

NEU: A Meta-Algorithm for Universal UAP-Invariant Feature Representation

Anastasis Kratsios

Department of Mathematics

Eidgenössische Technische Hochschule Zürich (ETH)

Rämistrasse 101, 8092 Zürich, ZH, Switzerland

ANASTASIS.KRATSIOS@MATH.ETHZ.CH

Cody Hyndman

Department of Mathematics and Statistics

Concordia University

1455 boulevard de Maisonneuve Ouest, Montréal, Québec, H3G 1M8, Canada

CODY.HYNDMAN@CONCORDIA.CA

Editor: George Konidaris

Abstract

Effective feature representation is key to the predictive performance of any algorithm. This paper introduces a meta-procedure, called Non-Euclidean Upgrading (NEU), which learns feature maps that are expressive enough to embed the universal approximation property (UAP) into most model classes while only outputting feature maps that preserve any model class's UAP. We show that NEU can learn any feature map with these two properties if that feature map is asymptotically deformable into the identity. We also find that the feature-representations learned by NEU are always submanifolds of the feature space. NEU's properties are derived from a new deep neural model that is universal amongst all orientation-preserving homeomorphisms on the input space. We derive qualitative and quantitative approximation guarantees for this architecture. We quantify the number of parameters required for this new architecture to memorize any set of input-output pairs while simultaneously fixing every point of the input space lying outside some compact set, and we quantify the size of this set as a function of our model's depth. Moreover, we show that deep feed-forward networks with most commonly used activation functions typically do not have all these properties. NEU's performance is evaluated against competing machine learning methods on various regression and dimension reduction tasks both with financial and simulated data.

Keywords: Geometric Deep Learning, Universal Feature Maps, Reconfiguration Networks, Pre-Processing, Homeomorphism Learning.

1. Introduction

The training phase of most learning problems seeks to identify a model \hat{f} belonging to a model class \mathcal{F} , which best approximates an unknown function f , as given by:

$$\min_{\hat{f} \in \mathcal{F}} \sum_{n \leq N} \frac{1}{N} L(f(x_n), \hat{f}(x_n), x_n) + P(\hat{f}), \quad (\text{L})$$

where $\{x_n\}_{n \leq N}$ is a given set of training data, L is a loss-function, and P is a penalty which encodes regularity into the model \hat{f} . The effectiveness of the learning task (L) often hinges on the appropriateness of the input data's representation. Where by a *representation* of the input space \mathbb{R}^d is the subset $\phi(\mathbb{R}^d) \subseteq \mathbb{R}^F$ and where $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^F$ is a *feature map* mapping into the *feature space* \mathbb{R}^F . Following Micchelli et al. (2006), we define feature maps as being continuous, and by Brouwer (1911) we observe that the feature space's dimension must be at-least that of \mathbb{R}^d .

Two popular but differing approaches to representation learning are offered by kernel methods and by deep learning. Introduced by Boser et al. (1992), the former of the two implicitly embeds \mathbb{R}^d into a high, and often infinite, dimensional linear space H by using the correspondence between feature maps and kernels identified in Aronszajn (1950); Argyriou et al. (2009). However, the effectiveness of these methods hinges on the appropriateness of the specified kernel; see Kanagawa et al. (2020) for example. In contrast, the deep learning paradigm offers a non-parametric approach to representation learning. This is because any deep feed-forward network (DNN) \hat{f} is necessarily of the form $\hat{f}(x) = W \circ \phi(x)$ where W is an affine function on \mathbb{R}^F and $\phi : \mathbb{R}^d \rightarrow \mathbb{R}^F$ is a feature map generated by iteratively applying feed-forward layers to the input space.

This paper introduces Non-Euclidean Upgrading (NEU), a meta-algorithm that incorporates a linearizing preprocessing step into (L) as summarized in Meta-Algorithm 1. During this preprocessing step, NEU generates a feature map that both increases the expressiveness of \mathcal{F} and preserves its approximation capabilities by learning a topological embedding of the input space \mathbb{R}^d into a low-dimensional feature space \mathbb{R}^F . NEU does this by training a new deep neural model type, called the *reconfiguration network* and denoted by $\Phi_{*,d}$, whose members form a universal class of regular feature maps.

NEU balances the newly found flexibility, which $\Phi_{*,d}$ embeds into \mathcal{F} by optimally re-weighting the relative impact of each training data-point $\{x_n\}_{n \leq N}$ in (L) so as to minimize the gap between the trained model's training and testing performance. We denote these new, data-dependent, weights by $\{w_n^{*,\lambda}\}_{n \leq N}$, where $\lambda > 0$ is a hyper-parameter.

Meta-Algorithm 1: Non-Euclidean Upgrading (NEU)

- input** : Hypothesis class \mathcal{F} , loss-function L , penalty function P ,
 Training Data $\{x_n\}_{n \leq N}$
 Feature map's depth J
 Robustness Hyper-parameter $\lambda > 0$
- output:** NEU-model $f^{NEU} \triangleq \hat{f} \circ \hat{\phi}_J$.
- 1 $\hat{\phi} \in \operatorname{argmin}_{\phi \in \Phi_{*,d}} \sum_{n \leq N} w_n^{*,\lambda} L(f(x_n), A\phi(x_n) + b, x_n) + P(A\phi + b)$; ▷ **Get Feature Map**
 - 2 $\hat{f} \in \operatorname{argmin}_{\hat{f} \in \mathcal{F}} \sum_{n \leq N} w_n^{*,\lambda} L(f(x_n), \hat{f} \circ \hat{\phi}(x_n) + b, x_n) + P(\hat{f} \circ \hat{\phi})$; ▷ **Get NEU-Model**
-

We motivate NEU through its properties. Many feature maps can impede the universal approximation property (UAP) of \mathcal{F} . Naturally, we require that any feature map generated by NEU, satisfies the following *UAP-invariance property*, which is introduced and characterized in Kratsios and Bilokopytov (2020):¹

1. The authors find that (P-i) holds exactly when ϕ is injective. Consequentially, Brouwer (1911) implies any UAP-invariant feature map must map into a feature space \mathbb{R}^F of dimension at-least d .

(P-i) If \mathcal{F} is a universal in $C(\mathbb{R}^d, \mathbb{R}^D)$ then so is $\mathcal{F} \circ \phi$.

Next, we require NEU should be able to learn the identity continuously. Thus, it should be capable of not imposing any additional unnecessary structure if, and once, the input space is sufficiently well-represented. Mathematically, we require that the collection of feature maps that NEU can generate, denoted for the moment by Φ , satisfy:

(P-ii) Any $\phi \in \Phi$ can be parameterized to continuously learn the identity; i.e., there is a continuous map $\phi_\alpha : [0, 1] \times \mathbb{R}^d \rightarrow \mathbb{R}^F$ such that

$$\phi_1(x) = \phi(x), \quad \phi_0((x_1, \dots, x_n)) = (x_1, \dots, x_d, 0, \dots, 0), \quad \phi_\alpha \text{ satisfies (P-i)}.$$

Property (P-ii) is critical when members of Φ are built by repeatedly composing many layers since failing (P-i) forces all deeper layers to simultaneously learn the target function and compensate the mistakes of erroneously applied earlier layers. As discussed in, Hardt and Ma (2016), property (P-ii) is core to the success of the batch normalization algorithm of Ioffe and Szegedy (2015), among other recent deep learning paradigms.

NEU is designed to exclusively generate feature maps satisfying both (P-i) and (P-ii). Our main universal approximation results will show that the feature maps generated by NEU are universal amongst all those satisfying both properties (P-i) and (P-ii). In contrast, we show that typically DNNs with the ReLU non-linearity of Hahnloser et al. (2000) fail to satisfy (P-i).

Together, properties (P-i) and (P-ii) only guarantee that a class of feature maps Φ does not disrupt the representation of any input space. However, we are most interested in identifying a feature map class Φ , which additionally improves the expressiveness of any model class possessing a basic level of expressiveness. By this, we mean that NEU should imbue most learning models with the universal approximation property:

(P-iii) If \mathcal{F} contains all linear maps, then $\{\mathcal{F} \circ \phi\}_{\phi \in \Phi}$ is universal.

Property (P-iii) is "asymptotic" since it guarantees that any function can eventually be approximated if a sufficiently complex feature map is used. We complement it with the following, non-asymptotic, refined *memorization* property:

(P-iv) If $n = d$ and \mathcal{F} contains all linear maps, then given any input-output pairs $\{x_n\}_{n \leq N}$ and $\{y_n\}_{n \leq N}$ and any tolerance $\delta > 0$, some feature map $\phi \in \Phi$ satisfies

$$\phi(x_n) = y_n \text{ and } \mu\left(\{x \in \mathbb{R}^d : \phi(x) \neq x\}\right) < \delta;$$

for every $n \leq N$; where μ is the Lebesgue measure on \mathbb{R}^d .

Property (P-iv) is a refinement of the arbitrary *memory capacity* of feed-forward networks studied in Jiang et al. (2009), which simultaneously asks that NEU be able to leave most of the unseen data unimpacted. We show that NEU generates feature maps satisfying properties (P-iii) and (P-iv) and that DNNs with commonly-used non-ReLU activation functions, such as the Swish non-linearity of Ramachandran et al. (2018), the Gaussian Error Linear Unit of Hendrycks and Gimpel (2016), the Soft-Plus activation of Glorot et al. (2011), and tanh activation functions all fail (P-iv). Analogously to Yarotsky (2018), Bölskei et al. (2019), Lu et al. (2020), and Kratsios and Papon (2021) our approximation guarantees are *quantitative* and analogously to Jiang et al. (2009), Yun et al. (2019), and Vershynin (2020a) our memorization guarantees are also quantitative.

Outline of the Paper

Section 2 begins by covering the topological background required for the framing of our main results and ends with the precise description of the deep neural model which NEU trains. Section 3 contains the paper’s main theoretical contributions. These include various universal approximation results, guarantees on the reconfiguration network’s memory capacity, and guarantees that the reconfiguration network can approximate at the desired optimal rates. The implications of these results are then unpacked in the context of NEU. Section 4 evaluates the predictive gain obtained by applying NEU across various regression and dimension reduction problems. Our implementations focus on financial data analysis. The performance of NEU regression methods is subsequently evaluated on simulated data to understand its implications in a fully controlled environment. Specifically, NEU is stress-tested using various pathological regression challenges. All proofs and any additional topological background is available in the supplementary material.

Notation

The following notation is maintained throughout this paper. We denote the set of continuous functions from \mathbb{R}^d to \mathbb{R}^D by $C(\mathbb{R}^d, \mathbb{R}^D)$. The set of DNNs from \mathbb{R}^d to \mathbb{R}^D with activation function $\sigma \in C(\mathbb{R})$ and at-least one hidden layer is denoted by $\mathcal{NN}_{d,D}^\sigma$.

2. Preliminaries

This section covers the background and definition required in the remainder of this paper.

2.1 Background

2.1.1 CONTINUOUS FUNCTIONS

We denote the Euclidean norm by $\|\cdot\|$ on \mathbb{R}^d (resp \mathbb{R}^D). Following Hornik et al. (1989), we view $C(\mathbb{R}^d, \mathbb{R}^D)$ as a metric space, with metric d_{ucc} defined for $f, g \in C(\mathbb{R}^d, \mathbb{R}^D)$ by

$$d_{ucc}(f, g) \triangleq \sum_{k \in \mathbb{N}_+} \frac{\sup_{\|x\| \leq k} \|f(x) - g(x)\|}{2^k \left(1 + \sup_{\|x\| \leq k} \|f(x) - g(x)\|\right)}. \quad (1)$$

This metric describes the uniform convergence on compacts topology standard in the universal approximation literature such as Leshno et al. (1993) and Kidger and Lyons (2020).

Analogously to Yarotsky (2018) our quantitative approximation results depend on the regularity of the unknown target function. The regularity of any $f \in C(\mathbb{R}^d, \mathbb{R}^D)$ is quantified by its (optimal) *modulus of continuity*, denoted by ω_f ,² measures the input’s space’s distortion upon applying f and it is defined by $\omega_f(\delta) \triangleq \sup_{x, y \in \mathbb{R}^d, \|x - y\| \leq \delta} \|f(x) - f(y)\|$.

2.1.2 ORIENTATION-PRESERVING HOMEOMORPHISMS

In Kratsios and Bilokopytov (2020), it was shown that a feature map has the UAP-invariant property if and only if it is injective. Geometrically, this is because any injective feature

2. By the Heine-Cantor Theorem (Munkres, 2000, Theorem 27.6), any continuous functions on a compact subset of its input space, such as $[-M, M]^d \subset \mathbb{R}^d$, has a well-defined modulus of continuity.

map is by definition, injective and continuous, and thus, as discussed in Kratsios and Papon (2021), it preserves all the topological information of any compact subset of \mathbb{R}^d .

This perfect preservation of topological information is formalized by *topological embeddings*. Topological embeddings are continuous bijections which have a continuous inverse defined on their image. These are closely related to *homeomorphisms*; where by a homeomorphism on \mathbb{R}^d we mean a bijection $\phi \in C(\mathbb{R}^d, \mathbb{R}^d)$ having a continuous inverse.

Throughout this paper, we focus on the subset $\mathcal{H}(\mathbb{R}^d) \subseteq C(\mathbb{R}^d, \mathbb{R}^d)$ consisting of homeomorphisms ϕ which preserve any the orientation of any basis of \mathbb{R}^d . For example, no reflection in \mathbb{R}^2 belongs to $\mathcal{H}(\mathbb{R}^2)$ but the map $x \mapsto 2x$ does not. This class is key to our analysis as it is interconnected with property (P-ii). This is because the central result of Kirby (1969) characterizes $\mathcal{H}(\mathbb{R}^d)$ as exactly describing the homeomorphisms on \mathbb{R}^d that can be continuously deformed into the identity; i.e.: there is a $\phi_\alpha \in C([0, 1] \times \mathbb{R}^d, \mathbb{R}^d)$ satisfying:

$$\phi_1(x) = \phi(x), \quad \phi_0(x) = x, \quad \phi_\alpha \text{ is a homeomorphism for each } \alpha \in [0, 1]. \quad (2)$$

Following Adams (2004), we refer to the function ϕ_α as an *ambient-isotopy*.

In Edwards and Kirby (1971), it is shown that the homeomorphisms in $\mathcal{H}(\mathbb{R}^d)$ are characterized by their *fragmentation property*. This means that, given any $\phi \in \mathcal{H}(\mathbb{R}^d)$, $\delta > 0$, and any $\{x_n\}_{n \leq N}$ for which $[-M, M]^d \subseteq \cup_{n \leq N} B(x_n, \delta)$ then there necessarily exist some $\{\phi_n\}_{n \leq N} \subseteq \mathcal{H}([-M, M]^d)$ satisfying

$$\phi = \phi_N \circ \dots \circ \phi_1 \text{ and } \phi_n(x - x_n) = x - x_n \quad (\text{whenever } \|x\| \geq \delta). \quad (3)$$

Our interest in the fragmentation property is that it allows us to quantify the complexity of a homeomorphisms; which we rely on for our quantitative results.

2.2 The Space of Rotation Matrices

Key to our analysis are the higher-dimensional rotation matrices, which have recently been connected to DNNs in Bansal et al. (2018), Jia et al. (2019), and Lezcano-Casado and Martínez-Rubio (2019). These matrices R are precisely those for which the map $x \mapsto Rx$ does not flip any basis of \mathbb{R}^d and it preserves the distances between any two vectors. These matrices are characterized by:

$$SO(d) \triangleq \left\{ R \in \text{Mat}_{d \times d} : R^\top R = RR^\top = I_d \text{ and } \det(R) = 1 \right\};$$

where I_d is the $d \times d$ identity matrix on \mathbb{R}^d and $\text{Mat}_{d \times d}$ denotes the set of $d \times d$ matrices.

Following Knapp (2002), every $R \in SO(d)$ can be expressed as the *matrix exponential* $\exp(A)$ of a $d \times d$ -skew symmetric matrix A ; where $\exp(A) \triangleq \sum_{k=0}^{\infty} \frac{1}{k!} A^k$. Analogously to Lezcano-Casado and Martínez-Rubio (2019), we identify the vector space of $d \times d$ -skew-symmetric matrices, denoted by \mathfrak{so}_d , with the Euclidean space of the same dimension. This identification is realized via the bijection $\text{Skw} : \mathbb{R}^{d(d-1)/2} \rightarrow \mathfrak{so}_d$ defined by

$$(x_{1,2}, \dots, x_{1,d}, \dots, x_{d-1,d}) \rightarrow \begin{pmatrix} 0 & x_{1,2} & \dots & x_{1,d} \\ -x_{1,2} & & & \\ \dots & & \ddots & x_{d-1,d} \\ -x_{1,d} & & -x_{d-1,d} & 0 \end{pmatrix}.$$

Next, we describe the deep neural models which NEU trains.

2.3 Reconfiguration Networks

Recall that the DNN architecture, originating in McCulloch and Pitts (1943), is built by repeatedly composing the following type of elementary functions $x \mapsto \sigma \bullet (Ax + b)$, where A is a $d_i \times d_{i+1}$ matrix ($d_i, d_{i+1} \in \mathbb{N}_+$), $b \in \mathbb{R}^{d_{i+1}}$, $\sigma : \mathbb{R} \rightarrow \mathbb{R}$ is a non-linear function which is fixed across each feed-forward layer, and \bullet denotes component-wise composition.

NEU trains a variant of the DNN architecture whose layers are constrained between \mathbb{R}^d , but with the key difference being that the *connection matrix* A is replaced by a specific $SO(d)$ -valued function called a *reconfiguration unit*. These units allow the network's connection to depend, in a highly structured and non-constant way, on spatial data with the added flexibility of being able to only locally manipulate input data. A visualization of reconfiguration units is found in Figure 1.

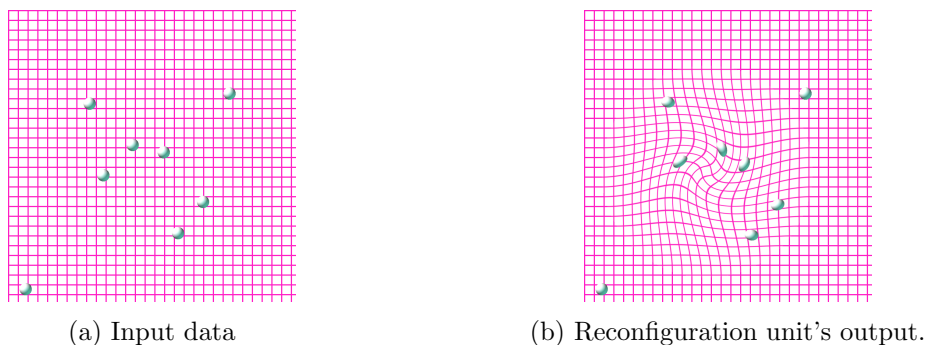


Figure 1: Visualization of reconfiguration unit's effect.

Definition 2.1 (Reconfiguration Unit) A reconfiguration unit is a matrix-valued function $A : \mathbb{R} \rightarrow \text{Mat}_{d \times d}$ with representation

$$\begin{aligned}
 A(x) &\triangleq \exp(M_1(\|x - c\|^2) + M_2(\|x - c\|^2)) \\
 M_1(z) &\triangleq \text{Skw}(W_{2,1} \circ \sigma_{\text{ReLU}} \bullet W_{1,1})(z) \\
 M_2(z) &\triangleq \text{Skw}(W_{2,2} \circ \sigma_{\text{ReLU}} \bullet W_{1,2})(z) \sqrt[8]{(z - \eta)(z + \eta)} I_{|z| < \eta},
 \end{aligned} \tag{4}$$

where, for $i = 0, 1$, $W_{1,i} : \mathbb{R} \rightarrow \mathbb{R}^{d(d-1)/2}$ and $W_{2,i} : \mathbb{R}^{d(d-1)/2} \rightarrow \mathbb{R}^{d(d-1)/2}$ are affine functions, $c \in \mathbb{R}^d$, $\eta > 0$, and where $\sigma_{\text{ReLU}}(x) \triangleq \max\{0, x\}$.

Remark 1 (Reconfiguration unit parameters) The map M_2 controls the local behaviour of A and the map M_1 controls its global behaviour. By setting $W_{2,2} = 0$, the reconfiguration unit A becomes the identity outside of the ball $\{z \in \mathbb{R}^d : \|z - c\| \leq \eta\}$.

