Polynomial Width is Sufficient for Set Representation with High-dimensional Features

Anonymous Author(s)
Affiliation
Address
email

Abstract

Set representation has become ubiquitous in deep learning for modeling the inductive bias of neural networks that are insensitive to the input order. DeepSets is the most widely used neural network architecture for set representation. It involves embedding each set element into a latent space with dimension L, followed by a sum pooling to obtain a whole-set embedding, and finally mapping the whole-set embedding to the output. In this work, we investigate the impact of the dimension L on the expressive power of DeepSets. Previous analyses either oversimplified high-dimensional features to be one-dimensional features or were limited to analytic activations, thereby diverging from practical use or resulting in L that grows exponentially with the set size N and feature dimension D. To investigate the minimal value of L that achieves sufficient expressive power, we present two set-element embedding layers: (a) linear + power activation (LP) and (b) logarithm + linear + exponential activations (LLE). We demonstrate that L being $\mathrm{poly}(N, D)$ is sufficient for set representation using both embedding layers. We also provide a lower bound of L for the LP embedding layer. Furthermore, we extend our results to permutation-equivariant set functions and the complex field.

1 Introduction

Enforcing invariance into neural network architectures has become a widely-used principle to design deep learning models [1–7]. In particular, when a task is to learn a function with a set as the input, the architecture enforces permutation invariance that asks the output to be invariant to the permutation of the input set elements [8, 9]. Neural networks to learn a set function have found a variety of applications in particle physics [10, 11], computer vision [12, 13] and population statistics [14–16], and have recently become a fundamental module (the aggregation operation of neighbors' features in a graph [17–19]) in graph neural networks (GNNs) [20, 21] that show even broader applications.

Previous works have studied the expressive power of neural network architectures to represent set functions [8, 9, 22–26]. Formally, a set with N elements can be represented as $\mathcal{S} = \{\boldsymbol{x}^{(1)}, \cdots, \boldsymbol{x}^{(N)}\}$ where $\boldsymbol{x}^{(i)}$ is in a feature space $\mathcal{X}$, typically $\mathcal{X} = \mathbb{R}^D$. To represent a set function that takes $\mathcal{S}$ and outputs a real value, the most widely used architecture DeepSets [9] follows Eq. (1).

$$f(\mathcal{S}) = \rho\left(\sum_{i=1}^{N} \phi(\boldsymbol{x}^{(i)})\right), \text{where } \phi : \mathcal{X} \to \mathbb{R}^L \text{ and } \rho : \mathbb{R}^L \to \mathbb{R} \text{ are continuous functions.} \quad (1)$$

DeepSets encodes each set element individually via ϕ, and then maps the encoded vectors after sum pooling to the output via ρ. The continuity of ϕ and ρ ensure that they can be well approximated by fully-connected neural networks [27, 28], which has practical implications. DeepSets enforces permutation invariance because of the sum pooling, as shuffling the order of $\boldsymbol{x}^{(i)}$ does not change

Submitted to 37th Conference on Neural Information Processing Systems (NeurIPS 2023). Do not distribute.

Table 1: A comprehensive comparison among all prior works on expressiveness analysis with L. Our results achieve the tightest bound on L while being able to analyze high-dimensional set features and extend to the equivariance case.

Prior Arts	L	$D > 1$	**Exact Rep.**	**Equivariance**
DeepSets [9]	$D+1$	✗	✓	✓
Wagstaff et al. [23]	D	✗	✓	✓
Segol et al. [25]	$\binom{N+D}{N}-1$	✓	✗	✓
Zweig & Bruna [26]	$\exp(\min\{\sqrt{N}, D\})$	✓	✗	✗
Our results	$\mathrm{poly}(N, D)$	✓	✓	✓

the output. However, the sum pooling compresses the whole set into an L-dimension vector, which places an information bottleneck in the middle of the architecture. Therefore, a core question on using DeepSets for set function representation is that given the input feature dimension D and the set size N, what the minimal L is needed so that the architecture Eq. (1) can represent/universally approximate any continuous set functions. The question has attracted attention in many previous works [9, 23–26] and is the focus of the present work.

An extensive understanding has been achieved for the case with one-dimensional features ($D = 1$). Zaheer et al. [9] proved that this architecture with bottleneck dimension $L = N$ suffices to *accurately* represent any continuous set functions when $D = 1$. Later, Wagstaff et al. proved that accurate representations cannot be achieved when $L < N$ [23] and further strengthened the statement to *a failure in approximation* to arbitrary precision in the infinity norm when $L < N$ [24].

However, for the case with high-dimensional features ($D > 1$), the characterization of the minimal possible L is still missing. Most of previous works [9, 25, 29] proposed to generate multi-symmetric polynomials to approximate permutation invariant functions [30]. As the algebraic basis of multi-symmetric polynomials is of size $L^* = \binom{N+D}{N} - 1$ [31] (exponential in $\min\{D, N\}$), these works by default claim that if $L \geq L^*$, f in Eq. 1 can approximate any continuous set functions, while they do not check the possibility of using a smaller L. Zweig and Bruna [26] constructed a set function that f requires bottleneck dimension $L > N^{-2}\exp(O(\min\{D, \sqrt{N}\}))$ (still exponential in $\min\{D, \sqrt{N}\}$) to approximate while it relies on the condition that ϕ, ρ only adopt analytic activations. This condition is overly strict, as most of the practical neural networks allow using non-analytic activations, such as ReLU. Zweig and Bruna thus left an open question *whether the exponential dependence on N or D of L is still necessary if ϕ, ρ allow using non-analytic activations.*

Present work. The main contribution of this work is to confirm a negative response to the above question. Specfically, we present the first theoretical justification that L being *polynomial* in N and D is sufficient for DeepSets (Eq. (1)) like architecture to represent any *continuous* set functions with *high-dimensional* features ($D > 1$). To mitigate the gap to the practical use, we consider two architectures to implement feature embedding ϕ (in Eq. 1) and specify the bounds on L accordingly:

- ϕ adopts *a linear layer with power mapping*: The minimal L holds a lower bound and an upper bound, which is $N(D + 1) \leq L < N^5 D^2$.
- Constrained on the entry-wise positive input space $\mathbb{R}^{N\times D}_{>0}$, ϕ adopts *two layers with logarithmic and exponential activations respectively*: The minimal L holds a tighter upper bound $L \leq 2N^2D^2$.

We prove that if the function ρ could be any continuous function, the above two architectures reproduce the precise construction of any set functions for high-dimensional features $D > 1$, akin to the result in [9] for $D = 1$. This result contrasts with [25, 26] which only present approximating representations. If ρ adopts a fully-connected neural network that allows approximation of any continuous functions on a bounded input space [27, 28], then the DeepSets architecture $f(\cdot)$ can approximate any set functions universally on that bounded input space. Moreover, our theory can be easily extended to permutation-equivariant functions and complex set functions, where the minimal L shares the same bounds up to some multiplicative constants.

Another comment on our contributions is that Zweig and Bruna [26] use difference in the needed dimension L to illustrate the gap between DeepSets [9] and Relational Network [32] in their expressive powers, where the latter encodes set elements in a pairwise manner rather than in a separate manner. The gap well explains the empirical observation that Relational Network achieves better expressive power with smaller L [23, 33]. Our theory does not violate such an observation while it shows that the gap can be reduced from an exponential order in N and D to a polynomial order. Moreover, many real-world applications have computation constraints where only DeepSets instead of Relational Network can be used, e.g., the neighbor aggregation operation in GNN being applied to large networks [21], and the hypergraph neural diffusion operation in hypergraph neural networks [7]. Our theory points out that in this case, it is sufficient to use polynomial L dimension to embed each element, while one needs to adopt a function ρ with non-analytic activitions.

2 Preliminaries

2.1 Notations and Problem Setup

We are interested in the approximation and representation of functions defined over sets [1]. In convention, an N-sized set $\mathcal{S} = \{\boldsymbol{x}^{(1)}, \cdots, \boldsymbol{x}^{(N)}\}$, where $\boldsymbol{x}^{(i)} \in \mathbb{R}^D, \forall i \in [N] (\triangleq \{1, 2, ..., N\})$, can be denoted by a data matrix $\boldsymbol{X} = \begin{bmatrix} \boldsymbol{x}^{(1)} & \cdots & \boldsymbol{x}^{(N)} \end{bmatrix}^\top \in \mathbb{R}^{N \times D}$. Note that we use the superscript (i) to denote the i-th set element and the subscript i to denote the i-th column/feature channel of $\boldsymbol{X}$, i.e., $\boldsymbol{x}_i = \begin{bmatrix} x_i^{(1)} & \cdots & x_i^{(N)} \end{bmatrix}^\top$. Let $\Pi(N)$ denote the set of all N-by-N permutation matrices. To characterize the unorderedness of a set, we define an equivalence class over $\mathbb{R}^{N \times D}$:

Definition 2.1 (Equivalence Class)**.** If matrices $\boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N \times D}$ represent the same set $\mathcal{X}$, then they are called equivalent up a row permutation, denoted as $\boldsymbol{X} \sim \boldsymbol{X'}$. Or equivalently, $\boldsymbol{X} \sim \boldsymbol{X'}$ if and only if there exists a matrix $\boldsymbol{P} \in \Pi(N)$ such that $\boldsymbol{X} = \boldsymbol{P}\boldsymbol{X'}$.

Set functions can be in general considered as permutation-invariant or permutation-equivariant functions, which process the input matrices regardless of the order by which rows are organized. The formal definitions of permutation-invariant/equivariant functions are presented as below:

Definition 2.2. (Permutation Invariance) A function $f : \mathbb{R}^{N \times D} \to \mathbb{R}^{D'}$ is called permutation-invariant if $f(\boldsymbol{P}\boldsymbol{X}) = f(\boldsymbol{X})$ for any $\boldsymbol{P} \in \Pi(N)$.

Definition 2.3. (Permutation Equivariance) A function $f : \mathbb{R}^{N \times D} \to \mathbb{R}^{N \times D'}$ is called permutation-equivariant if $f(\boldsymbol{P}\boldsymbol{X}) = \boldsymbol{P}f(\boldsymbol{X})$ for any $\boldsymbol{P} \in \Pi(N)$.

In this paper, we investigate the approach to design a neural network architecture with permutation invariance/equivariance. Below we will first focus on permutation-invariant functions $f : \mathbb{R}^{N \times D} \to \mathbb{R}$. Then, in Sec. 5, we show that we can easily extend the established results to permutation-equivariant functions through the results provided in [7, 34] and to the complex field. The obtained results for $D' = 1$ can also be easily extended to $D' > 1$ as otherwise f can be written as $\begin{bmatrix} f_1 & \cdots & f_{D'} \end{bmatrix}^\top$ and each f_i has single output feature channel.

2.2 DeepSets and The Difficulty in the High-Dimensional Case $D > 1$

The seminal work [9] establishes the following result which induces a neural network architecture for permutation-invariant functions.

Theorem 2.4 (DeepSets [9], $D = 1$)**.** *A continuous function $f : \mathbb{R}^N \to \mathbb{R}$ is permutation-invariant (i.e., a set function) if and only if there exists continuous functions $\phi : \mathbb{R} \to \mathbb{R}^L$ and $\rho : \mathbb{R}^L \to \mathbb{R}$ such that $f(\boldsymbol{X}) = \rho\left(\sum_{i=1}^N \phi(x^{(i)})\right)$, where L can be as small as N. Note that, here $x^{(i)} \in \mathbb{R}$.*

Remark 2.5. The original result presented in [9] states the latent dimension should be as large as $N + 1$. [23] tighten this dimension to exactly N.

Theorem 2.4 implies that as long as the latent space dimension $L \geq N$, any permutation-invariant functions can be implemented by a unified manner as DeepSets (Eq.(1)). Furthermore, DeepSets suggests a useful architecture for ϕ at the analysis convenience and empirical utility, which is formally defined below ($\phi = \psi_L$):

Definition 2.6 (Power mapping)**.** A power mapping of degree K is a function $\psi_K : \mathbb{R} \to \mathbb{R}^K$ which transforms a scalar to a power series: $\psi_K(z) = \begin{bmatrix} z & z^2 & \cdots & z^K \end{bmatrix}^\top$.

[1] In fact, we allow repeating elements in $\mathcal{S}$, therefore, $\mathcal{S}$ should be more precisely called multiset. With a slight abuse of terminology, we interchangeably use terms multiset and set throughout the whole paper.

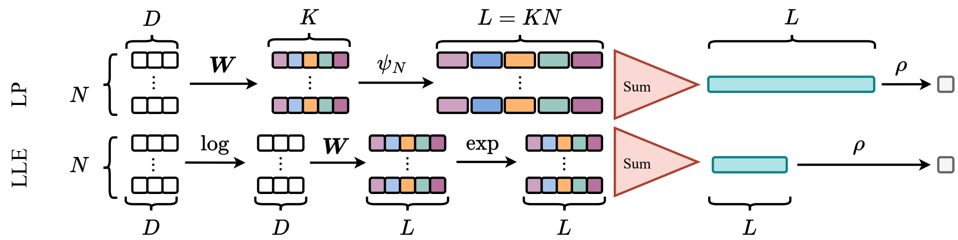

Figure 1: Illustration of the proposed linear + power mapping embedding layer (LP) and logarithm activation + linear + exponential activation embedding layer (LLE).

However, DeepSets [9] focuses on the case that the feature dimension of each set element is one (i.e., $D = 1$). To demonstrate the difficulty extending Theorem 2.4 to high-dimensional features, we reproduce the proof next, which simultaneously reveals its significance and limitation. Some intermediate results and mathematical tools will be recalled along the way later in our proof.

We begin by defining sum-of-power mapping (of degree K) $\Psi_K(\boldsymbol{X}) = \sum_{i=1}^{N} \psi_K(x_i)$, where ψ_K is the power mapping following Definition 2.6. Afterwards, we reveal that sum-of-power mapping $\Psi_K(\boldsymbol{X})$ has a continuous inverse. Before stating the formal argument, we formally define the injectivity of permutation-invariant mappings:

Definition 2.7 (Injectivity)**.** A set function $h : \mathbb{R}^{N\times D} \to \mathbb{R}^L$ is injective if there exists a function $g : \mathbb{R}^L \to \mathbb{R}^{N\times D}$ such that for any $\boldsymbol{X} \in \mathbb{R}^{N\times D}$, we have $g \circ f(\boldsymbol{X}) \sim \boldsymbol{X}$. Then g is an inverse of f.

