

Any-Dimensional Invariant Universality

Shengtai Yao*

Eitan Levin[†]

Mateo Díaz*

Abstract

Several machine learning models are defined for inputs of any size, such as graphs with different numbers of nodes and point clouds containing varying numbers of points. The universality properties of such any-dimensional models remain poorly understood, as universality is traditionally studied for models accepting inputs of a fixed size, defined on a compact subset of their domain. In sharp contrast, any-dimensional models can be viewed as sequences of functions defined on growing-sized inputs, and it is not clear in which sense they can be universal. We develop a systematic approach to establish any-dimensional universality, by identifying any-dimensional functions with a unique function taking inputs in a suitable infinite-dimensional limit space containing inputs of all finite sizes as well as their limits. Using the symmetries of these inputs and relations between inputs of different sizes, we show that this limit space admits a natural topology with rich families of compact sets on which any-dimensional universality can be established. We illustrate our approach by showing that several existing architectures fail to be universal, and we propose simple modifications that restore universality.

1 Introduction

Traditional supervised learning aims to learn mappings defined on fixed, finite-dimensional spaces by using training examples lying in those same spaces. In contrast, many modern learning tasks involve maps defined on inputs of varying dimensions. For instance, graph parameters, particle system dynamics, games, and regularizers remain meaningful regardless of the number of nodes, particles, players, or variables. There are a number of architectures in the literature that are defined on inputs of any dimension, e.g., DeepSets [1] for sets, Graph Neural Networks (GNN) [2] for graphs, and PointNet [3] for point clouds. Effectively, these any-dimensional architectures finitely parametrize an infinite sequence of invariant functions $(\hat{f}_n: V_n \rightarrow \mathbb{R})_n$ with increasingly larger input dimensions.¹ Despite the proliferation of any-dimensional architectures, their expressivity remains relatively unexplored. Specifically, the universality of most of the any-dimensional architectures above has mostly been studied in a fixed dimension or in finitely many dimensions [4, 5, 6]. This motivates our main question

Can any-dimensional models universally approximate functions across dimensions?

As stated, this question is ill-posed. To understand why, recall that classical universality concerns approximating a *single* continuous function $f: K \rightarrow \mathbb{R}$ on a fixed compact set K by a

*Department of Applied Mathematics and Statistics, Johns Hopkins University, Baltimore, MD 21218, USA. MD is partially supported by NSF awards CCF 2442615 and DMS 2502377, and a Sloan Research Fellowship.

[†]Department of Computing and Mathematical Sciences, Caltech, Pasadena, CA 91125, USA. EL is partially supported by AFOSR FA9550-23-1-0070 and FA9550-23-1-0204

¹By ‘invariant’ we mean that there is a sequence of groups (G_n) , such as permutations or rotations, such that for each n the group G_n acts on V_n and for all $g \in G_n$ and $v \in V_n$, we have $f_n(g \cdot v) = f_n(v)$.

model \widehat{f} uniformly on K . Any-dimensional models, however, produce a *sequence* of maps with varying domains. To resolve this mismatch, we adopt the framework of [7]. Specifically, we view the input spaces as a nested sequence $V_1 \subseteq V_2 \subseteq \dots$ of increasing-dimensional vector spaces, and embed them all into an infinite-dimensional limit space V_∞ we construct containing each V_n as a subspace. Under a mild compatibility condition, an any-dimensional model (\widehat{f}_n) admits a unique extension $\widehat{f}_\infty: V_\infty \rightarrow \mathbb{R}$ satisfying $\widehat{f}_\infty(x) = \widehat{f}_n(x)$ for any $x \in V_n$ and n . This identification lets us pose “universality across dimensions” as standard universality on an application-specific infinite-dimensional space.

Main contributions. Let us summarize our contributions.

(A recipe) We develop a general strategy for proving any-dimensional invariant universality. The strategy highlights two principles: (i) exploiting symmetries by passing to orbit spaces can yield rich families of compact sets, even in infinite dimensions, on which universality can be proved; and (ii) unlike in finite dimensions—where all norms induce the same topology—the choice of norm on V_∞ is essential, as continuity (and hence admissible activations and architectural primitives) depends on it.

(Instantiations) We apply this recipe to three representative domains: sets, graphs, and point clouds. We show that several widely used architectures yield extensions \widehat{f}_∞ that are either discontinuous in the natural topology induced by the learning task or fail to be universal. We then propose modifications—and, when necessary, new architectures—that restore continuity and achieve universality.

Outline. The remainder of this section reviews related work. Section 2 reviews the mathematical preliminaries of any dimensional learning necessary for the rest of the paper. Section 3 formalizes the type of any-dimensional universality that we wish to establish and outlines a general recipe for establishing universality. Section 4 presents concrete instantiations of this recipe for models that apply to sets, graphs, and point clouds. Finally, Section 5 concludes the paper with limitations and opportunities for future work.

1.1 Related work

Dimension-free learning. Many modern architectures are any-dimensional. Convolutional neural networks apply the same translation-invariant filters to images of arbitrary resolution [8, 9]. Recurrent neural networks and transformers handle variable-length sequences via recurrence or attention [10, 11, 12]. Neural operators such as DeepONet and FNO learn maps between (infinite-dimensional) function spaces and can be evaluated on grids of different resolutions [13, 14]. In geometric deep learning, several models have been proposed to handle sets of arbitrary cardinality [1, 3], graphs with varying numbers of nodes and edges [2, 15], and point clouds with any size [16, 17]. Recently, [18, 19, 20] leveraged *representation stability* [21, 22] to derive general-purpose dimension-free equivariant and invariant model classes, encoding neural networks, convex sets, and kernel machines. However, the ability to accept variable sizes does not guarantee consistent (or transferable) behavior across sizes. The study of transferability originated in the GNN literature [23] and has since been explored extensively in that context [24, 25, 26, 27, 28, 29]. More recently, [7] proposed a unified framework—beyond graphs—for formulating and certifying transferability across dimensions.

Universality. Universality is one of the pillars of deep learning theory. Let us start by commenting on finite-dimensional results. Classical universal approximation theorems show that fully connected networks approximate continuous functions on compact sets [30, 31, 32, 33], with recent extensions beyond compact domains under suitable growth and activation assumptions [34, 35, 36]. The expressivity of symmetric models has only recently been studied. For graphs, expressive power is subtle: message-passing GNNs are known to have intrinsic limitations [37], commonly formalized via the Weisfeiler–Lehman (WL) test [38], and many variants have been analyzed through this lens [37, 39, 40]. Alternative GNN based on tensor layers [4, 5], and those based on homomorphism density functions [41, 42] do achieve universality. For permutation-invariant problems, DeepSets is universal [1, 43]. Beyond these specific examples, [6] established universality of polynomial-invariant-based models that are symmetric under the action of any compact group. Nonetheless, these results are largely formulated for a fixed input dimension. By contrast, universality for any-dimensional models remains comparatively underdeveloped. The closest related results are [44, 45] for sets and [5, 46] for graphs. Our results and strategy highlight a more general perspective based on the transferability framework developed in [7], applicable beyond these concrete scenarios. Let us briefly comment on more specific differences. First, [45] studies lower bounds on approximation for certain RKHS classes; these classes are not our focus. Our DeepSets results extend [44] by covering a broader class of compact sets. On the graph side, [25] proves universality of graphon neural networks in simplified regimes, while [46] establishes universality with respect to the δ_p metric. In contrast, we work with the arguably more natural and widely studied cut metric δ_\square , which induces a coarser topology.

2 Preliminaries

In this section, we introduce the notation and technical background necessary for the paper.

Notation. We use \mathbb{R} and \mathbb{N} to denote the reals and positive integers, respectively. Define $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$ and $[n] := \{1, \dots, n\}$ for $n \in \mathbb{N}$. We write $O(k)$ for the orthogonal group in dimension k and S_n for the symmetric group on n letters. For a given metric space (\mathcal{X}, d) , we denote the open ball centered at x of radius η by $B(x, \eta) := \{u \in \mathcal{X} \mid d(u, x) < \eta\}$. Let $C(\mathcal{X}, \mathcal{Y})$ denote the space of continuous maps from a topological space \mathcal{X} to \mathcal{Y} , and write $C(\mathcal{X})$ when $\mathcal{Y} = \mathbb{R}$. Let $\mathcal{B}(\mathcal{X}, \mathcal{Y})$ be the space of bounded linear operators between normed spaces \mathcal{X} and \mathcal{Y} . Given an arbitrary norm $\|\cdot\|_{\mathbb{R}^k}$ on \mathbb{R}^k , the symbol $\ell_p(\mathbb{R}^k)$ denotes the space of sequences $X = (X_j \in \mathbb{R}^k)_{j \in \mathbb{N}}$ such that $\sum_{j=1}^{\infty} \|X_j\|_{\mathbb{R}^k}^p < \infty$. Similarly, for a given measure space (\mathbb{R}^m, μ) , let $L^p(\mathbb{R}^m; \mathbb{R}^k)$ be the space of measurable functions $X: \mathbb{R}^m \rightarrow \mathbb{R}^k$ with $\|X\|_p := \left(\int \|X(t)\|_{\mathbb{R}^k}^p d\mu(t)\right)^{1/p} < \infty$. The symbol $\|\cdot\|_p$ refers to either the ℓ^p - or L^p -norms, as appropriate. Finally, let $\mathcal{P}(\mathbb{R}^k)$ be the set of probability measures on \mathbb{R}^k , and let $\mathcal{P}_p(\mathbb{R}^k) \subseteq \mathcal{P}(\mathbb{R}^k)$ denote those with finite p th moments. We use the symbol W_p to denote the Wasserstein- p distance.

Any-dimensional learning. Next, we recall a few basic notions pertaining to any-dimensional models and their generalization across dimensions. Our exposition here follows [7]; we refer the interested reader to that reference for additional details. We focus exclusively on real-valued functions, and accordingly streamline several definitions, though all notions in this section extend to more general codomains. The key idea of this section is that a sequence of functions $(f_n: V_n \rightarrow \mathbb{R})$ with growing-sized inputs can be identified with a single extension $f_\infty: \overline{V_\infty} \rightarrow \mathbb{R}$, defined on an infinite-dimensional space $\overline{V_\infty}$ that ‘contains’ each V_n as a subspace. This viewpoint lets us pose

questions such as universality on one common domain, rather than across a sequence of domains. To state these ideas precisely, we start from a family of vector spaces encoding inputs of growing size, together with maps and relations linking them.

Definition 2.1. A *consistent sequence* is a triple $\mathbb{V} = \{(V_n)_{n \in \mathbb{N}}, (\varphi_{N,n})_{n \preceq N}, (G_n)_{n \in \mathbb{N}}\}$ indexed by a directed poset (\mathbb{N}, \preceq) ,² of finite dimensional vector spaces V_n , maps $\varphi_{N,n}$ and groups G_n acting linearly on V_n such that, for all $n \preceq N$, (i) the group G_n is embedded into G_N and (ii) $\varphi_{N,n}: V_n \hookrightarrow V_N$ is a linear, G_n -equivariant embedding.

The elements of different vector spaces in a consistent sequence represent objects of different sizes, while the groups and embeddings between them represent their symmetries and relations between objects of different sizes, respectively. To crystallize this notion, we provide two examples of consistent sequences that will play a crucial role in this work.

Example 2.2. Let $V_n = \mathbb{R}^n$ and $G_n = S_n$ be the symmetric group acting by coordinate permutation. One might endow this sequence with two different types of orderings and embeddings.

1. (**Zero-padding embedding**) Index by \mathbb{N} with the standard ordering \leq . Define $\varphi_{N,n}: V_n \hookrightarrow V_N$ by

$$\varphi_{N,n}(x_1, \dots, x_n) := (x_1, \dots, x_n, \underbrace{0, \dots, 0}_{(N-n) \text{ zeros}}).$$

We use the symbol \mathbb{V}_{zero} to denote this sequence.

2. (**Duplication embedding**) Index by \mathbb{N} with $n \preceq N$ if n divides N . Define $\varphi_{N,n}: V_n \hookrightarrow V_N$ by repeating each coordinate N/n times,

$$\varphi_{N,n}(x_1, \dots, x_n) = (\underbrace{x_1, \dots, x_1}_{N/n \text{ copies}}, \dots, \underbrace{x_n, \dots, x_n}_{N/n \text{ copies}}).$$

Similarly, we use \mathbb{V}_{dup} to denote this sequence.

To define and study continuous any-dimensional models, we further define a single, infinite-dimensional space containing inputs of all finite sizes. This is done in the following definition.

Definition 2.3. Define V_∞ as the disjoint union $\bigsqcup V_n$ modulo an equivalence relation

$$V_\infty := \bigsqcup_n V_n / \sim,$$

where $v \sim \varphi_{N,n}(v)$ whenever $n \preceq N$, and denote by $[v] \in V_\infty$ the equivalence class of $v \in V_n$. The limiting group G_∞ and its equivalence classes $[g]$ are defined analogously.

In turn, this allows us to identify any finite-dimensional element $v \in V_n$ with its infinite-dimensional counterpart $[v] \in V_\infty$. The next definition will allow us to extend a sequence of functions to this infinite dimensional space.

Definition 2.4. Let $\mathbb{V} = \{(V_n), (\varphi_{N,n}), (G_n)\}$ be a consistent sequence indexed by (\mathbb{N}, \preceq) . A sequence $(f_n: V_n \rightarrow \mathbb{R})$ is **compatible** with respect to \mathbb{V} if $f_N \circ \varphi_{N,n} = f_n$ for all $n \preceq N$, and each f_n is G_n -equivariant.

²A directed poset is a partial order \preceq on \mathbb{N} such that every two elements have a common upper bound.

In other words, a compatible sequence of functions is one that takes the same value on “equivalent” inputs, i.e., those related by symmetries and embeddings. A sequence of functions (f_n) is compatible if and only if there exists a unique G_∞ -equivariant map $f_\infty: V_\infty \rightarrow \mathbb{R}$ such that $f_\infty|_{V_n} = f_n$ for all $n \in \mathbb{N}$ [7]. Thus, we can identify any compatible sequence of functions with a unique extension that matches the functions when restricted to finite dimensions.

In order to reason about continuity of f_∞ , we will endow V_∞ with a metric. To do so, we consider a sequence of *compatible invariant norms* $(\|\cdot\|_{V_n}: V_n \rightarrow \mathbb{R})$. Compatibility yields the existence of a G_∞ -invariant norm $\|\cdot\|_{V_\infty}$ on V_∞ . This construction allows us to define a limit space that contains not only (equivalent classes of) finite-dimensional objects, but also their limits.

Definition 2.5. *The **limit space** is the pair $(\overline{V_\infty}, G_\infty)$ where $\overline{V_\infty}$ denotes the completion of V_∞ with respect to $\|\cdot\|_{V_\infty}$, endowed with the symmetrized metric*

$$\overline{d}(x, y) := \inf_{g \in G_\infty} \|g \cdot x - y\|_{V_\infty} \quad \text{for } x, y \in \overline{V_\infty}. \quad (1)$$

The **orbit space** is the metric space $(\overline{V_\infty}/G_\infty, \overline{d})$ with

$$\overline{V_\infty}/G_\infty := \overline{\{G_\infty \cdot x \mid x \in \overline{V_\infty}\}},$$

where the outer completion is taken with respect to \overline{d} .

This symmetrized metric is a pseudometric on $\overline{V_\infty}$ and a metric on the space of orbit closures of $\overline{V_\infty}$ under the action of G_∞ [7]. We slightly abuse notation, by denoting this set $\overline{V_\infty}/G_\infty$ and calling it “orbit space.” Two elements $x, y \in \overline{V_\infty}$ are in the same orbit (closure) if, and only if, $\overline{d}(x, y) = 0$. Let us provide a quick example that will be used recurrently.

Example 2.6 (Duplication sequence with ℓ_p norms). *Consider the duplication embedding from Example 2.2. Endow each $V_n = \mathbb{R}^n$ with the normalized ℓ_p norm for $p \in [1, \infty)$, i.e., $\|x\|_{V_n} = (\frac{1}{n} \sum_i x_i^p)^{1/p}$. Then, the limit space V_∞ can be identified with $L^p([0, 1])$, with permutations S_n acting by permuting the n consecutive intervals of length $1/n$ in the domain $[0, 1]$ of functions in $L^p([0, 1])$. The orbit space $\overline{V_\infty}/G_\infty$ corresponds to the space of probability distributions $\mathcal{P}_p(\mathbb{R})$. Furthermore, the symmetrized distance \overline{d} coincides with the Wasserstein- p metric. Details appear in Appendix C.1.2.*

Armed with these notions, we are now ready to define continuity across dimensions.

Definition 2.7. *Let \mathbb{V} be a consistent sequence endowed with a compatible invariant norm $(\|\cdot\|_{V_n})$. A sequence of invariant functions $(f_n: V_n \rightarrow \mathbb{R})$ is **continuously transferable** if (i) there exists an extension $f: \overline{V_\infty} \rightarrow \mathbb{R}$, i.e., $f_n = f_\infty|_{V_n}$ for all n , and (ii) the extension is continuous with respect to $\|\cdot\|_{V_\infty}$.*

We remark that if (f_n) is continuously transferable, then it must be compatible. Further, an extension f_∞ is continuous in the limit space with respect to $\|\cdot\|_{V_\infty}$ if, and only if, it is continuous in the orbit space $\overline{V_\infty}/G_\infty$ with respect to \overline{d} ; see Lemma A.4 in Appendix A. The term ‘transferable’ is motivated by the results in [7], which show that if a high-dimensional model is transferable, then it can be trained on low-dimensional inputs and generalize to higher-dimensional ones.