We combine reconfiguration units, biases, and activation functions to build complex deep neural models. However, again unlike DNNs, we use an activation function, which is always a homeomorphism on \mathbb{R} . Analogously to He et al. (2015) and Ramachandran et al. (2018) we allow the activation function to depend on an additional parameter that can be used to turn the activation function into the identity map.

Definition 2.2 (Reconfiguration Network) A reconfiguration network is a function $\phi \in C(\mathbb{R}^d, \mathbb{R}^d)$ with representation $\phi(x) = \phi_J(x)$, where ϕ_J is defined iteratively via

$$\phi_{n+1}(x) = \sigma_{\alpha_n}(A_n(\phi_n(x))(\phi_n(x) - c_n)) + b_n \quad \phi_0(x) = x, \quad (5)$$

for every $x \in \mathbb{R}^d$, some $J \in \mathbb{N}_+$, some $\{c_n, b_n\}_{n=1}^N \in \mathbb{R}^d$, reconfiguration units $\{A_n\}_{n=1}^N$, and some $\{\alpha_n\}_{n=1}^N$ in $[0, \infty)$, where

$$\sigma_\alpha(x) \triangleq x + \tanh(\alpha x). \quad (6)$$

The set of all reconfiguration networks is denoted by $\Phi_{*,d}$. J is called the depth of ϕ .

3. Main Results

This section contains the paper's main theoretical contributions. We begin by outlining the structured approximation capabilities of the reconfiguration networks, before describing their implications for NEU. The section closes upon examining NEU's robustification of (L).

3.1 Universal Orientation-Preserving Homeomorphisms

We find that reconfiguration networks can approximate any homeomorphism in $\mathcal{H}(\mathbb{R}^d)$.

Theorem 2 (Reconfiguration Networks are Universal in $\mathcal{H}(\mathbb{R}^d)$) Let $d \in \mathbb{N}_+$ with $d \geq 1$, $\phi \in \mathcal{H}(\mathbb{R}^d)$, K be a non-empty compact subset of \mathbb{R}^d , and let $\epsilon > 0$. There exist $\phi^\epsilon \in \Phi_{*,d}$ satisfying

$$(i) \sup_{x \in K} \|\phi(x) - \phi^\epsilon(x)\| < \epsilon.$$

$$(ii) \phi^\epsilon \in \mathcal{H}(\mathbb{R}^d).$$

Theorem 2 is a *qualitative* universal approximation result for DNN in $\mathcal{H}(\mathbb{R}^d)$. However, analogously to Barron (1993) and Siegel and Xu (2020), by assuming some additional regularity of the homeomorphism being approximated, we may obtain a *quantitative* approximation result describing the complexity of the reconfiguration network required to approximate a target homeomorphism.

Analogously to Yarotsky (2018); Kratsios and Papon (2021) the complexity of a reconfiguration network is quantified by its *depth*. Since homeomorphisms are more complex and structured objects than simple continuous functions, our rates depend both on the target homeomorphism's modulus of continuity, as in Yarotsky (2018), and its best behaviour on a fragmentation; in the sense of (3).

Let $M, \delta > 0$ and let $\omega : [0, \infty) \rightarrow [0, \infty)$ be a modulus of continuity. We establish bounds for the class $\mathcal{H}_{M,\omega}^\delta(\mathbb{R}^d) \subseteq \mathcal{H}(\mathbb{R}^d)$ consisting of all homeomorphisms $\phi \in \mathcal{H}(\mathbb{R}^d)$ mapping $[-M, M]^d$ into itself satisfying a refinement of (3). Here, we additionally require that there is at-least one pair $\{x_n\}_{n \leq N}$ and $\{\varphi_n\}_{n \leq N}$ satisfying (3) which also satisfies:

$$\|\varphi_n(x) - x_n\| = \|x - x_n\| \quad (7)$$

$$\|\varphi_n(x - x_n) - \varphi_n(y - x_n)\| = \|y - x\| \quad (\text{if } \|x - x_n\| = \|y - x_n\|) \quad (8)$$

$$\|x_n - x_m\| > \frac{\delta}{2} \quad (\text{if } n \neq m, n, m \leq N) \quad (9)$$

$$\max_{n \leq N} \omega_{\varphi_n}(u) \leq \omega(u) \quad (\forall u \geq 0). \quad (10)$$

Theorem 3 (Quantitative Approximation Rates) Fix $\delta, M > 0$, $d \in \mathbb{N}_+$, and a modulus of continuity ω . For any $\phi \in \mathcal{H}_{M,\omega}^\delta(\mathbb{R}^d)$, any $\nu \in \mathbb{N}_+$ there is a $C \geq 0$, not depending on d, D, δ or ω , and a reconfiguration network $\hat{\phi}$ of depth at-most $\nu \left\lceil \frac{(8M)^d d^{d-1}}{\delta^d (d-1)} \right\rceil$, such that:

$$\sup_{\|x\| \leq M} \left\| \phi(x) - \hat{\phi}(x) \right\| \leq \omega_{\nu,\delta,M,d,\omega}; \quad (11)$$

where $\omega_{\nu,\delta,M,d,\omega} \triangleq \omega \left\lceil \frac{(8dM)^d}{\delta^d} \right\rceil$ and $\omega \left\lceil \frac{(8dM)^d}{\delta^d} \right\rceil$ is determined recursively by:

$$\omega_{n+1} = 2dMC\omega \left(\left(2^{-1\nu} \left\lceil \frac{\delta^d}{(8dM)^d} \right\rceil d(d-1) \right)^{\frac{-1}{d}} \right) + \omega(\omega_n), \quad \omega_0 = 0. \quad (12)$$

Homeomorphisms depending on arbitrary moduli of continuity, as in Theorem 3, may have arbitrarily poor behaviour. However, in many applications one has differentiability of the unknown homeomorphism ϕ and therefore ϕ is Lipschitz on $[-M, M]^d$ by the Mean-Value Theorem; i.e.: $\omega(t) = L|t|$ for some $L > 0$. In this case, (11) implies the following.

Example 1 (Simple Bounds for Lipschitz Homeomorphisms) In the case that $\phi \in \mathcal{H}_{M,L|\cdot|}^\delta$, then recursive upper-bound simplifies to $\omega_{\nu,\delta,M,d,\omega} = \nu^{\frac{-1}{d}} \left(2dMCL \frac{1-L}{1-L} \frac{\delta^{d^2} M}{1-L} \right)$.

Reconfiguration networks can memorize arbitrarily many input-output pairs without much guessing. Analogously to Jiang et al. (2009), we provide an upper bound on the reconfiguration network's depth, trained on the memorization task.

3.2 Memorization without Guessing

We quantify the size of the compact subset $K \subseteq \mathbb{R}^d$ on which the reconfiguration network guesses, by the number of d -dimensional balls of radius $\delta > 0$ required to cover K . This quantity is known as the δ -external covering number of K , (see (Mohri et al., 2018, Chapter 3.5)), and it is denoted by $N_\delta(K)$. Its advantage over the Lebesgue measure of K , is that it does not ignore sets of Lebesgue measure zero.³

Theorem 4 (Quantitative Memory Capacity Bounds with Guessing Control) Let $d \geq 2$, $N \in \mathbb{N}_+$, $M \in \mathbb{N}$, and $\{x_i\}_{i=1}^N$, $\{y_m\}_{m=1}^M$, and $\{z_i\}_{i=1}^N$ be sets of distinct points in \mathbb{R}^d for which $x_i \neq z_i$, for $i = 1, \dots, N$ and such that $y_m \notin \{x_i, z_i\}_{i=1}^N$ for each $m \leq M$. Define $\Delta \triangleq \frac{1}{2} \min \{2, \min \|u - v\| : u \neq v, u, v \in \{x_i, z_i\}_{i=1}^N \cup \{y_m\}_{m=1}^M\}$. For any $0 < \delta < \Delta$ there exists a reconfiguration network ϕ of depth J and a compact subset $K \subseteq \mathbb{R}^d$ satisfying:

- (i) $\phi(x_n) = z_n$ for each $n = 1, \dots, N$.
- (ii) $\phi(y_m) = y_m$ for every $m = 1, \dots, M$.

3. The two quantities are related, since the volume of a d -dimensional of radius $\delta > 0$ is $\frac{\delta^d \pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)}$; hence

$$N_\delta(K) \leq N \text{ implies } \mu(K) \leq \frac{N \delta^d \pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2}+1)}.$$

(iii) $\phi(x) = x$ for every $x \in \mathbb{R}^d - K$.

Furthermore, the following upper-bounds on J and K hold:

$$N_\delta(K) \leq J \leq \left\lceil \frac{N\pi}{2(\min\{2\delta, 1\})} \right\rceil. \quad (13)$$

In contrast, DNNs with analytic activation function always fail to have the memorization without guessing property (P-iv).

Proposition 5 (No Memorization without Guessing) *Let $d \geq 1$. If σ is analytic then every $f \in \mathcal{NN}_{d,d}^\sigma$ fails (P-iv).*

DNNs with ReLU networks are not included in Proposition 5. However, the difference between these architectures and reconfiguration networks is addressed later in the paper.

3.3 Universal Approximation via Topological Embeddings

By precomposing with an injective linear map, the universal approximation capabilities of reconfiguration networks can be extended to universal topological embeddings. In turn, post-composing with a linear map we can approximate any continuous function.

Theorem 6 (Topologically Regular Universal Approximation) *Fix $f \in C(\mathbb{R}^d, \mathbb{R}^D)$, a $(d+D) \times d$ -matrix A , and $M > 0$. There exists a reconfiguration network $\hat{\phi} \in \mathcal{H}(\mathbb{R}^{d+D})$ and a $D \times (d+D)$ -matrix B , such that $\mathcal{M}_{\hat{\phi}, M, \epsilon} \triangleq \phi \circ (I \oplus A) \cdot [-M, M]^d \subseteq \mathbb{R}^{d+D}$ satisfies:*

(i) **Embedding:** $\hat{\phi} \circ (I_d \oplus A \cdot)$ is a homeomorphism from $[-M, M]^d$ onto $\mathcal{M}_{\hat{\phi}, M, \epsilon}$ and it is an isometry when $\mathcal{M}_{\hat{\phi}, M, \epsilon}$ is equipped with the metric:

$$d_{\hat{\phi}, M, \epsilon}(z_1, z_2) \triangleq \|(I_d \oplus A)^\dagger \phi^{-1}(z_1) - (I_d \oplus A)^\dagger \phi^{-1}(z_2)\|,$$

where $(I_d \oplus A)^\dagger \triangleq ((I_d \oplus A)^\top (I_d \oplus A))^{-1} (I_d \oplus A)$ and is a left-inverse of $I_d \oplus A$.

(ii) **Regular Feature Space:** $\mathcal{M}_{\hat{\phi}, M, \epsilon}$ is a topological submanifold of \mathbb{R}^{d+D} with boundary,

(iii) **Sparsity:** $L(x) = Bx$ and B has exactly D non-zero entries and,

(iv) **Universal Approximation:** $\sup_{\|x\| \leq M} \|f(x) - B\hat{\phi}(x, Ax)\| < \epsilon$.

Let us compare the topological regular approximation results of this section with the popular approach of generating ϕ with a DNN with ReLU activation function. We frame our result in the setting of generalized-ReLU networks, as defined in Gribonval Rémi et al. (2021).

Proposition 7 *Let $\sigma_{r-\text{ReLU}}(x) = \max\{0, x^r\}$ and $0 < d$. If $\hat{f} \in \mathcal{NN}_{d,D}^{\sigma_{r-\text{ReLU}}}$ and \hat{f} is UAP-invariant, then for every affine function $W : \mathbb{R}^d \rightarrow \mathbb{R}^d$ the deep ReLU network:*

$$\tilde{f} = \hat{f} \circ \sigma_{r-\text{ReLU}} \bullet W,$$

is not UAP-invariant. In particular, it is not a topological embedding.

Proposition 7 highlights the main geometric difference between DNNs and reconfiguration networks. Namely, the latter always represents the input space as an embedded topological submanifold of the feature space \mathbb{R}^F whereas the former typically does not.

Remark 8 (Discussion: Comparison with Kernel Methods) *Kernel methods implicitly linearize functions in $C(\mathbb{R}^d, \mathbb{R}^D)$ by representing them within a high-dimension space. In contrast, Theorem 6 guarantees that reconfiguration networks can perform this linearizing in a $d + D$ -dimensional space.*

To see this, we consider a familiar example. Consider the kernel $K_\psi(x, y) \triangleq \sum_{n=1}^{\infty} \frac{1}{2^{2n}} x^n y^n$ on $[0, 1]$. This Kernel is universal, in the sense that:

$$\left\{ \sum_{n=0}^F \beta_n K_\psi(\cdot, y_n) : F \in \mathbb{N}, y_1, \dots, y_F \in [0, 1], \beta_1, \dots, \beta_F \in \mathbb{R} \right\},$$

is a dense subset of $C([0, 1], \mathbb{R})^4$. Moreover, as discussed in (Micchelli et al., 2006, Section 3) the feature map $\tilde{\phi}$ associated to K_ψ is maps any $x \in [0, 1]$ to the sequence $(\frac{x^n}{2^n})_{n=0}^{\infty} \in \ell^2(\mathbb{N})$. Thus, for any $f \in C([0, 1], \mathbb{R})$, there exists a linear map $\tilde{B} : \ell^2 \rightarrow \mathbb{R}$ such that

$$\max_{x \in [0, 1]} \|\tilde{B} \circ \tilde{\phi}(x) - f(x)\| < \epsilon. \tag{14}$$

The contrast between equation (14) and Theorem 6 (iv) is that, for every $\epsilon > 0$, together the maps B and $\hat{\phi}(I_d \oplus A)$ only require 2 dimensions in order to approximately linearize the function f ; whereas together \tilde{B} and $\tilde{\phi}$ need infinitely many dimensions to do so. Therefore, amongst other things, Theorem 6 can be interpreted as an explicit low-dimensional analogue of kernel methods which are implicit and high-dimensional.

3.4 Non-Euclidean Upgrading

We close the approximation-theoretic portion of this paper by relating our results on reconfiguration networks back to the learning problem (L) and to NEU. We present two results, each offering a different perspective on the improvement which can be gained by NEU. Both qualitative and quantitative results are provided. We operate under the following assumptions; typical in non-convex optimization (see Dal Maso (1993)).

Assumption 3.1 *We assume the following regularity of L and P defining Problem (L):*

- (i) $L : \mathbb{R}^D \times \mathbb{R}^D \times \mathbb{R}^d \rightarrow \mathbb{R}$ is continuous, and bounded-below.
- (ii) $P : C(\mathbb{R}^d, \mathbb{R}^D) \rightarrow \mathbb{R}$ is continuous, bounded-below, and coercive; i.e.: for every $t \in \mathbb{R}$ the sub-level set:

$$\left\{ P(f) \leq t : f \in C(\mathbb{R}^d, \mathbb{R}^D) \right\},$$

is compact in $C(\mathbb{R}^d, \mathbb{R}^D)$.

4. This is a direct consequence of (Micchelli et al., 2006, Theorem 7) and the Weierstraß Approximation Theorem.

Remark 9 (Why not lower semi-continuity of L and of P ?) *In general non-convex settings, lower semi-continuous (lsc) objective functions are typically considered instead of continuous ones. However, László (2017) shows that if our objective function is not continuous but is lsc and, in addition, if we do not optimize this objective function over the entire space but instead only optimize it over a proper dense subset (such as $\{f \circ \phi : f \in \mathcal{F}, \phi \in \Phi_{\star, d+D}\}$ or \mathcal{F} if it is a universal model class), then the global optimum is typically unobtainable. However, in (László, 2017, Corollary 3.4) the authors shows that this is never a theoretical issue for continuous objective functions.*

The following result says that given a family of models \mathcal{F} which is at-least able to express linear functions, there must be a UAP-invariant feature map in $\Phi_{\star, d+D}$ which "upgrades" \mathcal{F} until it approximately achieves the optimal value of the learning problem (L). Moreover, the representation learned by the reconfiguration network never needs to be in dimension above $d + D$. Furthermore, the representation produced by the reconfiguration networks is an embedded topological submanifold of the feature space $\mathbb{R}^F = \mathbb{R}^{d+D}$ and the feature map is a topological embedding.

Corollary 10 (Non-Euclidean Upgrading I) *Let $\{x_n\}_{n=1}^N$ be a subset of \mathbb{R}^d and suppose that L and P satisfy Assumption 3.1. Suppose also that $\mathcal{F} \subseteq C(\mathbb{R}^{d+D}, \mathbb{R}^D)$ contains all linear maps from \mathbb{R}^{d+D} to \mathbb{R}^D ; i.e.:*

$$\{x \mapsto Ax + b : b \in \mathbb{R}^D, A \in \text{Mat}_{D \times d+D}\}.$$

Then, for every $\epsilon, M > 0$ and every full-rank matrix $A \in \text{Mat}_{d, d+D}$, there exists some $\hat{f}^\epsilon \in \mathcal{F}$ and some $\hat{\phi}^\epsilon \in \Phi_{\star, d+D}$ and some $f^\epsilon \in \mathcal{F}$ such that:

(i) *The ϵ -optimality criterion holds:*

$$\begin{aligned} & \sum_{n=1}^N L\left(f(x_n), f^\epsilon \circ \hat{\phi}^\epsilon((I_d \oplus A)x_n), x_n\right) + P\left(f^\epsilon \circ \hat{\phi}^\epsilon((I_d \oplus A)\cdot)\right) \\ & < \epsilon + \inf_{g \in C(X, \mathbb{R}^D)} \sum_{n=1}^N L(f(x_n), g(x_n), x_n) + P(g(x_n)). \end{aligned} \tag{15}$$

(ii) *$\hat{\phi}^\epsilon \circ (I_d \oplus A)\cdot$ is a homeomorphism from $[-M, M]^d$ onto $\mathcal{M}_{\hat{\phi}^\epsilon, M, \epsilon}$ and it is an isometry when $\mathcal{M}_{\hat{\phi}^\epsilon, M, \epsilon}$ is equipped with the metric:*

$$d_{\hat{\phi}^\epsilon, M, \epsilon}(z_1, z_2) \triangleq \|(I_d \oplus A)^\dagger \hat{\phi}^{-1}(z_1) - (I_d \oplus A)^\dagger \hat{\phi}^{-1}(z_2)\|.$$

(iii) *$\mathcal{M}_{\hat{\phi}^\epsilon, M, \epsilon}$ is a topological submanifold of \mathbb{R}^{d+D} with boundary.*

Remark 11 *In particular, Corollary 10 (i), implies that the "upgraded model": $f^\epsilon \hat{\phi}^\epsilon((I_d \oplus A)\cdot)$ must achieve a lower value of the objective function of the training problem (L).*

The following result says that reconfiguration networks can modify models in \mathcal{F} to match the optimizer of Problem (L) at every given data-point while leaving almost all of their other input-output pairs unaltered. The result is quantitative in the number of points modified, and the external-covering number of the set of input-output pairs left unaltered.

Corollary 12 (Non-Euclidean Upgrading II) *Let $\{x_n\}_{n=1}^N$ belong to a compact subset X of \mathbb{R}^d and suppose that L and P satisfy Assumption 3.1. Suppose also that $\mathcal{F} \subseteq C(\mathbb{R}^{d+D}, \mathbb{R}^D)$ contains all linear maps from \mathbb{R}^{d+D} to \mathbb{R}^D ; i.e.:*

$$\{x \mapsto Ax + b : b \in \mathbb{R}^D, A \in \text{Mat}_{D \times d+D}\}.$$

Then, there exists an $f^* \in C(X, \mathbb{R}^D)$ satisfying

$$f^* \in \operatorname{argmin}_{g \in C(X, \mathbb{R}^D)} \sum_{n=1}^N L(f(x_n), g(x_n), x_n) + P(g).$$

Moreover, if $f^*(x_n) \neq f^*(x_m)$ whenever $n \neq m$ and $x_n \neq 0$ for $n \leq N$, then there exists some $\hat{f} \in \mathcal{F}$ such that for every

$$0 < \delta < \frac{1}{2} \min \{\|u - v\| : u \neq v \text{ and } u, v \in \{f^*(x_n), x_m\}_{n,m=1}^N\}, \quad (16)$$

there exists a reconfiguration network $\hat{\phi}^\delta \in \Phi_{*,d+D}$ satisfying:

- (i) $\hat{f} \circ \hat{\phi}^\delta(Ax_n) = f^*(x_n)$ for every $n \leq N$,
- (ii) $N_\delta \left(\left\{ (x, \hat{f}(Ax)) \in \mathbb{R}^{d+D} : (x, \hat{f}(Ax)) \neq \hat{\phi}^\delta(x, \hat{f}(Ax)) \right\} \right) \leq \left\lceil \frac{N\pi}{2(\min\{2\delta, 1\})} \right\rceil$,
- (iii) $\hat{\phi}^\delta$ has depth at-most $\left\lceil \frac{N\pi}{2(\min\{2\delta, 1\})} \right\rceil$.