And we summarize the existence of continuous inverse of $\Psi_K(\boldsymbol{x})$ into the following lemma shown by [9] and improved by [23]. This result comes from homeomorphism between roots and coefficients of monic polynomials [35].

Lemma 2.8 (Existence of Continuous Inverse of Sum-of-Power [9,23])**.** *$\Psi_N : \mathbb{R}^N \to \mathbb{R}^N$ is injective, thus the inverse $\Psi_N^{-1} : \mathbb{R}^N \to \mathbb{R}^N$ exists. Moreover, Ψ_N^{-1} is continuous.*

Now we are ready to prove necessity in Theorem 2.4 as sufficiency is easy to check. By choosing $\phi = \psi_N : \mathbb{R} \to \mathbb{R}^N$ to be the power mapping (cf. Definition 2.6), and $\rho = f \circ \Psi_N^{-1}$. For any scalar-valued set $\boldsymbol{X} = \begin{bmatrix} x^{(1)} & \cdots & x^{(N)} \end{bmatrix}^\top$, $\rho\left(\sum_{i=1}^{N} \phi(x^{(i)})\right) = f \circ \Psi_N^{-1} \circ \Psi_N(\boldsymbol{x}) = f(\boldsymbol{P}\boldsymbol{X}) = f(\boldsymbol{X})$ for some $\boldsymbol{P} \in \Pi(N)$. The existence and continuity of Ψ_N^{-1} are due to Lemma 2.8.

Theorem 2.4 gives the *exact decomposable form* [23] for permutation-invariant functions, which is stricter than approximation error based expressiveness analysis. In summary, the key idea is to establish a mapping ϕ whose element-wise sum-pooling has a continuous inverse.

Curse of High-dimensional Features. We argue that the proof of Theorem 2.4 is not applicable to high-dimensional set features ($D \geq 2$). The main reason is that power mapping defined in Definition 2.6 only receives scalar input. It remains elusive how to extend it to a multivariate version that admits injectivity and a continuous inverse. A plausible idea seems to be applying power mapping for each channel $\boldsymbol{x}_i$ independently, and due to the injectivity of sum-of-power mapping Ψ_N, each channel can be uniquely recovered individually via the inverse Ψ_N^{-1}. However, we point out that each recovered feature channel $\boldsymbol{x'}_i \sim \boldsymbol{x}_i, \forall i \in [D]$, does not imply $\begin{bmatrix} \boldsymbol{x'}_1 & \cdots & \boldsymbol{x'}_D \end{bmatrix} \sim \boldsymbol{X}$, where the alignment of features across channels gets lost. Hence, channel-wise power encoding no more composes an injective mapping. Zaheer et al. [9] proposed to adopt multivariate polynomials as ϕ for high-dimensional case, which leverages the fact that multivariate symmetric polynomials are dense in the space of permutation invariant functions (akin to Stone-Wasserstein theorem) [30]. This idea later got formalized in [25] by setting $\phi(\boldsymbol{x}^{(i)}) = \begin{bmatrix} \cdots & \prod_{j\in[D]} (x_j^{(i)})^{\alpha_j} & \cdots \end{bmatrix}$ where $\boldsymbol{\alpha} \in \mathbb{N}^D$ traverses all $\sum_{j\in[D]} \alpha_j \leq n$ and extended to permutation equivariant functions. Nevertheless, the dimension $L = \binom{N+D}{D}$, i.e., exponential in $\min\{N, D\}$ in this case, and unlike DeepSets [9] which exactly recovers f for $D = 1$, the architecture in [9, 25] can only approximate the desired function.

3 Main Results

In this section, we present our main result which extends Theorem 2.4 to high-dimensional features. Our conclusion is that to universally represent a set function on sets of length N and feature dimension

D with the DeepSets architecture [9] (Eq. (1)), a dimension L at most polynomial in N and D is needed for expressing the intermediate embedding space.

Formally, we summarize our main result in the following theorem.

Theorem 3.1 (The main result)**.** *Suppose $D \geq 2$. For any continuous permutation-invariant function $f : \mathcal{K}^{N\times D} \to \mathbb{R}$, $\mathcal{K} \subseteq \mathbb{R}$, there exists two continuous mappings $\phi : \mathbb{R}^D \to \mathbb{R}^L$ and $\rho : \mathbb{R}^L \to \mathbb{R}$ such that for every $\boldsymbol{X} \in \mathcal{K}^{N\times D}$, $f(\boldsymbol{X}) = \rho\left(\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})\right)$ where*

- *For some $L \in [N(D+1), N^5D^2]$ when ϕ admits **linear layer + power mapping (LP)** architecture:*

$$\phi(\boldsymbol{x}) = \begin{bmatrix}\psi_N(\boldsymbol{w}_1\boldsymbol{x})^\top & \cdots & \psi_N(\boldsymbol{w}_K\boldsymbol{x})^\top\end{bmatrix} \tag{2}$$

  for some $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_K \in \mathbb{R}^D$, and $K = L/N$.

- *For some $L \in [ND, 2N^2D^2]$ when ϕ admits **logarithm activations + linear layer + exponential activations (LLE)** architecture:*

$$\phi(\boldsymbol{x}) = [\exp(\boldsymbol{w}_1 \log(\boldsymbol{x})) \quad \cdots \quad \exp(\boldsymbol{w}_L \log(\boldsymbol{x}))] \tag{3}$$

  for some $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_L \in \mathbb{R}^D$ and $\mathcal{K} \subseteq \mathbb{R}_{>0}$.

The bounds of L depend on the choice of the architecture of ϕ, which are illustrated in Fig. 1. In the LP setting, we adopt a linear layer that maps each set element into K dimension. Then we apply a channel-wise power mapping that separately transforms each value in the feature vector into an N-order power series, and concatenates all the activations together, resulting in a KN dimension feature. The LP architecture is closer to DeepSets [9] as they share the power mapping as the main component. Theorem 3.1 guarantees the existence of ρ and ϕ (in the form of Eq. (2)) which satisfy Eq. (1) without the need to set K larger than N^4D^2 while $K \geq D+1$ is necessary. Therefore, the total embedding size $L = KN$ is bounded by N^5D^2 above and $N(D+1)$ below. Note that this lower bound is not trivial as ND is the degree of freedom of the input $\boldsymbol{X}$. No matter how $\boldsymbol{w}_1, ..., \boldsymbol{w}_K$ are adopted, one cannot achieve an injective mapping by just using ND dimension.

In the LLE architecture, we investigate the utilization of logarithmic and exponential activations in set representation, which are also valid activations to build deep neural networks [36, 37]. Each set entry will be squashed by a element-wise logarithm first, then linearly embedded into an L-dimensional space via a group of weights, and finally transformed by an element-wise exponential activation. Essentially, each $\exp(\boldsymbol{w}_i \log(\boldsymbol{x})), i \in [L]$ gives a monomial of $\boldsymbol{x}$. The LLE architecture requires the feature space constrained on the positive orthant to ensure logarithmic operations are feasible. But the advantage is that the upper bound of L is improved to be $2N^2D^2$. The lower bound ND for the LLE architecture is a trivial bound due to the degree of freedom of the input $\boldsymbol{X}$. Note that the constraint on the positive orthant $\mathbb{R}_{>0}$ is not essential. If we are able to use monomial activations to process a vector $\boldsymbol{x}$ as used in [25, 26], then, the constraint on the positive orthant can be removed.

Remark 3.2. The bounds in Theorem 3.1 are non-asymptotic. This implies the latent dimensions specified by the corresponding architectures are precisely sufficient for expressing the input.

Remark 3.3. Unlike ϕ, the form of ρ cannot be explicitly specified, as it depends on the desired function f. The complexity of ρ remains unexplored in this paper, which may be high in practice.

Importance of Continuity. We argue that the requirements of continuity on ρ and ϕ are essential for our discussion. First, practical neural networks can only provably approximate continuous functions [27, 28]. Moreover, set representation without such requirements can be straightforward (but likely meaningless in practice). This is due to the following lemma.

Lemma 3.4 ([38])**.** *There exists a discontinuous bijective mapping between $\mathbb{R}^D$ and $\mathbb{R}$ if $D \geq 2$.*

By Lemma 3.4, we can define a bijective mapping $r : \mathbb{R}^D \to \mathbb{R}$ which maps the high-dimensional features to scalars, and its inverse exists. Then, the same proof of Theorem 2.4 goes through by letting $\phi = \psi_N \circ r$ and $\rho = f \circ r^{-1} \circ \Psi_N^{-1}$. However, we note both ρ and ϕ lose continuity.

Comparison with Prior Arts. Below we highlight the significance of Theorem 3.1 in contrast to the existing literature. A quick overview is listed in Tab. 1 for illustration. The lower bound in Theorem 3.1 corrects a natural misconception that the degree of freedom (i.e., $L = ND$ for multi-channel cases) is not enough for representing the embedding space. Fortunately, the upper bound in Theorem 3.1 shows the complexity of representing vector-valued sets is still manageable as it merely scales polynomially in N and D. Compared with Zweig and Bruna's finding [26], our result significantly improves this bound on L from exponential to polynomial by allowing non-analytic functions to amortize the expressiveness. Besides, Zweig and Bruna's work [26] is hard to be applied to the real domain, while ours are extensible to complex numbers and equivariant functions.

4 Proof Sketch

In this section, we introduce the proof techniques of Theorem 3.1, while deferring a full version and all missing proofs to the supplementary materials.

The proof of Theorem 3.1 mainly consists of two steps below, which is completely constructive:

1. For the LP architecture, we construct a group of K linear weights $\boldsymbol{w}_1 \cdots, \boldsymbol{w}_K$ with $K \le N^4 D^2$ such that the summation over the associated LP embedding (Eq. (2)): $\Psi(\boldsymbol{X}) = \sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})$ is injective and has a continuous inverse. Moreover, if $K \le D$, such weights do not exist, which induces the lower bound.
2. Similarly, for the LLE architecture, we construct a group of L linear weights $\boldsymbol{w}_1 \cdots, \boldsymbol{w}_L$ with $L \le 2N^2D^2$ such that the summation over the associated LLE embedding (Eq. (3)) is injective and has a continuous inverse. Trivially, if $L < ND$, such weights do not exist, which induces the lower bound.
3. Then the proof of upper bounds can be concluded for both settings by letting $\rho = f \circ \Psi^{-1}$ since $\rho\left(\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})\right) = f \circ \Psi^{-1} \circ \Psi(\boldsymbol{X}) = f(\boldsymbol{P}\boldsymbol{X}) = f(\boldsymbol{X})$ for some $\boldsymbol{P} \in \Pi(N)$.

Next, we elaborate on the construction idea which yields injectivity for both embedding layers in Sec. 4.1 and 4.2, respectively. To show injectivity, it is equivalent to establish the following statement for both Eq. (2) and Eq. (3), respectively:

$$\forall \boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N \times D}, \sum_{i=1}^N \phi(\boldsymbol{x}^{(i)}) = \sum_{i=1}^N \phi(\boldsymbol{x'}^{(i)}) \Rightarrow \boldsymbol{X} \sim \boldsymbol{X'} \tag{4}$$

In Sec. 4.3, we prove the continuity of the inverse map for LP and LLE via arguments similar to [35].

4.1 Injectivity of LP

In this section, we consider ϕ follows the definition in Eq. (2), which amounts to first linearly transforming each set element and then applying channel-wise power mapping. This is, we seek a group of linear transformations $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_K$ such that $\boldsymbol{X} \sim \boldsymbol{X'}$ can be induced from $\boldsymbol{X}\boldsymbol{w}_i \sim \boldsymbol{X'}\boldsymbol{w}_i, \forall i \in [K]$ for some K larger than N while being polynomial in N and D. The intuition is that linear mixing among each channel can encode relative positional information. Only if $\boldsymbol{X} \sim \boldsymbol{X'}$, the mixing information can be reproduced.

Formally, the first step accords to the property of power mapping (cf. Lemma 2.8), and we can obtain:

$$\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)}) = \sum_{i=1}^N \phi(\boldsymbol{x'}^{(i)}) \Rightarrow \boldsymbol{X}\boldsymbol{w}_i \sim \boldsymbol{X'}\boldsymbol{w}_i, \forall i \in [K]. \tag{5}$$

To induce $\boldsymbol{X} \sim \boldsymbol{X'}$ from $\boldsymbol{X}\boldsymbol{w}_i \sim \boldsymbol{X'}\boldsymbol{w}_i, \forall i \in [K]$, our construction divides the weights $\{\boldsymbol{w}_i, i \in [K]\}$ into three groups: $\{\boldsymbol{w}_i^{(1)} : i \in [D]\}$, $\{\boldsymbol{w}_j^{(2)} : j \in [K_1]\}$, and $\{\boldsymbol{w}_{i,j,k}^{(3)} : i \in [D], j \in [K_1], k \in [K_2]\}$. Each block is outlined as below:

1. Let the first group of weights $\boldsymbol{w}_1^{(1)} = \boldsymbol{e}_1, \cdots, \boldsymbol{w}_D^{(1)} = \boldsymbol{e}_D$ to buffer the original features, where $\boldsymbol{e}_i$ is the i-th canonical basis.
2. Design the second group of linear weights, $\boldsymbol{w}_1^{(2)}, \cdots, \boldsymbol{w}_{K_1}^{(2)}$ for K_1 as large as $N(N-1)(D-1)/2+1$, which, by Lemma 4.4 latter, guarantees at least one of $\boldsymbol{X}\boldsymbol{w}_j^{(2)}, j \in [K_1]$ forms an anchor defined below:

Definition 4.1 (Anchor)**.** Consider the data matrix $\boldsymbol{X} \in \mathbb{R}^{N\times D}$, then $\boldsymbol{a} \in \mathbb{R}^N$ is called an anchor of $\boldsymbol{X}$ if $\boldsymbol{a}_i \neq \boldsymbol{a}_j$ for any $i, j \in [N]$ such that $\boldsymbol{x}^{(i)} \neq \boldsymbol{x}^{(j)}$.