3 How to establish any-dimensional invariant universality?

In this section, we outline a general recipe for establishing any-dimensional invariant universality. We begin by formally stating our goal. Consider a continuously transferable, consistent sequence of invariant functions $(f_n: V_n \rightarrow \mathbb{R})$. Building upon the foundational concepts from Section 2, our

primary focus is on approximating its G_∞ -invariant extension f_∞ . To approximate this function we consider a class of continuously transferable any-dimensional models \mathcal{F} . We identify each of these models with its limiting extension \hat{f}_∞ and our goal is to show that: for any fixed accuracy $\varepsilon > 0$ and compact set K in the domain of f_∞ , there exists $\hat{f}_\infty \in \mathcal{F}$ such that

$$\sup_{x \in K} |f_\infty(x) - \hat{f}_\infty(x)| \leq \varepsilon. \quad (2)$$

We make two remarks concerning domains and topology.

Domain of the extension. Since invariant functions are constant on orbits, they admit two equivalent viewpoints: as functions on the limit space $f_\infty: \overline{V_\infty} \rightarrow \mathbb{R}$, or as functions on the orbit space $f_\infty: \overline{V_\infty}/G_\infty \rightarrow \mathbb{R}$. Thus, the question of approximation can be posed over any of these domains. Any consistent norm $\|\cdot\|_{V_\infty}$ induces a topology over the limit space, the corresponding topology on the orbit space is given by the symmetrized metric (1). As it turns out, compact sets in the orbit space are ‘richer’ than those in the limit space. The underlying reason for this phenomenon is that by taking orbits, we collapse ‘big’ subsets (potentially noncompact) into single orbits. To illustrate this point, consider the setting in Example 2.6 where $\overline{V_\infty} = L^p([0, 1])$ are integrable functions and $\overline{V_\infty}/G_\infty = \mathcal{P}_p(\mathbb{R})$ are probability distributions with p th moments. While the set $\widetilde{K} = \{X \in L^p([0, 1]) \mid \text{range}(X) \subseteq [0, 1]\}$ is not compact in $L^p([0, 1])$, the collection of its equivalence classes $K = \{\mu \in \mathcal{P}_p(\mathbb{R}) \mid \text{supp}(\mu) \subseteq [0, 1]\}$ is compact in $\mathcal{P}_p(\mathbb{R})$; we defer details to Appendix B. In Section 4, we will see more examples of this phenomena. In what follows, we focus on universality on the orbit space equipped with the symmetrized metric, i.e., we prove (2) for compact subsets K of the orbit space.

Induced topology and activation functions. In finite dimensional settings, all norms induce equivalent topologies, which means that continuity with respect to one norm guarantees continuity under any other. In sharp contrast, the topology of the limit space $\overline{V_\infty}$ (and correspondingly that of the orbit space) is heavily dependent on the choice of the consistent norm $\|\cdot\|_{V_\infty}$. This choice is often dictated by the learning task, i.e., the map f_∞ we aim to approximate (or in practice learn) is naturally continuous with respect to certain norms but not others. To illustrate this point, consider again the setting in Example 2.6 with $p = 2$, and the sequence of ‘second moment’ functions

$$f_n(x) = \frac{1}{n} \sum_{j=1}^n x_j^2 \quad \text{with extension} \quad f_\infty(\mu_X) = \mathbb{E}_{Y \sim \mu_X}[Y^2]$$

where $\mu_X \in \mathcal{P}_2(\mathbb{R})$ denotes the probability law associated with $X \in L^2([0, 1])$. In this case, f_∞ is continuous with respect to W_2 , but discontinuous with respect to W_1 . As we will see in Section 4, several existing any-dimensional architectures are not continuous in the natural topologies induced by norms of interest. A recurring theme is that the chosen topology constrains which activation functions and architectural primitives can be used.

Desiderata and a recipe. We now outline our strategy for proving universality. Our proof strategy is based on the Stone–Weierstrass theorem, which we recall for convenience; see [47][Theorem 5.7] for a proof.

Theorem 3.1 (Stone-Weierstrass). *Let $C(K)$ be the set of continuous real-valued functions on a compact metric space K endowed with the norm $\|f\|_\infty = \sup_{x \in K} |f(x)|$. Let \mathcal{F} be a subalgebra of $C(K)$ which contains a non-zero constant function.³ Then \mathcal{F} is dense in $C(K)$ if, and only if, it*

³A subalgebra of functions is a set closed under addition, scalar multiplication, and products.

separates points.⁴

Density in the sup-norm is equivalent to (2). This discussion suggests two desiderata for practical model classes \mathcal{F} : (i) functions in \mathcal{F} should be continuous with respect to the natural topology induced by the learning task; and (ii) the class \mathcal{F} should form a subalgebra that separates points. Somewhat surprisingly, many existing architectures fail to meet one (or both) of these requirements, likely because most prior work studies approximation on a fixed dimension V_n rather than on the orbit space. Guided by these desiderata, we propose the following three-step recipe.

Step 1 (Compact sets). Given a consistent norm $\|\cdot\|_\infty$ dictated by the learning task, characterize compact sets in the orbit space $\overline{V_\infty}/G_\infty$.

Step 2 (Continuity). Construct a collection \mathcal{F} of invariant any-dimensional models that are continuously transferable with respect to the symmetrized metric (1) on the orbit space $\overline{V_\infty}/G_\infty$.

Step 3 (Universality). Establish universality of \mathcal{F} via the Stone–Weierstrass theorem by showing that \mathcal{F} forms a subalgebra that separates points.

4 Instantiations of the recipe

In this section, we instantiate our general recipe for any-dimensional models on three representative domains: sets (Section 4.1), graphs (Section 4.2), and point clouds (Section 4.3). For brevity, we work directly with the corresponding infinite-dimensional extensions of our models, rather than the underlying sequences. Across all domains, we find that standard architectures are either discontinuous, or fail to be universal. We therefore either introduce minor, principled modifications to these existing architectures or propose new architectures, and we prove that the resulting models are both continuous and universal.

4.1 Set functions

We start with models for sets of any size. Since fully connected feedforward neural networks appear repeatedly throughout, we define the family

$$\text{NN}_{k,m}^\phi := \bigcup_{r=1}^{\infty} \{L_2 \circ \phi \circ L_1 \mid L_1 \in \mathcal{L}_{r,m}, L_2 \in \mathcal{L}_{k,r}\}, \quad (3)$$

where $\mathcal{L}_{m,n} := \{x \mapsto Wx + \theta \mid W \in \mathbb{R}^{m \times n}, \theta \in \mathbb{R}^m\}$ is the collection of affine maps from \mathbb{R}^n to \mathbb{R}^m and $\phi: \mathbb{R} \rightarrow \mathbb{R}$ is an activation function applied component-wise. Observe that these networks can be arbitrarily wide.

We consider two models whose input spaces parallel the consistent sequences in Example 2.2. Although both sequences share the same finite-dimensional spaces $V_n = \mathbb{R}^{n \times k}$, they use different embeddings and consistent norms, and therefore induce substantially different orbit spaces. In both settings we start from the DeepSets architecture [1], which maps

$$\text{DeepSets}^{\rho,\sigma}(X) = \sigma \left(\sum_{i=1}^{\infty} \rho(X_i) \right), \quad (4)$$

where $\rho \in \text{NN}_{r,k}^\phi$, $\sigma \in \text{NN}_{1,r}^\phi$, $r \in \mathbb{N}$. For each of the two orbit spaces we consider, we introduce a mild modification of DeepSets that restores continuity and yields universality.

⁴The function class \mathcal{F} separates points on K if for all $x, y \in K$ and $x \neq y$, there is an $f \in \mathcal{F}$ such that $f(x) \neq f(y)$.

4.1.1 Universality over sequences

Additional details about the following constructions can be found in Appendix C.1.1. For the first example, we consider the k -fold direct sum of the zero-padding consistent sequence $\mathbb{V}_{\text{zero}}^{\oplus k} = \{(\mathbb{R}^{n \times k}), (\varphi_{N,n}^{\oplus k}), (S_n)\}$. Given an arbitrary norm $\|\cdot\|_{\mathbb{R}^k}$ in \mathbb{R}^k , and a real number $p \in [1, \infty)$, we endow $\mathbb{V}_{\text{zero}}^{\oplus k}$ with the compatible sequence of ℓ_p norm $\|X\|_{V_n} = (\sum_{i=1}^n \|X_{i:}\|_{\mathbb{R}^k}^p)^{1/p}$, which are unchanged under zero-padding. The resulting limit space is the space of sequences $\overline{V_\infty} = \ell_p(\mathbb{R}^k)$, with the group G_∞ acting by permuting finitely many entries of these sequences.

Step 1 (Compact sets). For this example, we consider orbit closures of compact sets $K \subseteq \ell_p(\mathbb{R}^k)$, which we denote by K/G_∞ . These are compact subsets of $\overline{V_\infty}/G_\infty$ by the definition of the quotient topology on $\overline{V_\infty}/G_\infty$. Compact sets in $\ell_p(\mathbb{R}^k)$ are well-understood.

Proposition 4.1 (Compactness in $\ell_p(\mathbb{R}^k)$, [48]). *For $p \in [1, \infty)$, a set $K \subseteq \ell_p(\mathbb{R}^k)$ is compact if, and only if, it satisfies that (i) it is closed; (ii) it is bounded, i.e., $\sup_{X \in K} \|X\|_p < \infty$; and (iii) it exhibits uniform tail decay, i.e., $\lim_{N \rightarrow \infty} \sup_{X \in K} (\sum_{i \geq N} \|X_{i:}\|_{\mathbb{R}^k}^p)^{1/p} = 0$.*

For instance, the set $\{(x_i)_{i=1}^\infty : |x_i| \leq 2^{-i} \text{ for all } i \in \mathbb{N}\}$ is compact in $\ell_p(\mathbb{R})$ for all $p \in [1, \infty)$. We now turn to constructing a family of continuous models on $\overline{V_\infty}/G_\infty$.

Step 2 (Continuity). Our architecture is based on DeepSets (4). To ensure that our architecture is compatible with respect to our consistent sequence structure, i.e., that it is unchanged under zero-padding, we assume that $\rho(\mathbf{0}_k) = \mathbf{0}_r$. In general, (4) is not a continuous model as the infinite sum may diverge since the function ρ can decay arbitrarily slowly. Motivated by this observation, we propose

$$\text{DeepSets}_\infty^{\rho, \sigma}(X) = \sigma \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \cdot \rho(X_{i:}) \right), \quad (5)$$

where $\rho \in \text{NN}_{r,k}^\phi$ and $\sigma \in \text{NN}_{1,r}^\phi$ with $r \in \mathbb{N}$. The following proposition shows that $\text{DeepSets}_\infty^{\rho, \sigma}$ is indeed continuous on the space $\overline{V_\infty}/G_\infty$ of orbit closures when both ρ and σ are continuous. The proof is deferred to Appendix C.2.

Theorem 4.2 (Continuity on sequences). *Suppose $\rho \in C(\mathbb{R}^k, \mathbb{R}^r)$ and $\sigma \in C(\mathbb{R}^r, \mathbb{R})$ are continuous functions. Then, the map $\text{DeepSets}_\infty^{\rho, \sigma} : \overline{V_\infty}/G_\infty \rightarrow \mathbb{R}$ defined in (5) is continuous with respect to the symmetrized metric (1).*

As an immediate consequence, we get that all functions in

$$\mathcal{F}_{\text{DS}} := \left\{ \text{DeepSets}_\infty^{\rho, \sigma} \mid r \in \mathbb{N}, \rho \in \text{NN}_{r,k}^\phi, \sigma \in \text{NN}_{1,r}^\phi \right\} \text{ are continuous.}$$

Step 3 (Universality). Finally, we prove that the above variation of DeepSets is universal for sequences. The proof of the following result is deferred to Appendix C.3.

Theorem 4.3 (Universality of DeepSets_∞). *Suppose that the activation ϕ is continuous and non-polynomial. Let $K \subseteq \ell_p(\mathbb{R}^k)$ be any compact set. Then, \mathcal{F}_{DS} is dense in $C(K/G_\infty)$.*

4.1.2 Universality over measures

Additional details about the constructions in this section are deferred to Appendix C.1.2. We consider the k -fold direct sum of the duplication embedding consistent sequence $\mathbb{V}_{\text{dup}}^{\oplus k} := \{(\mathbb{R}^{n \times k}), (\varphi_{N,n}^{\oplus k}), (S_n)\}$

defined as in Example 2.2. Given an arbitrary norm $\|\cdot\|_{\mathbb{R}^k}$ on \mathbb{R}^k , and a real number $p \in [1, \infty)$, we endow each $\mathbb{R}^{n \times k}$ with the normalized ℓ_p norm, i.e., $\|X\|_{\bar{p}} = (\frac{1}{n} \sum_{i=1}^n \|X_i\|_{\mathbb{R}^k}^p)^{1/p}$. As explained in Example 2.6, the resulting limit space is the space of functions $L^p([0, 1], \mathbb{R}^k)$, and the space of orbit closures $\overline{V_\infty}/G_\infty$ can be identified with the Wasserstein space $\mathcal{P}_p(\mathbb{R}^k)$ equipped with the Wasserstein- p distance W_p .

Step 1 (Compact sets). We begin with a characterization of compact sets in the orbit space.

Proposition 4.4 (Compacta in $\mathcal{P}_p(\mathbb{R}^k)$ [49]). *For $p \in [1, \infty)$, a subset $Q \subseteq \mathcal{P}_p(\mathbb{R}^k)$ is compact if, and only if, it is (i) closed; (ii) tight, i.e., for any $\varepsilon > 0$, there is a compact $K_\varepsilon \subseteq \mathbb{R}^k$ such that $\sup_{\mu \in Q} \mu(\mathbb{R}^k \setminus K_\varepsilon) \leq \varepsilon$; and (iii) p -uniformly integrable, i.e.,*

$$\lim_{R \rightarrow \infty} \sup_{\mu \in Q} \int_{\|x\|_{\mathbb{R}^k} > R} \|x\|_{\mathbb{R}^k}^p d\mu(x) = 0.$$

For example, the collection of all measures supported on a closed ball of radius $R > 0$, $Q = \{\mu \in \mathcal{P}_p(\mathbb{R}^k) \mid \text{supp}(\mu) \subseteq \overline{B(0, R)}\}$, is compact in this space. We now turn to constructing a family of continuous models on $\mathcal{P}_p(\mathbb{R}^k)$ that can approximate any continuous function on a compact subset to arbitrary precision.

Step 2 (Continuity). We consider the normalized DeepSets architecture proposed in [44],

$$\overline{\text{DeepSets}}_\infty^{\rho, \sigma}(\mu) = \sigma\left(\int \rho d\mu\right), \quad (6)$$

where $\rho \in \text{NN}_{r, k}^\phi$ and $\sigma \in \text{NN}_{1, r}^\phi$, with $r \in \mathbb{N}$. To establish the continuity of this architecture we require that the entries of ρ do not grow too fast. In particular, we denote by $\mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^r)$ the class of continuous functions $\rho = (\rho_1, \dots, \rho_r)^\top$ such that each component ρ_j satisfies a p -th order growth condition; specifically, for each $j \in [r]$, there exists a constant $M_j > 0$ such that

$$|\rho_j(x)| \leq M_j(1 + \|x\|_{\mathbb{R}^k}^p), \quad \forall x \in \mathbb{R}^k.$$

The following theorem establishes the continuity of $\overline{\text{DeepSets}}_\infty^{\rho, \sigma}$ under this growth condition. We defer its proof to Appendix C.4.

Theorem 4.5 (Continuity on probability measures). *Suppose $\rho \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^r)$ and $\sigma \in C(\mathbb{R}^r, \mathbb{R})$ are continuous functions. Then, the map $\overline{\text{DeepSets}}_\infty^{\rho, \sigma} : \mathcal{P}_p(\mathbb{R}^k) \rightarrow \mathbb{R}$, defined in (6), is continuous with respect to W_p .*

As an immediate consequence of this theorem, we obtain that all functions in

$$\mathcal{F}_{\overline{\text{DS}}} := \left\{ \overline{\text{DeepSets}}_\infty^{\rho, \sigma} \mid r \in \mathbb{N}, \rho \in \text{NN}_{r, k}^\phi, \sigma \in \text{NN}_{1, r}^\phi \right\}$$

are continuous provided that ϕ is Lipschitz, since it ensures that $\text{NN}_{r, k}^\phi \subseteq \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^r)$, for $p \in [1, \infty)$.

Step 3 (Universality) Finally, we establish the universality of the normalized DeepSets architecture in (6). We defer the proof of this result to Appendix C.5.

Theorem 4.6 (Universality of $\overline{\text{DeepSets}}_\infty$). *Suppose that ϕ is a Lipschitz continuous, non-polynomial activation that is asymptotically polynomial at $\pm\infty$.⁵ Let $Q \subseteq \mathcal{P}_p(\mathbb{R}^k)$ be a compact set. Then, $\mathcal{F}_{\overline{\text{DS}}}$ is dense in $C(Q)$.*

⁵We say that $\phi(t)$ is asymptotically polynomial at $\pm\infty$ if there exist polynomials P_+ and P_- (possibly constant) such that $\phi(t)/P_\pm(t) \rightarrow 1$ as $t \rightarrow \pm\infty$, respectively.