Remark 13 *Corollary 12 (ii) highlights the need to quantify the smallness of the set K in Theorem 4 with its external covering number instead of its Lebesgue measure. This is because, the set being described in Corollary 12 (ii) is a d -dimensional subset of \mathbb{R}^{d+D} ; therefore, it is of Lebesgue measure 0. However, within the context of Corollary 12, this set is not negligible as it describes the set of input-output pairs which are transformed by the feature map $\hat{\phi}^\delta$.*

These two results show that NEU embeds a great deal of flexibility into any model class \mathcal{F} with a basic level of expressibility. The next section describes how to counter-balance this flexibility by robustifying the learning problem (L).

3.5 Robustification of Loss-Function for Improved Generalization

Fix $M > 0$, a non-empty training set $\{x_n\}_{n \leq N} \subseteq [-M, M]^d$, and a model \hat{f} for (L). We can evaluate its generalizability on $[-M, M]^d$ outside the training data $\{x_n\}_{n \leq N}$ by the gap in the error on the training data $\sum_{n \leq N} \frac{L(f(x_n), \hat{f}(x_n), x_n)}{N} + P(f)$ and the worst-case scenario error on $[-M, M]^d$. We define this gap by

$$\sup_{\|x\| \leq M} L(f(x_n), \hat{f}(x_n), x_n) + P(f) - \left[\frac{1}{N} \sum_{n \leq N} \frac{L(f(x_n), \hat{f}(x_n), x_n)}{N} + P(f) \right]. \quad (17)$$

A-priori it seems that, for any given \hat{f} , all the quantities in (17) are fixed by the learning problem. In fact, this is not the case here as we have implicitly made the assumption that the weight of each training data-point pulls equal weight on the left-hand side of (17).

Accordingly, we re-weight the training objective function to $\sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n) + P(f)$ with new weights $\{w_n\}_{n \leq N}$ in $[0, 1]$ summing to 1. Thus, we improve the generalizability of our model \hat{f} by extending (L) by coupling it with the following extension of (17)

$$\begin{aligned} & \min_{\hat{f} \in \mathcal{F}} \sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n) + P(\hat{f}) \\ \text{where: } & \operatorname{argmin}_{\substack{\sum_{n \leq N} w_n = 1 \\ w_n \in [0, 1]}} \sup_{\|x\| \leq M} L(f(x_n), \hat{f}(x_n), x_n) + P(f) \\ & - \left[\sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n) + P(f) \right]. \end{aligned} \quad (18)$$

The multi-function argmin is invariant under addition. Therefore, we may simplify the constraint in (18). Thus, we are interested in the following equivalent optimization problem:

$$\operatorname{argmax}_{\substack{\sum_{n \leq N} w_n = 1 \\ w_n \in [0, 1]}} \sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n). \quad (19)$$

The key advantage of (19) over (18) is that it is completely independent of the behaviour of our model's test-set performance quantified by $\sup_{\|x\| \leq M} L(f(x_n), \hat{f}(x_n), x_n) + P(f)$. Thus, any optimizer of (19) can be computed independently of any test-set information.

Nevertheless, problem (19) is generally ill-posed. In order to identify a good set of weights $\{w_n\}_{n \leq N}$, we interpret any set of weights $\{w_n\}_{n \leq N}$ as describing a discrete probability measure on $\{1, \dots, N\}$. Hence, by adding the following Kullback-Leibler divergence between the discrete probability measures implicitly defined by $\{w_n\}_{n \leq N}$ and the uniform probability measure on $\{1, \dots, N\}$ implicitly specified by the naive weighting scheme $\{\frac{1}{N}\}_{n \leq N}$ we obtain the following well-posed variant of (19), with hyper-parameter $\lambda > 0$

$$\operatorname{argmax}_{\substack{\sum_{n \leq N} w_n = 1 \\ w_n \in [0, 1]}} \sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n) - \lambda \sum_{n \leq N} w_n \log \left(\frac{w_n}{N} \right). \quad (20)$$

Theorem 14 (Optimal Robust Weights for (20)) *Let $\{x_n\}_{n \leq N}$ be a non-empty training dataset in \mathbb{R}^d and L be continuous. Then $\{w_n^{\lambda, \hat{f}}\}_{n=1}^N$ belongs to (20), where*

$$w_n^{\lambda, \hat{f}} \triangleq \frac{e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}{\sum_{n \leq N} e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}.$$

Moreover, the robust learning Problem (18) is equal to

$$\min_{\hat{f} \in \mathcal{F}} \sum_{n \leq N} \frac{e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)} L(f(x_n), \hat{f}(x_n), x_n)}{\sum_{n \leq N} e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}} + P(\hat{f}). \quad (21)$$

By the representation (18), the learning problem (21) necessarily yields better generalizability than its naive counterpart (L). This generalization improvement is quantified by the gap between (17) and the value of the constraint of (18). Given a model $\hat{f} \in \mathcal{F}$ and weights $\{w_n\}_{n \leq N}$ in $[0, 1]$ summing to 1 we define

$$\text{Gen}_M \left(\hat{f} \middle| \{w_n\}_{n \leq N} \right) \triangleq \sup_{\|x\| \leq M} L(f(x_n), \hat{f}(x_n), x_n) + P(f) - \left[\sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n) + P(f) \right].$$

Corollary 15 (NEU’s Loss-Function Modification Improves in Generalizability)

Let $x_1, \dots, x_N \in \mathbb{R}^d$, $\hat{f} \in \mathcal{F}$, and $M > 0$. Then $\text{Gen}_M \left(\hat{f} \middle| \left\{ \frac{1}{N} \right\}_{n \leq N} \right) - \text{Gen}_M \left(\hat{f} \middle| \{w_n^{\lambda, \hat{f}}\}_{n \leq N} \right)$ equals to:

$$\sum_{n \leq N} \left(\frac{e^{-\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}{\sum_{n \leq N} e^{-\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}} - \frac{1}{N} \right) L(f(x_n), \hat{f}(x_n), x_n) \geq 0.$$

We end this portion of the paper with the following observation.

Training Very Deep UAP-invariant Feature Maps

From the computational standpoint, property (P-i) can be used to to subdivide step 1 of Meta-Algorithm 1 into an incremental procedure, analogously to Bengio et al. (2007); Larochelle et al. (2009), allowing for the handling of extremely deep feature maps without negatively impacting the model’s UAP. This incremental procedure, summarized by sub-routine 2, views the feature map $\hat{\phi}$ Meta-Algorithm 1 step 1 as a composition $\hat{\phi} = \hat{\phi}_I \circ \dots \circ \hat{\phi}_1$ of deep reconfiguration networks $\{\hat{\phi}_i\}_{i \leq I}$ trained in a loop.

Sub-Routine 2: Incremental training of very deep UAP preserving feature maps.

input : Loss-function L , penalty function P ,
 Training and Validation Data $\{x_n\}_{n \leq N} = \{x_{n,t}\}_{n \leq N_t} \cup \{x_{n,v}\}_{n \leq N_v}$
 Number of blocks I of Feature Map, J depth per block
 Robustness Hyper-parameter $\lambda > 0$, $F \in \mathbb{N}_+$

output: NEU-Feature Map $\hat{\phi} \triangleq \hat{\phi}_I \circ \dots \circ \hat{\phi}_1$.

- 1 **for** $i \leq I$ **do**
- 2 $\hat{\phi} \in \underset{\phi \in \Phi_{*,F}}{\text{argmin}} \sum_{n \leq N} w_n^{*,\lambda} L \left(f(x_n), A\phi \circ \hat{\phi}_{i-1}(x_n) + b, x_n \right) + P(A\phi \circ \hat{\phi}_{i-1} + b)$;
- 3 $\phi_i = \hat{\phi} \circ \hat{\phi}_i$;
- 4 $\hat{A}_i, \hat{b}_i \in \underset{A,b}{\text{argmin}} \sum_{n \leq N_t} w_n^{*,\lambda} L \left(f(x_{n,t}), A\hat{\phi}_i(x_{n,t}) + b, x_{n,t} \right) + P(A\hat{\phi}_i + b)$
- 5 **end**
- 6 $\hat{I} \in \underset{i \leq I}{\text{argmin}} \sum_{n \leq N_v} w_n^{*,\lambda} L \left(f(x_{n,v}), \hat{A}_i \hat{\phi}_i(x_{n,v}) + \hat{b}_i, x_{n,v} \right) + P(\hat{A}_i \hat{\phi}_i + \hat{b}_i)$

Sub-routine 2 would typically yield sub-optimal $\hat{\phi}$. However, by replacing step 1 of Meta-Algorithm 1 with sub-routine 2, we can train feature maps that are too deep to train on a

single machine while also guaranteeing that these maps have property (P-i). In contrast, Proposition 7 guarantees that DNNs with the ReLU activation function cannot be trained analogously without disrupting the DNN’s UAP.

4. Numerical Evaluation of NEU-OLS and NEU-PCA

Next, we evaluate the performance of NEU across various learning tasks. First, we investigate the performance of NEU in the chaotic environment provided by real-world financial data. Then, we stress test NEU’s behaviour within the controlled environment provided by simulation studies. The Tensorflow (v.2.4.1) code and data-sets for our implementations is available online at [?](#).

4.1 Financial Data Analysis

The performance of the NEU meta-algorithm will be investigated both on regression and dimension reduction tasks using financial data. We begin with the regression problem of constructing a stock return replication and then move to non-Euclidean yield-curve analysis.

4.1.1 REGRESSION ANALYSIS: APPLE STOCK TRACKER

Predicting the relationship between the price of a set of assets is central to many trading strategies. For example, strategies that rely on illiquid assets may create a portfolio comprised entirely of liquid assets that tracks the illiquid asset’s movements. In this example the technique is demonstrated using liquid stocks for both the target and the tracking portfolio so we can better evaluate performance, which would be more difficult with illiquid assets due to missing contemporaneous prices for the illiquid target.

We consider Apple’s stock as the target asset while the tracking portfolio is comprised of IBM, Google, Cisco Systems, Microsoft, Acacia Communications, NXP Semiconductors NV, Qualcomm, Analog Devices, Glu Mobile, Jabil, Micron, and STMicroelectronics NV. Thus, the tracking portfolio is comprised of the stock of major companies in the same industry and Apple’s supply chain (see Seth (2018) and Stoller (2018)).

We build a tracking portfolio using various linear, non-linear, and discontinuous regression model classes. These include Elastic Net regularization of Zou and Hastie (2005), generalizing the LASSO regressor of Tibshirani (1996) and Tychonov regularization of Tikhonov (1963) (Elastic Net), kernel ridge regression (Kernel), gradient boosted random forests (GBRF), and a DNN with ReLU activation function. Each of the hyper-parameters is selected by cross-validations and randomized search from a large grid, hyper-parameters include the choice of kernel. The NEU version of these models is also used considered as a general evaluation of the improvement capabilities of NEU.

We consider 2 years of closing stock prices, ending on September 25th 2020, to compute the regression weights. The models are trained on the first 75% of the data and the remaining 25% is used to evaluate the out-of-sample predictive performance of the trained models, and is illustrated in Figure 2.




Figure 2: Regression models' 150 day-ahead out-of-sample predictions.

Train	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-ENET	-0.057708	6.552474e-09	0.058055	0.466568	0.455442
ENET	-0.054082	3.295653e-18	0.054309	0.452092	0.413066
NEU-GBRF	-0.069073	0.000000e+00	0.068067	0.525591	0.647727
GBRF	-0.066513	2.966087e-17	0.066309	0.488983	0.622644
NEU-kRidge	-0.005229	-8.935916e-05	0.005009	0.044634	0.003772
kRidge	-0.038941	-5.949302e-04	0.037408	0.326579	0.200016
NEU-OLS	-0.046744	5.180375e-03	0.055844	0.424678	0.361532
NEU-DNN	-0.026579	2.685197e-02	0.081460	0.439256	0.411793
DNN	-0.021030	3.271767e-02	0.088175	0.445274	0.427985

Table 1: Train performance of fixed time-horizon problem.

Test	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-ENET	-0.159483	0.407651	1.004434	1.448866	5.369902
ENET	-0.166303	0.364576	0.929381	1.368901	4.903049
NEU-GBRF	-0.096924	0.492442	1.096471	1.587017	5.989319
GBRF	-0.168434	0.427108	1.052304	1.591232	5.971403
NEU-kRidge	-0.150280	0.411124	1.026502	1.429903	5.515611
kRidge	-0.169045	0.380833	0.958410	1.427682	5.110043
NEU-OLS	-0.167247	0.344157	0.929869	1.309060	4.816255
NEU-DNN	-0.062913	0.455043	1.034286	1.407882	4.948565
DNN	-0.079580	0.462161	1.049870	1.444588	5.112289

Table 2: Test performance of fixed time-horizon problem.

NEU-OLS and the DNN model both outperform each of the linear models. However, NEU-OLS’s out-performance of the DNN model is a joint effort between its representation properties and its robustness properties. The in-sample advantage can be explained by NEU’s memory capacity, as demonstrated by Theorem 4, and its expressibility improvement, as demonstrated by Theorem 6. The out-of-sample performance, described in Table 2, has also benefited from the robustness of the NEU weights, described by Theorem 14.

4.1.2 DIMENSION REDUCTION: US-BOND YIELD CURVE

Principal component analysis (PCA) is commonly used in finance to reduce the effective dimension of data and a classical application is for representing the yield curve corresponding to zero-coupon bond prices. Denote by $B(t, T)$ the price at time t of a zero-coupon bond that pays the face value, by assumption \$1, at maturity T . The bond’s yield, denoted $y(t, T)$, is the continuously compounded interest rate at which an investment of $B(t, T)$ would accumulate to the face value. That is, $y(t, T) \triangleq -\frac{\ln(B(t, T))}{T-t}$. The yield curve is the map of a bond’s yields as a function of time to maturity, $(T - t) \rightarrow y(t, T)$. Since the bond prices for all maturity dates $T \geq t$ are not observed it is an important problem for a variety of financial applications to construct the curve using the available observed bond prices at a given time. We benchmark NEU against auto-encoders (AE) with bottleneck dimension equal to the number of principal components (or factors) and against kernel-PCA (kPCA), two popular non-linear alternatives to classical PCA. NEU variants of both these methods are also considered.

The daily bond data considered in this example consists of 6385 consecutive instances of stripped US government bond prices between June 8th 1990 to April 2^{rst} 2017. Each instance records the value of zero-coupon bonds with 1, 3, and 6 month, and 1, 2, 3, 5, 7, 10, 20, 30 year maturities.

We shall consider the performance of PCA, kPCA, NEU-PCA, NEU-kPCA, a deep auto-encoder (AE), and NEU-AE. The test set consists on instance ahead yield curves, and thus it measures the robustness to the dimension reduced yield-curves factor models to market movements.

Figures 3-5 show that NEU-AE’s performs best from all the proposed models both in and out-of-sample when using one, two, and three factors. NEU-PCA only becomes

competitive, both on the training and testing sets, when three factors are utilised. This observation highlights the importance of NEU’s UAP-invariance property (P-i) as, in this case, NEU is able to maintain and improve the expressiveness of the auto-encoder model.

	Test-MAE	Test-MSE	Train-MAE	Train-MSE
NEU-PCA	3.214483	15.518402	3.406530	16.814745
PCA	3.218449	15.302093	2.972712	13.817181
NEU-AE	2.715108	11.043391	3.010322	13.117202
AE	3.172362	14.792091	2.950363	13.495282
NEU-kPCA	3.554145	18.568798	3.687914	19.654074
kPCA	3.253660	15.626541	3.067545	14.513802

Table 3: Performance of reconstructed factor models - 1 Factor.

	Test-MAE	Test-MSE	Train-MAE	Train-MSE
NEU-PCA	3.189997	14.913756	3.227090	15.089537
PCA	3.195051	15.143381	2.955438	13.675315
NEU-AE	2.623852	10.425360	2.685428	10.647568
AE	2.827589	11.930328	2.703830	11.108966
NEU-kPCA	3.842316	21.080980	3.697985	20.525725
kPCA	3.252147	15.614535	3.064285	14.491074

Table 4: Performance of reconstructed factor models - 2 Factors.

	Test-MAE	Test-MSE	Train-MAE	Train-MSE
NEU-PCA	3.241300	14.439484	2.755066	11.301304
PCA	3.197890	15.176116	2.933380	13.557278
NEU-AE	3.093146	13.579660	2.613067	10.845548
AE	3.162107	14.688698	2.940275	13.391631
NEU-kPCA	3.883233	20.833542	3.464134	17.237757
kPCA	3.252588	15.178260	3.064387	14.493373

Table 5: Performance of reconstructed factor models - 3 Factors.

In each case, NEU-PCA and NEU-AE reconstructs the yield curve more accurately from a small number of learned driving factors. We find that the in and out of sample explanatory capabilities of NEU-PCA surpass even the auto-encoder. As expected, NEU-AE offers the best performance amongst all the models, however, the advantages over NEU-PCA is nevertheless marginal. As with the regression tasks, the kPCA’s rigid feature map negatively interacts with NEU’s feature map causing instability.

Remark 16 (Clashing NEU Features and Kernel Features) *This last point is a recurrent theme throughout our experiments; namely that the kernel methods such as $kPCA$ and $kRidge$'s features tend to clash with the features learned by NEU. At times, they harmonize and the Non-Euclidean Upgraded kernel model offers astounding performance, however, at other times the performance deteriorates. This unstable behaviour is not observed in the other non-Euclidean upgraded methods and this is because the other methods either do not impose any additional features (such as OLS, PCA, or Elastic Net) or are flexible enough to blend their feature representation with NEU's (as for GBRF, AE, or DNN).*

4.2 Simulated Experiments

Next, we unpack and understand the detailed behaviour of NEU in the controlled environment offered by simulation studies. We consider a series of regression problems. In each situation, the data is generated using to the non-linear regression model with additive and multiplicative noises

$$y = U_\delta m(x) + \sigma Z, \quad (22)$$

where $Z \sim N(0, 1)$ and $U_\delta \sim U(1-\delta, 1+\delta)$, $0 < \delta < 1$, $\sigma > 0$, and m is a non-linear function. The multiplicative noise U_δ encapsulates model misspecification as it discontinuously (in x) distorts the shape of the unknown function m , and the additive noise σZ quantifies the noise distorting the signal, as in classical regression problem formulations.

We consider four challenging non-linear functions, each exhibiting a distinct pathology. The first, is comprised of several distinct local sub-patterns. The second exhibits aperiodic oscillations. The third, is split by a sharp jump discontinuity. The last pattern is highly discontinuous and we use it to evaluate each model's ability to discern between a sharp irregular signal and varying levels of noise.

The NEU-OLS and NEU-DNN models will be benchmarked against three standard non-parametric regression algorithms, penalized smoothing splines regression (p -splines), locally weighted scatterplot smoothing (LOESS), kernel ridge regression (Ker-Ridge), and feed-forward artificial neural networks (DNN). Other than the DNNs which were discussed thoroughly in the paper's introductory section, we review the benchmark models here.

In each of our experiments, we visualize the feature representation learned by NEU by plotting each of the coordinates of $\phi(x)$. These plots are given in Figures 3, 5, and 7, respectively for each experiment. Essentially, these can be interpreted as the *features* learned by NEU, which are then fed into the upgraded model. In particular, when the model is linear, the target function is approximately expressible as a linear combination of these features.

We see that the target function is reflected by each of the feature maps learned by NEU. For example, in the first implementation, NEU's feature representation illustrated in Figure 5 has a dramatic change at the precise point where the two sub-patterns deviate from one another. In the second experiment, NEU's produces a feature map, illustrated by Figure 3, whose coordinates represent oscillations happening at different rates; these are then combined by the linear model being upgraded to produce the correct osculating pattern. In the final experiment, NEU's features are illustrated in Figure 7, and draw out two distinct and relatively flat heaps. This reflects the sharp discontinuity separating the two otherwise constant parts of the target function.

