Cho, and Y. Bengio, “How to construct deep recurrent neural networks,” *arXiv preprint arXiv:1312.6026*, 2013.
- [25] A. N. Kolmogorov and V. M. Tikhomirov, “ ε -entropy and ε -capacity of sets in function spaces,” *Uspekhi Matematicheskikh Nauk*, vol. 14, no. 2, pp. 3–86, 1959.
- [26] M. J. Wainwright, *High-dimensional Statistics: A Non-asymptotic Viewpoint*. Cambridge University Press, 2019, vol. 48.
- [27] G. Zames, “On the metric complexity of causal linear systems: ϵ -entropy and ϵ -dimension for continuous time,” *IEEE Transactions on Automatic Control*, vol. 24, no. 2, pp. 222–230, 1979.
- [28] V. Lakshmikantham, *Theory of Integro-differential Equations*. CRC press, 1995, vol. 1.

- [29] G. Alves, J. De Araújo, J. Cressoni, L. da Silva, M. da Silva, and G. Viswanathan, “Superdiffusion driven by exponentially decaying memory,” *Journal of Statistical Mechanics: Theory and Experiment*, vol. 2014, no. 4, p. P04026, 2014.
- [30] T. R. Moura, G. Viswanathan, M. da Silva, J. Cressoni, and L. da Silva, “Transient superdiffusion in random walks with a q-exponentially decaying memory profile,” *Physica A: Statistical Mechanics and its Applications*, vol. 453, pp. 259–263, 2016.
- [31] S. Karnik, R. Wang, and M. Iwen, “Neural network approximation of continuous functions in high dimensions with applications to inverse problems,” *Journal of Computational and Applied Mathematics*, vol. 438, p. 115557, 2024.
- [32] D. Yarotsky, “Optimal approximation of continuous functions by very deep ReLU networks,” in *Conference on Learning Theory*. PMLR, 2018, pp. 639–649.
- [33] E. Mainini and G. Mola, “Exponential and polynomial decay for first order linear Volterra evolution equations,” *Quarterly of Applied Mathematics*, vol. 67, no. 1, pp. 93–111, 2009.
- [34] M. Fabrizio and S. Polidoro, “Asymptotic decay for some differential systems with fading memory,” *Applicable Analysis*, vol. 81, no. 6, pp. 1245–1264, 2002.
- [35] P. Grohs, “Optimally sparse data representations,” *Harmonic and Applied Analysis: From Groups to Signals*, pp. 199–248, 2015.