The special case of Theorem 4.6 with Q consisting of all measures supported on a fixed compact set was proved in [44]. The requirement that ϕ is asymptotically polynomial at $\pm\infty$ guarantees the universality of neural networks on non-compact domains [34]. Representative activation functions meeting these conditions include ReLU, Softplus, and Sigmoid, among others.

4.2 Graph functions

Additional details about the constructions in this section can be found in Appendix D.1. Recall the divisibility ordering from the duplication embedding in Example 2.2. We consider a consistent sequence indexed with this ordering and given by $\mathbb{V}_{\text{dup}}^G = \{(\mathbb{R}_{\text{sym}}^{n \times n}), (\varphi_{N,n}), (S_n)\}$, where $\mathbb{R}_{\text{sym}}^{n \times n}$ denotes the space of $n \times n$ symmetric matrices with real entries, and $\varphi_{N,n}$ converts each entry into $N/n \times N/n$ blocks with the same entry. We endow $\mathbb{R}_{\text{sym}}^{n \times n}$ with the cut norm, $\|A\|_{\square} = \frac{1}{n^2} \max_{I, J \subseteq [n]} |\sum_{i \in I, j \in J} A_{ij}|$. The limit space $\overline{\mathbb{V}_{\infty}}$ can be identified with the space of kernels \mathcal{W} , i.e., measurable symmetric functions $W: [0, 1]^2 \rightarrow \mathbb{R}$ equipped with the cut norm $\|\cdot\|_{\square}$. The symmetrized metric (1) corresponds to the so-called cut metric δ_{\square} . This limit space has been widely studied in the literature on large graphs or graphons [50]. Intuitively, the functions W model adjacency matrices of infinite-node graphs.

Step 1 (Compact sets). As is standard in the graphon literature, we consider graphs with bounded edge weights; in particular, we restrict them to the unit interval $[0, 1]$. With this restriction, we focus on the subset of orbits corresponding to graphons, that is,

$$\mathcal{W}_0 := \{W \in \mathcal{W} \mid \text{range}(W) \subseteq [0, 1]\} / \sim, \quad (7)$$

where $W_1 \sim W_2$ whenever $\delta_{\square}(W_1, W_2) = 0$ [50]. This is the collection of all limits of simple graphs, i.e., graphs with edge weights in $\{0, 1\}$. This entire graphon space is compact.

Proposition 4.7 (Compactness of \mathcal{W}_0 , [50]). *The space \mathcal{W}_0 endowed with the cut metric δ_{\square} is compact.*

Step 2 (Continuity). Achieving continuity in the cut metric is subtle. In particular, pointwise nonlinearities are not continuous in cut metric: for any pointwise map ϕ , the operator $W \mapsto \phi(W)$ is continuous with respect to $\|\cdot\|_{\square}$ if and only if ϕ is linear [46]. Since continuity with respect to $\|\cdot\|_{\square}$ on the limit space is equivalent to continuity with respect to δ_{\square} on the orbit space (Lemma A.4), this precludes using pointwise activations. As it turns out, many existing graph architectures are discontinuous in $(\mathcal{W}_0, \delta_{\square})$ [4, 5]. One way to circumvent this issue is to coarsen the topology to make pointwise nonlinearities continuous, as in [46]. In contrast, we work directly with the cut distance δ_{\square} , which metrizes a standard notion of graph limits [50].

The continuity with respect to this distance is closely tied to *homomorphism densities*. Intuitively, these are functionals that quantify how frequently a fixed motif (a finite graph) appears in a graphon. Formally, given a motif $F = (V, E)$ and a graphon $W \in \mathcal{W}_0$, the homomorphism density is

$$t(F, W) = \int_{[0,1]^{|V|}} \prod_{ij \in E} W(x_i, x_j) \prod_{i \in V} dx_i. \quad (8)$$

In turn, the linear span of homomorphism densities of simple graphs, $\mathcal{HD} := \text{span}\{t(F, \cdot) \mid F \text{ is simple}\}$, is dense in the space of continuous functions $C(\mathcal{W}_0, \delta_{\square})$ [50]. Thus, it suffices to find a model class that can be expressed as linear combinations of homomorphism densities of simple graphs. Importantly, the desired model class must also not contain any homomorphism densities of non-simple graphs, as these are discontinuous in the cut metric [50]. With these considerations in

mind, we propose the following family of models

$$h_m^{a,b}(W) = \int_{[0,1]^m} \prod_{1 \leq i < j \leq m} (a_{ij}W(x_i, x_j) + b_{ij}) \prod_{\ell=1}^m dx_\ell, \quad (9)$$

where $a_{ij}, b_{ij} \in \mathbb{R}$ are parameters. In particular, the restriction $i < j$ ensures that only simple motifs can be parameterized. Based on this, we define the function class

$$\mathcal{F}_{\mathcal{W}} := \text{span}\{h_m^{a,b}(\cdot) \mid m \in \mathbb{N}, m \geq 2, a, b \in \mathbb{R}_{\text{sym}}^{m \times m}\}. \quad (10)$$

As we detail in Appendix D.4, this family can be represented by a neural-network-like deep architecture, making it directly amenable to gradient-based optimization. This property is informally stated as follows.

Theorem 4.8 (Informal). *The family of functions $\mathcal{F}_{\mathcal{W}}$ can be parameterized using a neural network architecture.*

The neural network architecture in this result is somewhat involved; to streamline the presentation, we defer to Appendix D.4. We mention in passing that the nonlinearities correspond to tensor contractions, which have been proposed in the past to enhance the expressivity of invariant architectures [51]. Our discussion thus far yields the following result; its formal proof appears in Appendix D.2.

Theorem 4.9 (Continuity of $\mathcal{F}_{\mathcal{W}}$). *All functions in $\mathcal{F}_{\mathcal{W}}$ are continuous with respect to δ_{\square} .*

Step 3 (Universality). Separating points in the graphon space $(\mathcal{W}_0, \delta_{\square})$ corresponds to distinguishing graphs up to weak isomorphism—a fundamental task in the study of graph neural networks. Prior work on the expressive power of standard GNN architectures [37, 39, 40] has linked their expressivity to the Weisfeiler–Lehman (WL) graph isomorphism test [38], which fails to distinguish certain non-isomorphic graphs. In contrast, the following theorem establishes that our model class (10) is universal; its proof is deferred to Appendix D.3.

Theorem 4.10 (Universality of $\mathcal{F}_{\mathcal{W}}$). *The class of functions $\mathcal{F}_{\mathcal{W}}$ in (10) is dense in $C(\mathcal{W}_0, \delta_{\square})$.*

The proof of Theorems 4.9 and 4.10 essentially proceed by expanding out the product in the definition of our model (9) and observing that it is a linear combination of homomorphism densities. By appropriately choosing the parameters a, b in (9), we further show that all homomorphism densities lie within our parametric family $\mathcal{F}_{\mathcal{W}}$.

We emphasize at this point that there are two key advantages to working with the parametrization (9) over directly learning coefficients in a linear combination of homomorphism densities, as done in [41, 42] for example. The first advantage pertains to the number of parameters needed to represent all homomorphism densities of simple graphs on m nodes. While the number of such densities scales as $\mathcal{O}\left(2^{\binom{m}{2}}/m!\right)$ and these densities are all linearly independent, our model is able to represent all of them using only $m(m-1)$ parameters. The second advantage pertains to the implementation and training of our model. Learning coefficients in a linear combination of homomorphism densities requires computing each homomorphism density separately for every motif and every input graph, which is costly. Our parameterization avoids this by encoding graph patterns implicitly through the learnable coefficients a, b in (9). Moreover, our model is directly amenable to gradient-based optimization thanks to its neural network-like decomposition in Appendix D.4.

The key enabler of universality here is the use of high-order tensors as hidden layers; the order of these tensors—as well as the width and depth of our networks—grows with m . Whether universal approximation is achievable with fixed width or depth remains an interesting open question.

4.3 Point cloud functions

Details of the constructions in this section appear in Appendix E.1. We consider a consistent sequence analogous to the one used for sets (measures), namely $\mathbb{V}_{\text{dup}}^P = \{(\mathbb{R}^{n \times k}), (\varphi_{N,n}^{\oplus k}), (S_n \times O(k))\}$, where $\mathbb{R}^{n \times k}$ represents sets of n points in \mathbb{R}^k for fixed $k \in \mathbb{N}$, and the maps $\varphi_{N,n}$ are duplication embeddings. The key difference is that the group is now $S_n \times O(k)$, where S_n is the permutation group and $O(k)$ is the orthogonal group, acting via $(g, h) \cdot X = gXh^\top$. We endow \mathbb{R}^k with the Euclidean norm $\|\cdot\|_{\mathbb{R}^k}$, and equip each V_n with the normalized ℓ_2 norm; the limit space can then be identified with $\overline{V_\infty} = L^2([0, 1]; \mathbb{R}^k)$.

Since equivariant operators appear repeatedly throughout, we define the family

$$\mathbf{LE}_{k \rightarrow l} := \left\{ L \in \mathbb{B}(k, l) \mid L \text{ is } S_{[0,1]}\text{-equivariant} \right\}, \quad (11)$$

where $\mathbb{B}(k, l) = \mathcal{B}\left(L^2([0, 1]^k), L^2([0, 1]^l)\right)$ denotes the space of bounded linear operators and $S_{[0,1]}$ is the group of measure-preserving bijections $\varphi: [0, 1] \rightarrow [0, 1]$. We say that L is $S_{[0,1]}$ -equivariant if, for every $\varphi \in S_{[0,1]}$ and every $W \in L^2([0, 1]^k)$, one has $L(W^\varphi) = L(W)^\varphi$ almost everywhere, where $W^\varphi(x_1, \dots, x_k) := W(\varphi(x_1), \dots, \varphi(x_k))$.

Step 1 (Compact sets). For this setting, we consider the set of orbit closures arising from a set of bounded functions in the limit space $\overline{V_\infty}$, namely

$$K_R := \{X \in L^2([0, 1]; \mathbb{R}^k) \mid \|X\|_{L^\infty} \leq R\},$$

where $R > 0$ is a positive constant. The next proposition establishes that the space of orbit closures of K_R , which we denote by K_R/G_∞ , is compact in $\overline{V_\infty}/G_\infty$. The proof is deferred to Appendix E.2.

Proposition 4.11 (Compactness of K_R/G_∞). *The set of orbit closures K_R/G_∞ is a compact subset of $\overline{V_\infty}/G_\infty$.*

Step 2 (Continuity). Next, we introduce a class of continuous models on the space of orbit closures. The architecture decomposes into two stages: given a point cloud X , we first form its Gram (inner-product) matrix $X^\top X$, and then apply a graph neural network (GNN) to this matrix. Intuitively, the Gram map enforces invariance to $O(k)$, while the GNN provides invariance to permutations S_n . Closely related architectures were first proposed in [16].

Let us formally define the architecture in the orbit space. For the first stage, we map each $X \in K_R/G_\infty$ to a graphon $W_X: [0, 1]^2 \rightarrow [0, 1]$ defined by

$$W_X(x, y) := \frac{1}{2R^2} \langle X(x), X(y) \rangle + \frac{1}{2}, \quad (12)$$

where we interpret X as any representative in its equivalence class. Notice that W_X is a symmetric measurable function, and so can be seen as a graphon. We endow the space of graphons with the δ_2 metric given by

$$\delta_2(W, \widetilde{W}) := \inf_{\varphi \in S_{[0,1]}} \|W - \widetilde{W}^\varphi\|_2.$$

For the second stage, we use Invariant Graphon Networks (IGNs) [46], i.e., functions $I_{\varrho, M, L, b}: \mathcal{W}_0 \rightarrow \mathbb{R}$ of the form

$$I_{\varrho, M, L, b}(W) := \sum_{m=1}^M L_m^{(2)}\left(\varrho\left(L_m^{(1)}(W) + b_m^{(1)}\right)\right) + b^{(2)}, \quad (13)$$

where $M \in \mathbb{N}_0$, $L_m^{(1)} \in \mathbf{LE}_{2 \rightarrow k_m}$, $L_m^{(2)} \in \mathbf{LE}_{k_m \rightarrow 0}$, $b_m^{(1)}, b_m^{(2)} \in \mathbb{R}$, and $k_m \in \mathbb{N}$ for each $m \in \{1, \dots, M\}$. Here $\varrho: \mathbb{R} \rightarrow \mathbb{R}$ is a nonlinearity acting pointwise.

Altogether, we consider the family of models given by

$$\mathcal{F}_{PC}^\varrho := \{X \mapsto I_{\varrho, M, L, b}(W_X)\}. \quad (14)$$

Here the architectural hyperparameters and parameters (i.e., M , the operators L , and the biases b) range over the admissible sets described after (13); we omit the full specification for brevity. The next theorem shows that this architecture is indeed continuous with respect to the symmetrized distance. The proof is in Appendix E.3.

Theorem 4.12 (Continuity of \mathcal{F}_{PC}^ϱ). *Let ϱ be continuous. Then, all models in \mathcal{F}_{PC}^ϱ are continuous with respect to the symmetrized metric.*

Step 3 (Universality). To close, we show that our model class for point clouds (14) is universal; the proof is deferred to Appendix E.4.

Theorem 4.13 (Universality of \mathcal{F}_{PC}^ϱ). *Let ϱ be continuous and non-polynomial. Then, \mathcal{F}_{PC}^ϱ is dense in $C(K_R/G_\infty)$.*

The above theorem gives universality in the space of continuous functions of point clouds that are invariant under rotations of the cloud. If we center the input point cloud before applying our proposed architecture, we also obtain universality in the smaller space of continuous functions invariant under all rigid motions of their inputs.

5 Conclusions and future work

To summarize, we have (i) formalized universality for any-dimensional invariant architectures, (ii) proposed a general recipe for constructing and certifying any-dimensional universal model classes, and (iii) used this recipe to prove universality for several concrete families. A key limitation of our framework is that it applies only to scalar-valued outputs. Many practically relevant any-dimensional architectures have both input and output sizes growing and are naturally equivariant as opposed to invariant. In this setting, the problem no longer reduces to studying functionals on orbit spaces. Developing tools that exploit such equivariance to obtain universality results for more general architectures is an important direction for future work. Another limitation is that we only focus on dense graphs in our current instantiations. There are several notions of limits for sparse graphs including [52, 53, 54], and applying our framework to these limits is another interesting direction.

References

- [1] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. *Advances in neural information processing systems*, 30, 2017.
- [2] Marco Gori, Gabriele Monfardini, and Franco Scarselli. A new model for learning in graph domains. In *Proceedings. 2005 IEEE international joint conference on neural networks, 2005.*, volume 2, pages 729–734. IEEE, 2005.
- [3] Charles R Qi, Hao Su, Kaichun Mo, and Leonidas J Guibas. Pointnet: Deep learning on point sets for 3d classification and segmentation. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 652–660, 2017.

- [4] Haggai Maron, Ethan Fetaya, Nimrod Segol, and Yaron Lipman. On the universality of invariant networks. In *International conference on machine learning*, pages 4363–4371. PMLR, 2019.
- [5] Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. *Advances in neural information processing systems*, 32, 2019.
- [6] Dmitry Yarotsky. Universal approximations of invariant maps by neural networks. *Constructive Approximation*, 55(1):407–474, 2022.
- [7] Eitan Levin, Yuxin Ma, Mateo Díaz, and Soledad Villar. On transferring transferability: Towards a theory for size generalization. *arXiv preprint arXiv:2505.23599*, 2025.
- [8] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 2002.
- [9] Mahdi Hashemi. Enlarging smaller images before inputting into convolutional neural network: zero-padding vs. interpolation. *Journal of Big Data*, 6(1):1–13, 2019.
- [10] Sepp Hochreiter and Jürgen Schmidhuber. Long short-term memory. *Neural computation*, 9(8):1735–1780, 1997.
- [11] Richard Socher, Cliff C Lin, Chris Manning, and Andrew Y Ng. Parsing natural scenes and natural language with recursive neural networks. In *Proceedings of the 28th international conference on machine learning (ICML-11)*, pages 129–136, 2011.
- [12] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [13] Lu Lu, Pengzhan Jin, Guofei Pang, Zhongqiang Zhang, and George Em Karniadakis. Learning nonlinear operators via deepnet based on the universal approximation theorem of operators. *Nature machine intelligence*, 3(3):218–229, 2021.
- [14] Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier neural operator for parametric partial differential equations. *arXiv preprint arXiv:2010.08895*, 2020.
- [15] Franco Scarselli, Marco Gori, Ah Chung Tsoi, Markus Hagenbuchner, and Gabriele Monfardini. The graph neural network model. *IEEE transactions on neural networks*, 20(1):61–80, 2008.
- [16] Soledad Villar, David W Hogg, Kate Storey-Fisher, Weichi Yao, and Ben Blum-Smith. Scalars are universal: Equivariant machine learning, structured like classical physics. *Advances in neural information processing systems*, 34:28848–28863, 2021.
- [17] Ben Blum-Smith, Ningyuan Huang, Marco Cuturi, and Soledad Villar. Functions on symmetric matrices and point clouds via lightweight invariant features from galois theory. *SIAM Journal on Applied Algebra and Geometry*, 9(4):902–938, 2025.
- [18] Eitan Levin and Venkat Chandrasekaran. Free descriptions of convex sets. *arXiv preprint arXiv:2307.04230*, 2023.
- [19] Eitan Levin and Mateo Díaz. Any-dimensional equivariant neural networks. In *International Conference on Artificial Intelligence and Statistics*, pages 2773–2781. PMLR, 2024.