For each simulation, 10^4 observations are generated on the interval $[-1, 1]$; the data is then normalized to the unit square for uniformity between the three examples. The models' tuning-parameters are then estimated by cross-validation.

4.2.1 APERIODIC OSCILLATIONS

We begin by evaluating each model's ability to handle aperiodic oscillations. To this end, we simulate from the unknown function $m_2(x) = \cos(e^{2+x})$.




Figure 3: NEU's features for: $\min(e^{\frac{-1}{(1+x)^2}}, x + \cos(x))$.

Test	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-OLS	-0.021313	-2.570136e-09	0.020961	0.435468	0.292991
Smoothing Splines	-0.006834	1.428117e-02	0.035403	0.435920	0.295227
LOESS	-0.012301	1.862346e-02	0.048980	0.630788	0.613180
ENET	-0.034831	-1.136868e-17	0.035714	0.740855	0.805618
NEU-GBRF	-0.023963	-5.684342e-17	0.024021	0.494789	0.382412
GBRF	-0.024555	-5.684342e-18	0.024455	0.499496	0.387694
NEU-kRidge	-0.020963	-2.570866e-04	0.020747	0.429108	0.284365
kRidge	-0.020745	-3.927500e-05	0.020682	0.434514	0.291221
NEU-DNN	0.004625	2.622699e-02	0.047207	0.435748	0.296543
DNN	-0.021901	-1.453060e-04	0.021675	0.441955	0.304498

Table 6: A-Periodic oscillations - $m_2(x) = \cos(e^{2+x})$: $\sigma = \delta = 0.5$

Figure 6 highlights the clash between the rigid structure imposed by the Kernel regression's implicit feature map and NEU's feature map. Since NEU's feature map is designed


Train	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-OLS	-0.007318	-0.005953	-0.004600	0.042497	0.003580
Smoothing Splines	-0.016907	-0.006644	0.002417	0.070279	0.180917
LOESS	0.016649	0.029347	0.042376	0.467625	0.321612
ENET	-0.007673	0.009254	0.025702	0.656024	0.526628
NEU-GBRF	0.000899	0.008715	0.016585	0.308246	0.120689
GBRF	-0.004417	0.003626	0.011627	0.313712	0.126322
NEU-kRidge	-0.008430	-0.006731	-0.005058	0.056449	0.005649
kRidge	-0.007045	-0.005678	-0.004326	0.046354	0.003695
DNN	-0.011299	-0.008462	-0.005682	0.102038	0.015454
NEU-DNN	0.016808	0.018696	0.020518	0.055287	0.007063

Table 7: A-Periodic Oscillations - $m_2(x) = \cos(e^{2+x})$: $\sigma = \delta = 0.5$


for models that are either linear or can efficiently approximate linear maps, then kernel regression’s feature map can, and in this case, it does, interfere with the representation learned by NEU. However, as is also reflected in Table 8, this only happens with the Kernel regression method and not with the GBRF, linear regression, or DNN methods.

4.2.2 FUNCTIONS WITH LOCAL BEHAVIOUR


Next, we compare each model’s abilities to learn from functions determined by several, exclusively local, sub-patterns. Thus, the unknown function m of (22) is taken to be $m(x) \triangleq \min(e^{-\frac{1}{(x+1)^2}}, x + \cos(x))$. The underlying pattern is therefore generated from two distinct sub-patterns $e^{-\frac{1}{(x+1)^2}}$ and $x + \cos(x)$, with the change between the two occurring every time the condition $e^{-\frac{1}{(x+1)^2}} < x + \cos(x)$ either holds or fails.




(a) NEU-DNN vs. DNN



(b) NEU-GBRF vs. GBRF



(c) NEU-OLS vs. OLS



(d) NEU-kRidge vs. Kernel

Figure 4: Error Distribution Comparisons for: $\cos(e^{2+x})$.




Figure 5: NEU’s features for: $\min(e^{\frac{-1}{(1+x)^2}}, x + \cos(x))$.

Test	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-OLS	-0.019358	6.980434e-04	0.021225	0.418056	0.266197
Smoothing Splines	-0.020310	4.915925e-11	0.019642	0.416499	0.265408
LOESS	-0.017103	3.094336e-03	0.023561	0.418556	0.268401
ENET	-0.020520	4.760636e-17	0.020322	0.419717	0.269394
NEU-GBRF	-0.020503	-1.705303e-16	0.020621	0.430601	0.284261
GBRF	-0.020795	7.958079e-17	0.020671	0.432727	0.286976
NEU-kRidge	-0.020580	1.588384e-05	0.020460	0.418158	0.267961
kRidge	-0.019606	1.823117e-06	0.020202	0.418136	0.267707
NEU-DNN	0.001613	2.200638e-02	0.041955	0.418093	0.269347
DNN	-0.012107	8.756267e-03	0.029376	0.418196	0.268941

Table 8: Non-Local function - $\min(e^{\frac{-1}{(1+x)^2}}, x + \cos(x))$: $\delta = \sigma = 0.5$.

Figure 8 show that NEU-OLS and NEU-DNN still offer the best out-of-sample performance amongst the DNN, LOESS, ENET, GBRF, kRidge, and DNN Models. However, this implementation suggests that smoothing splines may are better suited to locally-determined target functions. This is not surprising since NEU performs any localization after representing the pattern in a higher-dimensional space, whereas smoothing splines can locally approximate any function directly.

Figure 9 shows that in-sample, NEU-OLS offers provides a better fit than the Smoothing Splines, LOESS, ENET, GBRF, and the DNN model. However, kRidge seems best suited to the in-sample fitting of this type of pattern. We also note that, though the UAP-invariance property (P-i) guarantees that NEU-DNN is universal, since DNN was, Figure 9

Train	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-OLS	-0.003219	-0.001356	0.000763	0.029811	0.007868
Smoothing Splines	-0.005625	-0.004594	-0.003564	0.035477	0.002062
LOESS	-0.003280	-0.001950	-0.000587	0.037035	0.003622
ENET	-0.007302	-0.005885	-0.004524	0.047878	0.003902
NEU-GBRF	-0.008818	-0.004893	-0.000868	0.151281	0.031597
GBRF	-0.008343	-0.004443	-0.000530	0.148070	0.029684
NEU-kRidge	-0.005622	-0.004717	-0.003816	0.026037	0.001607
kRidge	-0.004352	-0.003566	-0.002819	0.024096	0.001135
NEU-DNN	0.014463	0.015718	0.016956	0.040284	0.003335
DNN	0.001173	0.002462	0.003720	0.035320	0.003244

Table 9: Non-Local function - $\min(e^{\frac{-1}{(1+x)^2}}, x + \cos(x)) : \delta = \sigma = 0.5$.


shows that at-times NEU-OLS can fit finite data-sets better than NEU-DNN does. This shows that though the DNN and the reconfiguration networks have arbitrarily large memory capacities, these two may at-times unfavorably interact since both memorize input-output pairs differently (which we see by comparing the proofs of Theorem 4 and the central result of Vershynin (2020b)). This interaction effect is less likely with NEU-OLS than NEU-DNN, since the any memorization only occurs in final layer in the in the former case.

Figure 6 reflect the bias reduction obtained by NEU’s feature map, which is most significant when applied to the OLS and DNN models. Here, NEU showcases the benefit of it being able to only locally modify a pattern, which is especially important in this case since two unrelated local sub-patterns determine m .


Tables 9 and 8 show that the NEU models achieve an improved performance both in-sample and out-of sample over their classical variants. The feature maps were only trained once for the OLS model and then used in the remaining models. Thus, the non-linear feature presentation must be correct as it transfers its improvement to each of the benchmarked regression models. We see that NEU-OLS offers the best accuracy, and most stable in-to-out of sample performance. This is likely due to the learned linearizing feature map not having to conflict with any other assumed feature map, as is the case for NEU-kRidge, NEU-DNN, and NEU-GBRF.

4.2.3 JUMP DISCONTINUITIES


The last simulation experiment explores a situation with discontinuities that is outside the scope of the standard p-spline, LOESS, and DNN methods. However, the NEU-OLS is able to perform well when the data exhibits these jump discontinuities. The function m in equation (22) is assumed to be $m(x) = I_{(-\infty, \frac{1}{2}]}(x)$. As reflected by Figure 7, NEU behaves the same when the unknown noisy function being approximated has a jump discontinuity as when distinct, locally-determined functions determined it; as in Figure 6. In both situations, NEU can learn a feature map that is determined only by local data, so it implicitly separates the behavior of the model on each part of the jump discontinuity, just as it did on the different sub-patterns in Figure 6.




(a) NEU-DNN vs. DNN



(b) NEU-GBRF vs. GBRF



(c) NEU-OLS vs. OLS



(d) NEU-kRidge vs. Kernel

Figure 6: Error Distribution Comparisons for: $\min(e^{\frac{-1}{(1+x)^2}}, x + \cos(x))$.





Figure 7: NEU’s features for: $I_{(-\infty,5]}$.

Test	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-OLS	-0.007420	-1.978159e-03	0.003309	0.219500	0.075624
Smoothing Splines	-0.005347	-1.902232e-12	0.005499	0.221883	0.077588
LOESS	-0.003747	2.843149e-03	0.009276	0.261553	0.109594
ENET	-0.007296	-5.968559e-17	0.007283	0.293706	0.134424
NEU-GBRF	-0.007103	3.467449e-16	0.007091	0.292915	0.128904
GBRF	-0.006864	3.240075e-16	0.007085	0.292915	0.128904
NEU-kRidge	-0.004780	4.746673e-04	0.005816	0.217793	0.074313
kRidge	-0.005810	1.327000e-06	0.005720	0.237115	0.088674
NEU-DNN	-0.017917	-1.253940e-02	-0.007353	0.218370	0.074929
DNN	0.004379	9.808546e-03	0.015200	0.218675	0.075156


Table 10: Discontinuous Function - $m(x) = I_{(-\infty, \frac{1}{2}]}(x)$. : $\delta = \epsilon = 0.5$

The presented simulations studies, highlight the strengths and weaknesses of NEU. In nearly every case, the Non-Euclidean Upgraded model outperforms its classical counterpart. However, we typically find that the more linear the original model, the more reliable NEU’s performance will be. This is likely due to conflicts between the assumed feature map, especially in NEU-kRidge, and the feature map being learned. This is because, assuming a feature map forces NEU to simultaneously learn the target while undoing any feature map misspecification.


Each of the simulated pathological regression problems, illustrated in this section showed that NEU is not only capable of out-performing each of the benchmark models regardless of how badly behaved the unknown target function is. Moreover, this performance improvement was maintained in the face of high amounts of multiplicative and additive noise.




(a) NEU-DNN vs. DNN



(b) NEU-GBRF vs. GBRF



(c) NEU-OLS vs. OLS



(d) NEU-kRidge vs. Kernel

Figure 8: Error Distribution Comparisons for: $\min(e^{\frac{-1}{(1+x)^2}}, x + \cos(x))$.

Train	Er. 95L	Er. Mean	Er. 95U	MAE	MSE
NEU-OLS	0.002959	0.003741	0.004552	0.011879	0.001689
Smoothing Splines	0.001218	0.002504	0.003796	0.022726	0.004295
LOESS	0.005587	0.009368	0.013222	0.135893	0.038387
ENET	0.002787	0.007589	0.012474	0.210154	0.063623
NEU-kRidge	0.001912	0.002406	0.002911	0.011165	0.000659
kRidge	-0.003193	-0.000655	0.001772	0.098367	0.015523
NEU-GBRF	-0.001760	0.002934	0.007482	0.238186	0.056756
GBRF	-0.001773	0.002934	0.007643	0.238186	0.056756
NEU-DNN	-0.010775	-0.010128	-0.009464	0.021392	0.001207
DNN	0.011372	0.012157	0.012948	0.019651	0.001737

Table 11: Discontinuous Function - $m(x) = I_{(-\infty, \frac{1}{2}]}(x)$. : $\delta = \epsilon = 0.5$.

These results mirror our theoretical findings and support the hypothesis that the performance improvement observed by using NEU in the context of quantitative finance, was not a singular instance but rather part of a general theoretically-founded trend.

5. Conclusion

This paper introduced the first generic algorithmic procedure for learning any feature map with the invariance properties (P-i) and (P-ii) while simultaneously guaranteeing the performance enhancement of properties (P-iii) and (P-v). From the perspective of Kernel methods, NEU is a universal procedure for learning a low-dimensional generic feature map with many desirable properties. From the standpoint of geometric deep learning, NEU also provides an answer to the recent research problem initiated in Kratsios and Bilokopytov (2020) and in Kratsios and Papon (2021), of how to generically learn optimal UAP-invariant feature maps from \mathbb{R}^d in cases where a UAP-invariant feature map is not already explicitly provided.

From the manifold learning perspective, reconfiguration networks are the first provably universal and computationally tractable class of topological embeddings. As a meta-algorithm, NEU can generically learn the optimal linearizing pre-processing step to nearly any model class \mathcal{F} , provided that \mathcal{F} at-least contains all linear maps. NEU’s theoretical properties were also supported experimentally.

NEU successfully introduced tools from geometric deep learning into financial data-analysis. NEU was found to outperform the current leading machine learning methods for non-parametric dimension reduction of yield curves and produce a competitive performance in the non-parametric stock-returns replication problem.

Acknowledgments

This research was supported by the ETH Zürich Foundation and the Natural Sciences and Engineering Research Council of Canada (NSERC). The authors thank Alina Stancu (Con-

cordia University) for helpful discussions, Josef Teichmann and the entire working group at ETH for their helpful feedback, and Behnoosh Zamanlooy for the Python guidance.

References

- Colin C. Adams. *The knot book*. American Mathematical Society, Providence, RI, 2004. ISBN 0-8218-3678-1. An elementary introduction to the mathematical theory of knots, Revised reprint of the 1994 original.
- Richard D. Anderson. On topological infinite deficiency. *Michigan Math. J.*, 14:365–383, 1967. ISSN 0026-2285.
- Andreas Argyriou, Charles A. Micchelli, and Massimiliano Pontil. When is there a representer theorem? Vector versus matrix regularizers. *J. Mach. Learn. Res.*, 10:2507–2529, 2009.
- Nachman Aronszajn. Theory of reproducing kernels. *Trans. Amer. Math. Soc.*, 68:337–404, 1950.
- Nitin Bansal, Xiaohan Chen, and Zhangyang Wang. Can we gain more from orthogonality regularizations in training deep networks? *Advances in Neural Information Processing Systems*, 31:4261–4271, 2018.
- Andrew R Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.
- Yoshua Bengio, Pascal Lamblin, Dan Popovici, and Hugo Larochelle. Greedy layer-wise training of deep networks. In *Advances in neural information processing systems*, pages 153–160, 2007.
- Helmut Bölcskei, Philipp Grohs, Gitta Kutyniok, and Philipp Petersen. Optimal approximation with sparsely connected deep neural networks. *SIAM J. Math. Data Sci.*, 1(1): 8–45, 2019.
- Bernhard E Boser, Isabelle M Guyon, and Vladimir N Vapnik. A training algorithm for optimal margin classifiers. In *Proceedings of the fifth annual workshop on Computational learning theory*, pages 144–152, 1992.
- L. E. J. Brouwer. Beweis der invarianz des n -dimensionalen gebiets. *Math. Ann.*, 71(3): 305–313, 1911. ISSN 0025-5831.
- G. Cybenko. Approximation by superpositions of a sigmoidal function. *Math. Control Signals Systems*, 2(4):303–314, 1989.
- Gianni Dal Maso. *An introduction to Γ -convergence*, volume 8 of *Progress in Nonlinear Differential Equations and their Applications*. Birkhäuser Boston, Inc., Boston, MA, 1993.
- Robert D. Edwards and Robion C. Kirby. Deformations of spaces of imbeddings. *Ann. of Math. (2)*, 93:63–88, 1971. ISSN 0003-486X. doi: 10.2307/1970753.

- Paul Embrechts and Marius Hofert. A note on generalized inverses. *Mathematical Methods of Operations Research*, 77(3):423–432, 2013.
- P. Thomas Fletcher. Geodesic regression and the theory of least squares on Riemannian manifolds. *Int. J. Comput. Vis.*, 105(2):171–185, 2013.
- Xavier Glorot, Antoine Bordes, and Yoshua Bengio. Deep sparse rectifier neural networks. In *Proceedings of the fourteenth international conference on artificial intelligence and statistics*, pages 315–323, 2011.
- Gribonval Rémi, Kutyniok Gitta, Nielsen Morten, and Voigtlaender Felix. Approximation spaces of deep neural networks. *Constructive Approximation*, 05 2021.
- Richard HR Hahnloser, Rahul Sarpeshkar, Misha A Mahowald, Rodney J Douglas, and H Sebastian Seung. Digital selection and analogue amplification coexist in a cortex-inspired silicon circuit. *Nature*, 405(6789):947–951, 2000.
- Moritz Hardt and Tengyu Ma. Identity matters in deep learning. *arXiv preprint arXiv:1611.04231*, 2016.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Delving deep into rectifiers: Surpassing human-level performance on imagenet classification. In *Proceedings of the IEEE international conference on computer vision*, pages 1026–1034, 2015.
- Dan Hendrycks and Kevin Gimpel. Gaussian error linear units (GELUs). *arXiv preprint arXiv:1606.08415*, 2016.
- Rodrigo Hernandez-Gutierrez. Countable dense homogeneity of function spaces. *Topology Proc.*, 56:125–146, 2020. ISSN 0146-4124.
- Heiko Hoffmann. On the continuity of the inverses of strictly monotonic functions. *Irish Math. Soc. Bull.*, 75(1):45–57, 2015.
- Kurt Hornik, Maxwell Stinchcombe, Halbert White, et al. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*, 2015.
- Kui Jia, Shuai Li, Yuxin Wen, Tongliang Liu, and Dacheng Tao. Orthogonal deep neural networks. *IEEE transactions on pattern analysis and machine intelligence*, 2019.
- N. Jiang, Z. Zhang, J. Wang, and X. Ma. The upper bound on the number of hidden neurons in multi-valued multi-threshold neural networks. In *2009 International Workshop on Intelligent Systems and Applications*, pages 1–4, 2009. doi: 10.1109/IWISA.2009.5073217.
- Motonobu Kanagawa, Bharath K. Sriperumbudur, and Kenji Fukumizu. Convergence analysis of deterministic kernel-based quadrature rules in misspecified settings. *Found. Comput. Math.*, 20(1):155–194, 2020.

- Patrick Kidger and Terry Lyons. Universal approximation with deep narrow networks. In *Conference on Learning Theory*, pages 2306–2327, 2020.
- Robion C. Kirby. Stable homeomorphisms and the annulus conjecture. *Ann. of Math. (2)*, 89:575–582, 1969. ISSN 0003-486X. doi: 10.2307/1970652.
- Anthony W Knapp. *Lie groups beyond an introduction*, volume 140. Springer, 2002.
- Anastasis Kratsios. The universal approximation property. *Annals of Mathematics and Artificial Intelligence*, 22 2021. ISSN 1573-7470. doi: <https://doi.org/10.1007/s10472-020-09723-1>.
- Anastasis Kratsios and Ievgen Bilokopytov. Non-euclidean universal approximation. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 10635–10646. Curran Associates, Inc., 2020.
- Anastasis Kratsios and Leonie Papon. Quantitative rates and fundamental obstructions to non-euclidean universal approximation with deep narrow feed-forward networks, 2021.
- Anastasis Kratsios and Behnoosh Zamanlooy. Learning sub-patterns in piece-wise continuous functions. *arXiv*, 2010.15571, 2021.
- Hugo Larochelle, Yoshua Bengio, Jérôme Louradour, and Pascal Lamblin. Exploring strategies for training deep neural networks. *Journal of Machine Learning Research*, 10(1):1–40, 2009.
- Szilárd László. Minimax results on dense sets and dense families of functionals. *SIAM J. Optim.*, 27(2):661–685, 2017. ISSN 1052-6234. doi: 10.1137/16M1092714.
- Moshe Leshno, Vladimir Ya Lin, Allan Pinkus, and Shimon Schocken. Multilayer feedforward networks with a nonpolynomial activation function can approximate any function. *Neural networks*, 6(6):861–867, 1993.
- Mario Lezcano-Casado and David Martínez-Rubio. Cheap orthogonal constraints in neural networks: A simple parametrization of the orthogonal and unitary group. *arXiv preprint arXiv:1901.08428*, 2019.
- Jianfeng Lu, Zuwei Shen, Haizhao Yang, and Shijun Zhang. Deep network approximation for smooth functions. *arXiv preprint arXiv:2001.03040*, 2020.
- David J. C. MacKay. *Information theory, inference and learning algorithms*. Cambridge University Press, New York, 2003.
- Warren S. McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *Bull. Math. Biophys.*, 5:115–133, 1943. ISSN 0007-4985.
- Charles A. Micchelli, Yuesheng Xu, and Haizhang Zhang. Universal kernels. *J. Mach. Learn. Res.*, 7:2651–2667, 2006. ISSN 1532-4435.