And suppose $\boldsymbol{a} = \boldsymbol{X}\boldsymbol{w}_{j^*}^{(2)}$ is an anchor of $\boldsymbol{X}$ for some $j^* \in [K_1]$ and $\boldsymbol{a'} = \boldsymbol{X'}\boldsymbol{w}_{j^*}^{(2)}$, then we show the following statement is true by Lemma 4.3 latter:

$$[\boldsymbol{a} \quad \boldsymbol{x}_i] \sim [\boldsymbol{a'} \quad \boldsymbol{x'}_i], \forall i \in [D] \Rightarrow \boldsymbol{X} \sim \boldsymbol{X'}. \tag{6}$$

3. Design a group of weights $\boldsymbol{w}_{i,j,k}^{(3)}$ for $i \in [D], j \in [K_1], k \in [K_2]$ with $K_2 = N(N-1)+1$ that mixes each original channel $\boldsymbol{x}_i$ with each $\boldsymbol{X}\boldsymbol{w}_j^{(2)}, j \in [K_1]$ by $\boldsymbol{w}_{i,j,k}^{(3)} = \boldsymbol{e}_i - \gamma_k \boldsymbol{w}_j^{(2)}$. Then we show in Lemma 4.5 that:

$$\boldsymbol{X}\boldsymbol{w}_i \sim \boldsymbol{X'}\boldsymbol{w}_i, \forall i \in [K] \Rightarrow \left[\boldsymbol{X}\boldsymbol{w}_j^{(2)} \quad \boldsymbol{x}_i\right] \sim \left[\boldsymbol{X'}\boldsymbol{w}_j^{(2)} \quad \boldsymbol{x'}_i\right], \forall i \in [D], j \in [K_1] \tag{7}$$

With such configuration, injectivity can be concluded by the entailment along Eq. (5), (7), (6): Eq. (5) guarantees the RHS of Eq. (7); The existence of the anchor in Lemma 4.4 paired with Eq. (6) guarantees $\boldsymbol{X} \sim \boldsymbol{X'}$. The total required number of weights $K = D + K_1 + DK_1K_2 \leq N^4 D^2$.

Below we provides a series of lemmas that demonstrate the desirable properties of anchors and elaborate on the construction complexity. Detailed proofs are left in Appendix. In plain language, by Definition 4.1, two entries in the anchor must be distinctive if the set elements at the corresponding indices are not equal. As a consequence, we derive the following property of anchors:

Lemma 4.2. *Consider the data matrix $\boldsymbol{X} \in \mathbb{R}^{N\times D}$ and $\boldsymbol{a} \in \mathbb{R}^N$ an anchor of $\boldsymbol{X}$. Then if there exists $\boldsymbol{P} \in \Pi(N)$ such that $\boldsymbol{P}\boldsymbol{a} = \boldsymbol{a}$ then $\boldsymbol{P}\boldsymbol{x}_i = \boldsymbol{x}_i$ for every $i \in [D]$.*

With the above property, anchors defined in Definition 4.1 indeed have the entailment in Eq. (6):

Lemma 4.3 (Union Alignment based on Anchor Alignment)**.** *Consider the data matrix $\boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N\times D}$, $\boldsymbol{a} \in \mathbb{R}^N$ is an anchor of $\boldsymbol{X}$ and $\boldsymbol{a'} \in \mathbb{R}^N$ is an arbitrary vector. If $[\boldsymbol{a} \quad \boldsymbol{x}_i] \sim [\boldsymbol{a'} \quad \boldsymbol{x'}_i]$ for every $i \in [D]$, then $\boldsymbol{X} \sim \boldsymbol{X'}$.*

However, the anchor $\boldsymbol{a}$ is required to be generated from $\boldsymbol{X}$ via a point-wise linear transformation. The strategy to generate an anchor is to enumerate as many linear weights as needs, so that for any $\boldsymbol{X}$, at least one j such that $\boldsymbol{X}\boldsymbol{w}_j^{(2)}$ becomes an anchor. We show that at most $N(N-1)(D-1)/2+1$ linear weights are enough to guarantee the existence of an anchor for any $\boldsymbol{X}$:

Lemma 4.4 (Anchor Construction)**.** *There exists a set of weights $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_{K_1}$ where $K_1 = N(N-1)(D-1)/2+1$ such that for every data matrix $\boldsymbol{X} \in \mathbb{R}^{N\times D}$, there exists $j \in [K_1]$, $\boldsymbol{X}\boldsymbol{w}_j$ is an anchor of $\boldsymbol{X}$.*

We wrap off the proof by presenting the following lemma which is applied to prove Eq. (7) by fixing any $i \in [D], j \in [K_1]$ in Eq. (7) while checking the condition for all $k \in [K_2]$:

Lemma 4.5 (Anchor Matching)**.** *There exists a group of coefficients $\gamma_1, \cdots, \gamma_{K_2}$ where $K_2 = N(N-1)+1$ such that the following statement holds: Given any $\boldsymbol{x}, \boldsymbol{x'}, \boldsymbol{y}, \boldsymbol{y'} \in \mathbb{R}^N$ such that $\boldsymbol{x} \sim \boldsymbol{x'}$ and $\boldsymbol{y} \sim \boldsymbol{y'}$, if $(\boldsymbol{x} - \gamma_k \boldsymbol{y}) \sim (\boldsymbol{x'} - \gamma_k \boldsymbol{y'})$ for every $k \in [K_2]$, then $[\boldsymbol{x} \quad \boldsymbol{y}] \sim [\boldsymbol{x'} \quad \boldsymbol{y'}]$.*

For completeness, we add the following lemma which implies LP-induced sum-pooling cannot be injective if $K \leq ND$, when $D \geq 2$.

Theorem 4.6 (Lower Bound)**.** *Consider data matrices $\boldsymbol{X} \in \mathbb{R}^{N\times D}$ where $D \geq 2$. If $K \leq D$, then for every $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_K$, there exists $\boldsymbol{X'} \in \mathbb{R}^{N\times D}$ such that $\boldsymbol{X} \not\sim \boldsymbol{X'}$ but $\boldsymbol{X}\boldsymbol{w}_i \sim \boldsymbol{X'}\boldsymbol{w}_i$ for every $i \in [K]$.*

Remark 4.7. Theorem 4.6 is significant in that with high-dimensional features, the injectivity is provably not satisfied when the embedding space has dimension equal to the degree of freedom.

4.2 Injectivity of LLE

In this section, we consider ϕ follows the definition in Eq. (3). First of all, we note that each term in the RHS of Eq. (3) can be rewritten as a monomial as shown in Eq. (8). Suppose we are able to use monomial activations to process a vector $\boldsymbol{x}^{(i)}$. Then, the constraint on the positive orthant $\mathbb{R}_{>0}$ in our main result Theorem 3.1 can be even removed.

$$\phi(\boldsymbol{x}) = [\cdots \quad \exp(\boldsymbol{w}_i \log(\boldsymbol{x})) \quad \cdots] = \left[\cdots \quad \prod_{j=1}^{D} \boldsymbol{x}_j^{\boldsymbol{w}_{i,j}} \quad \cdots\right] \tag{8}$$

Then, the assignment of $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_L$ amounts to specifying the exponents for D power functions within the product. Next, we prepare our construction with the following two lemmas:

Lemma 4.8. *For any pair of vectors $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^N, \boldsymbol{y}_1, \boldsymbol{y}_2 \in \mathbb{R}^N$, if $\sum_{i\in[N]} \boldsymbol{x}_{1,i}^{l-k}\boldsymbol{x}_{2,i}^{k} = \sum_{i\in[N]} \boldsymbol{y}_{1,i}^{l-k}\boldsymbol{y}_{2,i}^{k}$ for every $l, k \in [N]$ such that $0 \le k \le l$, then $[\boldsymbol{x}_1 \quad \boldsymbol{x}_2] \sim [\boldsymbol{y}_1 \quad \boldsymbol{y}_2]$.*

The above lemma is to show that we may use summations of monic bivariate monomials to align every two feature columns. The next lemma shows that such pairwise alignment yields union alignment.

Lemma 4.9 (Union Alignment based on Pairwise Alignment)**.** *Consider data matrices $\boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N\times D}$. If $[\boldsymbol{x}_i \quad \boldsymbol{x}_j] \sim [\boldsymbol{x'}_i \quad \boldsymbol{x'}_j]$ for every $i, j \in [D]$, then $\boldsymbol{X} \sim \boldsymbol{X'}$.*

Then the construction idea of $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_L$ can be drawn from Lemma 4.8 and 4.9:

1. Lemma 4.8 indicates if the weights in Eq. (8) enumerate all the monic bivariate monomials in each pair of channels with degrees less or equal to N, i.e., $\boldsymbol{x}_i^p\boldsymbol{x}_j^q$ for all $i, j \in [D]$ and $p + q \le N$, then we can yield:

$$\sum_{i=1}^{N} \phi(\boldsymbol{x}^{(i)}) = \sum_{i=1}^{N} \phi(\boldsymbol{x'}^{(i)}) \Rightarrow [\boldsymbol{x}_i \quad \boldsymbol{x}_j] \sim [\boldsymbol{x'}_i \quad \boldsymbol{x'}_j], \forall i, j \in [D]. \tag{9}$$

2. The next step is to invoke Lemma 4.9 which implies if every pair of feature channels is aligned, then we can conclude all the channels are aligned with each other as well.

$$[\boldsymbol{x}_i \quad \boldsymbol{x}_j] \sim [\boldsymbol{x'}_i \quad \boldsymbol{x'}_j], \forall i, j \in [D] \Rightarrow \boldsymbol{X} \sim \boldsymbol{X'}. \tag{10}$$

Based on these motivations, we assign the weights that induce all bivariate monic monomials with the degree no more than N. First of all, we reindex $\{\boldsymbol{w}_i, i \in [L]\}$ as $\{\boldsymbol{w}_{i,j,p,q}, i \in [D], j \in [D], p \in [N], q \in [p+1]\}$. Then weights can be explicitly specified as $\boldsymbol{w}_{i,j,p,q} = (q-1)\boldsymbol{e}_i + (p-q+1)\boldsymbol{e}_j$, where $\boldsymbol{e}_i$ is the i-th canonical basis. With such weights, injectivity can be concluded by entailment along Eq. (9) and (10). Moreover, the total number of linear weights is $L = D^2(N+3)N/2 \le 2N^2D^2$, as desired.

4.3 Continuous Lemma

In this section, we show that the LP and LLE induced sum-pooling are both homeomorphic. We note that it is intractable to obtain the closed form of their inverse maps. Notably, the following remarkable result can get rid of inversing a functions explicitly by merely examining the topological relationship between the domain and image space.

Lemma 4.10. *(Theorem 1.2 [35]) Let $(\mathcal{X}, d_{\mathcal{X}})$ and $(\mathcal{Y}, d_{\mathcal{Y}})$ be two metric spaces and $f : \mathcal{X} \to \mathcal{Y}$ is a bijection such that **(a)** each bounded and closed subset of $\mathcal{X}$ is compact, **(b)** f is continuous, **(c)** f^{-1} maps each bounded set in $\mathcal{Y}$ into a bounded set in $\mathcal{X}$. Then f^{-1} is continuous.*

Subsequently, we show the continuity in an informal but more intuitive way while deferring a rigorous version to the supplementary materials. Denote $\Psi(\boldsymbol{X}) = \sum_{i\in[N]} \phi(\boldsymbol{x}^{(i)})$. To begin with, we set $\mathcal{X} = \mathbb{R}^{N\times D}/\sim$ with metric $d_{\mathcal{X}}(\boldsymbol{X}, \boldsymbol{X}') = \min_{\boldsymbol{P}\in\Pi(N)} \|\boldsymbol{X} - \boldsymbol{P}\boldsymbol{X}'\|_1$ and $\mathcal{Y} = \{\Psi(\boldsymbol{X}) | \boldsymbol{X} \in \mathcal{X}\} \subseteq \mathbb{R}^L$ with metric $d_{\mathcal{Y}}(\boldsymbol{y}, \boldsymbol{y}') = \|\boldsymbol{y} - \boldsymbol{y}'\|_\infty$. It is easy to show that $\mathcal{X}$ satisfies the conditions **(a)** and $\Psi(\boldsymbol{X})$ satisfies **(b)** for both LP and LLE embedding layers. Then it remains to conclude the proof by verifying the condition **(c)** for the mapping $\mathcal{Y} \to \mathcal{X}$, i.e., the inverse of $\Psi(\boldsymbol{X})$. We visualize this mapping following our arguments on injectivity:

$$\begin{array}{lccccc} (LP) & \Psi(\boldsymbol{X}) & \xrightarrow{\text{Eq. (5)}} & [\cdots \quad \boldsymbol{P}_i\boldsymbol{X}\boldsymbol{w}_i \quad \cdots], i \in [K] & \xrightarrow{\text{Eqs. (6) + (7)}} & \boldsymbol{P}\boldsymbol{X} \\ (LLE) & \underbrace{\Psi(\boldsymbol{X})}_{\mathcal{Y}} & \xrightarrow{\text{Eq. (9)}} & \underbrace{[\cdots \quad \boldsymbol{Q}_{i,j}\boldsymbol{x}_i \quad \boldsymbol{Q}_{i,j}\boldsymbol{x}_j \quad \cdots], i, j \in [D]}_{\mathcal{Z}} & \xrightarrow{\text{Eq. (10)}} & \underbrace{\boldsymbol{Q}\boldsymbol{X}}_{\mathcal{X}} \end{array},$$

for some $\boldsymbol{X}$ dependent $\boldsymbol{P}$, $\boldsymbol{Q}$. Here, $\boldsymbol{P}_i, i \in [K]$ and $\boldsymbol{Q}_{i,j}, i, j \in [D] \in \Pi(N)$. According to homeomorphism between polynomial coefficients and roots (Corollary 3.2 in [35]), any bounded set in $\mathcal{Y}$ will induce a bound set in $\mathcal{Z}$. Moreover, since elements in $\mathcal{Z}$ contains all the columns of $\mathcal{X}$ (up to some changes of the entry orders), a bounded set in $\mathcal{Z}$ also corresponds to a bounded set in $\mathcal{X}$. Through this line of arguments, we conclude the proof.

5 Extensions

In this section, we discuss two extensions to Theorem 3.1, which strengthen our main result.