- [20] Mateo Díaz, Dmitriy Drusvyatskiy, Jack Kendrick, and Rekha R Thomas. Invariant kernels: Rank stabilization and generalization across dimensions. *arXiv preprint arXiv:2502.01886*, 2025.
- [21] Thomas Church and Benson Farb. Representation theory and homological stability. *Advances in Mathematics*, 245:250–314, 2013.
- [22] Thomas Church, Jordan S Ellenberg, and Benson Farb. Fi-modules and stability for representations of symmetric groups. *Duke Mathematical Journal*, 164(9), 2015.
- [23] Luana Ruiz, Luiz Chamon, and Alejandro Ribeiro. Graphon neural networks and the transferability of graph neural networks. *Advances in Neural Information Processing Systems*, 33:1702–1712, 2020.
- [24] Ron Levie, Wei Huang, Lorenzo Bucci, Michael Bronstein, and Gitta Kutyniok. Transferability of spectral graph convolutional neural networks. *Journal of Machine Learning Research*, 22(272):1–59, 2021.
- [25] Nicolas Keriven, Alberto Bietti, and Samuel Vaiter. Convergence and stability of graph convolutional networks on large random graphs. *Advances in Neural Information Processing Systems*, 33:21512–21523, 2020.
- [26] Luana Ruiz, Luiz FO Chamon, and Alejandro Ribeiro. Transferability properties of graph neural networks. *IEEE Transactions on Signal Processing*, 71:3474–3489, 2023.
- [27] Sohir Maskey, Ron Levie, and Gitta Kutyniok. Transferability of graph neural networks: an extended graphon approach. *Applied and Computational Harmonic Analysis*, 63:48–83, 2023.
- [28] Matthieu Cordonnier, Nicolas Keriven, Nicolas Tremblay, and Samuel Vaiter. Convergence of message passing graph neural networks with generic aggregation on random graphs. In *GSP 2023-6th Graph Signal Processing workshop*, pages 1–3, 2023.
- [29] Chen Cai and Yusu Wang. Convergence of invariant graph networks. In *International Conference on Machine Learning*, pages 2457–2484. PMLR, 2022.
- [30] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.
- [31] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. Multilayer feedforward networks are universal approximators. *Neural networks*, 2(5):359–366, 1989.
- [32] 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.
- [33] Kurt Hornik. Approximation capabilities of multilayer feedforward networks. *Neural networks*, 4(2):251–257, 1991.
- [34] Teun DH van Nuland. Noncompact uniform universal approximation. *Neural Networks*, 173:106181, 2024.
- [35] Ariel Neufeld and Philipp Schmocker. Universal approximation results for neural networks with non-polynomial activation function over non-compact domains. *arXiv preprint arXiv:2410.14759*, 2024.

- [36] Ahmed Abdeljawad and Thomas Dittrich. Weighted sobolev approximation rates for neural networks on unbounded domains. *arXiv preprint arXiv:2411.04108*, 2024.
- [37] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? *arXiv preprint arXiv:1810.00826*, 2018.
- [38] Andrei Leman and Boris Weisfeiler. A reduction of a graph to a canonical form and an algebra arising during this reduction. *Nauchno-Technicheskaya Informatsiya*, 2(9):12–16, 1968.
- [39] Haggai Maron, Heli Ben-Hamu, Hadar Serviansky, and Yaron Lipman. Provably powerful graph networks. *Advances in neural information processing systems*, 32, 2019.
- [40] Christopher Morris, Martin Ritzert, Matthias Fey, William L Hamilton, Jan Eric Lenssen, Gaurav Rattan, and Martin Grohe. Weisfeiler and leman go neural: Higher-order graph neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 4602–4609, 2019.
- [41] Takanori Maehara and Hoang NT. A simple proof of the universality of invariant/equivariant graph neural networks. *arXiv preprint arXiv:1910.03802*, 2019.
- [42] Hoang Nguyen and Takanori Maehara. Graph homomorphism convolution. In *International Conference on Machine Learning*, pages 7306–7316. PMLR, 2020.
- [43] Akiyoshi Sannai, Yuuki Takai, and Matthieu Cordonnier. Universal approximations of permutation invariant/equivariant functions by deep neural networks. *arXiv preprint arXiv:1903.01939*, 2019.
- [44] Christian Bueno and Alan Hylton. On the representation power of set pooling networks. *Advances in Neural Information Processing Systems*, 34:17170–17182, 2021.
- [45] Aaron Zweig and Joan Bruna. A functional perspective on learning symmetric functions with neural networks. In *International Conference on Machine Learning*, pages 13023–13032. PMLR, 2021.
- [46] Daniel Herbst and Stefanie Jegelka. Higher-order graphon neural networks: Approximation and cut distance. *arXiv preprint arXiv:2503.14338*, 2025.
- [47] Walter Rudin. *Functional Analysis*. International Series in Pure and Applied Mathematics. McGraw-Hill, New York, 2 edition, 1991.
- [48] Joseph Diestel. *Sequences and series in Banach spaces*, volume 92. Springer Science & Business Media, 2012.
- [49] Luigi Ambrosio, Nicola Gigli, and Giuseppe Savaré. *Gradient flows: in metric spaces and in the space of probability measures*. Springer, 2005.
- [50] László Lovász. *Large networks and graph limits*. American Mathematical Soc., 2012.
- [51] Marc Finzi, Max Welling, and Andrew Gordon Wilson. A practical method for constructing equivariant multilayer perceptrons for arbitrary matrix groups. In *International conference on machine learning*, pages 3318–3328. PMLR, 2021.

- [52] Christian Borgs, Jennifer Chayes, Henry Cohn, and Yufei Zhao. An L^p theory of sparse graph convergence I: Limits, sparse random graph models, and power law distributions. *Transactions of the American Mathematical Society*, 372(5):3019–3062, 2019.
- [53] Christian Borgs, Jennifer T. Chayes, Henry Cohn, and Nina Holden. Sparse exchangeable graphs and their limits via graphon processes. *Journal of Machine Learning Research*, 18(210):1–71, 2018.
- [54] Ágnes Backhausz and Balázs Szegedy. Action convergence of operators and graphs. *Canadian Journal of Mathematics*, 74(1):72–121, 2022.
- [55] Stephen Willard. *General topology*. Courier Corporation, 2012.
- [56] Walter Rudin. Principles of mathematical analysis. *3rd ed.*, 1976.
- [57] Cédric Villani et al. *Optimal transport: old and new*, volume 338. Springer, 2008.
- [58] Vladimir I Bogachev. *Measure theory*. Springer, 2007.
- [59] C. Borgs, J.T. Chayes, L. Lovász, V.T. Sós, and K. Vesztegombi. Convergent sequences of dense graphs I: Subgraph frequencies, metric properties and testing. *Advances in Mathematics*, 219(6):1801–1851, 2008.
- [60] Peter Diao, Dominique Guillot, Apoorva Khare, and Bala Rajaratnam. Differential calculus on graphon space. *Journal of Combinatorial Theory, Series A*, 133:183–227, 2015.
- [61] Sashi Mohan Srivastava. *A course on Borel sets*. Springer, 1998.
- [62] Stephen H Friedberg, Arnold J Insel, and Lawrence E Spence. *Linear algebra*. Prentice Hall, 1997.

A Missing details from Section 2

In this section, we present missing details from Section 2. We start with a more detailed definition of a consistent sequence.

Definition A.1 (Detailed version of Def. 2.1). *A **consistent sequence** of group representations over directed poset (\mathbb{N}, \preceq) is a triple $\mathbb{V} = \{(V_n)_{n \in \mathbb{N}}, (\varphi_{N,n})_{n \preceq N}, (G_n)_{n \in \mathbb{N}}\}$ consisting of the following elements.*

1. (**Groups**) *A sequence of groups (G_n) indexed by \mathbb{N} that embed into each other. Specifically, whenever $n \preceq N$ there is an injective group homomorphism $\theta_{N,n}: G_n \rightarrow G_N$ with*

$$\begin{aligned} \theta_{i,i} &= \text{id}_{G_i} && \text{for all } i \in \mathbb{N}, \\ \theta_{k,j} \circ \theta_{j,i} &= \theta_{k,i} && \text{whenever } i \preceq j \preceq k \text{ in } \mathbb{N}. \end{aligned}$$

2. (**Vector spaces**) *A sequence of finite-dimensional, real vector spaces (V_n) indexed by \mathbb{N} such that each V_n is a G_n -representation.*
3. (**Embeddings**) *A collection of embeddings $(\varphi_{N,n}: V_n \hookrightarrow V_N)_{n \preceq N}$ such that $\varphi_{N,n}$ is G_n -equivariant, i.e.,*

$$\varphi_{N,n}(g \cdot v) = \theta_{N,n}(g) \cdot \varphi_{N,n}(v) \text{ for all } g \in G_n, v \in V_n.$$

and such that

$$\begin{aligned}\varphi_{i,i} &= \text{id}_{V_i} && \text{for all } i \in \mathbb{N}, \\ \varphi_{k,j} \circ \varphi_{j,i} &= \varphi_{k,i} && \text{whenever } i \preceq j \preceq k \text{ in } \mathbb{N}.\end{aligned}$$

We can take direct sums of consistent sequences to obtain richer consistent sequences, as done in several of the examples of Section 4.

Definition A.2. *The k -fold direct sum of \mathbb{V} is defined as*

$$\mathbb{V}^{\oplus k} := \{(V_n^{\oplus k}), (\varphi_{N,n}^{\oplus k}), (G_n)\}$$

where $V_n^{\oplus d}$ denotes the direct sum of d copies of V_n and $\varphi_{N,n}^{\oplus d} : V_n^{\oplus d} \rightarrow V_n^{\oplus d}$ is defined by applying $\varphi_{N,n}$ to each component. The group G_n acts on $V_n^{\oplus d}$ by simultaneously acting on every copy of V_n , i.e. $g \cdot (v_1, \dots, v_d) := (g \cdot v_1, \dots, g \cdot v_d)$.

The next two propositions show that compatibility (Definition 2.4) is equivalent to the existence of a limiting extension.

Proposition A.3 ([7, Prop. B.6]). *Let $\mathbb{V} = \{(V_n), (\varphi_{N,n}), (G_n)\}$ and $\mathbb{U} = \{(U_n), (\psi_{N,n}), (G_n)\}$ be consistent sequences. A sequence of maps $(f_n : V_n \rightarrow U_n)$ is compatible if, and only if, it extends to the limit, i.e., there exists a G_∞ -equivariant map $f_\infty : V_\infty \rightarrow U_\infty$ such that $f_n = f_\infty|_{V_n}$ for all n .*

In particular, a sequence of norms $\|\cdot\|_{V_n}$ is compatible if and only if they extend to a norm $\|\cdot\|_{V_\infty}$ preserved by the action of G_∞ .

Finally, we show that continuity of invariant functions in the limit space is equivalent to continuity of the corresponding induced function on the orbit space.

Lemma A.4. *Let $\|\cdot\|_{V_\infty}$ be the extension of a compatible sequence of norms $\|\cdot\|_{V_n}$. Let $f_\infty : \overline{V_\infty} \rightarrow \mathbb{R}$ be G_∞ -invariant function, and let \bar{f}_∞ be the induced function on the space of orbit closures $\overline{V_\infty}/G_\infty$. Then, f_∞ is continuous on $(\overline{V_\infty}, \|\cdot\|_{V_\infty})$ if, and only if, \bar{f}_∞ is continuous on $(\overline{V_\infty}/G_\infty, \bar{d})$.*

Proof of Lemma A.4. Let $\pi : \overline{V_\infty} \rightarrow \overline{V_\infty}/G_\infty$ be the quotient map $\pi(x) = [x]$. Because f_∞ is G_∞ -invariant, we can write $\bar{f}_\infty \circ \pi$ for some $\bar{f}_\infty : \overline{V_\infty}/G_\infty \rightarrow \mathbb{R}$. We prove the two implications.

(\Leftarrow) Suppose $\bar{f}_\infty : \overline{V_\infty}/G_\infty \rightarrow \mathbb{R}$ is continuous with respect to \bar{d} . Note that π is continuous with respect to the norm $\|\cdot\|_{V_\infty}$ on $\overline{V_\infty}$ and symmetrized metric \bar{d} on $\overline{V_\infty}/G_\infty$, since $\bar{d}(\pi(x), \pi(y)) = \inf_{g \in G_\infty} \|g \cdot x - y\|_{V_\infty} \leq \|x - y\|_{V_\infty}$. Therefore, the composition $f_\infty = \bar{f}_\infty \circ \pi$ is continuous with respect to $\|\cdot\|_{V_\infty}$.

(\Rightarrow) Suppose $f_\infty : \overline{V_\infty} \rightarrow \mathbb{R}$ is continuous with respect to $\|\cdot\|_{V_\infty}$. Then for any $\varepsilon > 0$ and any $[x] \in \overline{V_\infty}/G_\infty$ there exists $\delta > 0$ such that any $y \in \overline{V_\infty}$ with $\|x - y\|_{V_\infty} < \delta$ satisfies $|f_\infty(x) - f_\infty(y)| < \varepsilon$. Now, consider any $[y'] \in \overline{V_\infty}/G_\infty$ satisfying $\bar{d}([x], [y']) < \delta$, then by the definition of \bar{d} , there exists an element $g_0 \in G_\infty$ such that $\|g_0 \cdot x - y'\|_{V_\infty} < \delta$. Since for any $g \in G_\infty$, the map $x \mapsto g \cdot x$ is an isometry on $\overline{V_\infty}$, we have $\|x - g_0^{-1} \cdot y'\|_{V_\infty} = \|g_0 \cdot x - y'\|_{V_\infty} < \delta$, which implies $|f_\infty(x) - f_\infty(g_0^{-1} \cdot y')| < \varepsilon$. Since f_∞ is G_∞ -invariant, we obtain $|\bar{f}_\infty([x]) - \bar{f}_\infty([y'])| = |f_\infty(x) - f_\infty(g_0^{-1} \cdot y')| < \varepsilon$, and conclude that \bar{f}_∞ is continuous. \square

B Missing details from Section 3

In this section, we establish the claim that the set of functions in $L^p[0, 1]$ with uniformly bounded range is not compact. This is a well known fact, but we prove it for completeness.

Lemma B.1. *The set $\widetilde{K} = \{X: [0, 1] \rightarrow [0, 1]\}$ is not compact in $L^p([0, 1])$ for any $p \in [1, \infty]$.*

Proof of Lemma B.1. It suffices to exhibit an infinite sequence $(X_n) \subseteq \widetilde{K}$ that has no convergent subsequence. To this end, let $X_n(t) = 1$ if $t \in \bigcup_{i=1}^{2^{n-1}} \left[\frac{2i-1}{2^n}, \frac{2i}{2^n} \right]$ and $X_n(t) = 0$ otherwise. Note that $\|X_n - X_m\|_p = \frac{1}{4^{1/p}}$ for all $n \neq m$, so (X_n) cannot have a convergent subsequence in L^p . \square

C Missing proofs and additional details from Section 4.1

In this section, we present the proofs and additional details about the constructions for set problems.

C.1 Consistent sequences on sets

We start by elaborating on the consistent sequences from Section 4.1. These were studied in detail in [7, Appx. F], and we refer the reader there for more details and proofs.

C.1.1 Zero-padding consistent sequence \mathbb{V}_{zero} with ℓ_p norm

The zero-padding consistent sequence $\mathbb{V}_{\text{zero}} = \{(V_n), (\varphi_{N,n}), (G_n)\}$ is defined as Example 2.2. Here we set $G_n = S_n$ to be the permutation group acting on \mathbb{R}^n by permuting coordinates, i.e., $(g \cdot x)_i = x_{g^{-1}(i)}$ for $g \in S_n$. For $n \leq N$, the embedding of groups $\theta_{N,n}: S_n \rightarrow S_N$ is given by

$$\theta_{N,n}(g) = \begin{bmatrix} g & 0 \\ 0 & I_{N-n} \end{bmatrix} \quad \text{for } g \in S_n.$$

For $p \in [1, \infty)$, each V_n is equipped with the ℓ_p -norms $\|x\|_p = (\sum_{i=1}^n |x_i|^p)^{1/p}$ which are permutation-invariant. By Proposition A.3, this induces a norm on V_∞ , also denoted as $\|\cdot\|_p$. Consequently, the limit space is identified with the classical sequence space

$$\overline{V_\infty} = \ell_p = \left\{ x = (x_i)_{i=1}^\infty : \|x\|_p = \left(\sum_{i=1}^\infty |x_i|^p \right)^{1/p} < \infty \right\}.$$

The symmetrized distance $\bar{d}_p(x, y) = \inf_{g \in G_\infty} \|g \cdot x - y\|_p$ defines a metric on the space of orbit closures $\overline{V_\infty}/G_\infty$.