- Mehryar Mohri, Afshin Rostamizadeh, and Ameet Talwalkar. *Foundations of machine learning*. Adaptive Computation and Machine Learning. MIT Press, Cambridge, MA, 2018.
- James R. Munkres. *Topology*. Prentice Hall, Inc., Upper Saddle River, NJ, 2000. ISBN 0-13-181629-2.
- NA. NIST digital library of mathematical functions, 2013. URL <http://dlmf.nist.gov/5.19#E4>.
- Prajit Ramachandran, Barret Zoph, and Quoc V. Le. Searching for activation functions. *ICLR 2018*, 2018.
- Ramona-Andreea Rohan. Some remarks on the exponential map on the groups $SO(n)$ and $SE(n)$. In *Geometry, integrability and quantization XIV*, pages 160–175. Avangard Prima, Sofia, 2013. doi: 10.7546/giq-14-2013-160-175.
- Denis Serre. *Supplementary Exercises to: Matrices*, volume 216 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2002. ISBN 0-387-95460-0.
- Shobhit Seth. 10 major companies tied to the Apple supply chain. *Investopedia.com*, July 2018. URL <https://www.investopedia.com/articles/investing/090315/10-major-companies-tied-apple-supply-chain.asp>. Accessed: 2018-07-25.
- Shai Shalev-Shwartz and Shai Ben-David. *Understanding machine learning: From theory to algorithms*. Cambridge University Press, 2014. ISBN 9781107057135.
- Jonathan W Siegel and Jinchao Xu. Approximation rates for neural networks with general activation functions. *Neural Networks*, 2020.
- Shlomo Sternberg. Lie algebras. *Online*, 2004. URL http://people.math.harvard.edu/~shlomo/docs/lie_algebras.pdf.
- Kristin Stoller. The world’s largest tech companies 2018: Apple, Samsung take top spots again. *Forbes*, June 2018.
- Robert Tibshirani. Regression shrinkage and selection via the lasso. *J. R. Stat. Soc. Series B. Stat. Methodol*, pages 267–288, 1996.
- Andrey Nikolayevich Tikhonov. Solution of incorrectly formulated problems and the regularization method. *Dokl. Akad. Nauk. SSSR*, 151:501–504, 1963.
- Jan van Mill. *The infinite-dimensional topology of function spaces*, volume 64 of *North-Holland Mathematical Library*. North-Holland Publishing Co., Amsterdam, 2001.
- Roman Vershynin. Memory capacity of neural networks with threshold and ReLU activations. *CoRR*, abs/2001.06938, 2020a.
- Roman Vershynin. Memory Capacity of Neural Networks with Threshold and Rectified Linear Unit Activations. *SIAM J. Math. Data Sci.*, 2(4):1004–1033, 2020b.

Renjie Wang, Cody Hyndman, and Anastasis Kratsios. The entropic measure transform. *Canad. J. Statist.*, 48(1):97–129, 2020. ISSN 0319-5724. doi: 10.1002/cjs.11537.

Dmitry Yarotsky. Optimal approximation of continuous functions by very deep relu networks. In Sébastien Bubeck, Vianney Perchet, and Philippe Rigollet, editors, *Proceedings of the 31st Conference On Learning Theory*, volume 75 of *Proceedings of Machine Learning Research*, pages 639–649, 2018.

Chulhee Yun, Suvrit Sra, and Ali Jadbabaie. Small relu networks are powerful memorizers: a tight analysis of memorization capacity. In *Advances in Neural Information Processing Systems*, pages 15558–15569, 2019.

Hui Zou and Trevor Hastie. Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Series B. Stat. Methodol.*, 67(2):301–320, 2005.

6. Proofs

This appendix contains proofs of the paper’s main results and some auxiliary technical lemmas. We draw the reader’s attention to the fact that many of the paper’s results are heavily interdependent and that this sequential dependence is different from the order giving the cleanest presentation, which we chose for the paper’s main body. Accordingly, proofs will be presented in their logical order, even if it differs from the paper’s main exposé.

6.1 Technical Lemmas

This section contains some technical lemmas which we often refer to throughout the paper’s proofs.

Lemma 17 (Properties of Reconfiguration Units/Network) *Every reconfiguration unit is a reconfiguration network. Moreover, the following hold:*

(i) *Every $\phi \in \Phi_{\star;d}$ is a homeomorphism in $\mathcal{H}(\mathbb{R}^d)$,*

(ii) *For every $\phi \in \Phi_{\star;d}$, there exist $(\phi_\alpha)_{\alpha \in [0,1]} \subset \mathcal{H}(\mathbb{R}^d)$ such that for each $x \in \mathbb{R}^d$:*

$$\phi_0(x) = x, \quad \xi_\alpha \in \Phi_{\star;d}, \quad \phi_1(x) = \phi(x), \quad (23)$$

for every $\alpha \in [0, 1]$. Moreover, for every $x \in \mathbb{R}^d$, the map $\alpha \mapsto \phi_\alpha(x)$ is continuous.

(iii) *For every $x, y \in \mathbb{R}^d$ and every $\epsilon > 0$ there exists a reconfiguration unit ϕ such that*

$$\phi(x) = y \text{ and } \phi(z) = z,$$

for every $z \in \mathbb{R}^d$ satisfying $\|x - z\| > \epsilon + \|x - y\|$,

(iv) *If $\phi_1, \phi_2 \in \Phi_{\star;d}$, then $\phi_2 \circ \phi_1 \in \Phi_{\star;d}$.*

Proof [Proof of Lemma 17] First, we observe that (iv) holds by construction and the fact that any two reconfiguration units are composable since they map to and from \mathbb{R}^d .

By (iv) and the fact that the composition of homeomorphisms is again a homeomorphism, then it is enough to establish (ii) on a single map of the form

$$\sigma_{\alpha_n} (A(x)(x - c)) + b. \quad (24)$$

First, observe that for every $\alpha \in [0, 1]$, the map $x \mapsto \sigma_\alpha(x)$ is monotonically increasing, continuous, and surjective; thus, by (Hoffmann, 2015) implies that σ_α is a homeomorphism from \mathbb{R} onto itself. Since the d -fold Cartesian products of homeomorphisms is again a homeomorphism, then the map $x \mapsto \sigma_\alpha \bullet x$ is a homeomorphism from \mathbb{R}^d onto itself. Moreover, for any $b \in \mathbb{R}^d$, the maps $x \mapsto x + b$ and $x \mapsto x - c$ are homeomorphisms; thus the map of (24) is a homeomorphism if each reconfiguration unit $A(x)$ is a homeomorphism. Hence, we show that $A \in \mathcal{H}(\mathbb{R}^d)$.

For notational simplicity, we let $F(x) \triangleq \text{Skw}(f_0)(x) + A_1 \text{Skw}(f_1)(x) L_\eta(x)$, $f_0, f_1 \in \mathcal{NN}_{d,d(d-1)/2}^{\sigma_{\text{ReLU}}}$ with 1 hidden unit, and we observe that the reconfiguration unit $A(x)$ may be written as $A(x) = \exp(F(\|x\|))x$. Define the map $B(y) \triangleq \exp(-F(\|y\|))y$. Since \exp is continuous, matrix multiplication is continuous, and F is continuous then both A and B are continuous. Thus, if B is a two-sided inverse of A then A is a homeomorphism; we show this now. First, observe that for every $z \in \mathbb{R}$, $F(z)$ is a $d \times d$ skew-symmetric matrix and therefore, by (Rohan, 2013, Section 4), for every $x \in \mathbb{R}^d$, $\exp(F(\|x\|))$ is an orthogonal matrix with determinant $\det(\exp(F(\|x\|))) = 1$. Thus, $\exp(F(\|x\|))$ is an isometry fixing the origin; hence

$$\|\exp(F(\|x\|))x\| = \|x\|. \quad (25)$$

Therefore, (25) implies that

$$\|y\| = \|\exp(F(\|x\|))x\| = \|x\|,$$

where $y = F(\|x\|)x$. Hence, we compute that

$$\begin{aligned} B(y) &= \exp(-F(\|y\|))y \\ &= \exp(-F(\|x\|))y \\ &= \exp(-F(\|x\|))\exp(F(\|x\|))x \\ &= \exp(-F(\|x\|) + F(\|x\|))x \\ &= \exp(0)x \\ &= I_d x \\ &= x; \end{aligned} \quad (26)$$

where I_d is the $d \times d$ identity matrix. Mutatis mundais, the computation of the right-inverse is analogous. Therefore, B is the two-sided inverse of A and thus A is a homeomorphism. This gives (i).

For (ii), since the composition of identity maps is again the identity map then it is enough to demonstrate that any one reconfiguration unit may be parameterized to be the identity. Indeed, setting

$$\phi_\alpha(x) \triangleq \sigma_{\alpha\alpha_n} \exp(\alpha \text{Skw}(f_0)(\|x - c\|^2) + \alpha \text{Skw}(f_1)(x) L_\eta(\|x - c\|^2)); \quad \alpha \in [0, 1],$$

this gives the result since $\exp(0) = I_d$.

Next, we complete the proof by showing (iii). By (ii) and (iv) any reconfiguration unit $A(x) \in \Phi_{*,d}$. Let $x, y, z \in \mathbb{R}^d$ be such that $\|x - y\| < \|x - z\|$. Then, the barycenter $c = 2^{-1}(x + y)$ satisfies $\|x - c\| = \|y - c\| < \|c - z\|$. We set $\eta \triangleq 2^{-1}(\|c - z\| - \|c - x\|)$ and $f_0 = 0$.

Now, since the subset of $\mathcal{NN}_{d,d(d-1)/2}^{\sigma_{\text{ReLU}}}$ with one hidden unit, contains all constant functions, Skw is a bijection from \mathfrak{so}_d onto $\mathbb{R}^{d(d-1)/2}$, and since any matrix $B \in \exp(\mathfrak{so}_d)$ satisfies $B^\top B = I_d$ and $\det(B) = 1$ then for any $B \in \mathfrak{so}_d$ the constant function $f_1^B \triangleq \text{Skw}^{-1}(\log(B))$ belongs to $\mathcal{NN}_{d,d(d-1)/2}^{\sigma_{\text{ReLU}}}$. Moreover, the reconfiguration unit $A^{c,B,\eta}(u) \triangleq \exp(f_1^B L_\eta(\|u - c\|^2))(u - c)$, belongs to $\Phi_{*,d}$ and by construction $A^{c,B,\eta}(z) = z$ since $L_\eta(\|z - c\|^2) = 0$. Therefore $\exp(f_1^B L_\eta(\|z - c\|^2))z = \exp(0)z = z$. Furthermore, $A^{c,B,\eta}(x) \neq x$ and $A^{c,B,\eta}(y) \neq y$ whenever $B \neq 0$.

At this point, we seek a matrix B such that $Bx = y$, $B^\top B = I_d$, and $\det(B) = 1$. This matrix is explicitly computed to be

$$B = \left(\frac{y - c}{\|x - c\|} - \frac{x - c}{\|x - c\|} \right) \frac{x - c}{\|x - c\|}^\top + (\tilde{y} - \tilde{x}) \tilde{x}^\top, \quad (27)$$

$$\tilde{x} \triangleq \frac{y - (y^\top x)x}{\|y - (y^\top x)x\|}, \quad \tilde{y} \triangleq \frac{-x + (x^\top y)y}{\|x - (x^\top y)y\|}, \quad c \triangleq \frac{x + y}{2}.$$

Since $L_\eta(u) > 0$ while $|u| < \eta$ then $L_\eta(\|x - c\|)^{-1}$ is a well-defined strictly positive number. Setting $f_1^B \triangleq L_\eta(\|x - c\|)^{-1} \log(B)$, where B is given by (27), gives the conclusion. \blacksquare

Proof [Proof of Proposition 5] We argue analogously to the proof of (Kratsios and Zamanlooy, 2021, Lemma A.1). If σ is analytic, then so is the function $\Sigma : \mathbb{R}^d \rightarrow \mathbb{R}^d$ obtained by component-wise application of σ . Since the composition of analytic functions is again analytic, and since every affine function is analytic, then every $f \in \mathcal{NN}_{d,d}^\sigma$ is analytic. Since the difference of two analytic functions is analytic, then for every $f \in \mathcal{NN}_{d,d}^\sigma$, the function $F(x) \triangleq f(x) - x$ is analytic from \mathbb{R}^d to \mathbb{R}^d .

Suppose that property (P-iv) (from page 3) holds. Denote the zero-set of F by $Z_F \triangleq \{x \in \mathbb{R}^d : \phi(x) = x\}$ and note that by hypothesis we have that $\mu(Z_F) > 0$. Since Z_F is a set of positive Lebesgue measure then it must have a accumulation point and therefore it is identically 0 by the Principle of Permanence⁵. Hence, F must be identically 0 on all of \mathbb{R}^d and therefore $f(x) = x$ for every $x \in \mathbb{R}^d$. However, this contradicts the hypothesis that $f(x_1) = y_1$ for some $x_1 \neq y_1$; both in \mathbb{R}^d . Thus property (P-iv) never holds if σ is analytic. \blacksquare

The proof of Proposition 7 uses the following observation.

Lemma 18 *Let $h : X \rightarrow Y$ and $g : Y \rightarrow Z$, where X, Y, Z are non-empty sets. Then, $g \circ h$ is injective only if h is injective.*

Proof [Proof of Lemma 18] Let $g \circ h : X \rightarrow Z$ be injective. For a contradiction, assume that h is not injective. Then there would exist distinct $x_1, x_2 \in X$ for which $h(x_1) = h(x_2)$. Therefore,

5. The Principle of Permanence is sometimes called the *Uniqueness Theorem of Analytic Functions*

$g \circ h(x_1) = g(f(x_1)) = g(f(x_2)) = g \circ h(x_2)$ which contradicts the assumed injectivity of $g \circ h$. Hence, h must be injective. \blacksquare

Proof [Proof of Proposition 7] Let $f \in \mathcal{NN}_{d,D}^\sigma$. Clearly, if $d = 0$ then f is injective since there is only one point in its domain. Assume that $d \geq 1$. Let $\hat{f} \in \mathcal{NN}_{d,D}^{\sigma_{r-\text{ReLU}}}$ be UAP-invariant and therefore by (Kratsios and Bilokopytov, 2020, Theorem 3.4) is must be injective.

Define the function $g_W \in C(\mathbb{R}^d, \mathbb{R}^d)$ by $g_W \triangleq \sigma_{r-\text{ReLU}} \bullet W$, where W is an affine function from \mathbb{R}^d to itself. By Lemma 18, $\hat{f} \circ g_W$ is injective only if g_W is injective.

We show that this is never the case. To see this, we note that $\sigma_{r-\text{ReLU}} \bullet W$ is injective only if $\text{Im}(W) = [0, \infty)^d$; where \mathbb{R}^d is the co-domain of W . By Lemma 18, $\sigma_{r-\text{ReLU}} \bullet W_1$ is injective only if W is also injective.

We show that these two conditions cannot hold simultaneously. Suppose that W is injective and takes values in $[0, \infty)^d$. Since W is affine then $W(x) = Ax + b$ for some $d \times d$ -dimensional matrix A and some $b \in \mathbb{R}^d$. Since W is injective then $A \neq 0$. Let $e_i \in \mathbb{R}^d$ be the vector with entry 1 in its i^{th} component and zero otherwise; that is, $\{e_i\}_{i=1}^d$ is the standard orthonormal basis of \mathbb{R}^d . Therefore, $W((-b_j - A_{i,j})e_i)_j = -1$ for every $1 \leq j \leq d$ and every $1 \leq i \leq d$. Thus, $\text{Im}(W) \cap (-\infty, 0)^d \neq \emptyset$, which is a contradiction. In turn, this contradicts the assumption that $\hat{f} \circ \sigma_{r-\text{ReLU}} \bullet W$ is not injective. By (Kratsios and Bilokopytov, 2020, Theorem 3.4), $\hat{f} \circ \sigma_{r-\text{ReLU}} \bullet W$ is not UAP-invariant. Moreover, by definition it is not an embedding on $[0, 1]^D$, since all topological embeddings are injective. Therefore, if $\hat{f} \in \mathcal{NN}_{d,D}^{\sigma_{r-\text{ReLU}}}$ is UAP-invariant, then for every affine map $W : \mathbb{R}^d \rightarrow \mathbb{R}^d$, the DNN $\hat{f} \circ \sigma_{r-\text{ReLU}} \bullet W$ is not; nor is it a topological embedding on $[0, 1]^d$. \blacksquare

6.2 Proof of Theorem 4

The proof of Theorem 4 depends on the following lemma. Recall the length of a piece-wise smooth curve $\gamma : [0, 1] \rightarrow \mathbb{R}^d$, denoted by $\text{Len}(\gamma)$, is defined by:

$$\text{Len}(\gamma) = \int_{t \in D} \|\dot{\gamma}(t)\| dt = \int_0^1 \|\dot{\gamma}(t)\| dt,$$

where D is a dense subset of $[0, 1]$ for which $\mu(D) = 1$ and on which γ is differentiable, and where $\dot{\gamma}$ is the derivative of γ on D .

Lemma 19 *Let $d, N \in \mathbb{N}_+$ with $d > 1$ and x_1, \dots, x_N, x, z be distinct points in \mathbb{R}^d . Fix $0 < \delta < \Delta$ where Δ is given by:*

$$\Delta \triangleq \frac{1}{2} \min \left\{ 2, \min_{i,j=1,\dots,N, i \neq j} \|x_i - x_j\|, \min_{i=1,\dots,N} \|x_i - z\|, \min_{i=1,\dots,N} \|x_i - x\| \right\}.$$

There exists a piece-wise smooth curve $\gamma_{[x;z]} : [0, 1] \rightarrow \mathbb{R}^d$ satisfying

$$(a) \min_{\substack{j=1,\dots,n \\ t \in [0,1]}} \|\gamma_{[x;z]}(t) - x_j\| > 0,$$

$$(b) \text{Len}(\gamma_{[x:z]}) \leq \frac{\pi}{2\Delta},$$

$$(c) \gamma_{[x:z]}(0) = x \text{ and } \gamma_{[x:z]}(1) = z.$$

In particular, there exists reconfiguration units A_1, \dots, A_J such that

$$(i) A_J \circ \dots \circ A_1(x) = z,$$

$$(ii) A_J \circ \dots \circ A_1(x_i) = x_i,$$

(iii) There exists a compact subset $K \subseteq \mathbb{R}^d$ for which $A_J \circ \dots \circ A_1(x) = x$ for every $x \notin K$.