Permutation Equivariance. Permutation-equivariant functions (cf. Definition 2.3) are considered as a more general family of set functions. Our main result does not lose generality to this class of functions. By Lemma 2 of [7], Theorem 3.1 can be directly extended to permutation-equivariant functions with *the same lower and upper bounds*, stated as follows:

Theorem 5.1 (Extension to Equivariance)**.** *For any permutation-equivariant function $f : \mathcal{K}^{N\times D} \to \mathbb{R}^N$, $\mathcal{K} \subseteq \mathbb{R}$, there exists continuous functions $\phi : \mathbb{R}^D \to \mathbb{R}^L$ and $\rho : \mathbb{R}^D \times \mathbb{R}^L \to \mathbb{R}$ such that $f(\boldsymbol{X})_j = \rho\left(\boldsymbol{x}^{(j)}, \sum_{i\in[N]} \phi(\boldsymbol{x}^{(i)})\right)$ for every $j \in [N]$, where $L \in [N(D+1), N^5D^2]$ when ϕ admits LP architecture, and $L \in [ND, 2N^2D^2]$ when ϕ admits LLE architecture ($\mathcal{K} \in \mathbb{R}_{>0}$).*

Complex Domain. The upper bounds in Theorem 3.1 is also true to complex features up to a constant scale (i.e., $\mathcal{K} \subseteq \mathbb{C}$). When features are defined over $\mathbb{C}^{N\times D}$, our primary idea is to divide each channel into two real feature vectors, and recall Theorem 3.1 to conclude the arguments on an $\mathbb{R}^{N\times 2D}$ input. All of our proof strategies are still applied. This result directly contrasts to Zweig and Bruna's work [26] whose main arguments were established on complex numbers. We show that even moving to the complex domain, polynomial length of L is still sufficient for the DeepSets architecture [9]. We state a formal version of the theorem in the supplementary material.

6 Related Work

Works on neural networks to represent set functions have been discussed extensively in the Sec. 1. Here, we review other related works on the expressive power analysis of neural networks.

Early works studied the expressive power of feed-forward neural networks with different activations [27,28]. Recent works focused on characterizing the benefits of the expressive power of deep architectures to explain their empirical success [39–43]. Modern neural networks often enforce some invariance properties into their architectures such as CNNs that capture spatial translation invariance. The expressive power of invariant neural networks has been analyzed recently [22,44,45].

The architectures studied in the above works allow universal approximation of continuous functions defined on their inputs. However, the family of practically useful architectures that enforce permutation invariance often fail in achieving universal approximation. Graph Neural Networks (GNNs) enforce permutation invariance and can be viewed as an extension of set neural networks to encode a set of pair-wise relations instead of a set of individual elements [20,21,46,47]. GNNs suffer from limited expressive power [5, 17, 18] unless they adopt exponential-order tensors [48]. Hence, previous studies often characterized GNNs' expressive power based on their capability of distinguishing non-isomorphic graphs. Only a few works have ever discussed the function approximation property of GNNs [49–51] while these works still miss characterizing such dependence on the depth and width of the architectures [52]. As practical GNNs commonly adopt the architectures that combine feed-forward neural networks with set operations (neighborhood aggregation), we believe the characterization of the needed size for set function approximation studied in [26] and this work may provide useful tools to study finer-grained characterizations of the expressive power of GNNs.

7 Conclusion

This work investigates how many neurons are needed to model the embedding space for set representation learning with the DeepSets architecture [9]. Our paper provides an affirmative answer that polynomial many neurons in the set size and feature dimension are sufficient. Compared with prior arts, our theory takes high-dimensional features into consideration while significantly advancing the state-of-the-art results from exponential to polynomial.

Limitations. The tightness of our bounds is not examined in this paper, and the complexity of ρ is uninvestigated and left for future exploration. Besides, deriving an embedding layer agnostic lower bound for the embedding space remains another widely open question.

A Formal Definitions

In this section, we begin by providing rigorous definitions to specify the topology of the input space of permutation-invariant functions.

Definition A.1. Equipped $\mathcal{K}^{N\times D}$ with the equivalence relation $\sim$ (cf. Definition 2.1), we define metric space $(\mathcal{K}^{N\times D}/\sim, d_F)$, where $d_F : (\mathcal{K}^{N\times D}/\sim) \times (\mathcal{K}^{N\times D}/\sim) \to \mathbb{R}_{\geq 0}$ is the optimal transport distance:

$$d_F(\boldsymbol{X}, \boldsymbol{X'}) = \min_{\boldsymbol{P}\in\Pi(N)} \|\boldsymbol{PX} - \boldsymbol{X'}\|_{\infty,\infty}, \tag{11}$$

and $\mathcal{K}$ can be either $\mathbb{R}$ or $\mathbb{C}$.

Remark A.2. The $\|\cdot\|_{\infty,\infty}$ norm takes the absolute value of the maximal entry: $\max_{i\in[N],j\in[D]} |\boldsymbol{X}_{i,j}|$. Other topologically equivalent matrix norms also apply.

Lemma A.3. *The function $d_F : (\mathcal{K}^{N\times D}/\sim) \times (\mathcal{K}^{N\times D}/\sim) \to \mathbb{R}_{\geq 0}$ is a distance metric on $\mathcal{K}^{N\times D}/\sim$.*

Proof. Identity, positivity, and symmetry trivially hold for d_F. It remains to show the triangle inequality as below: for arbitrary $\boldsymbol{X}, \boldsymbol{X'}, \boldsymbol{X''} \in (\mathcal{K}^{N\times D}/\sim, d_F)$,

$$\begin{aligned} d_F(\boldsymbol{X}, \boldsymbol{X''}) &= \min_{\boldsymbol{P}\in\Pi(N)} \|\boldsymbol{PX} - \boldsymbol{X''}\|_{\infty,\infty} \leq \min_{\boldsymbol{P}\in\Pi(N)} \left(\|\boldsymbol{PX} - \boldsymbol{Q}^*\boldsymbol{X'}\|_{\infty,\infty} + \|\boldsymbol{Q}^*\boldsymbol{X'} - \boldsymbol{X''}\|_{\infty,\infty} \right) \\ &= \min_{\boldsymbol{P}\in\Pi(N)} \|\boldsymbol{PX} - \boldsymbol{Q}^*\boldsymbol{X'}\|_{\infty,\infty} + \|\boldsymbol{Q}^*\boldsymbol{X'} - \boldsymbol{X''}\|_{\infty,\infty} \\ &= d_F(\boldsymbol{X}, \boldsymbol{X'}) + d_F(\boldsymbol{X}, \boldsymbol{X''}), \end{aligned}$$

where $\boldsymbol{Q}^* = \operatorname{argmin}_{\boldsymbol{Q}\in\Pi(N)} \|\boldsymbol{QX'} - \boldsymbol{X''}\|_{\infty,\infty}$. □

Also we reveal a topological property for $(\mathcal{K}^{N\times D}/\sim, d_F)$ which is essential to show continuity later.

Lemma A.4. *Each bounded and closed subset of $(\mathcal{K}^{N\times D}/\sim, d_F)$ is compact.*

Proof. Without loss of generality, the proof is done by extending Theorem 2.4 in [35] to high-dimensional set elements. □

Then we can rephrase the definition of permutation invariant function as a proper function mapping between the two metric spaces: $f : (\mathcal{K}^{N\times D}/\sim, d_F) \to (\mathcal{K}^{D'}, d_\infty)$, where $d_\infty : \mathcal{K}^{D'} \times \mathcal{K}^{D'} \to \mathbb{R}_{\geq 0}$ is the ℓ_∞-norm induced distance metric.

We also recall the definition of injectivity for permutation-invariant functions:

Definition A.5 (Injectivity)**.** The following statements are equivalent:

1. A permutation-invariant function $f : (\mathcal{K}^{N\times D}/\sim, d_F) \to (\mathcal{K}^{D'}, d_\infty)$ is injective.
2. There exists a function $g : (\mathcal{K}^{D'}, d_\infty) \to (\mathcal{K}^{N\times D}/\sim, d_F)$ such that for every $\boldsymbol{X} \in \mathcal{K}^{N\times D}$, $g \circ f(\boldsymbol{X}) \sim \boldsymbol{X}$.
3. For every $\boldsymbol{X}, \boldsymbol{X'} \in \mathcal{K}^{N\times D}$ such that $f(\boldsymbol{X}) = f(\boldsymbol{X'})$, then $\boldsymbol{X} \sim \boldsymbol{X'}$.

We give an intuitive definition of continuity for permutation-invariant functions via the epsilon-delta statement:

Definition A.6 (Continuity)**.** A permutation-invariant function $f : (\mathcal{K}^{N\times D}/\sim, d_F) \to (\mathcal{K}, d_\infty)$ is continuous if for arbitrary $\boldsymbol{X} \in \mathcal{K}^{N\times D}$ and $\epsilon > 0$, there exists $\delta > 0$ such that for every $\boldsymbol{X'} \in \mathcal{K}^{N\times D}$, $d_F(\boldsymbol{X}, \boldsymbol{X'}) < \delta$ then $d_\infty(f(\boldsymbol{X}) - f(\boldsymbol{X'})) < \epsilon$.

Remark A.7. Since d_F is a distance metric, other equivalent definitions of continuity still applies.

B Sum-of-Power Mapping

In this section, we extend Definition 2.6 and Lemma 2.8 to the complex version, which provides the mathematical tools for our later proof.

Definition B.1. Define (complex) power mapping: $\psi_N : \mathbb{C} \to \mathbb{C}^N$, $\psi_N(z) = \begin{bmatrix} z & z^2 & \cdots & z^N \end{bmatrix}^\top$ and (complex) sum-of-power mapping $\psi_N : (\mathbb{C}^N / \sim, d_F) \to (\mathbb{C}^N, d_\infty)$, $\psi_N(\boldsymbol{z}) = \sum_{i=1}^N \psi_N(z_i)$.

Lemma B.2 (Existence of Continuous Inverse of Complex Sum-of-Power [35])**.** *ψ_N is injective, thus the inverse $\psi_N^{-1} : (\mathbb{C}^N, d_\infty) \to (\mathbb{C}^N / \sim, d_F)$ exists. Moreover, ψ_N^{-1} is continuous.*

Lemma B.3 (Corollary 3.2 [35])**.** *Consider a function $\zeta : (\mathbb{C}^N, d_\infty) \to (\mathbb{C}^N / \sim, d_F)$ that maps the coefficients of a polynomial to its root multi-set. Then for any bounded subset $\mathcal{U} \subset (\mathbb{C}^N, d_\infty)$, the image $\zeta(\mathcal{U}) = \{\zeta(\boldsymbol{z}) : \boldsymbol{z} \in \mathcal{U}\}$ is also bounded.*

Remark B.4. Lemma B.3 is also true for real numbers when we constrain the domain of ζ to be real coefficients such that the corresponding polynomial can fully split over the real domain.

Lemma B.5. *Consider the N-degree sum-of-power mapping: $\psi_N : (\mathbb{C}^{\mathbb{C},N} / \sim, d_F) \to (\mathbb{C}^N, d_\infty)$, where $\psi_N(\boldsymbol{x}) = \sum_{i=1}^N \psi_N(x_i)$. Denote the range of ψ_N as $\mathcal{Z} \subseteq \mathbb{C}^N$ and its inverse mapping $\psi_N^{-1} : (\mathcal{Z}, d_\infty) \to (\mathbb{C}^N / \sim, d_F)$ (existence guaranteed by Lemma B.2). Then for every bounded set $\mathcal{U} \subset (\mathcal{Z}, d_\infty)$, the image $\Psi_N^{-1}(\mathcal{U}) = \{\Psi_N^{-1}(\boldsymbol{z}) : \boldsymbol{z} \in \mathcal{U}\}$ is also bounded.*

Proof. We borrow the proof technique from [9] to reveal a polynomial mapping between $(\mathcal{Z}, d_\infty)$ and $(\mathbb{C}^N / \sim, d_F)$. For every $\boldsymbol{\xi} \in (\mathbb{C}^N / \sim, d_F)$, let $\boldsymbol{z} = \psi_N(\boldsymbol{\xi})$ and construct a polynomial:

$$P_{\boldsymbol{\xi}}(x) = \prod_{i=1}^N (x - \xi_i) = x^N - a_1 x^{N-1} + \cdots + (-1)^{N-1} a_{N-1} x + (-1)^N a_N, \tag{12}$$

where $\boldsymbol{\xi}$ are the roots of $P_{\boldsymbol{\xi}}(x)$ and the coefficients can be written as elementary symmetric polynomials, i.e.,

$$a_n = \sum_{1 \le j_1 \le j_2 \le \cdots \le j_n \le N} \xi_{j_1} \xi_{j_2} \cdots \xi_{j_n}, \forall n \in [N]. \tag{13}$$

On the other hand, the elementary symmetric polynomials can be uniquely expressed as a function of $\boldsymbol{z}$ by Newton-Girard formula:

$$a_n = \frac{1}{n} \det \begin{bmatrix} z_1 & 1 & 0 & 0 & \cdots & 0 \\ z_2 & z_1 & 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots \\ z_{n-1} & z_{n-2} & z_{n-3} & z_{n-4} & \cdots & 1 \\ z_n & z_{n-1} & z_{n-2} & z_{n-3} & \cdots & 1 \end{bmatrix} := Q(\boldsymbol{z}), \forall n \in [N] \tag{14}$$

where the determinant $Q(\boldsymbol{z})$ is also a polynomial in $\boldsymbol{z}$. Then the proof proceeds by observing that for any bounded subset $\mathcal{U} \subseteq (\mathcal{Z}, d_\infty)$, the resulting $\mathcal{A} = Q(\mathcal{U})$ is also bounded in $(\mathbb{C}^N, d_\infty)$. Therefore, by Lemma B.3, any bounded coefficient set $\mathcal{A}$ will produce a bounded root multi-set in $(\mathbb{C}^N / \sim, d_F)$. $\square$

Remark B.6. Lemma B.3 is also true for real numbers. By Remark B.4, we can constrain the ambient space of $\mathcal{A}$ in Lemma B.3 to be real coefficients whose corresponding polynomial can split over real numbers, and the same proof proceeds.