Next, consider the k -fold direct sum $\mathbb{V}_{\text{zero}}^{\oplus k} = \{(\mathbb{R}^{n \times k}), (\varphi_{N,n}^{\oplus k}), (S_n)\}$ of \mathbb{V}_{zero} , as defined in Definition A.2. Fix an arbitrary norm $\|\cdot\|_{\mathbb{R}^k}$ on \mathbb{R}^k , and define the ℓ_p -norm on $\mathbb{R}^{n \times d}$ as

$$\|X\|_p = \left(\sum_{i=1}^n \|X_{i:}\|_{\mathbb{R}^k}^p \right)^{1/p},$$

where $X_{i:}$ denotes the i -th row of X . The corresponding limit space can be represented as:

$$\overline{V_\infty} = \ell_p(\mathbb{R}^k) = \left\{ X = (X_{i:})_{i=1}^\infty : \|X\|_p = \left(\sum_{i=1}^\infty \|X_{i:}\|_{\mathbb{R}^k}^p \right)^{1/p} < \infty \right\}.$$

C.1.2 Duplication consistent sequence \mathbb{V}_{dup} with normalized ℓ_p -norms

The duplication embedding consistent sequence $\mathbb{V}_{\text{dup}} = \{(V_n), (\varphi_{N,n}), (G_n)\}$ is defined as Example 2.2. The group $G_n = S_n$ again acts on \mathbb{R}^n by permuting coordinates. For $n|N$, the embedding of groups $\theta_{N,n}: S_n \rightarrow S_N$ is given by $\theta_{N,n}(g) = g \otimes I_{N/n}$, where \otimes denotes the Kronecker product.

The space V_∞ can be identified with the space of step functions, whose discontinuity points are rational. More precisely, each $x \in \mathbb{R}^n$ corresponds to a step function $f_x : [0, 1] \rightarrow \mathbb{R}$ given by

$$f_x(t) = x_i \quad \text{for } t \in \left(\frac{i-1}{n}, \frac{i}{n} \right], i \in [n],$$

and $f_x(0) = x_1$. For $p \in [1, \infty)$, each V_n is equipped with the normalized ℓ_p -norms $\|x\|_{\bar{p}} = \left(\frac{1}{n} \sum_{i=1}^n |x_i|^p \right)^{1/p}$, which are compatible. Under the identification with step functions, this corresponds to the L^p norm of functions, given by $\|f_x\|_p = \left(\int_0^1 |f_x(t)|^p dt \right)^{1/p}$. By Proposition A.3, this induces a norm on V_∞ . Consequently, the limit space can be identified with the L^p space, specifically,

$$\overline{V_\infty} \cong L^p([0, 1]) = \left\{ f : [0, 1] \rightarrow \mathbb{R} \text{ measurable} : \int_0^1 |f(t)|^p dt < \infty \right\}.$$

The permutations in S_n act on functions on $L^p([0, 1])$ by permuting consecutive intervals of length $1/n$. Formally, each $\sigma \in S_n$ defines a measure-preserving bijection $\tilde{\sigma} : [0, 1] \rightarrow [0, 1]$ via $\tilde{\sigma}((i-1)/n + t) = (\sigma(i) - 1)/n + t$ for $t \in [0, 1/n)$ and $i \in [n]$ (and $\tilde{\sigma}(1) = 1$ for simplicity).

A function $f \in \overline{V_\infty}$ gives rise to a probability measure μ_f defined as the distribution of $f(T)$ for T uniformly sampled from $[0, 1]$. Equivalently, $\mu_f = f_{\#}\lambda$ is the pushforward under f of the Lebesgue measure λ on $[0, 1]$. All elements in the orbit of f correspond to the same measure μ_f , and conversely, two functions are in the orbit-closures of each other if and only if they correspond to the same measure. Thus, the orbit space $\overline{V_\infty}/G_\infty$ can be identified with $\mathcal{P}_p(\mathbb{R})$. Furthermore, the symmetrized distance coincides with the Wasserstein- p distance in this case.

Likewise, for the k -fold direct sum of \mathbb{V}_{dup} we fix an arbitrary norm on \mathbb{R}^k . The limit space $\overline{V_\infty}$ can be identified with $L^p([0, 1]; \mathbb{R}^k)$, and the orbit space can be identified with $\mathcal{P}_p(\mathbb{R}^k)$ endowed with the Wasserstein p -distance with respect to $\|\cdot\|_{\mathbb{R}^k}$ in the same way. More details and proofs can be found in [7, Appx. F].

C.2 Proof of Theorem 4.2

Proof of Theorem 4.2. We show that the model $\text{DeepSets}_{\infty}^{\rho, \sigma}$ can be decomposed as $\text{DeepSets}_{\infty}^{\rho, \sigma} = \sigma \circ \tau$ with $\tau : \overline{V_\infty}/G_\infty \rightarrow \mathbb{R}^r$ continuous. Since $\overline{V_\infty}/G_\infty$ is a metric space equipped with the symmetrized distance, continuity on $\overline{V_\infty}/G_\infty$ is equivalent to continuity on each of its compact subset [55]. Finally, by definition, $\sigma : \mathbb{R}^r \rightarrow \mathbb{R}$ is a continuous function, this decomposition implies continuity of the entire model.

For any $n \in \mathbb{N}$, define the maps $\tau : V_\infty \rightarrow \mathbb{R}^r$ and $\tau_n : V_\infty \rightarrow \mathbb{R}^r$ by

$$\tau(X) = \sum_{i=1}^{\infty} \|X_i\|_{\mathbb{R}^k}^p \rho(X_i) \quad \text{and} \quad \tau_n(X) = \sum_{i=1}^{n-1} \|X_i\|_{\mathbb{R}^k}^p \rho(X_i),$$

respectively. We equip \mathbb{R}^r with a standard norm $\|\cdot\|_{\mathbb{R}^r}$. Fix an arbitrary compact subset $K \subseteq \ell_p(\mathbb{R}^k)$.

By Proposition 4.1, there exists $M > 0$ such that $\sup_{X \in K} \left(\sum_{j=1}^{\infty} \|X_j\|_{\mathbb{R}^k}^p \right)^{1/p} \leq M$. This implies $\sup_{X \in K} \sup_{i \in \mathbb{N}} \|X_i\|_{\mathbb{R}^k} \leq M$. Since $\rho : \mathbb{R}^k \rightarrow \mathbb{R}^r$ is continuous, it maps compact sets to compact sets; thus, there exists $M_\rho > 0$ such that $\sup_{X \in K} \sup_{i \in \mathbb{N}} \|\rho(X_i)\|_{\mathbb{R}^r} \leq M_\rho$. Each τ_n is continuous on K as it is a finite sum of continuous functions (coordinate projections and ρ). According to the uniform tail decay property from Proposition 4.1, we have

$$\lim_{n \rightarrow \infty} \sup_{X \in K} \|\tau_n(X) - \tau(X)\|_{\mathbb{R}^r} \leq M_\rho \lim_{n \rightarrow \infty} \sup_{X \in K} \sum_{i \geq n} \|X_i\|_{\mathbb{R}^k}^p = 0.$$

Hence, τ_n uniformly converges to τ on K . By invoking the uniform limit theorem [56], we conclude that τ is continuous on K . Since K was an arbitrary compact subset, we conclude that τ is continuous on all of $\overline{V_\infty}$. Furthermore, since τ is G_∞ -invariant, it induces a function on the orbit space $\overline{V_\infty}/G_\infty$, which we still denote by τ . By Lemma A.4, this induced map $\tau: \overline{V_\infty}/G_\infty \rightarrow \mathbb{R}^r$ is continuous, which implies that $\text{DeepSets}_\infty^{\rho, \sigma} = \sigma \circ \tau$ is continuous. \square

C.3 Proof of Theorem 4.3

To facilitate the proof, we introduce the set of invariant functions $\mathcal{F}_{\text{cont}}$ given by

$$\mathcal{F}_{\text{cont}} := \left\{ f: K/G_\infty \rightarrow \mathbb{R}, f(X) = \sigma \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho(X_{i:}) \right) \mid r \in \mathbb{N}, \rho \in C(\mathbb{R}^k, \mathbb{R}^r), \sigma \in C(\mathbb{R}^r, \mathbb{R}) \right\}.$$

This class serves as an intermediate bridge between the target space $C(K/G_\infty)$ and the neural network class \mathcal{F}_{DS} .

The proof of Theorem 4.3 is split into two main components. We first invoke the Stone-Weierstrass Theorem to prove that $\mathcal{F}_{\text{cont}}$ is dense in $C(K/G_\infty)$. Next, we show that any element in $\mathcal{F}_{\text{cont}}$ can be uniformly approximated by \mathcal{F}_{DS} through the universal approximation property of neural networks. Finally, the result follows by the transitivity of the density property. The details are organized into the three lemmas below, Lemma C.1, C.2, C.4.

Lemma C.1. *The function class $\mathcal{F}_{\text{cont}}$ is a subalgebra of $C(K/G_\infty)$.*

Proof of Lemma C.1. To establish that $\mathcal{F}_{\text{cont}}$ is a subalgebra of $C(K/G_\infty)$, we verify its closure under scalar multiplication, addition, and pointwise multiplication. Let $f_1, f_2 \in \mathcal{F}_{\text{cont}}$ with $f_j(X) = \sigma_j \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_j(X_{i:}) \right)$ for $j \in \{1, 2\}$, where $\rho_j \in C(\mathbb{R}^k, \mathbb{R}^{r_j})$ and $\sigma_j \in C(\mathbb{R}^{r_j}, \mathbb{R})$.

Scalar Multiplication. For any $\lambda \in \mathbb{R}$, $(\lambda f_1)(X) = \sigma_\lambda \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_1(X_{i:}) \right)$ where $\sigma_\lambda := \lambda \sigma_1$. Since σ_λ is also continuous, $\lambda f_1 \in \mathcal{F}_{\text{cont}}$.

Addition. Define the concatenated map $\rho_0: \mathbb{R}^k \rightarrow \mathbb{R}^{r_1+r_2}$ as $\rho_0(x) = (\rho_1(x), \rho_2(x))$. Since its components ρ_1 and ρ_2 are continuous, $\rho_0: \mathbb{R}^k \rightarrow \mathbb{R}^{r_1+r_2}$ is continuous by the property of product topologies. Define $\sigma_{\text{add}}: \mathbb{R}^{r_1+r_2} \rightarrow \mathbb{R}$ by $\sigma_{\text{add}}(u_1, u_2) = \sigma_1(u_1) + \sigma_2(u_2)$ for $u_1 \in \mathbb{R}^{r_1}, u_2 \in \mathbb{R}^{r_2}$. σ_{add} is also continuous since the addition operator on \mathbb{R} is continuous. Consequently,

$$(f_1 + f_2)(X) = \sigma_1 \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_1(X_{i:}) \right) + \sigma_2 \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_2(X_{i:}) \right) = \sigma_{\text{add}} \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_0(X_{i:}) \right).$$

Thus, $\mathcal{F}_{\text{cont}}$ is closed under addition.

Pointwise Multiplication. Similarly, define the concatenated map $\rho_0: \mathbb{R}^k \rightarrow \mathbb{R}^{r_1+r_2}$ as $\rho_0(V) = (\rho_1(V), \rho_2(V))$, which is continuous. Define $\sigma_{\text{mul}}: \mathbb{R}^{r_1+r_2} \rightarrow \mathbb{R}$ by $\sigma_{\text{mul}}(u_1, u_2) = \sigma_1(u_1) \cdot \sigma_2(u_2)$ for $u_1 \in \mathbb{R}^{r_1}, u_2 \in \mathbb{R}^{r_2}$, which is also continuous, because of the continuity of multiplication operator on \mathbb{R} . We then have

$$(f_1 \cdot f_2)(X) = \sigma_1 \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_1(X_{i:}) \right) \sigma_2 \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_2(X_{i:}) \right) = \sigma_{\text{mul}} \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho_0(X_{i:}) \right).$$

Thus, $\mathcal{F}_{\text{cont}}$ is closed under pointwise multiplication. \square

Lemma C.2. *$\mathcal{F}_{\text{cont}}$ separates points in K/G_∞ , and contains a nonzero constant function.*

To establish point separation on the quotient space K/G_∞ , we first characterize the conditions under which the metric $\bar{d}(X, Y)$ vanishes. Specifically, we identify the properties that distinguish points in distinct orbits. We claim that points in different orbits must differ in their first finite coordinate indices, see Claim C.3. Based on this, we can construct functions in $\mathcal{F}_{\text{cont}}$ to separate them.

Claim C.3. *For any compact $K \subseteq \ell_p(\mathbb{R}^k)$ and $X, Y \in K$, if the multisets $\{X_i : \|X_i\|_{\mathbb{R}^k} > \varepsilon\}$ and $\{Y_i : \|Y_i\|_{\mathbb{R}^k} > \varepsilon\}$ are equal for all $\varepsilon > 0$, then $\bar{d}(X, Y) = 0$.*

Proof of Claim C.3. Define the index sets by $I_x(\varepsilon) := \{i \in \mathbb{N} : \|X_i\|_{\mathbb{R}^k} > \varepsilon\}$, and $I_y(\varepsilon) := \{i \in \mathbb{N} : \|Y_i\|_{\mathbb{R}^k} > \varepsilon\}$, respectively. Since $X, Y \in K$, by Proposition 4.1, $\sup_{x \in K} \left(\sum_{i=1}^{\infty} \|X_i\|_{\mathbb{R}^k}^p \right) < \infty$. Therefore, $I_x(\varepsilon)$ and $I_y(\varepsilon)$ are finite and have equal cardinality. Thus, there exists $\sigma_\varepsilon \in G_\infty$, such that $X_{\sigma_\varepsilon^{-1}(i)} = Y_i$, for all $i \in I_y(\varepsilon)$. We then have

$$\bar{d}(X, Y)^p \leq \|\sigma_\varepsilon \cdot X - Y\|_p^p = \sum_{i \notin I_y(\varepsilon)} \|X_{\sigma_\varepsilon^{-1}(i)} - Y_i\|_{\mathbb{R}^k}^p.$$

For $i \notin I_y(\varepsilon)$, we have $\|Y_i\|_{\mathbb{R}^k} \leq \varepsilon$. Furthermore, since σ_ε is a bijection, the index set $\{\sigma_\varepsilon^{-1}(i) : i \notin I_y(\varepsilon)\}$ is exactly $\{j \in \mathbb{N} : j \notin I_x(\varepsilon)\}$, where $\|X_j\|_{\mathbb{R}^k} \leq \varepsilon$. Since $p \geq 1$, apply the inequality $\|a - b\|^p \leq 2^{p-1}(\|a\|^p + \|b\|^p)$ and obtain

$$\bar{d}(X, Y)^p \leq 2^{p-1} \left(\sum_{i \notin I_x(\varepsilon)} \|X_i\|_{\mathbb{R}^k}^p + \sum_{i \notin I_y(\varepsilon)} \|Y_i\|_{\mathbb{R}^k}^p \right).$$

Letting $\varepsilon \rightarrow 0$ shows that $\bar{d}(X, Y) = 0$, meaning X and Y represent the same point in K/G_∞ . \square

We are ready to prove separation of points.

Proof of Lemma C.2. First, $\mathcal{F}_{\text{cont}}$ contains a nonzero constant function by taking $\sigma \equiv 1$, which yields $f(X) = 1$ for all X .

To establish point separation, consider two distinct points in K/G_∞ , represented by $X, Y \in K$ such that $\bar{d}(X, Y) > 0$. By the contrapositive of Claim C.3, there exists $\varepsilon > 0$ such that the finite multisets $M_X := \{X_i : \|X_i\|_{\mathbb{R}^k} > \varepsilon\}$ and $M_Y := \{Y_i : \|Y_i\|_{\mathbb{R}^k} > \varepsilon\}$ are distinct. Thus, there exists some $v \in \mathbb{R}^k$ with $\|v\|_{\mathbb{R}^k} > \varepsilon$ such that their multiplicities differ: $m_X(v) \neq m_Y(v)$, where $m_X(v) := \#\{i \in \mathbb{N} : X_i = v, X_i \in M_X\}$ and $m_Y(v) := \#\{i \in \mathbb{N} : Y_i = v, Y_i \in M_Y\}$. Let $Q := (M_X \cup M_Y) \setminus \{v\}$, which is a finite set as both M_X and M_Y are finite. We define the separation radius

$$\eta := \min \left(\min_{u \in Q} \|u - v\|_{\mathbb{R}^k}, \|v\|_{\mathbb{R}^k} - \varepsilon \right).$$

Consider the ball $B(v, \eta)$ around v . For any $i \in \mathbb{N}$, we distinguish the following two cases.

Case 1. If $\|X_i\|_{\mathbb{R}^k} \leq \varepsilon$, then by the reverse triangle inequality, $\|X_i - v\|_{\mathbb{R}^k} \geq \|v\|_{\mathbb{R}^k} - \|X_i\|_{\mathbb{R}^k} \geq \|v\|_{\mathbb{R}^k} - \varepsilon \geq \eta$.

Case 2. If $X_i \in M_X$ and $X_i \neq v$, then by the definition of Q , $\|X_i - v\|_{\mathbb{R}^k} \geq \min_{u \in Q} \|u - v\|_{\mathbb{R}^k} \geq \eta$.

Combined, these cases imply that $X_i \in B(v, \eta)$ if and only if $X_i = v$. The same logic applies to the sequence Y . By Urysohn's Lemma [55], there exists a continuous function $\varphi : \mathbb{R}^k \rightarrow [0, 1]$, such that $\varphi\left(\overline{B(v, \frac{\eta}{2})}\right) = 1$ and $\varphi\left(\mathbb{R}^k \setminus B(v, \eta)\right) = 0$. Define $f \in \mathcal{F}_{\text{cont}}$ by setting $r = 1$, $\sigma(u) = u$, and

$\rho = \varphi$, then,

$$f(X) = \sigma \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho(X_{i:}) \right) = \|v\|_{\mathbb{R}^k}^p m_X(v) \neq \|v\|_{\mathbb{R}^k}^p m_Y(v) = \sigma \left(\sum_{i=1}^{\infty} \|Y_{i:}\|_{\mathbb{R}^k}^p \rho(Y_{i:}) \right) = f(Y).$$

Therefore, we conclude that $\mathcal{F}_{\text{cont}}$ separates points in K/G_{∞} . \square

Lemma C.4. \mathcal{F}_{DS} is dense in $\mathcal{F}_{\text{cont}}$ on K .