Furthermore, the following estimates hold:

$$J \leq \left\lceil \frac{\pi}{2(\min\{2\delta, 1\})} \right\rceil, \quad N_\delta(K) \leq J, \text{ and } \mu(K) \leq \frac{J\delta^2\pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2} + 1)}$$

Proof [Proof of Lemma 19] Let $n \in \mathbb{N}$, and x_1, \dots, x_n, x, z be distinct points \mathbb{R}^d . For any $y \in \mathbb{R}^d$ and $\epsilon > 0$ let

$$B_\epsilon(y) \triangleq \{p \in \mathbb{R}^d : \|p - y\| < \epsilon\}.$$

We denote the boundary of $B_\epsilon(y)$ by $\partial B_\epsilon(y)$. We proceed by induction. If $N = 0$ then $\delta > 0$ can be chosen arbitrarily. Let $l_{[x:z]} : [0, 1] \rightarrow \mathbb{R}^d$ be the straight line joining x to z . Since this function is piece-wise smooth its length, which we denote $\text{Len}(l_{[x:z]})$, is given by

$$\text{Len}(l_{[x:z]}) = \int_0^1 \left\| \dot{l}_{[x:z]}(t) \right\| dt = \|x - z\|. \quad (28)$$

Since $l_{[x:z]}$ and $B_\delta(x_i)$ is a convex body in \mathbb{R}^d then $l_{[x:z]}$ can each $\partial B_\delta(x_i)$ at-most two points. Without loss of generality, if $l_{[x:z]}([0, 1]) \cap B_\delta(x_i) = \emptyset$ for some $i = 1, \dots, N$ then $l_{[x:z]}$ does not need to be modified to avoid x_i ; thus we may assume that $l_{[x:z]}([0, 1]) \cap B_\delta(x_i) \neq \emptyset$ for each $i = 1, \dots, N$. For each $i = 1, \dots, N$ let $t_I^i < t_O^i$ be the respective first and final times where $l_{[x:z]}$ intersects $\partial B_\delta(x_i)$; these are given by

$$t_I^i \triangleq \min \{t \in [0, 1] : l_{[x:z]}(t) \in \partial B_\delta(x_i)\} \text{ and } t_O^i \triangleq \min \{t \in [t_I^i, 1] : l_{[x:z]}(t) \in \partial B_\delta(x_i)\}$$

Accordingly define the points $p_I^i \triangleq l_{[x:z]}(t_I^i)$ and $p_O^i \triangleq l_{[x:z]}(t_O^i)$. We modify $l_{[x:z]}$ to circumvent $B_\delta(x_i)$ and connect p_I^i to p_O^i about a minimal length curve on $\partial B_\delta(x_i)$. From Fletcher (2013) one finds that this is given by a great circle on the sphere $\partial B_\delta(x_i)$ given by the curve $\gamma^i : [0, 1] \rightarrow \partial B_\delta(x_i)$

$$\gamma^i(t) \triangleq \delta^2 \left(p_I^i \cos(t \|v^i\|) + \frac{\sin(t \|v^i\|)}{\|v^i\|} v^i \right),$$

where $v^i \triangleq \arccos(\langle p_I^i, p_O^i \rangle)(p_O^i - \langle p_I^i, p_O^i \rangle p_I^i)$ (here we have adjusted the formula in Fletcher (2013) to the case where $\delta \leq 1$ by using basic trigonometry). Accordingly, we modify $l_{[x:z]}$ to the following curve

$$\gamma_{[x:z]}(t) \triangleq \begin{cases} \gamma^i \left(\frac{(t-t_I^i)t_O^i}{(t_O^i-t_I^i)} + \frac{(t-t_O^i)t_I^i}{(t_I^i-t_O^i)} \right) & : t \in [t_I^i, t_O^i] \\ l_{[x:z]}(t) & : \text{else} \end{cases}, \quad (29)$$

this is indeed well-defined since $B_\delta(x_i) \cap B_\delta(x_j) = \emptyset$ for every $i \neq j$ by definition of Δ .

Since the length of γ^i is given by the geodesic distance on the sphere $\partial B_\delta(x^i)$ which was shown in Fletcher (2013) to be equal to

$$\delta^2 \arccos(\langle p_I^i, p_O^i \rangle) \leq \frac{\delta^2 \pi}{2}. \quad (30)$$

Then, since γ^i is piece-wise smooth then, by (28) and (30), it's length is computed to be

$$\text{Len}(\gamma_{[x:z]}) = \int_0^1 \|\dot{\gamma}(t)\| dt \leq \frac{\pi}{2} \|x - z\|. \quad (31)$$

Moreover, by construction $\gamma_{[x:z]}$ satisfies the bound

$$\delta < \inf_{t \in [0,1], x_1, \dots, x_N} \|\gamma_{[x:z]}(t) - x_i\|. \quad (32)$$

Thus, (a)-(c) hold.

Since $\text{Len}(\gamma_{[x:z]}) \leq \frac{\pi}{2\Delta} < \infty$ then we may pick $1 \leq t_1 < \dots < t_{\tilde{J}} \leq 1$ such that the length of the segment of $\gamma_{[x:z]}$ between $[t_i, t_{i+1}]$ is at-most 2δ ; i.e.:

$$\int_{t_i}^{t_{i+1}} \|\dot{\gamma}(t)\| dt \leq 2\delta \quad (\forall i = 1, \dots, \tilde{J}),$$

where $\tilde{J} \triangleq \left\lceil \frac{\text{Len}(\gamma_{[x:z]})}{\min\{2\delta, 1\}} \right\rceil$ balls of radius δ . In particular, by (31) we may take

$$\tilde{J} \leq \left\lceil \frac{\text{Len}(\gamma_{[x:z]})}{\min\{2\delta, 1\}} \right\rceil \leq \left\lceil \frac{\pi}{2(\min\{2\delta, 1\})} \right\rceil.$$

Furthermore, combining with (32), we observe that the collection of closed balls $\left\{ \overline{B_\delta(t_j)} \right\}_{j=1}^{\tilde{J}}$ satisfies:

$$\gamma_{[x:z]}([0, 1]) \subset \bigcup_{j=1}^{\tilde{J}} B_\delta(t_j) \text{ and } B_\delta(t_j) \cap \{x_i\}_{i=1}^n = \emptyset.$$

By Lemma 6.1 (iv) there exists reconfiguration units A_1, \dots, A_J such that for each $j = 1, \dots, \tilde{J} - 1$

$$A_j(\gamma_{[x:z]}(t_j)) = \gamma_{[x:z]}(t_{j+1}) \text{ and } A_j(p) = p; \quad (\forall p \in \mathbb{R}^d - \overline{B_\delta(\gamma_{[x:z]}(t_j))}). \quad (33)$$

Let $\tilde{K} \triangleq \bigcup_{j=1}^{\tilde{J}} \overline{B_\delta(\gamma_{[x:z]}(t_j))}$, by the Heine-Borel theorem each $\overline{B_\delta(\gamma_{[x:z]}(t_j))}$ is compact and since \tilde{J} is finite then \tilde{K} is compact.

Lastly, by the sub-additive of measure and by (NA, 2013, Equation 5.19 (iii)), we have the following estimate of $\mu(\tilde{K})$

$$\mu(\tilde{K}) \leq \sum_{j=1}^{\tilde{J}} \mu(\overline{B_\delta(\gamma_{[x:z]}(t_j))}) = \sum_{j=1}^{\tilde{J}} \frac{\delta^2 \pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2} + 1)} = \frac{J \delta^2 \pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2} + 1)}.$$

We have defined an open cover of K by the balls $\{B_\delta(\gamma_{[x:z]}(t_j))\}_{j=1}^{\tilde{J}}$. Therefore $N_\delta(\tilde{K}) \leq \tilde{J}$. \blacksquare

We will obtain the result now follows from Theorem 4 by repeatedly applying Lemma 19. **Proof** [Proof of Theorem 4] We proceed by induction on $N \in \mathbb{N}_+$. Suppose that $N = 1$, set $x = x_1$, $z = z_1$, and let $\{y_k\}_{k=1}^K$ be given. Since

$$\begin{aligned} \Delta &\leq \frac{1}{2} \min \left\{ 2, \min \|u - v\| : u \neq v, u, v \in \{x_i, z_i\}_{i=1}^N \cup \{y_k\}_{k=1}^K \right\} \\ &\leq \frac{1}{2} \min \left\{ 2, \min_{i,j=1,\dots,N, i \neq j} \|x_i - x_j\|, \min_{i=1,\dots,N} \|x_i - z\|, \min_{i=1,\dots,N} \|x_i - x\| \right\}, \end{aligned}$$

then Lemma 19 applies; hence, there exists some sequence of reconfiguration units $A_{1,1}, \dots, A_{1,N_1}$ satisfying Theorem 4. This yields the base case of the induction.

For the induction hypothesis, suppose that for any $N \in \mathbb{N}_+$ with $N \geq 1$, there exists some sequence $A_{1,N}, \dots, A_{J_N,N}$ and some non-empty compact subset $K_{N_n} \subseteq \mathbb{R}^d$ satisfying the conclusion of Theorem 4 for the given set $\{y_k\}_{k=1}^K$. Set $x = x_{N+1}$ and $z = z_{N+1}$. Our requirements on Δ imply that Lemma 19 applies; whence, there exists some sequence of reconfiguration units $\{A_{s,N+1}\}_{s=1}^{J_{N+1}}$ and some non-empty compact subset $K_{N+1} \subseteq \mathbb{R}^d$ such that

$$A_{J_{N+1}} \circ \dots \circ A_1(x) = z \tag{34}$$

$$A_{J_{N+1}} \circ \dots \circ A_1(\tilde{x}) = \tilde{x} \quad \forall \tilde{x} \in K_{N+1} \cup [\{y_m\}_{m=1}^M \cup \{x_n, z_n\}_{n=1}^N] \tag{35}$$

$$J_{N+1} \leq \left\lceil \frac{\pi}{2(\min\{2\delta, 1\})} \right\rceil \tag{36}$$

$$\mu(K_{N+1}) \leq \frac{J_{N+1} \delta^2 \pi^{\frac{d}{2}}}{\Gamma(\frac{d}{2} + 1)}, \tag{37}$$

$$N_\delta(K_{N+1}) \leq J_{N+1}. \tag{38}$$

Consider the sequence of reconfiguration units $A_{1,N}, \dots, A_{J_N,N}, A_{1,N+1}, \dots, A_{J_{N+1},N+1}$ and the set $K \triangleq K_N \cup K_{N+1}$. Note that this sequence is of length $J \triangleq J_{N+1} + J_N$. Note that the set K is compact since the finite union of compact subsets of \mathbb{R}^d is again compact, K_N is compact by induction hypothesis, and K_{N+1} is compact by Lemma 19.

Consider the reconfiguration network $\Phi \triangleq A_{J_{N+1},N+1} \circ \dots \circ A_{1,N+1} \circ \dots \circ A_{J_N,N} \circ \dots \circ A_{1,N} \circ \dots \circ A_1$. For notational simplicity let $\Phi_1 \triangleq A_{J_N,N} \circ \dots \circ A_{1,N}$, and $\Phi_2 \triangleq A_{J_{N+1},N+1} \circ \dots \circ A_{1,N+1}$. By the induction hypothesis, $\Phi_1(x) = x$ and $\Phi_1(z) = z$. Therefore, (34) implies that

$$\Phi(x_{N+1}) = \Phi(x) = \Phi_2 \circ \Phi_1(x) = \Phi_2(x) = z = z_{N+1}.$$

Again, by the induction hypothesis, we have that $\Phi_1(x_n) = z_n$ for each $n \leq N$. Hence, (35) implies that $\Phi(x_n) = \Phi_2 \circ \Phi_1(x) = \Phi_2(z_n) = z_n$, for each $n \leq N$. Thus (i) holds.

By the induction hypothesis, $\Phi_1(y_m) = y_m$ for each $m \leq M$. Thus, (35) yields

$$\Phi(y_m) = \Phi_2 \circ \Phi_1(y_m) = \Phi_2(y_m) = y_m,$$

for each $m \leq M$. Hence, (ii) holds.

Since, $K_N, K_{N+1} \subseteq K$ then $\mathbb{R}^d - K \subseteq [\mathbb{R}^d - K_N] \cap [\mathbb{R}^d - K_{N+1}]$. Thus, the induction hypothesis implies that $\Phi_1|_{\mathbb{R}^d - K_N}(p) = p$ for every $p \in \mathbb{R}^d - K$. Likewise, by definition of K_{N+1} , $\Phi_2|_{\mathbb{R}^d - K_N}(p) = p$ for every $p \in \mathbb{R}^d - K_{N+1}$. Therefore, for every $p \in \mathbb{R}^d - K$ we have that $\Phi(p) = \Phi_2 \circ \Phi_1(p) = \Phi_2(p) = p$. Therefore, (iii) holds.

Next, by the induction hypothesis, the definition of J , and by (36) we have that

$$J = J_N + J_{N+1} \leq \left\lceil \frac{N\pi}{2(\min\{2\delta, 1\})} \right\rceil + \left\lceil \frac{\pi}{2(\min\{2\delta, 1\})} \right\rceil \leq \left\lceil \frac{(N+1)\pi}{2(\min\{2\delta, 1\})} \right\rceil.$$

This gives the bound on J . Lastly, combining the induction hypothesis on the covering number of K_N with the bound of (38) we make the following computation

$$N_\delta(K) \leq N_\delta(K_N) + N_\delta(K_{N+1}) \leq J_N + J_{N+1} = J.$$

This completes the induction hypothesis and therefore the proof. \blacksquare

6.3 Proof of Theorem 3

Proof [Proof of Theorem 3] Let $\phi \in \mathcal{H}_{M,\omega}^\delta(\mathbb{R}^d)$, $M, \delta > 0$, and $d \in \mathbb{N}_+$. By definition, there exists $\{x_n\}_{n \leq N}$ in \mathbb{R}^d and $\{\varphi_n\}_{n \leq N}$ in $\mathcal{H}(\mathbb{R}^d)$ such that $[-M, M]^d \subseteq \cup_{n \leq N} B(x_n, \delta)$, (3), and (7)-(10) hold.

Step 1 - Upper-bound N : Let us begin by upper-bounding N . Observe that if $x \in [-M, M]^d$ then $\|x\| \leq 2\sqrt{d}M$. Therefore, (Shalev-Shwartz and Ben-David, 2014, page 337) implies that

$$\begin{aligned} N_{2^{-1}\delta}([-M, M]^d) &\leq \sup_{x \in [-M, M]^d} \left(\frac{2\|x\|\sqrt{d}}{2^{-1}\delta} \right)^d \\ &= \left(\frac{8\sqrt{d}M\sqrt{d}}{\delta} \right)^d. \end{aligned} \tag{39}$$

Next, $[-M, M]^d \subseteq \cup_{n \leq N} B(x_n, \delta)$ and (9), together with the bound on the packing number $N_{2^{-1}\delta}^{pack}([-M, M]^d)$ in (Shalev-Shwartz and Ben-David, 2014, Lemma 27.3) imply that

$$N \leq N_\delta([-M, M]^d) \leq N_{2^{-1}\delta}^{pack}([-M, M]^d) \leq N_{2^{-1}\delta}([-M, M]^d). \tag{40}$$

Combining (39) with (40) yields the following upper-bound on N

$$N \leq N_{2^{-1}\delta}([-M, M]^d) \leq \left(\frac{8dM}{\delta} \right)^d. \tag{41}$$

Step 2 - Representation of each φ_n : Next, let us determine the form of any $\phi \in \mathcal{H}_{M,\omega}^\delta(\mathbb{R}^d)$. By the fragmentation representation (3), we only need to describe a single φ_n for $n \leq N$. In what follows, for $x \in \mathbb{R}^d$ and $r \geq 0$, we use $S^{d-1}(x, r)$ to denote the $d-1$ dimensional sphere in \mathbb{R}^d with center x and radius r ; defined by $S^{d-1}(x, r) \triangleq \{z \in \mathbb{R}^d : \|x - z\| = r\}$. Together, properties (7) and (8) imply that φ_n maps $S^{d-1}(x_n, \delta)$ to itself and that $\varphi_n|_{S^{d-1}(x_n, r)} : S^{d-1}(x_n, \delta) \rightarrow S^{d-1}(x_n, r)$ is an isometry, for every $r \geq 0$. Since any isometry on \mathbb{R}^d belongs to the Euclidean group⁶ then it is of the form $\varphi_n|_{S^{d-1}(x_n, r)} =$

6. The Euclidean group is set of all isometries on \mathbb{R}^d . A map $f \in C(\mathbb{R}^d, \mathbb{R}^d)$ belongs to the Euclidean group if and only if it is an affine map of the form $f(x) = Ax + b$ where A is an orthogonal matrix and $b \in \mathbb{R}^d$.

$A_r x + b_r$ for some orthogonal matrix A_r and some $b_r \in \mathbb{R}^d$. However, since φ_n maps each $S^{d-1}(x_n, r)$ into itself then it must fix x_n and therefore $b_n = -A_n x_n$. Next, since $\varphi_n \in \mathcal{H}(\mathbb{R}^d)$ then it is orientation-preserving and therefore $A_r \in SO(d)$. Thus, for every $n \leq N$ and every $r \geq 0$

$$\varphi_n(x) = A_{\|x\|}(x - x_n). \quad (42)$$

The map $\exp : \mathfrak{so}_d \mapsto SO(d)$ is a continuous surjection with continuous right-inverse $\log : SO(d) \mapsto \mathfrak{so}_d$ (see (Rohan, 2013, Section 4)). Hence, the map $r \mapsto \tilde{R}_r \triangleq \log(A_r)$ from $[0, \infty)$ to \mathfrak{so}_d must be continuous and by (42) it satisfies

$$\varphi_n(x) = \exp\left(\tilde{R}_{\|x-x_n\|}\right)(x - x_n). \quad (43)$$

since $\text{Skw} : \mathbb{R}^{d(d-1)/2} \rightarrow \mathfrak{so}_d$ is a bijective isometry then there exists some function $R \in C([0, \infty), \mathbb{R}^{d(d-1)/2})$ satisfying $\text{Skw} \circ R = \tilde{R}$. Thus, (43) simplifies to

$$\varphi_n(x) = \exp\left(\text{Skw}R_{\|x-x_n\|}\right)(x - x_n). \quad (44)$$

Since $\varphi_n \in \mathcal{H}(\mathbb{R}^d)$ then representation (44), the continuity of \log on $SO(d)$, Skw^{-1} on \mathfrak{so}_d , the continuity of the Euclidean norm, the continuity of affine transformations, and the fact that the composition of continuous functions is again continuous all together imply that the map $x \mapsto \exp(\text{Skw}(R_{\|x-x_n\|}))(x - x_n)$ must be also be continuous.

Step 3 - Representation of Reconfiguration Network's Layers: For any given $\tilde{N}, n \in \mathbb{N}_+$ with $n \leq N$, whenever we set $\alpha_k = 0$, $\eta_k = 0$ for $1 \leq k \leq \tilde{N}_n$, $c_k = b_k = 0$ for $1 < k \leq \tilde{N}_n$, and $-c_1 = b_1 = x_n$, we can represent any such reconfiguration network $\phi^{(n)} \triangleq \phi_{\tilde{N}_n, n} \circ \dots \circ \phi_{1, n}$, with $\phi_{k, n}$ for the form:

$$\phi_{k, n}(x) = \prod_{k \leq \tilde{N}_n} \exp\left(\text{Skw}(W_{2, k: n} \circ \sigma_{\text{ReLU}} \bullet W_{1, k: n})(\|x - x_n\|^2)\right)(x - x_n). \quad (45)$$

Using the fact that $\exp(A + B) = \exp(A)\exp(B)$ for any $d \times d$ matrices A and B and the linearity of Skw , we rewrite (45) as

$$\phi_{k, n}(x) = \exp\left(\text{Skw}\left(\sum_{k \leq \tilde{N}_n} W_{2, k: n} \circ \sigma_{\text{ReLU}} \bullet W_{1, k: n}(\|x - x_n\|^2)\right)\right)(x - x_n). \quad (46)$$

Finally, by construction each $\phi_{k, n}$ maps each $S^{d-1}(x_n, r)$, for $r \geq 0$, isometrically into itself we have that $\|\phi_{k, n}(x) - x_n\| = \|x - x_n\|$, for every $x \in \mathbb{R}^d$ and every $k \leq \tilde{N}$, and therefore

$$\phi^{(n)}(x) = \phi_{\tilde{N}_n, n} \circ \dots \circ \phi_{1, n} = \exp\left(\text{Skw}(f_n(\|x - x_n\|^2))\right)(x - x_n), \quad (47)$$

where $f_n \in \mathcal{NN}_{1, d(d-1)/2}^{\sigma_{\text{ReLU}}}$ with 1 hidden layer and of width $\nu_{\tilde{N}_n} \triangleq \frac{\tilde{N}_n d(d-1)}{2}$.

Thus, in the next step, we consider reconfiguration networks of the form $\hat{\phi}(x) = \phi^{(N)} \circ \dots \circ \phi^{(1)}$; where each $\phi^{(n)}$ is represented via (47). Hence, ϕ will always be a reconfiguration network of depth $\sum_{n \leq N} \nu_{\tilde{N}_n}$.