C Proofs of LP Embedding Layer

In this section, we complete the proofs for the LP embedding layer (Eq. (2)). First we constructively show an upper bound that sufficiently achieves injectivity following our discussion in Sec. 4.1, and then prove Theorem 4.6 to reveal a lower bound that is necessary for injectivity. Finally, we show prove the continuity of the inverse of our constructed LP embedding layer with the techniques introduced in Sec. 4.3.

C.1 Upper Bound for Injectivity

To prove the upper bound, we construct an LP embedding layer with $L \le N^5 D^2$ output neurons such that its induced summation is injective. The main ingredient of our construction is anchor defined in Definition 4.1. Two key properties of anchors are stated in Lemma 4.2 and 4.3 (restated as follows) and proved below:

Lemma C.1. *Consider the data matrix $\boldsymbol{X} \in \mathbb{R}^{N\times D}$ and $\boldsymbol{a} \in \mathbb{R}^N$ an anchor of $\boldsymbol{X}$. Then if there exists $\boldsymbol{P} \in \Pi(N)$ such that $\boldsymbol{P}\boldsymbol{a} = \boldsymbol{a}$ then $\boldsymbol{P}\boldsymbol{x}_i = \boldsymbol{x}_i$ for every $i \in [D]$.*

Proof of Lemma 4.2. Prove by contradiction. Suppose $\boldsymbol{P}\boldsymbol{x}_i \neq \boldsymbol{x}_i$ for some $i \in [D]$, then there exist some $p, q \in [N]$ such that $\boldsymbol{x}_i^{(p)} \neq \boldsymbol{x}_i^{(q)}$ while $a_p = a_q$. However, this contradicts the definition of an anchor (cf. Definition 4.1). $\square$

Lemma C.2 (Union Alignment based on Anchor Alignment)**.** *Consider the data matrix $\boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N\times D}$, $\boldsymbol{a} \in \mathbb{R}^N$ is an anchor of $\boldsymbol{X}$ and $\boldsymbol{a'} \in \mathbb{R}^N$ is an arbitrary vector. If $[\boldsymbol{a} \quad \boldsymbol{x}_i] \sim [\boldsymbol{a'} \quad \boldsymbol{x'}_i]$ for every $i \in [D]$, then $\boldsymbol{X} \sim \boldsymbol{X'}$.*

Proof of Lemma 4.3. According to definition of equivalence, there exists $\boldsymbol{Q}_i \in \Pi(N)$ for every $i \in [D]$ such that $[\boldsymbol{a} \quad \boldsymbol{x}_i] = [\boldsymbol{Q}_i\boldsymbol{a'} \quad \boldsymbol{Q}_i\boldsymbol{x'}_i]$. Moreover, since $[\boldsymbol{a} \quad \boldsymbol{x}_i] \sim [\boldsymbol{a'} \quad \boldsymbol{x'}_i]$, it must hold that $\boldsymbol{a} \sim \boldsymbol{a'}$, i.e., there exists $\boldsymbol{P} \in \Pi_N$ such that $\boldsymbol{P}\boldsymbol{a} = \boldsymbol{a'}$. Combined together, we have that $\boldsymbol{Q}_i\boldsymbol{P}\boldsymbol{a} = \boldsymbol{a}$.

Next, we choose $\boldsymbol{Q'}_i = \boldsymbol{P}^\top \boldsymbol{Q}_i^\top$ so $\boldsymbol{Q'}_i\boldsymbol{a} = \boldsymbol{Q'}_i\boldsymbol{Q}_i\boldsymbol{P}\boldsymbol{a} = \boldsymbol{a}$. Due to the property of anchors (Lemma 4.2), we have $\boldsymbol{Q'}_i\boldsymbol{x}_i = \boldsymbol{x}_i$. Notice that $\boldsymbol{x}_i = \boldsymbol{Q'}_i\boldsymbol{x}_i = \boldsymbol{P}^\top\boldsymbol{Q}_i^\top\boldsymbol{Q}_i\boldsymbol{x'}_i = \boldsymbol{P}\boldsymbol{x'}_i$. Therefore, we can conclude the proof as we have found a permutation matrix $\boldsymbol{P}$ that simultaneously aligns $\boldsymbol{x}_i$ and $\boldsymbol{x'}_i$ for every $i \in [D]$, i.e., $\boldsymbol{X} = [\boldsymbol{x}_1 \quad \cdots \quad \boldsymbol{x}_D] = [\boldsymbol{P}\boldsymbol{x}_1 \quad \cdots \quad \boldsymbol{P}\boldsymbol{x}_D] = \boldsymbol{P}\boldsymbol{X'}$. $\square$

Next, we need to examine how many weights are needed to construct an anchor via linear combining all the existing channels. We restate Lemma 4.4 with more specifications as well as a mathematical device to prove it as below:

Lemma C.3. *Consider D linearly independent weight vectors $\{\boldsymbol{w}_1, \cdots, \boldsymbol{w}_D \in \mathbb{R}^D\}$. Then for every $p, q \in [N]$ such that $\boldsymbol{x}^{(p)} \neq \boldsymbol{x}^{(q)}$, there exists $\boldsymbol{w}_j, j \in [D]$, such that $\boldsymbol{x}^{(p)\top}\boldsymbol{w}_j \neq \boldsymbol{x}^{(q)\top}\boldsymbol{w}_j$.*

Proof. This is the simple fact of full-rank linear systems. Prove by contradiction. Suppose for $\forall j \in [D]$ we have $\boldsymbol{x}^{(p)\top}\boldsymbol{w}_j = \boldsymbol{x}^{(q)\top}\boldsymbol{w}_j$. Then we form a linear system: $\boldsymbol{x}^{(p)\top}[\boldsymbol{w}_1 \quad \cdots \quad \boldsymbol{w}_D] = \boldsymbol{x}^{(q)\top}[\boldsymbol{w}_1 \quad \cdots \quad \boldsymbol{w}_D]$. Since $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_D$ are linearly independent, it yields $\boldsymbol{x}^{(p)} = \boldsymbol{x}^{(q)}$, which reaches the contradiction. $\square$

Lemma C.4 (Anchor Construction)**.** *Consider a set of weight vectors $\{\boldsymbol{w}_1, \cdots, \boldsymbol{w}_{K_1} \in \mathbb{R}^D\}$ with $K_1 = N(N-1)(D-1)/2 + 1$, of which every D-length subset, i.e., $\{\boldsymbol{w}_j : \forall j \in \mathcal{J}\}, \forall \mathcal{J} \subseteq [K_1], |\mathcal{J}| = D$, is linearly independent, then for every data matrix $\boldsymbol{X} \in \mathbb{R}^{N\times D}$, there exists $j^* \in [K_1]$, $\boldsymbol{X}\boldsymbol{w}_{j^*}$ is an anchor of $\boldsymbol{X}$.*

Proof. Define a set of pairs which an anchor needs to distinguish: $\mathcal{I} = \{(p,q) : \boldsymbol{x}^{(p)} \neq \boldsymbol{x}^{(q)}\} \subseteq [N]^2$ Consider a D-length subset $\mathcal{J} \subseteq [K]$ with $|\mathcal{J}| = D$. Since $\{\boldsymbol{w}_j : \forall j \in \mathcal{J}\}$ are linear independent, we assert by Lemma C.3 that for every pair $(p,q) \in \mathcal{I}$, there exists $j \in \mathcal{J}$, $\boldsymbol{x}^{(p)\top}\boldsymbol{w}_j \neq \boldsymbol{x}^{(q)\top}\boldsymbol{w}_j$. It is equivalent to claim: for every pair $(p,q) \in \mathcal{I}$, at most $D-1$ many $\boldsymbol{w}_j, j \in [K_1]$ satisfy $\boldsymbol{x}^{(p)\top}\boldsymbol{w}_j = \boldsymbol{x}^{(q)\top}\boldsymbol{w}_j$. Based on pigeon-hold principle, as long as $K_1 \ge N(N-1)(D-1)/2 + 1 = (D-1)\binom{N}{2} + 1 \ge (D-1)|\mathcal{I}| + 1$, there must exist $j^* \in [K_1]$ such that $\boldsymbol{x}^{(p)\top}\boldsymbol{w}_{j^*} \neq \boldsymbol{x}^{(q)\top}\boldsymbol{w}_{j^*}$ for $\forall (p,q) \in \mathcal{I}$. By Definition 4.1, $\boldsymbol{X}\boldsymbol{w}_{j^*}$ generates an anchor. $\square$

Proposition C.5. *The linear independence condition in Lemma C.4 can be satisfied with probability one by drawing i.i.d. Gaussian random vectors $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_{K_1} \overset{i.i.d.}{\sim} \mathcal{N}(\boldsymbol{0}, \boldsymbol{I})$.*

Proof. We first note that generating a $D \times K_1$ Gaussian random matrix ($D \le K_1$) is equivalent to drawing a matrix with respect to a probability measure defined over $\mathcal{M} = \{\boldsymbol{X} \in \mathbb{R}^{D\times K} : \operatorname{rank}(\boldsymbol{X}) \le D\}$. Since rank-$D$ matrices are dense in $\mathcal{M}$ [53, 54], we can conclude that for $\forall \mathcal{J} \subseteq [K_1]$, $|\mathcal{J}| = D$, $\mathbb{P}(\{\boldsymbol{w}_j : j \in \mathcal{J}\}$ are linearly independent$) = 1$. By union bound, $\mathbb{P}(\{\boldsymbol{w}_j : j \in \mathcal{J}\}$ for all $\mathcal{J} \in [K]$, $|\mathcal{J}| = D$ are linearly independent$) = 1$. $\square$

We also restate Lemma 4.5 in the following lemma to demonstrate the weight construction for anchor matching:

Lemma C.6 (Anchor Matching)**.** *Consider a group of coefficients* $\Gamma = \{\gamma_1, \cdots, \gamma_{K_2} \in \mathbb{R}\}$ *with* $\gamma_i \neq 0, \forall i \in [K_2]$, $\gamma_i \neq \gamma_j, \forall i,j \in [K_2]$, *and* $K_2 = N(N-1)+1$ *such that for all 4-tuples* $(\gamma_i, \gamma_j, \gamma_k, \gamma_l) \subset \Gamma$, *if* $\gamma_i \neq \gamma_j, \gamma_i \neq \gamma_k$ *then* $\gamma_i/\gamma_j \neq \gamma_k/\gamma_l$. *It must hold that: Given any* $\boldsymbol{x}, \boldsymbol{x'}, \boldsymbol{y}, \boldsymbol{y'} \in \mathbb{R}^N$ *such that* $\boldsymbol{x} \sim \boldsymbol{x'}$ *and* $\boldsymbol{y} \sim \boldsymbol{y'}$, *if* $(\boldsymbol{x} - \gamma_k \boldsymbol{y}) \sim (\boldsymbol{x'} - \gamma_k \boldsymbol{y'})$ *for every* $k \in [K_2]$, *then* $[\boldsymbol{x} \quad \boldsymbol{y}] \sim [\boldsymbol{x'} \quad \boldsymbol{y'}]$.

Proof. We note that $\boldsymbol{x} \sim \boldsymbol{x'}$ and $\boldsymbol{y} \sim \boldsymbol{y'}$ imply that there exist $\boldsymbol{P}_x, \boldsymbol{P}_y \in \Pi(N)$ such that $\boldsymbol{P}_x \boldsymbol{x} = \boldsymbol{x'}$ and $\boldsymbol{P}_y \boldsymbol{y} = \boldsymbol{y'}$. Also $(\boldsymbol{x} - \gamma_k \boldsymbol{y}) \sim (\boldsymbol{x'} - \gamma_k \boldsymbol{y'}), \forall k \in [K_2]$ implies there exists $\boldsymbol{Q}_k \in \Pi(N), \forall k \in [K_2]$ such that $\boldsymbol{Q}_k(\boldsymbol{x} - \gamma_k \boldsymbol{y}) = \boldsymbol{x'} - \gamma_k \boldsymbol{y'}$. Substituting the former to the latter, we can obtain:

$$\left(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x\right) \boldsymbol{x} = \gamma_k \left(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x\right) \boldsymbol{y}, \quad \forall k \in [K_2], \tag{15}$$

where for each $k \in [K_2]$, Eq. (15) corresponds to N equalities as follows. Here, we let $(\boldsymbol{Z})_i$ denote the ith column of the matrix $\boldsymbol{Z}$.

$$\begin{gathered}(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_1^\top \boldsymbol{x} = \gamma_k (\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_1^\top \boldsymbol{y} \\ \vdots \\ (\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_N^\top \boldsymbol{x} = \gamma_k (\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_N^\top \boldsymbol{y}\end{gathered} \tag{16}$$

We compare and entries in $\boldsymbol{x} = [\cdots x_p \cdots]^\top$ and for each entry index $p \in [N]$, we define a set of non-zero pairwise differences between x_p and other entries in $\boldsymbol{x}$: $\mathcal{D}_x^{(p)} = \{x_p - x_q : q \in [N], x_p \neq x_q\}$. Similarly, for $\boldsymbol{y}$, we define $\mathcal{D}_y^{(p)} = \{y_p - y_q : q \in [N], y_p \neq y_q\}$. We note that for every $n \in [N]$, either $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x} = 0$ or $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x} \in \mathcal{D}_x^{(p)}$ for some $p \in [N]$ as $(\boldsymbol{Q}_k^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x}$ is some x_q. Then it is sufficient to show there must exist $k \in [K_2]$ such that none of equations in Eq. (16) can be induced by some elements in $\mathcal{D}_x^{(p)}$, i.e.,

$$\exists k^* \in [K_2], \forall p, n \in [N] \text{ such that } (\boldsymbol{I} - \boldsymbol{Q}_{k^*}^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x} \notin \mathcal{D}_x^{(p)}. \tag{17}$$

This is because Eq. (17) implies:

$$\begin{aligned} & (\boldsymbol{I} - \boldsymbol{Q}_{k^*} \boldsymbol{P}_x)^\top \boldsymbol{x} = \boldsymbol{0} \quad \Rightarrow \quad \boldsymbol{x} = \boldsymbol{Q}_{k^*}^\top \boldsymbol{P}_x \boldsymbol{x} = \boldsymbol{Q}_{k^*}^\top \boldsymbol{x'}, \\ \text{(Since } \gamma_k \neq 0, \forall k \in [K_2]) \quad & (\boldsymbol{I} - \boldsymbol{Q}_{k^*} \boldsymbol{P}_y)^\top \boldsymbol{y} = \boldsymbol{0} \quad \Rightarrow \quad \boldsymbol{y} = \boldsymbol{Q}_{k^*}^\top \boldsymbol{P}_y \boldsymbol{y} = \boldsymbol{Q}_{k^*}^\top \boldsymbol{y'}, \end{aligned}$$

which is $[\boldsymbol{x} \quad \boldsymbol{y}] = \boldsymbol{Q}_{k^*}^\top [\boldsymbol{x'} \quad \boldsymbol{y'}]$.