Proof of Lemma C.4. We invoke the Universal Approximation Theorem (UAT) [32]: for any continuous non-polynomial activation function ϕ , the set $\text{NN}_{r,k}^{\phi}$ is dense in $C(\mathbb{R}^k, \mathbb{R}^r)$ under the supremum norm on compact sets. More formally, fix a compact set $Q \subseteq \mathbb{R}^k$ and a norm $\|\cdot\|_{\mathbb{R}^r}$ on \mathbb{R}^r . For any continuous function $h \in C(\mathbb{R}^k, \mathbb{R}^r)$, and any approximation precision $\varepsilon > 0$, there exists $\hat{h} \in \text{NN}_{r,k}^{\phi}$ such that $\sup_{v \in Q} \|h(v) - \hat{h}(v)\|_{\mathbb{R}^r} \leq \varepsilon$.

Given any function $f \in \mathcal{F}_{\text{cont}}$, it admits the form $f(X) = \sigma \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho(X_{i:}) \right)$, for some $r \in \mathbb{N}$, $\sigma \in C(\mathbb{R}^r, \mathbb{R})$, $\rho \in C(\mathbb{R}^k, \mathbb{R}^r)$. Since $K \in \ell_p(\mathbb{R}^k)$ is compact, by Proposition 4.1, there exists $M > 0$ such that $\sup_{X \in K} \left(\sum_{j=1}^{\infty} \|X_{j:}\|_{\mathbb{R}^k}^p \right)^{1/p} \leq M$. This implies $\sup_{X \in K} \sup_{i \in \mathbb{N}} \|X_{i:}\|_{\mathbb{R}^k} \leq M$. Since $\rho \in C(\mathbb{R}^k, \mathbb{R}^r)$, it maps compact sets to compact sets, i.e., there exists $M_{\rho} > 0$ such that $\sup_{X \in K} \sup_{i \in \mathbb{N}} \|\rho(X_{i:})\|_{\mathbb{R}^r} \leq M_{\rho}$. Thus for any $X \in K$, $\left\| \sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho(X_{i:}) \right\|_{\mathbb{R}^r} \leq M^p M_{\rho}$. This means that the input domains of functions σ and ρ are both compact, which we denote by Q_{σ} and Q_{ρ} , respectively.

Approximating σ . For any $\varepsilon > 0$, UAT ensures there exists $\hat{\sigma} \in \text{NN}_{1,r}^{\phi}$ such that $\sup_{u \in Q_{\sigma}} |\sigma(u) - \hat{\sigma}(u)| \leq \frac{\varepsilon}{2}$. Moreover, since $\hat{\sigma}$ is uniformly continuous on compact sets, there exists $\delta > 0$ such that $\|u_1 - u_2\|_{\mathbb{R}^r} \leq \delta \Rightarrow |\hat{\sigma}(u_1) - \hat{\sigma}(u_2)| \leq \frac{\varepsilon}{2}$.

Approximating ρ . Applying UAT to ρ , there exists $\hat{\rho} \in \text{NN}_{r,k}^{\phi}$ such that $\sup_{v \in Q_{\rho}} \|\rho(v) - \hat{\rho}(v)\|_{\mathbb{R}^r} \leq \frac{\delta}{M^p}$. Consequently,

$$\sup_{X \in K} \left\| \sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p (\rho(X_{i:}) - \hat{\rho}(X_{i:})) \right\|_{\mathbb{R}^r} \leq \sup_{X \in K} \sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \sup_{X \in K} \|\rho(X_{i:}) - \hat{\rho}(X_{i:})\|_{\mathbb{R}^r} \leq M^p \frac{\delta}{M^p} = \delta.$$

We define the overall approximating function $\hat{f} \in \mathcal{F}_{\text{DS}}$ by $\hat{f}(X) = \hat{\sigma} \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \hat{\rho}(X_{i:}) \right)$, then

$$\begin{aligned} \sup_{X \in K} |f(X) - \hat{f}(X)| &\leq \sup_{X \in K} \left| \sigma \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho(X_{i:}) \right) - \hat{\sigma} \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho(X_{i:}) \right) \right| \\ &\quad + \left| \hat{\sigma} \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \rho(X_{i:}) \right) - \hat{\sigma} \left(\sum_{i=1}^{\infty} \|X_{i:}\|_{\mathbb{R}^k}^p \hat{\rho}(X_{i:}) \right) \right| \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon, \end{aligned}$$

which shows that \mathcal{F}_{DS} uniformly approximates every function in $\mathcal{F}_{\text{cont}}$. By the definition of both of the function classes, $\mathcal{F}_{\text{DS}} \subseteq \mathcal{F}_{\text{cont}}$, therefore, \mathcal{F}_{DS} is dense in $\mathcal{F}_{\text{cont}}$ on K with respect to the supremum norm. \square

Finally, combining the above lemmas, we prove Theorem 4.3.

Proof of Theorem 4.3. Invoking the continuity property of our model, Theorem 4.2, $\mathcal{F}_{\text{cont}} \subseteq C(K/G_{\infty})$. By Proposition 3.1, the Stone–Weierstrass Theorem, and Lemmas C.1 and C.2, we

conclude that $\mathcal{F}_{\text{cont}}$ is dense in $C(K/G_\infty)$ with respect to the supremum norm. Combining this with Lemma C.4 and the transitivity of density, it follows that \mathcal{F}_{DS} is dense in $C(K/G_\infty)$, completing the proof that the model is universal. \square

C.4 Proof of Theorem 4.5

Proof of Theorem 4.5. This theorem follows directly from the fact that W_p -convergence is equivalent to weak convergence in \mathcal{P}_p . Consider any sequence of measures $(\mu_i)_{i \in \mathbb{N}} \subseteq \mathcal{P}_p(\mathbb{R}^k)$ converging to μ in the Wasserstein p -metric, i.e., $W_p(\mu_i, \mu) \rightarrow 0$. According to the characterization of convergence in $\mathcal{P}_p(\mathbb{R}^k)$ [57], this is equivalent to $\int \varphi d\mu_i \rightarrow \int \varphi d\mu$ for all continuous functions $\varphi: \mathbb{R}^k \rightarrow \mathbb{R}$ satisfying the growth condition $|\varphi(x)| \leq C(1 + \|x\|_{\mathbb{R}^k}^p)$ for some $C > 0$.

In our framework, $\rho \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^r)$ is a continuous mapping whose components ρ_j satisfy $|\rho_j(x)| \leq M_j(1 + \|x\|_{\mathbb{R}^k}^p)$. It follows that $\int \rho d\mu_i \rightarrow \int \rho d\mu$ as $i \rightarrow \infty$, establishing the continuity of the map $\mu \mapsto \int \rho d\mu$. Since $\sigma: \mathbb{R}^r \rightarrow \mathbb{R}$ is continuous, the composition $\overline{\text{DeepSets}}_\infty^{\rho, \sigma}(\mu) = \sigma(\int \rho d\mu)$ is continuous with respect to the W_p topology. \square

C.5 Proof of Theorem 4.6

We first define an intermediate class of functions, $\mathcal{F}_{\text{cont}}^-$, to bridge $C(Q)$ and \mathcal{F}_{DS}

$$\mathcal{F}_{\text{cont}}^- := \left\{ f: Q \rightarrow \mathbb{R}, f(\mu) = \sigma \left(\int \rho d\mu \right) \mid r \in \mathbb{N}, \rho \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^r), \sigma \in C(\mathbb{R}^r, \mathbb{R}) \right\}.$$

The proof strategy of Theorem 4.6 is similar to Theorem 4.3 by showing that: (i) $\mathcal{F}_{\text{cont}}^-$ is dense in $C(Q)$ via the Stone-Weierstrass Theorem, and (ii) \mathcal{F}_{DS} is dense in $\mathcal{F}_{\text{cont}}^-$ via the universal approximation property. We decompose the proof into three auxiliary Lemmas: C.5, C.6, and C.7.

Lemma C.5. *The function class $\mathcal{F}_{\text{cont}}^-$ is a subalgebra of $C(Q)$.*

Proof of Lemma C.5. Invoking Theorem 4.5, $\mathcal{F}_{\text{cont}}^- \subseteq C(Q)$. To show $\mathcal{F}_{\text{cont}}^-$ is a subalgebra, it suffices to verify that it is closed under scalar multiplication, addition and pointwise multiplication. Let $f_1, f_2 \in \mathcal{F}_{\text{cont}}^-$ with $f_j(\mu) = \sigma_j(\int \rho_j d\mu)$ for $j \in \{1, 2\}$, where $\rho_j \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^{r_j})$ and $\sigma_j \in C(\mathbb{R}^{r_j}, \mathbb{R})$.

Scalar Multiplication. For any $\lambda \in \mathbb{R}$, $(\lambda f_1)(\mu) = \sigma_\lambda(\int \rho_1 d\mu)$ where $\sigma_\lambda := \lambda \sigma_1$. Since σ_λ is also continuous, $\lambda f_1 \in \mathcal{F}_{\text{cont}}^-$.

Addition. Define the concatenated map $\rho_0: \mathbb{R}^k \rightarrow \mathbb{R}^{r_1+r_2}$ as $\rho_0(x) = (\rho_1(x), \rho_2(x))$. Since its components ρ_1 and ρ_2 are continuous, $\rho_0: \mathbb{R}^k \rightarrow \mathbb{R}^{r_1+r_2}$ is continuous by the property of product topologies. Furthermore, since $\rho_j \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^{r_j})$ for $j \in \{1, 2\}$, each individual component function of $\rho_0 = (\rho_{1,1}, \dots, \rho_{1,r_1}, \rho_{2,1}, \dots, \rho_{2,r_2})^\top$ satisfies the p -th order growth condition. Consequently, we have $\rho_0 \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^{r_1+r_2})$. Define $\sigma_{\text{add}}: \mathbb{R}^{r_1+r_2} \rightarrow \mathbb{R}$ by $\sigma_{\text{add}}(u_1, u_2) = \sigma_1(u_1) + \sigma_2(u_2)$ for $u_1 \in \mathbb{R}^{r_1}, u_2 \in \mathbb{R}^{r_2}$. Observe that σ_{add} is also continuous since the addition operator on \mathbb{R} is continuous. Consequently,

$$(f_1 + f_2)(\mu) = \sigma_1 \left(\int \rho_1 d\mu \right) + \sigma_2 \left(\int \rho_2 d\mu \right) = \sigma_{\text{add}} \left(\left[\int \rho_1 d\mu, \int \rho_2 d\mu \right]^\top \right) = \sigma_{\text{add}} \left(\int \rho_0 d\mu \right).$$

Thus, $\mathcal{F}_{\text{cont}}^-$ is closed under addition.

Pointwise Multiplication. Similarly, define the concatenated map $\rho_0: \mathbb{R}^k \rightarrow \mathbb{R}^{r_1+r_2}$ as $\rho_0(V) = (\rho_1(V), \rho_2(V))$, which is continuous. Define $\sigma_{\text{mul}}: \mathbb{R}^{r_1+r_2} \rightarrow \mathbb{R}$ by $\sigma_{\text{mul}}(u_1, u_2) = \sigma_1(u_1) \times \sigma_2(u_2)$ for $u_1 \in \mathbb{R}^{r_1}, u_2 \in \mathbb{R}^{r_2}$, which is also continuous, because of the continuity of multiplication operator on \mathbb{R} . We then have

$$(f_1 \cdot f_2)(\mu) = \sigma_1 \left(\int \rho_1 d\mu \right) \sigma_2 \left(\int \rho_2 d\mu \right) = \sigma_{\text{mul}} \left(\left[\int \rho_1 d\mu, \int \rho_2 d\mu \right]^\top \right) = \sigma_{\text{mul}} \left(\int \rho_0 d\mu \right).$$

Thus, $\mathcal{F}_{\text{cont}}$ is closed under pointwise multiplication. \square

Lemma C.6. *The function class $\mathcal{F}_{\text{cont}}$ separates points in Q , and contains a nonzero constant function.*

Proof of Lemma C.6. The fact that $\mathcal{F}_{\text{cont}}$ contains a nonzero constant function is trivial. To show it separates points in Q , suppose $\mu, \nu \in Q \subseteq \mathcal{P}_p(\mathbb{R}^k)$, and $\mu \neq \nu$. There must exist a Borel set $B \subseteq \mathbb{R}^k$ such that $\mu(B) \neq \nu(B)$. Without loss of generality, assume $\mu(B) > \nu(B)$, and let $\delta := \mu(B) - \nu(B) > 0$. Since \mathbb{R}^k is a metric space, any measure in $\mathcal{P}_p(\mathbb{R}^k)$ is regular [58]. By the property of regularity, for any $\varepsilon > 0$, there exists a closed set $E \subseteq B$ and an open set $O \supset B$ such that

$$\mu(B) \leq \mu(E) + \varepsilon, \quad \nu(B) \geq \nu(O) - \varepsilon.$$

Since E and $\mathbb{R}^k \setminus O$ are disjoint closed sets in \mathbb{R}^k , by Urysohn's lemma [55], there exists a continuous function $\varphi: \mathbb{R}^k \rightarrow [0, 1]$ such that $\varphi(E) = 1$ and $\varphi(\mathbb{R}^k \setminus O) = 0$. It follows that

$$\int \varphi d\mu \geq \mu(E) \geq \mu(B) - \varepsilon, \quad \int \varphi d\nu \leq \nu(O) \leq \nu(B) + \varepsilon.$$

Set $\varepsilon = \frac{\delta}{3} > 0$, we obtain

$$\int \varphi d\mu - \int \varphi d\nu \geq \mu(B) - \nu(B) - 2\varepsilon = \frac{\delta}{3} > 0.$$

Recall that $\mathcal{F}_{\text{cont}} := \left\{ f: Q \rightarrow \mathbb{R}, f(\mu) = \sigma(\int \rho d\mu) \mid r \in \mathbb{N}, \rho \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^r), \sigma \in C(\mathbb{R}^r, \mathbb{R}) \right\}$. By choosing $r = 1$, $\sigma(u) = u$, and $\rho = \varphi$, we observe that φ is bounded and continuous, which implies $|\varphi(x)| \leq 1 \leq (1 + \|x\|_{\mathbb{R}^k}^p)$, and thus $\rho \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R})$. Therefore, it follows that for any $\mu, \nu \in Q \subseteq \mathcal{P}_p(\mathbb{R}^k)$ with $\mu \neq \nu$, there exists a function $f \in \mathcal{F}_{\text{cont}}$ such that $f(\mu) \neq f(\nu)$. \square

Lemma C.7. *The function class \mathcal{F}_{DS} is dense in $\mathcal{F}_{\text{cont}}$ on Q .*

Proof of Lemma C.7. The proof for this lemma proceeds by analogy with Lemma C.4 except for the following subtle issue: the compactness of $Q \subseteq \mathcal{P}_p(\mathbb{R}^k)$, as described in Proposition 4.4, does not ensure that the input domain of ρ is compact. Consequently, the standard Universal Approximation Theorem cannot be directly applied. To address this issue, we leverage the property that compactness in $\mathcal{P}_p(\mathbb{R}^k)$ guarantees the probability mass in the tails of the measures is negligible. This allows us to replace ρ with a function vanishing at infinity, which can then be uniformly approximated by neural networks on the entire space \mathbb{R}^k using recent results for non-compact domains [34].

For any function $f \in \mathcal{F}_{\text{cont}}$, we can represent it as $f(\mu) = \sigma(\int \rho d\mu)$, for some $r \in \mathbb{N}$, $\rho \in \mathcal{F}_p(\mathbb{R}^k, \mathbb{R}^r)$, and $\sigma \in C(\mathbb{R}^r, \mathbb{R})$. We fix a compact set $Q \subseteq \mathcal{P}_p(\mathbb{R}^k)$, and let \mathbb{R}^r be equipped with the Euclidean norm $\|\cdot\|_{\mathbb{R}^r}$. By the growth condition on ρ , there exist positive constants $M_j > 0$ such that for each component function ρ_j , we have $|\rho_j(x)| \leq M_j(1 + \|x\|_{\mathbb{R}^k}^p)$ for all $x \in \mathbb{R}^k$. Defining the vector $M := (M_1, \dots, M_r)^\top \in \mathbb{R}^r$, it follows that $\|\rho(x)\|_{\mathbb{R}^r} \leq \|M\|_{\mathbb{R}^r}(1 + \|x\|_{\mathbb{R}^k}^p)$, $\forall x \in \mathbb{R}^k$. Moreover, Theorem 4.5 ensures the continuity of the map $\mu \mapsto \int \rho d\mu$. Thus, the image of the compact set Q under this map, denoted by $Q_\sigma \subseteq \mathbb{R}^r$, is also compact, forming the input domain for σ .

Approximating σ . For any $\varepsilon > 0$, UAT ensures there exists $\hat{\sigma} \in \text{NN}_{1,r}^\phi$ such that $\sup_{u \in Q_\sigma} |\sigma(u) - \hat{\sigma}(u)| \leq \frac{\varepsilon}{2}$. Moreover, since $\hat{\sigma}$ is uniformly continuous on compact sets, there exists $\delta > 0$ such that $\|u_1 - u_2\|_{\mathbb{R}^r} \leq \delta \Rightarrow |\hat{\sigma}(u_1) - \hat{\sigma}(u_2)| \leq \frac{\varepsilon}{2}$.