Step 4 - Upper-bounding the Modulus of Continuity of $r \mapsto R_r$: We conclude this portion of the proof by computing an upper-bound of the modulus of continuity of

R . Indeed, since we are estimating on $[-M, M]^d$ and since the Euclidean norm maps $[-M, M]^d$ to the compact interval $[0, 2\sqrt{d}M]$ then we only need to concern ourselves with approximating the modulus of continuity of $r \mapsto R_r$ on $[0, 2\sqrt{d}M]$. Since $r \mapsto R_r$ is continuous and since $[0, 2\sqrt{d}M]$ is compact, then by the Heine-Cantor Theorem ((Munkres, 2000, Theorem 27.6)) every continuous function on a compact set is uniformly continuous. Hence, R genuinely admits an optimal modulus of continuity, which we denote by ω_R .

If f, g are composable with respective optimal moduli of continuity ω_f and ω_g then the optimal modulus of continuity of $g \circ f$, denoted $\omega_{g \circ f}$, satisfies the bound

$$\omega_{g \circ f}(t) \leq \omega_g \circ \omega_f(t), \quad (48)$$

for every $t \in [0, \infty)$. Since the Riemannian exponential map on $SO(d)$ coincides with its Lie exponential at the identity and therefore by (Sternberg, 2004, Equation 1.11), we know that for every pairs of $d \times d$ skew symmetric matrices X and Y we have that:

$$\begin{aligned} d \exp(Y) &= \exp(I_d) \int_0^1 \exp(-sX) Y \exp(sX) ds \\ &= \int_0^1 Y ds \\ &= Y; \end{aligned} \quad (49)$$

where we have use the fact that if $X \in \mathfrak{so}_d$, and similarly for sX for any $s \in [0, 1]$ since \mathfrak{so}_d is a vector space, then $\exp(X)$ is an isometry. Since \exp is smooth then the mean-value theorem and (49) imply that \exp is 1-Lipschitz from \mathfrak{so}_d with the Fröbenius norm $\|\cdot\|_F$ to $SO(d)$ with its Riemannian distance and, in particular, so for its Euclidean norm. Furthermore, the generalized right-inverse of ω_{\exp} , in the sense of Embrechts and Hofert (2013), is defined by

$$\omega_{\exp}^{-1}(t) \triangleq \inf\{s \in \mathbb{R} : \omega_{\exp}(s) \geq t\} = t$$

and it satisfies $t \leq \omega_{\exp} \circ \omega_{\exp}^{-1}(t)$ for every $t \in [0, \infty)$. These considerations, together with (Embrechts and Hofert, 2013, Proposition 2.3) and (10) imply that

$$\begin{aligned} \omega_R(t) &\leq \omega_R \circ \omega_{\exp} \circ \omega_{\exp}^{-1}(t) = \omega_{\exp \circ R} \circ \omega_{\exp}^{-1}(t) \\ &= \omega_{\phi_n} \circ \omega_{\exp}^{-1}(t) \leq \omega \circ \omega_{\exp}^{-1}(t) \\ &= \omega(t), \end{aligned} \quad (50)$$

for every $t \in [0, \infty)$.

Step 5 - Computing Bounds (Single block Case): By (Serre, 2002, Excercise 106), the map $\exp : \mathfrak{so}_d \rightarrow SO(d)$ is 1-Lipschitz when \mathfrak{so}_d is equipped with the operator norm $\|\cdot\|_{op}$. Combining the representations for ϕ_n and of $\hat{\phi}^{(n)}$ is steps 2 and 3, respectively,

we compute the following estimate.

$$\begin{aligned}
 \left\| \phi_n(x) - \phi^{(n)}(x) \right\| &= \left\| \exp(\text{Skw} \circ R_{\|x-x_n\|})(x-x_n) - \exp(\text{Skw}(f_n(\|x-x_n\|^2)))(x-x_n) \right\| \\
 &\leq \left\| \exp(\text{Skw} \circ R_{\|x-x_n\|}) - \exp(\text{Skw}(f_n(\|x-x_n\|^2))) \right\|_{op} \|x-x_n\| \\
 &\leq \left\| \exp(\text{Skw} \circ R_{\|x-x_n\|}) - \exp(\text{Skw}(f_n(\|x-x_n\|^2))) \right\|_{op} 2\sqrt{d}M \\
 &\leq \left\| \text{Skw} \circ R_{\|x-x_n\|} - \text{Skw} \circ f_n(\|x-x_n\|^2) \right\|_{op} 2\sqrt{d}M \\
 &= \sqrt{d} \left\| \text{Skw} \circ R_{\|x-x_n\|} - \text{Skw} \circ f_n(\|x-x_n\|^2) \right\|_F 2\sqrt{d}M \\
 &= \left\| R_{\|x-x_n\|} - f_n(\|x-x_n\|^2) \right\| 2dM;
 \end{aligned}$$

where we appealed the bound $\|\cdot\|_{op} \leq \sqrt{d}\|\cdot\|_F$ and the fact that Skw is an isometry between \mathfrak{so}_d with the Fröbenius norm, and $\mathbb{R}^{d(d-1)/2}$ with the Euclidean norm $\|\cdot\|$. Since (50) gives us a uniform upper-bound of the modulus of continuity of R , and since $f_n \in \mathcal{NN}_{1, \frac{d(d-1)}{2}}^{\sigma_{\text{ReLU}}}$ with (breadth) $\frac{\nu_{\tilde{N}_n} d(d-1)}{2}$ and one hidden layer, then we may apply (Yarotsky, 2018, Proposition 1) to the left-hand side of (6.3) to obtain the following bound

$$\left\| \phi_n(x) - \phi^{(n)}(x) \right\| = \left\| R_{\|x-x_n\|} - f_n(\|x-x_n\|^2) \right\| 2dM \leq 2dMC\omega \left(\left(\frac{\nu_{\tilde{N}_n} d(d-1)}{2} \right)^{\frac{-1}{d}} \right), \quad (51)$$

where $C > 0$ is a constant independent ν .

Step 5 - Computing Bounds (General Case): Suppose that N is known exactly. We make the following abbreviations $\omega_{N-1} \triangleq \sup_{\|x\| \leq M} \left\| \phi^{(N-1)} \circ \dots \circ \phi^{(1)}(x) - \varphi_{N-1} \circ \dots \circ \varphi_1(x) \right\|$, $\tilde{\varphi}_{N-1} \triangleq \varphi_{N-1} \circ \dots \circ \varphi_1$ and $\tilde{\phi}^{N-1} \triangleq \phi^{(N-1)} \circ \dots \circ \phi^{(1)}$. Then, using (10), the definition of the modulus of continuity of ϕ_N , and Step 4 we compute the following recursive-bound for any $x \in [-M, M]^d$.

$$\begin{aligned}
 \left\| \phi_n \circ \tilde{\phi}^{(n-1)}(x) - \varphi_n \circ \tilde{\varphi}_{n-1}(x) \right\| &= \left\| \phi_n \circ \tilde{\phi}^{(n-1)}(x) - \varphi_n \circ \tilde{\phi}^{(n-1)}(x) \right. \\
 &\quad \left. - \varphi_n \circ \tilde{\varphi}_{n-1}(x) + \varphi_n \circ \tilde{\phi}^{(n-1)}(x) \right\| \\
 &\leq \left\| \phi_n \circ \tilde{\phi}^{(n-1)}(x) - \varphi_n \circ \tilde{\phi}^{(n-1)}(x) \right\| \\
 &\quad + \left\| \varphi_n \circ \tilde{\phi}^{(n-1)}(x) - \varphi_n \circ \tilde{\varphi}_{n-1}(x) \right\| \\
 &\leq \left\| \phi_n \circ \tilde{\phi}^{(n-1)}(x) - \varphi_n \circ \tilde{\phi}^{(n-1)}(x) \right\| \quad (52) \\
 &\quad + \omega \left(\left\| \tilde{\phi}^{(n-1)}(x) - \tilde{\varphi}_{n-1}(x) \right\| \right) \\
 &\leq 2dMC\omega \left(\left(\frac{\nu_{\tilde{N}_n} d(d-1)}{2} \right)^{\frac{-1}{d}} \right) \\
 &\quad + \omega(\omega_{n-1})
 \end{aligned}$$

Set $\nu_{\tilde{N}_n} \triangleq \nu \left\lceil \frac{\delta^d}{(8dM)^d} \right\rceil$. By the bound (41), we know that $\nu_{\tilde{N}_n} \triangleq \nu \left\lceil \frac{\delta^d}{(8dM)^d} \right\rceil \leq \frac{\nu}{N}$ and by the monotonicity of modulus of continuity ω we have that for every $k \geq 0$

$$\omega\left(\nu_{\tilde{N}_n} k\right) \leq \omega\left(\frac{\nu}{N} k\right). \quad (53)$$

Therefore, we have the following approximation bound

$$\sup_{\|x\| \leq M} \left\| \phi_n \circ \tilde{\phi}^{N-1} - \hat{\phi}^{(N)} \circ \tilde{\phi}^{(N-1)}(x) \right\| \leq \omega_N,$$

where ω_N is determined by the recurrence relation:

$$\omega_{n+1} = 2dMC\omega\left(\left(\frac{\nu_{\tilde{N}_n} d(d-1)}{2}\right)^{\frac{-1}{d}}\right) + \omega(\omega_n), \quad \omega_0 = 0. \quad (54)$$

Since ω takes values in the non-negative integers; then, the sequence $\{\omega_n\}_{n \in \mathbb{N}}$ is monotonically increasing. Whence, $\omega_N \leq \omega \left\lceil \frac{(8dM)^d}{\delta^d} \right\rceil$. This gives us the desired bound. \blacksquare

6.4 Proof of Theorem 6

The proof of Theorem 6 relies on the observation that given any $(d+D) \times d$ matrix A , any $f \in C(\mathbb{R}^d, \mathbb{R}^D)$ can be expressed as $f = P \circ \Phi_f(x, Ax)$ where $P : \mathbb{R}^{d+D} \rightarrow \mathbb{R}^D$ is the (linear) orthogonal projection mapping $(x, y) \in \mathbb{R}^d \times \mathbb{R}^D$ to $y \in \mathbb{R}^D$ and $\Phi_f \in \mathcal{H}(\mathbb{R}^{d+D})$ is given by

$$\Phi_{f,A}(x, y) = (x, y + f(x) - Ax).$$

Since P is 1-Lipschitz, A is arbitrary, then we only need to approximate $\Phi_{f,A}$; which we know can always be done using Theorem 2.

Proof [Proof of Theorem 6] By Theorem 2 since $\Phi_{f,A} \in \mathcal{H}(\mathbb{R}^{d+D})$ there exists some $\hat{\phi} \in \mathcal{NN}_{\star; d+D}$ satisfying

$$\sup_{(x,y) \in \mathbb{R}^{d+D}; \|(x,y)\| \leq M} \left\| \Phi_{f,A}(x, y) - \hat{\phi}(x, y) \right\| < \epsilon. \quad (55)$$

Since the orthogonal projection $P(x, y) \triangleq y$ is 1-Lipschitz, then (55) yields (iv) via

$$\begin{aligned} \sup_{\|x\| \leq M} \left\| f(x) - P\hat{\phi}(x, Ax) \right\| &= \sup_{\|x\| \leq M} \left\| P \circ \Phi_{f,A}(x, Ax) - P\hat{\phi}(x, Ax) \right\| \\ &= \sup_{\|x\| \leq M} \left\| \Phi_{f,A}(x, Ax) - \hat{\phi}(x, Ax) \right\| \\ &< \epsilon. \end{aligned} \quad (56)$$

Setting $B \triangleq P$, yields (iii).

Since $x \mapsto I_d \oplus Ax$ is a linear map, it is continuous. By construction, $I_d \oplus A$ is of fullrank; whence the associated linear map (obtained by matrix-multiplication) is injective

and since $\hat{\phi} \in \mathcal{H}(\mathbb{R}^{d+D})$, then it is also continuous and injective. Hence, $\psi(x) \triangleq \hat{\phi}(I_d \oplus Ax)$ is a continuous injective map. Since ψ is continuous then the pre-image of any compact subset $K \subseteq \mathbb{R}^{d+D}$ is closed in \mathbb{R}^d and $\psi^{-1}[K] \cap [-M, M]^d$ is a closed subset of the compact set $[-M, M]^d$. Hence, $\psi^{-1}[K] \cap [-M, M]^d$ is compact in $[-M, M]^d$. Therefore, ψ is a continuous, injective, and by (Munkres, 2000, Exercises 26.6) it is a closed map also. Thus, it is a topological embedding; i.e. a homeomorphism onto its image with respect to the subspace topology. This gives (ii) and the statement that ψ is an isometry in (i).

Moreover, since ψ is a homeomorphism then it is a bijection with inverse $\phi^{-1} \circ A^\dagger$ where $I_d \oplus A^\dagger$ is the left-inverse of $I_d \oplus A$ given in the statement of (i). Since ϕ is a bijection then the metric

$$d_{\hat{\phi}, M, \epsilon}(z_1, z_2) \triangleq \|\psi^{-1}(z_1) - \phi^{-1}(z_2)\|$$

is well-defined on $\mathcal{M}_{\hat{\phi}, M, \epsilon}$ and by construction ψ is an isometry from $[-M, M]^d$ to $\mathcal{M}_{\hat{\phi}, M, \epsilon}$. This completes the proof of (i). \blacksquare

6.5 Proof of Theorem 2

6.5.1 COMMENTS ON THE METHOD OF PROOF FOR THEOREM 2

We begin by noting that the space $\mathcal{H}(\mathbb{R}^d)$ is highly non-linear. To see this, observe that the identity map $1_{\mathbb{R}^d}$ belongs to $\mathcal{H}(\mathbb{R}^d)$, however, $1_{\mathbb{R}^d} - k1_{\mathbb{R}^d} \notin \mathcal{H}(\mathbb{R}^d)$ for any $k > 0$.

This non-linearity unfortunately renders most classical tools used to establish universality, such as the Stone-Weierstraß theorem (used in Hornik et al. (1989)) from classical approximation theory, the Hahn-Banach theorem (used in Micchelli et al. (2006)) from functional analysis, the Wiener-Tauberian Theorem (used in Cybenko (1989)) from harmonic analysis, or hypercyclicity results (used in Kratsios (2021)) from linear dynamical systems, useless.

Nevertheless, we still have access to the less standard tools from infinite-dimensional topology; which were specifically built to handle limits of homeomorphisms in $\mathcal{H}(\mathbb{R}^d)$. Specifically, we appeal to *the inductive convergence criterion*, a less-known result which provides conditions on a sequence of homeomorphisms $(\Phi_k)_{k \in \mathbb{N}_+}$ in $\mathcal{H}(\mathbb{R}^d)$ such that the limit of $\Phi_k \circ \dots \circ \Phi_1$ in the topology of $\mathcal{H}(\mathbb{R}^d)$ exists. For a self-contained exposition, we briefly discuss state and discuss the result here before applying it.

The result relies on some notation which we now develop. Let $\epsilon > 0$, $K \subseteq \mathbb{R}^d$ be non-empty and compact. For any $\Phi \in \mathcal{H}(K)$ define the quantity:

$$\alpha(h, \epsilon) \triangleq \inf \{ \delta : \exists x, y \in K, \|x - y\| > \epsilon \text{ and } \|\Phi(x) - \Phi(y)\| = \delta \}.$$

As discussed on (Anderson, 1967, page 367), if $d \in \mathbb{N}_+$ then for any $\Phi \in \mathcal{H}(K)$ and any $\epsilon > 0$ we have that $\alpha(\Phi, \epsilon) > 0$. Several formulations of this lemma exist, see van Mill (2001); Hernandez-Gutierrez (2020); however, we prefer to use of the phrasing presented in Anderson (1967).

Lemma 20 (Inductive Convergence Criterion (Anderson, 1967, Lemma 2.1))

If $(\phi_k)_{k \in \mathbb{N}_+}$ is a sequence in $\mathcal{H}(K)$ and if

$$\sup_{x \in K} \|\phi_k(x) - x\| < 3^{-k} \min \left\{ 1, \alpha \left(\phi_{k-1} \circ \dots \circ \phi_1, 2^{-k} \right) \right\},$$

for each $k > 1$ then $\{\phi_k \circ \dots \circ \phi_1\}_{k \in \mathbb{N}_+}$ converges (to a homeomorphism) in $\mathcal{H}(K)$.

6.5.2 PROOF FOR THEOREM 2

Proof [Proof of Theorem 2] Let $d \in \mathbb{N}_+$ with $d \geq 2$, $K \subseteq \mathbb{R}^d$ be compact, and let $\phi \in \mathcal{H}(\mathbb{R}^d)$. We deduce some simplifications on K , on ϕ , and on their relationship before moving onto the main part of the proof. Since $\phi \in \mathcal{H}(\mathbb{R}^d)$ then it is isotopic to the identity; thus, the second Corollary⁷ to (Kirby, 1969, Theorem 2) implies that ϕ can be fragmented, i.e.: for any open cover $\mathcal{U} \triangleq \{U_n\}_{n=1}^N$ of \mathbb{R}^d there exists $\phi_{\mathcal{U}:1}, \dots, \phi_{\mathcal{U}:N} \in \mathcal{H}(\mathbb{R}^d)$ such that

$$\begin{aligned} \phi &= \phi_{\mathcal{U}:N} \circ \dots \circ \phi_{\mathcal{U}:1} \\ \phi_{\mathcal{U}:n}(x) &= x \quad (\forall x \notin U_n). \end{aligned} \tag{57}$$

In particular, (57) implies that $\phi_{\mathcal{U}:n}(U_n) = U_n$, for every $n \leq N$, since the value of ϕ_n is determined to be the identity outside of U_n . In particular, fix some $\delta > 0$ and define $K'_\delta \triangleq \{z \in \mathbb{R}^d : (\exists x \in K) \|z - x\| < \delta\}$; then for the open cover $\{K'_\delta, \mathbb{R}^d - K\}$ of \mathbb{R} , the fragmentation property of ϕ implies that there are $\{\phi_{\mathcal{U}:i}\}_{i=1}^2 \subset \mathcal{H}(\mathbb{R}^d)$ satisfying

$$\begin{aligned} \phi &= \phi_{\mathcal{U}:2} \circ \phi_{\mathcal{U}:1} \\ \phi_{\mathcal{U}:2}(x) &= x \quad (\forall x \in K) \\ \phi_{\mathcal{U}:1}(K) &= K'_\delta. \end{aligned} \tag{58}$$

In particular, (58) implies that for each $x \in \text{int}(K)$ we have that $\phi(x) = \phi_{\mathcal{U}:2} \circ \phi_{\mathcal{U}:1}(x) = \phi_{\mathcal{U}:1}(x)$. Therefore, for every $\hat{\phi} \in \Phi_{\star,d}$, the following holds

$$\sup_{x \in K'_\delta} \left\| \hat{\phi}(x) - \phi(x) \right\| = \sup_{x \in K'_\delta} \left\| \hat{\phi}(x) - \phi_1(x) \right\|. \tag{59}$$

Since K'_δ is, by definition, dense in its closure $\overline{K'_\delta}$ and since continuous functions are determined by their value on dense sets then

$$\max_{x \in K} \left\| \hat{\phi}(x) - \phi(x) \right\| \leq \sup_{x \in K'_\delta} \left\| \hat{\phi}(x) - \phi(x) \right\| = \max_{x \in \overline{K'_\delta}} \left\| \hat{\phi}(x) - \phi_1(x) \right\|. \tag{60}$$

Therefore, we only need to approximate ϕ_1 on $\overline{K'_\delta}$ in order to approximate ϕ on $K \subset \overline{K'_\delta}$. Thus, without loss of generality, we may assume that $K = \overline{\text{int}(K)}$ and that $\phi : K \rightarrow K$.

With these key simplifications, we proceed to the main part of the proof. Since \mathbb{R}^d is separable then so is K . Therefore, there exists a sequence $\{x_n\}_{n \in \mathbb{N}}$ contained in K such that $\{x_n\}_{n \in \mathbb{N}}$ is a dense subset of K . Since the value of any continuous function is determined by its value on any dense subset, and since ϕ is continuous, then we only need to approximate ϕ on $\{x_n\}_{n \in \mathbb{N}}$.