To show Eq. (17), we construct N bipartite graphs $\mathcal{G}^{(p)} = (\mathcal{D}_x^{(p)}, \mathcal{D}_y^{(p)}, \mathcal{E}^{(p)})$ for $p \in [N]$ where each $\alpha \in \mathcal{D}_x^{(p)}$ or each $\beta \in \mathcal{D}_y^{(p)}$ is viewed as a node and $(\alpha, \beta) \in \mathcal{E}^{(p)}$ gives an edge if $\alpha = \gamma_k \beta$ for some $k \in [K_2]$. Then we prove the existence of k^* via seeing a contradiction that does counting the number of connected pairs (α, β) from two perspectives.

Perspective of $\mathcal{D}_x^{(p)}$. We argue that for $\forall p \in [N]$ and arbitrary $\alpha_1, \alpha_2 \in \mathcal{D}_x^{(p)}$, $\alpha_1 \neq \alpha_2$, there exists at most one $\beta \in \mathcal{D}_y^{(p)}$ such that $(\alpha_1, \beta) \in \mathcal{E}^{(p)}$ and $(\alpha_2, \beta) \in \mathcal{E}^{(p)}$. Otherwise, suppose there exists $\beta' \in \mathcal{D}_y^{(p)}$, $\beta' \neq \beta$ such that $(\alpha_1, \beta') \in \mathcal{E}^{(p)}$ and $(\alpha_2, \beta') \in \mathcal{E}^{(p)}$. Then we have $\alpha_1 = \gamma_i \beta$, $\alpha_2 = \gamma_j \beta$, $\alpha_1 = \gamma_k \beta'$, and $\alpha_2 = \gamma_l \beta'$ for some $\gamma_i, \gamma_j, \gamma_k, \gamma_l \in \Gamma$, which is $\gamma_i/\gamma_k = \gamma_k/\gamma_l$. As $\alpha_1 \neq \alpha_2$, it is obvious that $\gamma_i \neq \gamma_j$. Similarly, we have $\gamma_i \neq \gamma_k$. Altogether, it contradicts our assumption on Γ. Therefore, $|\mathcal{E}^{(p)}| \leq 2\max\{|\mathcal{D}_x^{(p)}|, |\mathcal{D}_y^{(p)}|\} \leq 2(N-1)$. And the total edge number of all bipartite graphs should be less than $2N(N-1)$.

Perspective of Γ. We note that if for some $k \in [K_2]$ that makes $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x} \in \mathcal{D}_x^{(p)}$ for some $p, n \in [N]$, i.e., $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x} = \gamma_k (\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_y)_n^\top \boldsymbol{y} \neq 0$, this γ_k contributes at least two edges in the entire bipartite graph, i.e., there being another $n' \in [N]$, $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_{n'}^\top \boldsymbol{x} = \gamma_k (\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_y)_{n'}^\top \boldsymbol{y} \neq 0$. Otherwise, there exists a unique $n^* \in [N]$ such that $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_{n^*}^\top \boldsymbol{x} \in \mathcal{D}_x^{(p)} (\neq 0)$ and $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x} = 0$ for all $n \neq n^*$. This cannot be true because $\boldsymbol{1}^\top (\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)\boldsymbol{x} = 0$. By which, if $\forall k \in [K_2]$, $\exists p, n \in [N]$ such that $(\boldsymbol{I} - \boldsymbol{Q}_k^\top \boldsymbol{P}_x)_n^\top \boldsymbol{x} \in \mathcal{D}_x^{(p)}$ (i.e., Eq. (17) is always false), then the total number of edges is at least $2K_2 = 2N(N-1) + 2$.

Hereby, we conclude the proof by the contradiction, in which the minimal count of edges $2K_2$ by **Perspective of** Γ already surpasses the maximal number $2N(N-1)$ by **Perspective of** $\mathcal{D}_x^{(p)}$. $\square$

Remark C.7. A handy choice of Γ in Lemma C.6 are prime numbers, which are provably positive, infinitely many, and not divisible by each other.

We wrap off this section by formally stating and proving the injectivity statement of the LP embedding layer.

Theorem C.8. *Suppose $\phi : \mathbb{R}^D \to \mathbb{R}^L$ admits the form of Eq.* (2) *where $L = KN \le N^5 D^2$, $K = D + K_1 + DK_1K_2$ and $\boldsymbol{W} = \begin{bmatrix} \boldsymbol{w}_1^{(1)} & \cdots & \boldsymbol{w}_D^{(1)} & \boldsymbol{w}_1^{(2)} & \cdots & \boldsymbol{w}_{K_1}^{(2)} & \cdots & \boldsymbol{w}_{i,j,k}^{(3)} & \cdots \end{bmatrix}$ is constructed as follows:*

1. *Let the first group of weights $\boldsymbol{w}_1^{(1)} = \boldsymbol{e}_1, \cdots, \boldsymbol{w}_D^{(1)} = \boldsymbol{e}_D$ to buffer the original features, where $\boldsymbol{e}_i$ is the i-th canonical basis.*
2. *Choose the second group of linear weights, $\boldsymbol{w}_1^{(2)}, \cdots, \boldsymbol{w}_{K_1}^{(2)}$ for K_1 as large as $N(N-1)(D-1)/2 + 1$, such that the conditions in Lemma C.4 are satisfied.*
3. *Design the third group of weights $\boldsymbol{w}_{i,j,k}^{(3)}$ for $i \in [D], j \in [K_1], k \in [K_2]$ where $\boldsymbol{w}_{i,j,k}^{(3)} = \boldsymbol{e}_i - \gamma_k \boldsymbol{w}_j^{(2)}$, $K_2 = N(N-1) + 1$, and $\gamma_k, k \in [K_2]$ are chosen such that conditions in Lemma C.6 are satisfied.*

Then $\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})$ is injective (cf. Definition A.5).

Proof. Suppose $\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)}) = \sum_{i=1}^N \phi(\boldsymbol{x'}^{(i)})$ for some $\boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N \times D}$. Due to the property of power mapping (cf. Lemma 2.8):

$$\begin{aligned}\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)}) = \sum_{i=1}^N \phi(\boldsymbol{x'}^{(i)}) \Rightarrow & \boldsymbol{X}\boldsymbol{w}_i^{(1)} \sim \boldsymbol{X'}\boldsymbol{w}_i^{(1)}, \forall i \in [D], \boldsymbol{X}\boldsymbol{w}_i^{(2)} \sim \boldsymbol{X'}\boldsymbol{w}_i^{(2)}, \forall i \in [K_1], \\ & \boldsymbol{X}\boldsymbol{w}_{i,j,k}^{(3)} \sim \boldsymbol{X'}\boldsymbol{w}_{i,j,k}^{(3)}, \forall i \in [D], j \in [K_1], k \in [K_2].\end{aligned} \tag{18}$$

By Lemma C.4, it is guaranteed that there exists $j^* \in [K_1]$ such that $\boldsymbol{X}\boldsymbol{w}_{j^*}^{(2)}$ is an anchor, and according to Eq. (18), we have $\boldsymbol{X}\boldsymbol{w}_{j^*}^{(2)} \sim \boldsymbol{X'}\boldsymbol{w}_{j^*}^{(2)}$. By Lemma C.6, we induce:

$$\begin{aligned}& \boldsymbol{X}\boldsymbol{w}_i^{(1)} \sim \boldsymbol{X'}\boldsymbol{w}_i^{(1)}, \forall i \in [D], \boldsymbol{X}\boldsymbol{w}_{j^*}^{(2)} \sim \boldsymbol{X'}\boldsymbol{w}_{j^*}^{(2)}, \boldsymbol{X}\boldsymbol{w}_{i,j,k}^{(3)} \sim \boldsymbol{X'}\boldsymbol{w}_{i,j,k}^{(3)}, \forall i \in [D], j \in [K_1], k \in [K_2] \\ & \Rightarrow \begin{bmatrix} \boldsymbol{X}\boldsymbol{w}_{j^*}^{(2)} & \boldsymbol{x}_i \end{bmatrix} \sim \begin{bmatrix} \boldsymbol{X'}\boldsymbol{w}_{j^*}^{(2)} & \boldsymbol{x'}_i \end{bmatrix}, \forall i \in [D].\end{aligned} \tag{19}$$

Since $\boldsymbol{X}\boldsymbol{w}_{j^*}^{(2)}$ is an anchor, by union alignment (Lemma 4.3), we have:

$$\begin{bmatrix} \boldsymbol{X}\boldsymbol{w}_{j^*}^{(2)} & \boldsymbol{x}_i \end{bmatrix} \sim \begin{bmatrix} \boldsymbol{X'}\boldsymbol{w}_{j^*}^{(2)} & \boldsymbol{x'}_i \end{bmatrix}, \forall i \in [D] \Rightarrow \boldsymbol{X} \sim \boldsymbol{X'}. \tag{20}$$

Here $K = D + K_1 + DK_1K_2 \le N^4 D^2$, thus $L = KN \le N^5 D^2$, which concludes the proof. $\square$

C.2 Continuity

Next, we show that under the construction of Theorem C.8, the inverse of $\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})$ is continuous. The main idea is to check the three conditions provided in Lemma 4.10:

Corollary C.9. *Consider channel-wise high-dimensional sum-of-power* $\widehat{\Psi_N}(\boldsymbol{X}) : (\mathbb{R}^{N\times K}/\sim, d_F) \to (\mathbb{R}^{NK}, d_\infty)$ *defined as below:*

$$\widehat{\Psi_N}(\boldsymbol{X}) = \begin{bmatrix}\Psi_N(\boldsymbol{x}_1)^\top & \cdots & \Psi_N(\boldsymbol{x}_K)^\top\end{bmatrix}^\top \in (\mathbb{R}^{NK}, d_\infty), \tag{21}$$

and an associated mapping $\widehat{\Phi_N} : (\mathbb{R}^{NK}, d_\infty) \to (\mathbb{R}^N/\sim, d_F)^K$*:*

$$\widehat{\Phi_N}(\boldsymbol{Z}) = \begin{bmatrix}\Psi_N^{-1}(\boldsymbol{z}_1) & \cdots & \Psi_N^{-1}(\boldsymbol{z}_K)\end{bmatrix}, \tag{22}$$

where $\boldsymbol{Z} = \begin{bmatrix}\boldsymbol{z}_1^\top & \cdots & \boldsymbol{z}_K^\top\end{bmatrix}^\top$*,* $\boldsymbol{z}_i \in \mathbb{R}^N, \forall i \in [K]$*. We denote the induced product metric over* $(\mathbb{R}^N/\sim, d_F)^K$ *as* $d_F^K : (\mathbb{R}^N/\sim)^K \times (\mathbb{R}^N/\sim)^K \to \mathbb{R}_{\geq 0}$*:*

$$d_F^K(\boldsymbol{Z}, \boldsymbol{Z'}) = \max_{i\in[K]} d_F(\boldsymbol{z}_i, \boldsymbol{z'}_i). \tag{23}$$

Then the mapping $\widehat{\Phi_N}$ *maps any bounded set in* $(\mathbb{R}^{NK}, d_\infty)$ *to a bounded set in* $(\mathbb{R}^N/\sim, d_F)^K$*.*

Proof. Proved by noting that if $d_\infty(\boldsymbol{z}_i, \boldsymbol{z'}_i) \leq C_1$ for some $\boldsymbol{z}_i, \boldsymbol{z'}_i \in (\mathbb{R}^N, d_\infty), \forall i \in [K]$ and a constant $C_1 \geq 0$, then $d_F(\Psi_N^{-1}(\boldsymbol{z}_i), \Psi_N^{-1}(\boldsymbol{z'}_i)) \leq C_2, \forall i \in [K]$ for some constant $C_2 \geq 0$ by Lemma B.5 and Remark B.6. Finally, we have:

$$d_F^K(\widehat{\Phi_N}(\boldsymbol{Z}), \widehat{\Phi_N}(\boldsymbol{Z'})) = \max_{i\in[K]} d_F(\Psi_N^{-1}(\boldsymbol{z}_i), \Psi_N^{-1}(\boldsymbol{z'}_i)) \leq C_2,$$

which is also bounded above. □

Now we are ready to present and prove the continuity of the LP embedding layer.