Tail Control and Truncation of ρ . Since Q is compact, it is tight and p -uniformly integrable, by Proposition 4.4. Then for the constant $\frac{\delta}{4\|M\|_{\mathbb{R}^r}}$, there exists a positive constant $R > 0$ such that

$$\begin{aligned} \sup_{\mu \in Q} \mu \left(\mathbb{R}^k \setminus \overline{B(0, R)} \right) &\leq \frac{\delta}{4\|M\|_{\mathbb{R}^r}} && \text{(tightness),} \\ \sup_{\mu \in Q} \int_{\mathbb{R}^k \setminus \overline{B(0, R)}} \|x\|_{\mathbb{R}^k}^p d\mu &\leq \frac{\delta}{4\|M\|_{\mathbb{R}^k}} && (p\text{-uniformly integrability).} \end{aligned}$$

Using Urysohn's Lemma [55], there exists a continuous function $\varphi: \mathbb{R}^k \rightarrow [0, 1]$ such that $\varphi(\overline{B(0, R)}) = 1$ and $\varphi(\overline{B(0, 2R)}) = 0$. We truncate function ρ with φ by

$$\rho_\varphi(x) := \varphi(x)\rho(x), \quad \forall x \in \mathbb{R}^k.$$

Evidently, $\rho_\varphi \in C_0(\mathbb{R}^k, \mathbb{R}^r)$, the space of continuous functions vanishing at infinity. Furthermore, by the tail control of ρ , the discrepancy between ρ and its truncated counterpart ρ_φ can also be controlled on Q

$$\sup_{\mu \in Q} \int \|\rho_\varphi - \rho\|_{\mathbb{R}^r} d\mu \leq \sup_{\mu \in Q} \int_{\mathbb{R}^k \setminus \overline{B(0, R)}} \|\rho(x)\|_{\mathbb{R}^r} d\mu(x) \leq \|M\|_{\mathbb{R}^r} \sup_{\mu \in Q} \int_{\mathbb{R}^k \setminus \overline{B(0, R)}} (1 + \|x\|_{\mathbb{R}^k}^p) d\mu(x) \leq \frac{\delta}{2}.$$

Approximating ρ_φ . For the approximation of C_0 functions, we invoke a specialized UAT [34]: if the activation function ϕ is continuous, nonpolynomial, and asymptotically polynomial at $\pm\infty$, then any function in $C_0(\mathbb{R}^k, \mathbb{R})$ can be uniformly approximated by $\text{NN}_{1,k}^\phi$ on \mathbb{R}^k . This result extends naturally to the vector-valued space $C_0(\mathbb{R}^k, \mathbb{R}^r)$ and $\text{NN}_{r,k}^\phi$ by approximating each component function independently. Since $\rho_\varphi \in C_0(\mathbb{R}^k, \mathbb{R}^r)$, there exists $\hat{\rho} \in \text{NN}_{r,k}^\phi$ such that $\sup_{x \in \mathbb{R}^k} \|\hat{\rho}(x) - \rho_\varphi(x)\|_{\mathbb{R}^r} \leq \frac{\delta}{2}$. Thus,

$$\sup_{\mu \in Q} \left\| \int (\rho - \hat{\rho}) d\mu \right\|_{\mathbb{R}^r} \leq \sup_{\mu \in Q} \int \|\rho_\varphi - \rho\|_{\mathbb{R}^r} d\mu + \sup_{\mu \in Q} \int \|\rho_\varphi - \hat{\rho}\|_{\mathbb{R}^r} d\mu \leq \delta.$$

We define the overall architecture as $\hat{f} = \hat{\sigma}(\int \hat{\rho} d\mu) \in \mathcal{F}_{\overline{\text{DS}}}$, and estimate the approximation error

$$\begin{aligned} \sup_{\mu \in Q} |f(\mu) - \hat{f}(\mu)| &= \sup_{\mu \in Q} \left| \sigma \left(\int \rho d\mu \right) - \hat{\sigma} \left(\int \hat{\rho} d\mu \right) \right| \\ &\leq \sup_{\mu \in Q} \left| \sigma \left(\int \rho d\mu \right) - \hat{\sigma} \left(\int \rho d\mu \right) \right| + \sup_{\mu \in Q} \left| \hat{\sigma} \left(\int \rho d\mu \right) - \hat{\sigma} \left(\int \hat{\rho} d\mu \right) \right| \\ &\leq \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon, \end{aligned}$$

which concludes that $\mathcal{F}_{\overline{\text{DS}}}$ is dense in $\mathcal{F}_{\overline{\text{cont}}}$. \square

Proof of Theorem 4.6. By Proposition 3.1, the Stone–Weierstrass Theorem, and Lemmas C.5 and C.6, we conclude that $\mathcal{F}_{\overline{\text{cont}}}$ is dense in $C(Q)$. Combining this with Lemma C.7 and the transitivity of density, it follows that $\mathcal{F}_{\overline{\text{DS}}}$ is dense in $C(Q)$. \square

D Details and missing proofs from Section 4.2

In this section, we present additional details and proofs for our results on graphs.

D.1 Duplication consistent sequence for graphs

We start by elaborating on the consistent sequence for graphs used in Section 4.2. More details and proofs for the following assertions can be found in [7, Appdx. G].

The duplication embedding consistent sequence for graphs $\mathbb{V}_{\text{dup}}^G = \{(V_n), (\varphi_{N,n}), (G_n)\}$ is defined as follows. The index set $(\mathbb{N}, \cdot | \cdot)$ is the set of natural numbers with divisibility partial order, where $n \preceq N$ if and only if $n | N$. For each $n \in \mathbb{N}$, $V_n = \mathbb{R}_{\text{sym}}^{n \times n}$. For $n \preceq N$, the duplication embedding $\varphi_{N,n}: \mathbb{R}_{\text{sym}}^{n \times n} \hookrightarrow \mathbb{R}_{\text{sym}}^{N \times N}$ is given by $\varphi_{N,n}(A) = A \otimes \left(\mathbb{1}_{N/n} \mathbb{1}_{N/n}^\top \right)$, for $A \in \mathbb{R}_{\text{sym}}^{n \times n}$, where \otimes denotes the Kronecker product. The group is the symmetric group S_n and the group embedding is the same as the case in Appendix C.1.2. S_n acts on V_n via $g \cdot A = gAg^\top$.

The space V_∞ can be identified with the space of step graphons $W_A: [0, 1]^2 \rightarrow \mathbb{R}$ given by

$$W_A(x, y) = A_{ij} \quad \text{for } (x, y) \in \left(\frac{i-1}{n}, \frac{i}{n} \right] \times \left(\frac{j-1}{n}, \frac{j}{n} \right], i, j \in [n].$$

The symmetric group acts on the induced step graphon by

$$\sigma \cdot W_A = W_A^{\sigma^{-1}} := W_A(\sigma^{-1}(x), \sigma^{-1}(y)).$$

We endow each V_n with the cut norm, which is defined as

$$\|A\|_\square := \frac{1}{n^2} \max_{S \subseteq [n], T \subseteq [n]} \left| \sum_{i \in S, j \in T} A_{ij} \right| \quad \text{for } A \in \mathbb{R}_{\text{sym}}^{n \times n}.$$

This cut norm extends to a norm in V_∞ , which coincides with the cut norm on graphons, defined as:

$$\|W\|_\square := \sup_{S, T \subseteq [0, 1]} \left| \int_{S \times T} W(x, y) dx dy \right|.$$

The symmetrized distance coincides with the δ_\square distance, which is defined by

$$\bar{d}(W, U) = \delta_\square(W, U) := \inf_{\varphi \in S_{[0, 1]}} \| (U^\varphi - W) \|_\square,$$

where $S_{[0, 1]}$ denotes the set of all measure-preserving bijections $[0, 1] \rightarrow [0, 1]$; and $U^\varphi(x, y) := U(\varphi(x), \varphi(y))$. The orbit space can be identified with the space of symmetric measurable functions $[0, 1]^2 \rightarrow \mathbb{R}$, modulo the equivalence relation $W_1 \sim W_2$ whenever $\delta_\square(W_1, W_2) = 0$.

D.2 Proof of Theorem 4.9

Proof of Theorem 4.9. We expand the formulation (9) and get

$$h_m(W) = \int_{[0, 1]^m} \prod_{1 \leq i < j \leq m} (a_{ij} W(x_i, x_j) + b_{ij}) \prod_{\ell=1}^m dx_\ell = \sum_{S \subseteq E_m} \left(\prod_{e \in S} a_e \prod_{e \notin S} b_e \right) t(S; W),$$

where E_m denotes the set of edges of the m -node complete graph. Thus, each $h_m(\cdot)$ is a linear combination of homomorphism densities of simple graphs, which are continuous with respect to the cut distance δ_\square [59, Thm. 2.7], and hence $\mathcal{F}_W \subseteq C(\mathcal{W}_0, \delta_\square)$ as claimed. \square

D.3 Proof of Theorem 4.10

Proof of Theorem 4.10. To establish the density of \mathcal{F}_W , it suffices to show that $\mathcal{HD} \subseteq \mathcal{F}_W$, where $\mathcal{HD} := \text{span}\{t(F, \cdot) \mid F \text{ is a simple graph}\}$ is known to be dense in $C(\mathcal{W}_0, \delta_\square)$ by [60, Thm. 2.2].

For any simple graph $F = (V, E)$ with $|V| = k \leq m$, consider the functional $h_m(W)$ defined in

(9) with parameters

$$a_{ij} = \begin{cases} 1 & \text{if } \{i, j\} \in E \\ 0 & \text{if } \{i, j\} \notin E \end{cases}, \quad \text{and} \quad b_{ij} = \begin{cases} 0 & \text{if } \{i, j\} \in E \\ 1 & \text{if } \{i, j\} \notin E \end{cases},$$

and note that $h_m(W) = t(F, W)$, proving that $\mathcal{HD} \subseteq \mathcal{F}_W$ as desired. \square

D.4 A universal deep model for graphs

First, we introduce a tensor contraction operator that serves as the nonlinear activation function in the network. For $k \in \mathbb{N}_0$, define a class of functions $\mathcal{G}_k := \{G : [0, 1]^k \rightarrow \mathbb{R} \text{ measurable, bounded}\} / \sim_{\text{a.e.}}$ where $G_1 \sim_{\text{a.e.}} G_2$ if $G_1 = G_2$ almost everywhere, and set $\mathcal{G}_0 = \mathbb{R}$.

Definition D.1. We define the operator $T : \mathcal{G}_{k+1} \times \prod_{i=1}^k \mathcal{G}_2 \rightarrow \mathcal{G}_k$ as

$$(T(P, G_1, \dots, G_k))(y_1, \dots, y_k) = \int_{[0,1]} P(x, y_1, \dots, y_k) \prod_{i=1}^k G_i(x, y_i) dx.$$

For $k = 0$, the operator $T : \mathcal{G}_1 \rightarrow \mathbb{R}$ is given by $\int_{[0,1]} P(x) dx$.

We define the j -th linear layer as

$$\mathcal{L}_j(P_j, W) := \begin{pmatrix} P_j \\ a_j^1 W + b_j^1 \\ \vdots \\ a_j^{k_j} W + b_j^{k_j} \end{pmatrix},$$

where $P_j \in \mathcal{G}_{k_j+1}$, for some $k_j \in \mathbb{N}_0$; $a_j^i, b_j^i \in \mathbb{R}$, for $i \in [k_j]$; $W \in \mathcal{W}_0$. We denote the j -th linear layer map as $\tilde{\mathcal{L}}_j : \mathcal{G}_{k_j+1} \times \mathcal{W}_0 \rightarrow (\mathcal{G}_{k_j+1} \times \prod_{i=1}^{k_j} \mathcal{G}_2) \times \mathcal{W}_0$ given by $\tilde{\mathcal{L}}_j(P_j, W) = (\mathcal{L}_j(P_j, W), W)$. Furthermore, we denote the nonlinearity in the network as $\tilde{T} : (\mathcal{G}_{k+1} \times \prod_{i=1}^k \mathcal{G}_2) \times \mathcal{W}_0 \rightarrow \mathcal{G}_k \times \mathcal{W}_0$ given by $\tilde{T}((P, G_1, \dots, G_k), W) = (T(P, G_1, \dots, G_k), W)$, where T is the operator in Definition D.1.

We then contract the output of \mathcal{L}_j using T defined in Definition D.1 to get the input of $(j+1)$ -th linear layer

$$T \circ \mathcal{L}_j(P_j, W) := P_{j+1} \in \mathcal{G}_{k_j}.$$

Suppose we have $m \geq 1$ layers, set $P_1 \in \mathcal{G}_m, P_1(x_1, \dots, x_m) \equiv 1$, and let $k_j = m - j$. For the j -th tensor output, $1 \leq j \leq m + 1$, we have $P_j \in \mathcal{G}_{m+1-j}$. The overall architecture $h_D^m : \mathcal{W}_0 \rightarrow \mathbb{R}$ can be written as

$$(h_D^m(W), W) = \tilde{T} \circ \tilde{\mathcal{L}}_m \circ \dots \circ \tilde{T} \circ \tilde{\mathcal{L}}_1(P_1, W). \quad (15)$$

Theorem D.2. Let $\mathcal{F}_{\text{deep}}$ be the class of functions defined as $\mathcal{F}_{\text{deep}} := \text{span}\{h_D^m \mid m \in \mathbb{N}\}$, where each h_D^m has the form of (15). Then $\mathcal{F}_{\text{deep}}$ is dense in $C(\mathcal{W}_0, \delta_\square)$.

Proof of Theorem D.2. We establish the result by induction, demonstrating that the deep model form (15) recovers the structure of (9), and complete the proof by Theorem 4.10. We prove inductively that for any $r \in \mathbb{N}$ such that $2 \leq r \leq m + 1$, the output tensor $P_r \in \mathcal{G}_{m+1-r}$ of the r th layer can be represented as

$$P_r(\mathbf{x}_{\geq r}) = \int_{[0,1]^{r-1}} \prod_{1 \leq i < j \leq r-1} L_i^{j-i}(W(x_i, x_j)) \prod_{i=1}^{r-1} \prod_{j=r}^m L_i^{j-i}(W(x_i, x_j)) d\mathbf{x}_{< r}, \quad (16)$$

where $L_i^v(W) := a_i^v W + b_i^v$ with $a_i^v, b_i^v \in \mathbb{R}$, $\mathbf{x}_{\geq r} := (x_r, \dots, x_m)$, and $d\mathbf{x}_{< r} := dx_1 \dots dx_{r-1}$. Having done so, we set $r = m + 1$ and conclude that

$$P_{m+1} = \int_{[0,1]^m} \prod_{1 \leq i < j \leq m} L_i^{j-i}(W(x_i, x_j)) d\mathbf{x}_{< m+1} = \int_{[0,1]^m} \prod_{1 \leq i < j \leq m} (a_i^{j-i} W(x_i, x_j) + b_i^{j-i}) d\mathbf{x}_{< m+1}.$$

By setting $a_{ij} := a_i^{j-i}$ and $b_{ij} := b_i^{j-i}$, this expression exactly matches the form (9), completing the proof. Thus, we proceed to establish (16).

Base Case ($r = 2$). Directly applying Definition D.1, we have

$$P_2(x_2, \dots, x_m) = \int_{[0,1]} \prod_{u=1}^{m-1} (a_1^u W(x_1, x_{u+1}) + b_1^u) dx_1 = \int_{[0,1]} \prod_{j=2}^m L_1^{j-1}(W(x_1, x_j)) dx_1$$

which is consistent with the inductive hypothesis (16).

Inductive Step. Suppose the hypothesis holds for some $r \in \mathbb{N}$ ($2 \leq r \leq m$), such that $P_r \in \mathcal{G}_{m+1-r}$ takes the form of (16). Then, for $r + 1$, we obtain

$$\begin{aligned} & P_{r+1}(x_{r+1}, \dots, x_m) \\ &= \int_{[0,1]} P_r(x_r, x_{r+1}, \dots, x_m) \prod_{u=1}^{m-r} (a_r^u W(x_r, x_{r+u}) + b_r^u) dx_r \\ &= \int_{[0,1]} \left(\int_{[0,1]^{r-1}} \prod_{1 \leq i < j \leq r-1} L_i^{j-i}(W(x_i, x_j)) \prod_{i=1}^{r-1} \prod_{j=r}^m L_i^{j-i}(W(x_i, x_j)) d\mathbf{x}_{< r} \right) \prod_{j=r+1}^m L_r^{j-r}(W(x_r, x_j)) dx_r \\ &= \int_{[0,1]^r} \left(\prod_{1 \leq i < j \leq r-1} L_i^{j-i}(W(x_i, x_j)) \right) \left(\prod_{i=1}^{r-1} L_i^{r-i}(W(x_i, x_r)) \right) \left(\prod_{i=1}^r \prod_{j=r+1}^m L_i^{j-i}(W(x_i, x_j)) \right) d\mathbf{x}_{< r+1} \\ &= \int_{[0,1]^r} \prod_{1 \leq i < j \leq r} L_i^{j-i}(W(x_i, x_j)) \prod_{i=1}^r \prod_{j=r+1}^m L_i^{j-i}(W(x_i, x_j)) d\mathbf{x}_{< r+1}. \end{aligned}$$

This confirms that the hypothesis holds for $r + 1$. \square

We remark that the above architecture can be made more expressive by allowing general widths for the linear layers, and by using general linear equivariant maps [46]. The latter would require adapting the nonlinear operators in Definition D.1 appropriately, and we do not further explore these extensions.