Define the sequence $\{z_n\}_{n \in \mathbb{N}} \in \phi(K)$ by $z_n \triangleq \phi(x_n)$, for each $n \in \mathbb{N}$. For any $N \in \mathbb{N}_+$, Theorem 4 implies that there exists some $\hat{\phi} \in \Phi_{\star,d}$ satisfying $\hat{\phi}(x_n) = \phi(x_n)$ for every $n \leq N$. Set $\hat{\phi}_0 \triangleq \hat{\phi}$. We recursively generate a sequence of reconfiguration networks in $\Phi_{\star,d}$ converging to ϕ on $\{x_n\}_{n \in \mathbb{N}}$ and therefore on K .

7. The results are not numbered in Kirby (1969).

By Lemma 19, for every $n > N$ there exists some continuous curve with finite length $\gamma_{x_n:z_n} : [0, 1] \rightarrow \mathbb{R}^d$ satisfying $\gamma_{x_n:z_n}(0) = x_n$, $\gamma_{x_n:z_n}(1) = z_n$, and for which there is some $\Delta_n > 0$ satisfying

$$\Delta_n \triangleq \min_{k < n, t \in [0, 1]} \|\gamma_{[x_n:z_n]}(t) - z_k\| > 0.$$

. For any decreasing of positive numbers $\{\epsilon_{n:k}\}_{k \in \mathbb{N}}$ in $(0, \Delta_n)$ Since $\gamma_{[x_n:z_n]}$ has finite length then one can choose an increasing sequence $\{t_{n:k}\}_{k \in \mathbb{N}}$ in $[0, 1]$ such that $t_{n:0} = 0$ and $\gamma_{[x_n:z_n]}([0, 1]) \subseteq \bigcup_{k \in \mathbb{N}} \text{Ball}(\gamma_{[x_n:z_n]}(t_{n:k}), \epsilon_{n:k})$. Since $\gamma_{[x_n:z_n]}$ is continuous and $[0, 1]$ is compact, then (Munkres, 2000, Theorem 26.5) implies that $\gamma_{[x_n:z_n]}([0, 1])$ is compact and therefore there must exist a finite sub-collection $\{(\epsilon_{n:k_j}, t_{n:k_j})\}_{j=1}^{J_n}$ such that

$$\gamma_{[x_n:z_n]}([0, 1]) \subseteq \bigcup_{j=1}^{J_n} \text{Ball}(\gamma_{[x_n:z_n]}(t_{n:k_j}), \epsilon_{n:k_j}). \quad (61)$$

By Lemma 6.1 (ii), we choose reconfigurations networks $\hat{\phi}_{n:1}, \dots, \hat{\phi}_{n:J_n-1} \in \Phi_{\star:d}$ so that

$$\hat{\phi}_{n:j}(\gamma_{[x_n:z_n]}(t_{n:k_j})) = \gamma_{[x_n:z_n]}(t_{n:k_{j+1}}), \quad (62)$$

$$\sup_{u \in \mathbb{R}^d, \|u - \gamma_{[x_n:z_n]}(t_{n:k_j})\| \geq \epsilon_{n:k_j}} \|\hat{\phi}_{n:j}(u) - u\| = 0 \quad (63)$$

for $k \in \{1, \dots, J_n - 1\}$. In particular, $\sup_{u \in \mathbb{R}^d} \|\hat{\phi}_{n:j}(u) - u\| \leq 2\epsilon_{n:k_j}$.

Let us describe how to choose $\{\epsilon_{n:k}\}_{k \in \mathbb{N}}$. We define this sequence recursively according to

$$\epsilon_{n,k} \triangleq \min \left\{ 3^{-N(n,k)}, 3^{-N(n,k)} \alpha \left(\hat{\phi}_{n:k-1} \circ \dots \circ \hat{\phi}_{n:1} \circ \dots \circ \hat{\phi}_{1:1}, \frac{1}{2^{N(n,k)}} \right) \right\},$$

where $N(n, k)$ is the number of elements in $\epsilon_{1,k_1}, \dots, \epsilon_{1,k_{J_1}}, \dots, \epsilon_{n,1}, \dots, \epsilon_{n,k_{J_n-1}}$. Note that, as discussed in (Anderson, 1967, page 367 preceding Lemma 2.1), the quantity

$$\alpha \left(\hat{\phi}_{n:k-1} \circ \dots \circ \hat{\phi}_{n:1} \circ \dots \circ \hat{\phi}_{1:1}, \frac{1}{2^{N(n,k)}} \right)$$

is positive by the compactness of K . Hence, the sequence $\{\epsilon_{n,k}\}_{k \in \mathbb{N}}$ is decreasing and non-negative. Thus, by the argument preceding (62), there exist well-defined sequence $\hat{\phi}_{1,1}, \dots, \hat{\phi}_{1,J_1}, \dots, \hat{\phi}_{n,1}, \dots, \hat{\phi}_{n,J_n}, \dots$ satisfying

$$\hat{\phi}_{n,J_n} \circ \dots \circ \hat{\phi}_{n,1} \circ \dots \circ \hat{\phi}_{1,J_1} \circ \dots \circ \hat{\phi}_{1,1}(x_n) = \phi(x_n) \quad (64)$$

$$\sup_{u \in \mathbb{R}^d} \|\hat{\phi}_{n:j}(u) - u\| < \epsilon_{n,k}. \quad (65)$$

Hence, (65) and the Inductive Convergence Criterion imply that the limit

$$\lim_{n \uparrow \infty} \hat{\phi}_{n,J_n} \circ \dots \circ \hat{\phi}_{n,1} \circ \dots \circ \hat{\phi}_{1,J_1} \circ \dots \circ \hat{\phi}_{1,1}$$

converges in $\mathcal{H}(\mathbb{R}^d)$ (with respect to the topology of uniform convergence on compacts). Denote its limit by $\hat{\phi}_\infty$. Next, (64) implies that $\hat{\phi}_\infty(x_n) = \phi(x_n)$ for every $n \in \mathbb{N}$. Moreover, (Munkres, 2000, Theorem 46.5) guarantees that the uniform limit of continuous functions

is again continuous, therefore $\hat{\phi}_\infty$ and ϕ are both continuous. Since continuous functions agreeing on dense subsets are equal, then we conclude that $\hat{\phi}_\infty = \phi$. In particular, for the given value of $\epsilon > 0$, there must be some integer $N(n, k) \in \mathbb{N}_+$ such that for every $N \in \mathbb{N}$ with $N \geq N(n, k)$ we have that:

$$\sup_{x \in K} \left\| (\hat{\phi}_{N, J_N} \circ \cdots \circ \hat{\phi}_{N, 1} \circ \cdots \circ \hat{\phi}_{1, J_1} \circ \cdots \circ \hat{\phi}_{1, 1})(x) - \phi(x) \right\| < \epsilon.$$

This concludes the proof. ■

6.6 Proof of Corollary 10

Lemma 21 *Let $\{x_n\}_{n=1}^N$ be points in X , where X is a compact subset of \mathbb{R}^d , and suppose that L and P satisfy Assumption 3.1. Then:*

- (i) *The map $g \mapsto \sum_{n=1}^N L(f(x_n), g(x_n), x_n) + P(g)$ is continuous from $C(X, \mathbb{R}^D)$ to \mathbb{R} ,*
- (ii) *There exists some $f^* \in C(\mathbb{R}^d, \mathbb{R}^D)$ satisfying*

$$f^* \in \operatorname{argmin}_{g \in C(X, \mathbb{R}^D)} \sum_{n=1}^N L(f(x_n), g(x_n), x_n) + P(g). \quad (66)$$

Proof [Proof of Lemma 21] Since \mathbb{R}^D and \mathbb{R} are metric spaces then (Munkres, 2000, Theorem 46.8) implies that the topology of uniform convergence and the compact-open⁸ topologies coincide on $C(X, \mathbb{R}^D)$ and on $C(X, \mathbb{R}^d)$. Since L and f are continuous and since X is compact, then (Munkres, 2000, Theorem 46.10) implies that, for every $n \in \{1, \dots, N\}$ the evaluation map:

$$g \mapsto L(f(x_n), g(x_n), x_n),$$

is continuous from $C(X, \mathbb{R}^D)$ to $C(X, \mathbb{R})$. Since the sum of (finitely many) continuous functions is continuous then the map

$$g \mapsto \sum_{n=1}^N L(f(x_n), g(x_n), x_n) + P(g), \quad (67)$$

is continuous on $C(X, \mathbb{R}^D)$. This gives (i).

Continuing our argument, for every $g \in C(\mathbb{R}^d, \mathbb{R}^D)$ the quantity $\sum_{n=1}^N L(f(x_n), g(x_n), x_n) + P(g)$ is finite since L and P were assumed to be finite-valued. Since L was assumed to be bounded-below, then in particular:

$$\min_{n=1, \dots, N, g \in C(X, \mathbb{R}^D)} L(f(x_n), g(x_n), x_n) \geq \min_{y_1, y_2 \in \mathbb{R}^D, x \in \mathbb{R}^d} L(f(x_n), g(x_n), x_n) > -\infty.$$

Now since P was also assumed to be bounded-below on $C(\mathbb{R}^d, \mathbb{R}^D)$ then in particular it is bounded-below on $C(X, \mathbb{R}^D)$; hence the functional of (67) is continuous, bounded-below,

8. See (Munkres, 2000, page 285) for a definition of the compact-open topology; however, its definition is never explicitly needed in our argument; but rather only its properties.

and finite-valued on $C(X, \mathbb{R}^D)$. Therefore, (Dal Maso, 1993, Theorem 1.15) applies therefore there exists some $f^* \in C(X, \mathbb{R}^D)$ satisfying (66). This gives (ii). \blacksquare

Proof [Proof of Corollary 10] By Lemma 21 (ii) there exists some $f^* \in C(X, \mathbb{R}^D)$ satisfying (66). By Theorem 3.3, for every full-rank matrix $A \in \text{Mat}_{d,d+D}$ and every $\delta > 0$, there exists a matrix $B^\delta \in \text{Mat}_{d+D,D}$ and some $\hat{\phi}^\delta \in \Phi_{*,d+D}$ such that

$$\sup_{x \in X} \left\| f^*(x) - B^\delta \hat{\phi}^\delta((I_d \oplus A)x) \right\| < \delta. \quad (68)$$

By Lemma 21 (i) the map $F : g \mapsto \sum_{n=1}^N L(f(x_n), g(x_n), x_n) + P(g)$ is continuous and therefore (68) implies that for every $\epsilon > 0$ we may pick some $\delta > 0$ such that:

$$F\left(B^\delta \hat{\phi}^\delta((I_d \oplus A)\cdot)\right) \leq F(f^*) + \epsilon.$$

Relabeling, yields the conclusion. \blacksquare

6.7 Proof of Corollary 12

Proof [Proof of Corollary 12] The existence of f^* is guaranteed by Lemma 21. Define the matrices $B \in \text{Mat}_{D \times d+D}$ and $A \in \text{Mat}_{d+D,d}$, respectively, by:

$$A \triangleq (I_d | 0_D) \text{ and } B \triangleq (0_d | I_D), \quad (69)$$

where 0_D and 0_d are respectively the respective zero matrices of $\text{Mat}_{D \times D}$ and of $\text{Mat}_{d \times d}$. Set $\hat{f}(z) \triangleq Bz$ and observe that $\hat{f} \in \mathcal{F}$ since \mathcal{F} contains all linear maps.

For every $n \leq N$ set

$$\begin{aligned} z_n &\triangleq (0, \dots, 0, f^*(x_n)) \in \mathbb{R}^{d+D} \\ x'_n &\triangleq (x_n, 0, \dots, 0) \in \mathbb{R}^{d+D}. \end{aligned}$$

Since $f^*(x_n) \neq f^*(x_m)$ for every $n \neq m$, with $n, m \leq N$, then for any $n, m \leq N$ if $n \neq m$ then $z_n \neq z_m$ and $x'_n \neq x'_m$ (the latter observation follows from the fact that if x'_n did equal to x'_m then x_n would equal to x_m and then $f(x_n)$ would equal to $f(x_m)$ which is a contradiction of our assumption). Since $x_n \neq 0$ for any $n \leq N$ then $x'_n \neq z_m$ for any $n, m \leq N$ since these must differ their first coordinate. Hence, $\{z_n\}_{n=1}^N \cup \{x'_n\}_{n=1}^N$ is a set of $2N$ distinct points in \mathbb{R}^{d+D} and $d \geq 1$, $D \geq 1$, thus $d + D \geq 2$. Since Δ of Theorem 4 is at-most equal to the right-hand side of (16) then Theorem 4 applies; hence, there exists some $\hat{\phi} \in \Phi_{*,d+D}$ satisfying

$$\hat{\phi}(x'_n) = z_n, \quad (70)$$

for every $n \leq N$. In particular, (70) and the definitions of A and B imply that for every $n \leq N$

$$\hat{f} \circ \hat{\phi}^\delta(Ax_n) = B\hat{\phi}^\delta(Ax_n) = B\hat{\phi}^\delta(x'_n) = Bz_n = f^*(x_n). \quad (71)$$

Thus, (i) holds.

Moreover, Theorem 4 implies that $\hat{\phi}^\delta$ has depth at-most $\left\lceil \frac{N\pi}{2(\min\{2\delta, 1\})} \right\rceil$; depending on our choice of $\delta > 0$; sufficiently small; and in particular, as specified by condition (16). This gives (iii).

By construction we have the containment:

$$\begin{aligned} & \left\{ (x, \hat{f}(Ax)) \in \mathbb{R}^{d+D} : (x, \hat{f}(Ax)) \neq \hat{\phi}^\delta(x, \hat{f}(Ax)) \right\} \\ & \subseteq \left\{ (x, y) \in \mathbb{R}^{d+D} : (x, y) \neq \hat{\phi}^\delta(x, y) \right\}, \end{aligned} \quad (72)$$

and the bound on the external covering number of the right-hand side of (72) is also guaranteed by Theorem 4. This gives (ii) and completes the proof. \blacksquare

6.8 Proof of Theorem 14

Proof [Proof of Theorem 14] Observe that the term $\sup_{x \in K} L(f(x_n), \hat{f}(x_n), x_n)$ does not depend on $\{w_n\}_{n=1}^N \in (0, 1)^n$; hence:

$$\begin{aligned} & \operatorname{argmin}_{\substack{w_n \in (0, 1) \\ \{w_n\}_{n=1}^N w_n = 1}} \sup_{x \in K} L(f(x_n), \hat{f}(x_n), x_n) - \left[\sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n) - \lambda \sum_{n \leq N} \frac{1}{N} \log \left(\frac{w_n}{N} \right) \right] \\ = & \operatorname{argmin}_{\substack{w_n \in (0, 1) \\ \{w_n\}_{n=1}^N w_n = 1}} - \left[\sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n) - \lambda \sum_{n \leq N} \frac{1}{N} \log \left(\frac{w_n}{N} \right) \right]. \end{aligned} \quad (73)$$

Notice that the set $\left\{ \{w_n\}_{n=1}^N : w_n \in (0, 1), \sum_{n \leq N} w_n = 1 \right\}$ describes the collection of probability measures \mathbb{Q} which are equivalent to the uniform probability measure \mathbb{P} on $\{x_n\}_{n=1}^N$. Thus, $\sum_{n \leq N} w_n L(f(x_n), \hat{f}(x_n), x_n)$ is the expectation of $L(f(x_n), \hat{f}(x_n), x_n)$ for a random variable X taking values on $\{x_n\}_{n=1}^N$ and $\sum_{n \leq N} \frac{1}{N} \log \left(\frac{w_n}{N} \right)$ is the relative entropy of the measure \mathbb{Q} with respect to \mathbb{P} , in the sense of (Wang et al., 2020, Definition 3). Therefore (Wang et al., 2020, Proposition 1) applies, thus, the measure \mathbb{P}^* defined by

$$\mathbb{P}^*(X = x_i) = \frac{e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}{\sum_{n \leq N} e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}$$

optimizes the right-hand side of (73). \blacksquare

Proof [Proof of Corollary 15] Fix $\hat{f} \in C(\mathbb{R}^d, \mathbb{R}^D)$ and let $\{w_n^{\lambda, \hat{f}}\}_{n=1}^N$ be as in Theorem 14 and use $\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N}$ to denote the measure on $\{x_n\}_{n=1}^N$ taking values with probabilities $\{w_n^{\lambda, \hat{f}}\}_{n=1}^N$. Then by Theorem 14

$$\operatorname{Gen}_M \left(\hat{f} \middle| \{w_n^{\lambda, \hat{f}}\}_{n=1}^N \right) + \lambda \operatorname{D}_{KL} \left(\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N} \middle\| \mathbb{P} \right) \leq \operatorname{Gen}_M \left(\hat{f} \middle| \left\{ \frac{1}{N} \right\}_{n=1}^N \right) + \lambda \operatorname{D}_{KL} (\mathbb{P} \middle\| \mathbb{P}). \quad (74)$$

By Gibbs' inequality, see (MacKay, 2003, Exercise 2.26), we have that $D_{KL}(\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N} \parallel \mathbb{P}) \geq 0$ and $D_{KL}(\mathbb{P} \parallel \mathbb{P}) = 0$. Thus, since $\lambda > 0$ then (74) reduces to

$$\begin{aligned} \text{Gen}_M(\hat{f} \parallel \{w_n^{\lambda, \hat{f}}\}_{n=1}^N) &\leq \text{Gen}_M(\hat{f} \parallel \{w_n^{\lambda, \hat{f}}\}_{n=1}^N) + \lambda D_{KL}(\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N} \parallel \mathbb{P}) \\ &\leq \text{Gen}_M\left(\hat{f} \parallel \left\{\frac{1}{N}\right\}_{n=1}^N\right) + \lambda D_{KL}(\mathbb{P} \parallel \mathbb{P}) \\ &= \text{Gen}_M\left(\hat{f} \parallel \left\{\frac{1}{N}\right\}_{n=1}^N\right), \end{aligned} \quad (75)$$

with equality in the left-most inequality if and only if $\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N} = \mathbb{P}$. Hence, from (75) we have that

$$0 \leq \text{Gen}_M\left(\hat{f} \parallel \left\{\frac{1}{N}\right\}_{n=1}^N\right) - \text{Gen}_M(\hat{f} \parallel \{w_n^{\lambda, \hat{f}}\}_{n=1}^N), \quad (76)$$

with equality if and only if $\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N} = \mathbb{P}$. Applying Theorem 14 we have an explicit form for $\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N}$ and therefore, we compute the right-hand side of (76) as follows

$$\begin{aligned} 0 &\leq \text{Gen}_M\left(\hat{f} \parallel \left\{\frac{1}{N}\right\}_{n=1}^N\right) - \text{Gen}_M(\hat{f} \parallel \{w_n^{\lambda, \hat{f}}\}_{n=1}^N) \\ &= \left[\sup_{x \in K} L(f(x_n), \hat{f}(x_n), x_n) - \frac{1}{N} \sum_{n \leq N} L(f(x_n), \hat{f}(x_n), x_n) \right] \\ &\quad - \left[\sup_{x \in K} L(f(x_n), \hat{f}(x_n), x_n) - \sum_{n \leq N} w_n^{\lambda, \hat{f}} L(f(x_n), \hat{f}(x_n), x_n) \right] \\ &= \sum_{n \leq N} w_n^{\lambda, \hat{f}} L(f(x_n), \hat{f}(x_n), x_n) - \frac{1}{N} \sum_{n \leq N} L(f(x_n), \hat{f}(x_n), x_n) \\ &= \sum_{n \leq N} \left(\frac{e^{-\lambda^{-1} L(\theta, x_n)}}{\sum_{n \leq N} e^{-\lambda^{-1} L(\theta, x_n)}} - \frac{1}{N} \right) L(f(x_n), \hat{f}(x_n), x_n). \end{aligned} \quad (77)$$

Lastly, note that $\mathbb{Q}^{\{w_n^{\lambda, \hat{f}}\}_{n=1}^N} = \mathbb{P}$ if and only if for each $1 \leq n \leq N$ we have

$$\frac{e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}{\sum_{n \leq N} e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}} = \frac{1}{N}. \quad (78)$$

Thus, (78) holds if and only if for every $1 \leq n, m \leq N$

$$\frac{e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}{\sum_{n \leq N} e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}} = \frac{1}{N} = \frac{e^{\lambda^{-1} L(\theta, x_m)}}{\sum_{n \leq N} e^{\lambda^{-1} L(f(x_n), \hat{f}(x_n), x_n)}}. \quad (79)$$

■