Theorem C.10. *Suppose* ϕ *admits the form of Eq.* (2) *and follows the construction in Theorem C.8, then the inverse of LP-induced sum-pooling* $\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})$ *is continuous.*

Proof. The proof is done by invoking Lemma 4.10. First of all, the inverse of $\Psi(\boldsymbol{X}) = \sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})$, denoted as $\Psi^{-1} : (\mathbb{R}^{NK}, d_\infty) \to (\mathbb{R}^{N\times D}/\sim, d_F)$, exists due to Theorem C.8. By Lemma A.4, any closed and bounded subset of $(\mathbb{R}^{N\times D}/\sim, d_F)$ is compact. Trivially, $\Psi(\boldsymbol{X})$ is continuous. Then it remains to show the condition **(c)** in Lemma 4.10. We decompose Ψ^{-1} into two mappings following the similar idea of proving its existence:

$$(\mathbb{R}^{NK}, d_\infty) \quad \xrightarrow{\widehat{\Phi_N}} \quad (\mathbb{R}^N/\sim, d_F)^K \quad \xrightarrow{\pi} \quad (\mathbb{R}^{N\times D}/\sim, d_F)\ ,$$

where $\widehat{\Phi_N}$ is defined in Eq. (22) and π exists due to Eqs. (19) and (20) in Theorem C.8. Also according to our construction in Theorem C.8, the first identity block induces that: for any $\boldsymbol{Z} \in (\mathbb{R}^N/\sim, d_F)^K$, $\boldsymbol{z}_i \sim \pi(\boldsymbol{Z})_i$ for every $i \in [D]$. Therefore, $\forall \boldsymbol{Z}, \boldsymbol{Z'} \in (\mathbb{R}^N/\sim, d_F)^K$ such that $d_F^K(\boldsymbol{Z}, \boldsymbol{Z'}) \leq C$ for some constant $C > 0$, we have:

$$d_F(\pi(\boldsymbol{Z}), \pi(\boldsymbol{Z'})) \leq \max_{i\in[D]} d_F(\boldsymbol{z}_i, \boldsymbol{z'}_i) \leq d_F^K(\boldsymbol{Z}, \boldsymbol{Z'}) \leq C, \tag{24}$$

which implies π maps every bounded set in $(\mathbb{R}^N/\sim, d_F)^K$ to a bounded set in $(\mathbb{R}^{N\times D}/\sim, d_F)$. Now we conclude the proof by the following chain of argument:

$$\mathcal{Z} \subseteq (\mathbb{R}^{NK}, d_\infty) \text{ is bounded} \quad \xRightarrow{\text{Corollary C.9}} \quad \widehat{\Phi_N}(\mathcal{Z}) \text{ is bounded} \quad \xRightarrow{\text{Eq. (24)}} \quad \pi \circ \widehat{\Phi_N}(\mathcal{Z}) \text{ is bounded.}$$

□

C.3 Lower Bound for Injectivity

In this section, we prove Theorem 4.6 which shows that $K \geq D + 1$ is necessary for injectivity of LP-induced sum-pooling when $D \geq 2$. Our argument mainly generalizes Lemma 2 of [55] to our equivalence class. To proceed our argument, we define the linear subspace $\boldsymbol{V}$ by vectorizing $\begin{bmatrix}\boldsymbol{X}\boldsymbol{w}_1 & \cdots & \boldsymbol{X}\boldsymbol{w}_K\end{bmatrix}$ as below:

$$\mathcal{V} := \left\{ \begin{bmatrix}\boldsymbol{X}\boldsymbol{w}_1 \\ \vdots \\ \boldsymbol{X}\boldsymbol{w}_K\end{bmatrix} : \boldsymbol{X} \in \mathbb{R}^{N\times D} \right\} = \mathcal{R}\left(\begin{bmatrix}(\boldsymbol{w}_1 \otimes \boldsymbol{I}_N) \\ \vdots \\ (\boldsymbol{w}_K \otimes \boldsymbol{I}_N)\end{bmatrix} \right), \tag{25}$$

where $\mathcal{R}(\boldsymbol{Z})$ denotes the column space of $\boldsymbol{Z}$ and $\otimes$ is the Kronecker product. $\mathcal{V}$ is a linear subspace of $\mathbb{R}^{NK}$ with dimension at most $\mathbb{R}^{ND}$, characterized by $\boldsymbol{w}_1, \cdots, \boldsymbol{w}_K \in \mathbb{R}^D$. For the sake of notation simplicity, we denote $\Pi(N)^{\otimes K} = \{\operatorname{diag}(\boldsymbol{Q}_1, \cdots, \boldsymbol{Q}_K) : \forall \boldsymbol{Q}_1, \cdots, \boldsymbol{Q}_K \in \Pi(N)\}$, and $\boldsymbol{I}_K \otimes \Pi(N) = \{\boldsymbol{I}_K \otimes \boldsymbol{Q} : \forall \boldsymbol{Q} \in \Pi(N)\}$. Next, we define the notion of unique recoverability [55] as below:

Definition C.11 (Unique Recoverability)**.** The subspace $\mathcal{V}$ is called uniquely recoverable under $\boldsymbol{Q} \in \Pi(N)^{\otimes K}$ if whenever $\boldsymbol{x}, \boldsymbol{x'} \in \mathcal{V}$ satisfy $\boldsymbol{Q}\boldsymbol{x} = \boldsymbol{x'}$, there exists $\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(N)$, $\boldsymbol{P}\boldsymbol{x} = \boldsymbol{x'}$.

Subsequently, we derive a necessary condition for the unqiue recoverability:

Lemma C.12. *A linear subspace $\mathcal{V} \subseteq \mathbb{R}^{NK}$ is uniquely recoverable under $\boldsymbol{Q} \in \Pi(N)^{\otimes K}$ only if there exists $\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(N)$, $\boldsymbol{Q}(\mathcal{V}) \cap \mathcal{V} \subset \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda=1}$, where $\mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda}$ denotes the eigenspace corresponding to the eigenvalue λ.*

Proof. We first show that $\boldsymbol{Q}(\mathcal{V}) \cap \mathcal{V} \subseteq \bigcup_{\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(K)} \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda=1}$. Since the LHS is a subspace, while the RHS is a union of subspaces, there exists $\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(K)$ such that $\boldsymbol{Q}(\mathcal{V}) \cap \mathcal{V} \subseteq \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda=1}$. Then it remains to show $\boldsymbol{Q}(\mathcal{V}) \cap \mathcal{V} \subseteq \bigcup_{\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(K)} \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda=1}$.

Proved by contradiction. Suppose there exists $\boldsymbol{x} \in \boldsymbol{Q}(\mathcal{V}) \cap \mathcal{V}$ but $\boldsymbol{x} \notin \bigcup_{\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(K)} \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda=1}$. Or equivalently, there exists $\boldsymbol{x'} \in \mathcal{V}$ and $\boldsymbol{x} = \boldsymbol{Q}\boldsymbol{x'}$, while for $\forall \boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(N)$, $\boldsymbol{x} \neq \boldsymbol{Q}\boldsymbol{P}^\top \boldsymbol{x}$. This implies $\boldsymbol{Q}^\top \boldsymbol{x} = \boldsymbol{x'} \neq \boldsymbol{P}\boldsymbol{x}$ for $\forall \boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(N)$. However, this contradicts the fact that $\mathcal{V} \subseteq \mathbb{R}^{NK}$ is uniquely recoverable (cf. Definition C.11). ☐

We also introduce a useful Lemma C.13 that gets rid of the discussion on $\boldsymbol{Q}$ in the inclusion:

Lemma C.13. *Suppose $\mathcal{V} \subseteq \mathbb{R}^N$ is a linear subspace, and $\boldsymbol{A}$ is a linear mapping. $\boldsymbol{A}(\mathcal{V}) \cap \mathcal{V} \cap \mathcal{E}_{\boldsymbol{A}, \lambda} = \boldsymbol{0}$ if and only if $\mathcal{V} \cap \mathcal{E}_{\boldsymbol{A}, \lambda} = \boldsymbol{0}$.*

Proof. The sufficiency is straightforward. The necessity is shown by contradiction: Suppose $\mathcal{V} \cap \mathcal{E}_{\boldsymbol{A}, \lambda} \neq \boldsymbol{0}$, then there exists $\boldsymbol{x} \in \mathcal{V} \cap \mathcal{E}_{\boldsymbol{A}, \lambda}$ such that $\boldsymbol{x} \neq \boldsymbol{0}$. Then $\boldsymbol{A}\boldsymbol{x} = \lambda \boldsymbol{x}$ implies $\boldsymbol{x} \in \boldsymbol{A}(\mathcal{V})$. Hence, $\boldsymbol{x} \in \boldsymbol{A}(\mathcal{V}) \cap \mathcal{V} \cap \mathcal{E}_{\boldsymbol{A}, \lambda}$ which reaches the contradiction. ☐

Now we are ready to present the proof of Theorem 4.6:

Proof of Theorem 4.6. Proved by contrapositive. First notice that, $\forall \boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N \times D}, \boldsymbol{X}\boldsymbol{w}_i \sim \boldsymbol{X'}\boldsymbol{w}_i, \forall i \in [K] \Rightarrow \boldsymbol{X} \sim \boldsymbol{X'}$ holds if and only if $\dim \mathcal{V} = ND$ and $\mathcal{V}$ is uniquely recoverable under all possible $\boldsymbol{Q} \in \Pi(N)^{\otimes K}$. By Lemma C.12, for every $\boldsymbol{Q} \in \Pi(N)^{\otimes K}$, there exists $\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(N)$ such that $\boldsymbol{Q}(\mathcal{V}) \cap \mathcal{V} \subset \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda=1}$. This is $\boldsymbol{Q}(\mathcal{V}) \cap \mathcal{V} \cap \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda} = 0$ for all $\lambda \neq 1$. By Lemma C.13, we have $\mathcal{V} \cap \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda} = 0$ for all $\lambda \neq 1$. Then proof is concluded by discussing the dimension of ambient space $\mathbb{R}^{NK}$ such that an ND-dimensional subspace $\mathcal{V}$ can reside. To ensure $\mathcal{V} \cap \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda} = 0$ for all $\lambda \neq 1$, it is necessary that $\dim \mathcal{V} \leq \min_{\lambda \neq 1} \operatorname{codim} \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda}$ for every $\boldsymbol{Q} \in \Pi(N)^{\otimes K}$ and its associated $\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(N)$. Relaxing the dependence between $\boldsymbol{Q}$ and $\boldsymbol{P}$, we derive the inequality:

$$ND = \dim \mathcal{V} \leq \min_{\boldsymbol{Q} \in \Pi(N)^{\otimes K}} \max_{\boldsymbol{P} \in \boldsymbol{I}_K \otimes \Pi(N)} \min_{\lambda \neq 1} \operatorname{codim} \mathcal{E}_{\boldsymbol{Q}\boldsymbol{P}^\top, \lambda} \leq NK - 1, \tag{26}$$

where the last inequality considers the scenario where every non-one eigenspace is one-dimensional, which is achievable when $K \geq 2$. Hence, we can bound $K \geq D + 1/K$, i.e., $K \geq D + 1$. ☐

D Proofs for LLE Embedding Layer

In this section, we present the complete proof for the LLE embedding layer (Eq. (3)). Similar to the LP embedding layer, we construct an LLE whose induced sum-pooling is injective following arguments in Sec. 4.2 and has continuous inverse with the techniques introduced in Sec. 4.3.

D.1 Upper Bound for Injectivity

To prove the upper bound, we construct an LLE embedding layer with $L \leq 2N^2D^2$ output neurons such that its induced sum-pooling is injective. The main proof technique is to bind every pair of channel with complex numbers and invoke the injectiviy of sum-of-power mapping over the complex domain.

With Lemma B.2, we can prove Lemma 4.8 as below:

Proof of Lemma 4.8. If for any pair of vectors $\boldsymbol{x}_1, \boldsymbol{x}_2 \in \mathbb{R}^N, \boldsymbol{y}_1, \boldsymbol{y}_2 \in \mathbb{R}^N$ such that $\sum_{i\in[N]} \boldsymbol{x}_{1,i}^{l-k}\boldsymbol{x}_{2,i}^{k} = \sum_{i\in[N]} \boldsymbol{y}_{1,i}^{l-k}\boldsymbol{y}_{2,i}^{k}$ for every $l, k \in [N]$, $0 \leq k \leq l$, then for $\forall l \in [N]$,

$$\begin{aligned}\sum_{i=1}^{N}(\boldsymbol{x}_{1,i} + \boldsymbol{x}_{2,i}\sqrt{-1})^l &= \sum_{i=1}^{N}\sum_{k=0}^{l}(\sqrt{-1})^k \boldsymbol{x}_{1,i}^{l-k}\boldsymbol{x}_{2,i}^{k} = \sum_{k=0}^{l}(\sqrt{-1})^k \left(\sum_{i=1}^{N}\boldsymbol{x}_{1,i}^{l-k}\boldsymbol{x}_{2,i}^{k}\right) \\ &= \sum_{k=0}^{l}(\sqrt{-1})^k \left(\sum_{i=1}^{N}\boldsymbol{y}_{1,i}^{l-k}\boldsymbol{y}_{2,i}^{k}\right) \\ &= \sum_{i=1}^{N}(\boldsymbol{y}_{1,i} + \boldsymbol{y}_{2,i}\sqrt{-1})^l \\ &= \psi_N(\boldsymbol{y}_1 + \boldsymbol{y}_2\sqrt{-1})\end{aligned} \tag{27}$$

Then by Lemma B.2, we have $(\boldsymbol{x}_1 + \boldsymbol{x}_2\sqrt{-1}) \sim (\boldsymbol{y}_1 + \boldsymbol{y}_2\sqrt{-1})$, which is essentially $[\boldsymbol{x}_1 \quad \boldsymbol{x}_2] \sim [\boldsymbol{y}_1 \quad \boldsymbol{y}_2]$. ☐

Now we are ready to prove the injectiviy of the LLE layer.

Theorem D.1. *Suppose $\phi : \mathbb{R}^D \to \mathbb{R}^L$ admits the form of Eq.* (3) *and $\boldsymbol{W} = [\cdots \quad \boldsymbol{w}_{i,j,p,q} \quad \cdots] \in \mathbb{R}^{D\times L}, i \in [D], j \in [D], p \in [N], q \in [p+1]$ is constructed as follows:*

$$\boldsymbol{w}_{i,j,p,q} = (q-1)\boldsymbol{e}_i + (p-q+1)\boldsymbol{e}_j, \tag{28}$$

where $\boldsymbol{e}_i$ is the i-th canonical basis. Then $\sum_{i=1}^{N}\phi(\boldsymbol{x}^{(i)})$ is injective (cf. Definition A.5).

Proof. First of all, notice that $L = D^2\sum_{p=1}^{N}(p+1) = D^2(N+3)N/2 \leq 2N^2D^2$. According to Eq. 8, we can rewrite for $\forall i \in [D], j \in [D], p \in [N], q \in [p+1]$

$$\phi(\boldsymbol{x})_{i,j,p,q} = \exp(\boldsymbol{w}_{i,j,p,q}^\top \log(\boldsymbol{x})) = \prod_{k=1}^{D} \boldsymbol{x}_k^{\boldsymbol{w}_{i,j,p,q,k}} = \boldsymbol{x}_i^{q-1}\boldsymbol{x}_j^{p-q+1}, \tag{29}$$

Then for $\boldsymbol{X}, \boldsymbol{X'} \in \mathbb{R}^{N\times D}$, $\sum_{i\in[N]}\phi(\boldsymbol{x}^{(i)} = \sum_{i\in[N]}\phi(\boldsymbol{x'}^{(i)})$ implies $\sum_{i\in[N]}\boldsymbol{x}_i^{q-1}\boldsymbol{x}_j^{p-q+1} = \sum_{i\in[N]}\boldsymbol{x'}_i^{q-1}\boldsymbol{x'}_j^{p-q+1}$ for $\forall i \in [D], j \in [D], p \in [N], q \in [p+1]$. By Lemma 4.8, we have $[\boldsymbol{x}_i \quad \boldsymbol{x}_j] \sim [\boldsymbol{x'}_i \quad \boldsymbol{x'}_j]$ for $\forall i \in [D], j \in [D]$. Finally, Lemma 4.9 directly yields $\boldsymbol{X} \sim \boldsymbol{X'}$. ☐

D.2 Continuity

The proof idea of continuity for LLE layer shares the same outline with the LP layer.