E Details and missing proofs from Section 4.3

E.1 Duplication consistent sequence for point clouds

We elaborate on the consistent sequence used in Section 4.3. See [7, Appx. H] for more details and proofs of the following assertions.

Similar to the case of graphs, the duplication embedding consistent sequence for point clouds $\mathbb{V}_{\text{dup}}^P = \{(V_n), (\varphi_{N,n}), (G_n)\}$ is defined as follows. The index set $(\mathbb{N}, \cdot | \cdot)$ is the set of natural numbers with divisibility partial order, where $n \preceq N$ if and only if $n | N$. For each $n \in \mathbb{N}$, $V_n = \mathbb{R}^{n \times k}$, which represents sets of n points in \mathbb{R}^k , and k is fixed. The group is $G_n = S_n \times O(k)$, where S_n is the permutation group, and $O(k)$ is the orthogonal group. The group action on V_n is

defined as

$$(g, h) \cdot X = gXh^\top.$$

For $n \preceq N$, the duplication embedding $\varphi_{N,n}: \mathbb{R}^{n \times k} \hookrightarrow \mathbb{R}^{N \times k}$ is given by $\varphi_{N,n}(X) = X \otimes \mathbb{1}_{N/n}$, and the group embedding $\theta_{N,n}: S_n \times O(k) \hookrightarrow S_N \times O(k)$ is given by $\theta_{N,n}(g, h) = (g \otimes I_{N/n}, h)$. We consider the Euclidean norm on \mathbb{R}^k denoted by $\|\cdot\|_{\mathbb{R}^k}$, which corresponds to the inner product preserved by elements of $O(k)$. We equip each V_n with the normalized ℓ_2 norm:

$$\|X\|_{\bar{2}} = \left(\frac{1}{n} \sum_{i=1}^n \|X_{i\cdot}\|_{\mathbb{R}^k}^2 \right)^{1/2}.$$

Similarly to the case of sets, we can identify each matrix $X \in \mathbb{R}^{n \times k}$ with a step function $[0, 1] \rightarrow \mathbb{R}^k$. Then the limit space can be identified with $\overline{V_\infty} = L^2([0, 1]; \mathbb{R}^k)$, and the symmetrized metric can be written as:

$$\bar{d}(X, Y) = \inf_{h \in O(k)} \inf_{\varphi \in S_{[0,1]}} \left(\int_0^1 \|h(X(t)) - Y(\varphi(t))\|_{\mathbb{R}^k}^2 dt \right)^{1/2},$$

where $S_{[0,1]}$ is the set of measure preserving bijections.

We can further identify each matrix $X \in \mathbb{R}^{n \times k}$ with an empirical probability measure in \mathbb{R}^k . The space of orbit closures $\overline{V_\infty}/G_\infty$ can then be identified with the space of orbits of probability measures on \mathbb{R}^k under the $O(k)$ action by pushforwards $h \cdot \mu = h_\# \mu$ for $h \in O(k)$. From this perspective, we can further rewrite the symmetrized distance as

$$\bar{d}(X, Y) = \inf_{h \in O(k)} W_2(h \cdot \mu_X, \mu_Y) \quad \text{for } X, Y \in \overline{V_\infty},$$

where $\mu_X = X_\# \lambda$, λ is the Lebesgue measure on $[0, 1]$.

E.2 Proof of Proposition 4.11

Proof of Proposition 4.11. Consider the set of probability measures corresponding to the set K_R , denoted by $Q_R := \{X_\# \lambda \mid X \in K_R\}$. We show that Q_R is equal to the set of probability measures with R -bounded support $\mathcal{P}_2(\overline{B(0, R)}) := \{\mu \in \mathcal{P}_2(\mathbb{R}^k) \mid \text{supp } \mu \subseteq \overline{B(0, R)}\}$.

First, we show that $Q_R \subseteq \mathcal{P}_2(\overline{B(0, R)})$. For any $\mu \in Q_R$, there exists a function $X \in K_R$ such that $\mu = X_\# \lambda$. By the definition of the push-forward measure, for any Borel set $B \subseteq \mathbb{R}^k$, we have $\mu(B) = \lambda(X^{-1}(B))$. Since $X \in K_R$, for almost every $t \in [0, 1]$ we have $\|X(t)\|_{\mathbb{R}^k} \leq R$. Consequently, $\text{supp } \mu \subseteq \overline{B(0, R)}$, which implies $Q_R \subseteq \mathcal{P}_2(\overline{B(0, R)})$.

Conversely, we show that $\mathcal{P}_2(\overline{B(0, R)}) \subseteq Q_R$. For any $\mu \in \mathcal{P}_2(\overline{B(0, R)})$, there exists a Borel mapping $X: [0, 1] \rightarrow \overline{B(0, R)}$ such that $X_\# \lambda = \mu$ because $\overline{B(0, R)}$ is a standard Borel space [61, Thm. 3.3.13]. Since the image of X is contained within the ball, we have $\|X(t)\|_{\mathbb{R}^k} \leq R$ for all t . Moreover, since $\int_{[0,1]} \|X(t)\|_{\mathbb{R}^k}^2 d\lambda(t) \leq R^2 < \infty$, it follows that $X \in L^2([0, 1]; \mathbb{R}^k)$, and thus $X \in K_R$. This concludes that $\mathcal{P}_2(\overline{B(0, R)}) \subseteq Q_R$.

Consequently, the set $Q_R = \mathcal{P}_2(\overline{B(0, R)})$ is a compact subset of the Wasserstein space $\mathcal{P}_2(\mathbb{R}^k)$ with respect to the W_2 metric by Proposition 4.4. Since the canonical projection map $\pi: \mathcal{P}_2(\mathbb{R}^k) \rightarrow \mathcal{P}_2(\mathbb{R}^k)/O(k)$ is continuous (see the proof of Lemma A.4), we conclude that K_R/G_∞ , which can be identified with $Q_R/O(k)$, is compact in $\overline{V_\infty}/G_\infty$. \square

E.3 Proof of Theorem 4.12

Proof of Theorem 4.12. By [46], the IGN architecture in (13) is continuous with respect to δ_2 distance. Therefore, to establish the overall model is continuous, it suffices to show that $X \mapsto W_X$ is continuous with symmetrized distance \bar{d} in its input and δ_2 distance in its image. For $X, Y \in K_R = \left\{ X \in L^2([0, 1]; \mathbb{R}^k) : \|X\|_\infty \leq R \right\}$, the distance of their corresponding covariance functions can be bounded by

$$\begin{aligned} \delta_2(W_X, W_Y)^2 &= \frac{1}{4R^4} \inf_{\varphi \in \mathcal{S}_{[0,1]}} \left(\int_{[0,1]^2} |\langle X(x), X(y) \rangle - \langle Y(\varphi(x)), Y(\varphi(y)) \rangle|^2 dx dy \right) \\ &= \frac{1}{4R^4} \inf_{\varphi \in \mathcal{S}_{[0,1]}} \left(\int_{[0,1]^2} |\langle X(x) - Y(\varphi(x)), X(y) \rangle + \langle Y(\varphi(x)), X(y) - Y(\varphi(y)) \rangle|^2 dx dy \right) \\ &\leq \frac{1}{4R^2} \inf_{\varphi \in \mathcal{S}_{[0,1]}} \left(\int_{[0,1]^2} (\|X(x) - Y(\varphi(x))\|_{\mathbb{R}^k} + \|X(y) - Y(\varphi(y))\|_{\mathbb{R}^k})^2 dx dy \right). \end{aligned}$$

Applying Young's inequality, namely, $(a + b)^2 \leq 2(a^2 + b^2)$, we obtain

$$\begin{aligned} \delta_2(W_X, W_Y)^2 &\leq \frac{1}{2R^2} \inf_{\varphi \in \mathcal{S}_{[0,1]}} \left(\int_{[0,1]^2} (\|X(x) - Y(\varphi(x))\|_{\mathbb{R}^k}^2 + \|X(y) - Y(\varphi(y))\|_{\mathbb{R}^k}^2) dx dy \right) \\ &= \frac{1}{R^2} \inf_{\varphi \in \mathcal{S}_{[0,1]}} \left(\int_0^1 \|X(x) - Y(\varphi(x))\|_{\mathbb{R}^k}^2 dx \right). \end{aligned}$$

For any $h \in O(k)$, we have $W_{h \cdot X} = \langle h \cdot X(x), h \cdot X(y) \rangle = \langle X(x), X(y) \rangle = W_X$, so

$$\delta_2(W_X, W_Y)^2 = \delta_2(W_{h \cdot X}, W_Y)^2 \leq \frac{1}{R^2} \inf_{\varphi \in \mathcal{S}_{[0,1]}} \left(\int_0^1 \|h \cdot X(x) - Y(\varphi(x))\|_{\mathbb{R}^k}^2 dx \right).$$

Taking the infimum over the orthogonal group $O(k)$, the inequality still holds:

$$\delta_2(W_X, W_Y) \leq \frac{1}{R} \inf_{h \in O(k)} \inf_{\varphi \in \mathcal{S}_{[0,1]}} \left(\int_0^1 \|h \cdot X(x) - Y(\varphi(x))\|_{\mathbb{R}^k}^2 dx \right)^{1/2} = \frac{1}{R} \bar{d}(X, Y),$$

which implies that the map $X \mapsto W_X$ is Lipschitz continuous. \square

E.4 Proof of Theorem 4.13

The proof relies on the correspondence between K_R/G_∞ and the orbit space of $(\mathcal{W}_0, \delta_2)$ under the action of the orthogonal group $O(k)$, together with the universality of the model $I_{\varrho, M, L, b}^\varrho$ on compact subsets of $(\mathcal{W}_0, \delta_2)$ [46]. We first state and prove several auxiliary lemmas.

Lemma E.1. *For $X, Y \in K_R$, and W_X, W_Y are the corresponding covariance functions defined in (12), then $\bar{d}(X, Y) = 0$ if and only if $\delta_2(W_X, W_Y) = 0$.*

This lemma not only ensures that the architecture is well-defined, but also implies that point separation on K_R/G_∞ is equivalent to separate points on $(\mathcal{W}_0, \delta_2)/O(k)$. The proof is built on the following claim that if the covariance functions are equal almost everywhere, then there must exist an orthogonal transformation between the original functions. More formally,

Claim E.2. *Let $X, Y \in L^2([0, 1], \mathbb{R}^k)$. Suppose that*

$$\langle X(x), X(y) \rangle = \langle Y(x), Y(y) \rangle \quad \text{for almost every } x, y \in [0, 1].$$

Then there exists an orthogonal transformation $h \in O(k)$ such that $Y = hX$ almost everywhere.

Proof of Claim E.2. First, we can identify the space $L^2([0, 1], \mathbb{R}^k)$ with $\mathcal{B}(L^2([0, 1]), \mathbb{R}^k)$, which is the set of bounded linear operators between two Hilbert spaces $L^2([0, 1]) \rightarrow \mathbb{R}^k$, equipped with the Hilbert–Schmidt norm $\|\cdot\|_{HS}$. In more detail, each $X \in L^2([0, 1], \mathbb{R}^k)$ can be written as

$$X = (X_1, \dots, X_k)^\top \quad \text{with} \quad X_i \in L^2([0, 1]),$$

which defines a bounded linear map $\Phi: L^2([0, 1]) \rightarrow \mathbb{R}^k$ by

$$\Phi_X \varphi = (\langle X_1, \varphi \rangle, \dots, \langle X_k, \varphi \rangle)^\top.$$

Conversely, by Riesz Representation Theorem any bounded linear operator $\Phi \in \mathcal{B}(L^2([0, 1]), \mathbb{R}^k)$ can be written as this form for some $X_1, \dots, X_k \in L^2([0, 1])$.

Define $\mathcal{V}_m := \text{span}\{X_1, \dots, X_k\}$. Then Φ_X vanishes on the orthogonal complement \mathcal{V}_k^\perp . Since $\Phi_X: \mathcal{V}_k \rightarrow \mathbb{R}^k$ is a linear map between finite-dimensional vector spaces, it admits a singular value decomposition [62]. Thus, there exist non-negative singular values $\sigma_1, \dots, \sigma_k \in \mathbb{R}_{\geq 0}$, orthonormal bases $\{v_1, \dots, v_k\} \subseteq \mathbb{R}^k$, and $\{u_1, \dots, u_k\} \subseteq L^2([0, 1])$ such that

$$\Phi_X = \sum_{i=1}^k \sigma_i \langle u_i, \cdot \rangle v_i.$$

Moreover, for each $\sigma_i > 0$, the u_i is an eigenvector of the self-adjoint operator $\Phi_X^* \Phi_X$ corresponding to the eigenvalue σ_i^2 . For indices i with $\sigma_i = 0$, the vectors u_i can be chosen to complete the set $\{u_1, \dots, u_k\}$ into an orthonormal basis of \mathcal{V}_k .

The self-adjoint operator $\Phi_X^* \Phi_X$ evaluated on $\varphi \in L^2([0, 1])$ yields

$$(\Phi_X^* \Phi_X \varphi)(x) = \sum_{i=1}^k \langle X_i, \varphi \rangle X_i(x) = \sum_{i=1}^k X_i(x) \int_0^1 X_i(y) \varphi(y) dy = \int_0^1 \langle X(x), X(y) \rangle \varphi(y) dy.$$

Since $\langle X(x), X(y) \rangle = \langle Y(x), Y(y) \rangle$ almost everywhere, then $\Phi_X^* \Phi_X = \Phi_Y^* \Phi_Y$. Therefore, we can choose the same orthonormal bases $\{u_1, \dots, u_k\} \subseteq L^2([0, 1])$ for Φ_X and Φ_Y , such that

$$\Phi_X = \sum_{i=1}^k \sigma_i \langle u_i, \cdot \rangle v_i, \quad \Phi_Y = \sum_{i=1}^k \sigma_i \langle u_i, \cdot \rangle w_i,$$

where $\{v_1, \dots, v_k\}$ and $\{w_1, \dots, w_k\}$ are both the orthonormal bases in \mathbb{R}^k . Then there exists $h \in O(k)$ such that $h(v_i) = w_i$, which implies

$$h \Phi_X = \sum_{i=1}^k \sigma_i \langle u_i, \cdot \rangle (h v_i) = \sum_{i=1}^k \sigma_i \langle u_i, \cdot \rangle w_i = \Phi_Y.$$

For any $\omega \in \mathbb{R}^k$, $\langle Y(\cdot), \omega \rangle = \Phi_Y^* \omega = (h \Phi_X)^* \omega = \langle h X(\cdot), \omega \rangle$, which yields $Y = hX$ in $L^2([0, 1], \mathbb{R}^k)$. \square

Proof of Lemma E.1. First, by Theorem 4.12, if $\bar{d}(X, Y) = 0$, then $\delta_2(W_X, W_Y) = 0$. We only need to show that $\delta_2(W_X, W_Y) = 0$ implies $\bar{d}(X, Y) = 0$. Since $\delta_2(W_X, W_Y) = 0$, there exists measure-preserving maps $\varphi, \psi \in \bar{S}_{[0,1]}$, such that $W_X^\varphi = W_Y^\psi$ almost everywhere [50, Cor. 10.35], hence

$$\langle X(\varphi(x)), X(\varphi(y)) \rangle = \langle Y(\psi(x)), Y(\psi(y)) \rangle \quad \text{for almost every } x, y \in [0, 1].$$

By Claim E.2, there exists $h \in O(k)$ such that $hX^\varphi = Y^\psi$, where $X^\varphi(x) = X(\varphi(x))$. Therefore,

$$\begin{aligned} \bar{d}(X, Y) &= \inf_{h \in O(k)} \inf_{\varphi \in \bar{S}_{[0,1]}} \left(\int_0^1 \|h(X(t)) - Y(\varphi(t))\|_{\mathbb{R}^k}^2 dt \right)^{1/2} \\ &= \inf_{h \in O(k)} \inf_{\varphi, \psi \in \bar{S}_{[0,1]}} \left(\int_0^1 \|h(X(\varphi(t))) - Y(\psi(t))\|_{\mathbb{R}^k}^2 dt \right)^{1/2} \\ &\leq 0, \end{aligned}$$

where $\bar{S}_{[0,1]}$ denotes the set of measure-preserving maps $[0, 1] \rightarrow [0, 1]$. This proves that $\bar{d}(X, Y) = 0$, as desired. \square

Proof of Theorem 4.13. We define the class of functions \mathcal{F}_h as

$$\mathcal{F}_h := \text{span}\{X \mapsto t(F, W_X) \mid F \text{ is simple}\},$$

where t denotes the homomorphism density defined in (8). It is straightforward to see that \mathcal{F}_h is a subalgebra, since $t(F_1, \cdot) t(F_2, \cdot) = t(F_1 \sqcup F_2, \cdot)$. By Lemmas E.1 and [50, Cor. 10.34], \mathcal{F}_h separates points. Therefore, \mathcal{F}_h is dense in $C(K_R/G_\infty)$.

By Theorem 4.12, $X \mapsto W_X$ is a continuous mapping. Since the input K_R/G_∞ is compact by Proposition 4.11, the image of K_R/G_∞ is a compact subset of $(\mathcal{W}_0, \delta_2)$. The IGN model is universal on compact subsets of $(\mathcal{W}_0, \delta_2)$, and thus can approximate arbitrarily well any homomorphism density of simple graphs [46]. Consequently, \mathcal{F}_{PC}^g is dense in $C(K_R/G_\infty)$. \square