Corollary D.2. *Consider the mapping $\widetilde{\Phi_N} = \psi_N^{-1} \circ \tau : (\mathbb{R}^{N(N+3)/2}, d_\infty) \to (\mathbb{C}^N/\sim, d_F)$, where $\tau : (\mathbb{R}^{N(N+3)/2}, d_\infty) \to (\mathbb{C}^N, d_\infty)$ is a linear mapping that combines sum of monomials to polynomials following Eq.* (27)*. Construct the mapping $\widehat{\Phi_N} : (\mathbb{R}^L, d_\infty) \to (\mathbb{C}^N/\sim, d_F)^{D^2}$ based on $\widetilde{\Phi_N}$:*

$$\widehat{\Phi_N}(\boldsymbol{Z}) = \left[\widetilde{\Phi_N}(\boldsymbol{z}_1) \quad \cdots \quad \widetilde{\Phi_N}(\boldsymbol{z}_{D^2})\right], \tag{30}$$

where $\boldsymbol{Z} = \left[\boldsymbol{z}_1^\top \quad \cdots \quad \boldsymbol{z}_{D^2}^\top\right]^\top \in \mathbb{R}^L$, $\boldsymbol{z}_i \in \mathbb{R}^{(N(N+3)/2)}, \forall i \in [D^2]$. We denote the induced product metric over $(\mathbb{C}^N/\sim, d_F)^{D^2}$ as $d_F^{D^2} : (\mathbb{C}^N/\sim)^{D^2} \times (\mathbb{C}^N/\sim)^{D^2} \to \mathbb{R}_{\geq 0}$:

$$d_F^{D^2}(\boldsymbol{Z}, \boldsymbol{Z'}) = \max_{i\in[D^2]} d_F(\boldsymbol{z}_i, \boldsymbol{z'}_i). \tag{31}$$

Then the mapping $\widehat{\Phi_N}$ maps any bounded set in $(\mathbb{R}^L, d_\infty)$ to a bounded set in $(\mathbb{C}^N/\sim, d_F)^{D^2}$.

Proof. Since τ is a linear mapping, any $\boldsymbol{Z}, \boldsymbol{Z'} \in \mathbb{R}^L$ such that $d_\infty(\boldsymbol{z}_i, \boldsymbol{z'}_i) \le C_1, \forall i \in [D^2]$ for some constant $C_1 \ge 0$, then $d_F(\tau(\boldsymbol{z}_i), \tau(\boldsymbol{z'}_i)) \le C_2$ for some constant $C_2 \ge 0$. By Lemma B.3, ψ_N^{-1} maps any bounded set in $(\mathbb{C}^N, d_\infty)$ to a bounded set in $(\mathbb{C}^N / \sim, d_F)$. This is $d_F(\widetilde{\Phi_N}(\boldsymbol{z}_i), \widetilde{\Phi_N}(\boldsymbol{z'}_i)) \le C_2, \forall i \in [D^2]$ for some constant $C_2 \ge 0$. Finally, we have:

$$d_F^{D^2}(\widehat{\Phi_N}(\boldsymbol{Z}), \widehat{\Phi_N}(\boldsymbol{Z'}) = \max_{i \in [D^2]} d_F(\widetilde{\Phi_N}(\boldsymbol{z}_{i,j}), \widetilde{\Phi_N}(\boldsymbol{z'}_{i,j})) \le C_2,$$

which is also bounded above. □

Theorem D.3. *Suppose ϕ admits the form of Eq.* (3) *and follows the construction in Theorem D.1, then the inverse of LLE-induced sum-pooling $\sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})$ is continuous.*

Proof. It is sufficient to verify three conditions in Lemma 4.10. First of all, we denote the inverse of $\Psi(\boldsymbol{X}) = \sum_{i=1}^N \phi(\boldsymbol{x}^{(i)})$, denoted as $\Psi^{-1} : (\mathbb{R}^L, d_\infty) \to (\mathbb{R}^{N \times D} / \sim, d_F)$, which exists thanks to Theorem D.1. By Lemma A.4, any closed and bounded subset of $(\mathbb{R}^{N \times D} / \sim, d_F)$ is compact. Obviously, $\Psi(\boldsymbol{X})$ is continuous. Then it remains to show the condition **(c)** in Lemma 4.10. Similar to Theorem C.10, we decompose Ψ^{-1} into two mappings following the clue of proving its existence:

$$(\mathbb{R}^{NK}, d_\infty) \quad \xrightarrow{\widehat{\Phi_N}} \quad (\mathbb{C}^N / \sim, d_F)^{D^2} \quad \xrightarrow{\pi} \quad (\mathbb{R}^{N \times D} / \sim, d_F) \ ,$$

where $\widehat{\Phi_N}$ is defined in Eq. (30) and π exists due to Theorem D.1. Also according to our construction in Theorem D.1, for any $\boldsymbol{Z} \subset (\mathbb{C}^N / \sim, d_F)^{D^2}$ and $\forall i, j \in [D]$, there exists $k \in [D^2]$ such that $(\pi(\boldsymbol{z})_i + \pi(\boldsymbol{z})_j \sqrt{-1}) \sim \boldsymbol{z}_k$. Therefore, $\forall \boldsymbol{Z}, \boldsymbol{Z'} \in (\mathbb{C}^N / \sim, d_F)^{D^2}$ such that $d_F^{D^2}(\boldsymbol{Z}, \boldsymbol{Z'}) \le C$ for some constant $C > 0$, we have:

$$d_F(\pi(\boldsymbol{Z}), \pi(\boldsymbol{Z'})) \le \max_{i \in [D^2]} d_F(\boldsymbol{z}_i, \boldsymbol{z'}_i) \le d_F^{D^2}(\boldsymbol{Z}, \boldsymbol{Z'}) \le C, \tag{32}$$

which implies π maps every bounded set in $(\mathbb{C}^N / \sim, d_F)^K$ to a bounded set in $(\mathbb{R}^{N \times D} / \sim, d_F)$. Now we conclude the proof by the following chain of argument:

$$\mathcal{Z} \subseteq (\mathbb{R}^{NK}, d_\infty) \text{ is bounded} \quad \xRightarrow{\text{Corollary D.2}} \quad \widehat{\Phi_N}(\mathcal{Z}) \text{ is bounded} \quad \xRightarrow{\text{Eq. (32)}} \quad \pi \circ \widehat{\Phi_N}(\mathcal{Z}) \text{ is bounded.}$$

□

E Extension to Permutation Equivariance

In this section, we prove Theorem 5.1, the extension of Theorem 3.1 to equivariant functions, following the similar arguments with [7]:

Lemma E.1 ([7, 34]). *$f : \mathbb{R}^{N \times D} \to \mathbb{R}^N$ is a permutation-equivariant function if and only if there is a function $\rho : \mathbb{R}^{N \times D} \to \mathbb{R}$ that is permutation invariant to the last $N-1$ entries, such that $f(\boldsymbol{Z})_i = \rho(\boldsymbol{z}^{(i)}, \underbrace{\boldsymbol{z}^{(i+1)}, \cdots, \boldsymbol{z}^{(N)}, \cdots, \boldsymbol{z}^{(i-1)}}_{N-1})$ for any $i \in [N]$.*

Proof. (Sufficiency) Define $\pi : [N] \to [N]$ be an index mapping associated with the permutation matrix $\boldsymbol{P} \in \Pi(N)$ such that $\boldsymbol{P}\boldsymbol{Z} = \left[\boldsymbol{z}^{(\pi(1))}, \cdots, \boldsymbol{z}^{(\pi(N))}\right]^\top$. Then we have:

$$f\left(\boldsymbol{z}^{(\pi(1))}, \cdots, \boldsymbol{z}^{(\pi(N))}\right)_i = \rho\left(\boldsymbol{z}^{(\pi(i))}, \boldsymbol{z}^{(\pi(i+1))}, \cdots, \boldsymbol{z}^{(\pi(N))}, \cdots, \boldsymbol{z}^{(\pi(i-1))}\right).$$

Since $\rho(\cdot)$ is invariant to the last $N-1$ entries, it can shown that:

$$f(\boldsymbol{P}\boldsymbol{Z})_i = \rho\left(\boldsymbol{z}^{(\pi(i))}, \boldsymbol{z}^{(\pi(i+1))}, \cdots, \boldsymbol{z}^{(\pi(N))}, \cdots, \boldsymbol{z}^{(\pi(i-1))}\right) = f(\boldsymbol{Z})_{\pi(i)}.$$

(Necessity) Given a permutation-equivariant function $f : \mathbb{R}^{N \times D} \to \mathbb{R}^N$, we first expand it to the following form: $f(\boldsymbol{Z})_i = \rho_i(\boldsymbol{z}^{(1)}, \cdots, \boldsymbol{z}^{(N)})$. Permutation-equivariance means

$\rho_{\pi(i)}(\boldsymbol{z}^{(1)}, \cdots, \boldsymbol{z}^{(N)}) = \rho_i(\boldsymbol{z}^{\pi(1)}, \cdots, \boldsymbol{z}^{\pi(N)})$ for any permutation mapping π. Suppose given an index $i \in [N]$, consider any permutation $\pi : [N] \to [N]$ such that $\pi(i) = i$. Then, we have:

$$\rho_i \left(\boldsymbol{z}^{(1)}, \cdots, \boldsymbol{z}^{(i)}, \cdots, \boldsymbol{z}^{(N)}\right) = \rho_{\pi(i)} \left(\boldsymbol{z}^{(1)}, \cdots, \boldsymbol{z}^{(i)}, \cdots, \boldsymbol{z}^{(N)}\right) = \rho_i \left(\boldsymbol{z}^{(\pi(1))}, \cdots, \boldsymbol{z}_i, \cdots, \boldsymbol{z}^{(\pi(N))}\right),$$

which implies $\rho_i : \mathbb{R}^{N\times D} \to \mathbb{R}$ must be invariant to the $N-1$ elements other than the i-th element. Now, consider a permutation π where $\pi(1) = i$. Then we have:

$$\begin{aligned}\rho_i \left(\boldsymbol{z}^{(1)}, \boldsymbol{z}^{(2)}, \cdots, \boldsymbol{z}^{(N)}\right) &= \rho_{\pi(1)} \left(\boldsymbol{z}^{(1)}, \boldsymbol{z}^{(2)}, \cdots, \boldsymbol{z}^{(N)}\right) = \rho_1 \left(\boldsymbol{z}^{(\pi(1))}, \boldsymbol{z}^{(\pi(2))}, \cdots, \boldsymbol{z}^{(\pi(N))}\right) \\ &= \rho_1 \left(\boldsymbol{z}^{(i)}, \boldsymbol{z}^{(i+1)}, \cdots, \boldsymbol{z}^{(N)}, \cdots, \boldsymbol{z}^{(i-1)}\right),\end{aligned}$$

where the last equality is due to the invariance to $N-1$ elements, stated beforehand. This implies two results. First, for all i, $\rho_i(\boldsymbol{z}^{(1)}, \boldsymbol{z}^{(2)}, \cdots, \boldsymbol{z}^{(i)}, \cdots, \boldsymbol{z}^{(N)}), \forall i \in [N]$ should be written in terms of $\rho_1(\boldsymbol{z}^{(i)}, \boldsymbol{z}^{(i+1)}, \cdots, \boldsymbol{z}^{(N)}, \cdots, \boldsymbol{z}^{(i-1)})$. Moreover, ρ_1 is permutation invariant to its last $N-1$ entries. Therefore, we just need to set $\rho = \rho_1$ and broadcast it accordingly to all entries. We conclude the proof. $\square$

Proof of Theorem 5.1 [7]. Sufficiency can be shown by verifying the equivariance. We conclude the proof by showing the necessity with Lemma E.1. First we rewrite any permutation equivariant function $f(\boldsymbol{x}^{(1)}, \cdots, \boldsymbol{x}^{(N)}) : \mathbb{R}^{N\times D} \to \mathbb{R}^N$ as:

$$f\left(\boldsymbol{x}^{(1)}, \cdots, \boldsymbol{x}^{(N)}\right)_i = \tau\left(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(i+1)}, \cdots, \boldsymbol{x}^{(N)}, \cdots, \boldsymbol{x}^{(i-1)}\right), \tag{33}$$

where π is invariant to the lask $N-1$ elements, according to Lemma E.1. Given ϕ with either LP or LLE architectures, $\Psi(\boldsymbol{X}) = \sum_{i=1}^{N} \phi(\boldsymbol{x}^{(i)})$ is injective and has continuous inverse if:

- $L \in [N(D+1), N^5 D^2]$ when ϕ admits LP architecture. (By Theorem C.8 and C.10).
- $L \in [ND, 2N^2 D^2]$ when ϕ admits LLE architecture. (By Theorem D.1 and D.3).

The proof proceeds by letting $\rho : \mathbb{R}^D \times \mathbb{R}^L \to \mathbb{R}$ take the form $\rho(\boldsymbol{x}, \boldsymbol{z}) = \tau(\boldsymbol{x}, \Phi^{-1}(\boldsymbol{z} - \phi(\boldsymbol{x})))$, and observe that:

$$\begin{aligned}\tau\left(\boldsymbol{x}^{(i)}, \boldsymbol{x}^{(i+1)}, \cdots, \boldsymbol{x}^{(N)}, \cdots, \boldsymbol{x}^{(i-1)}\right) &= \tau\left(\boldsymbol{x}^{(i)}, \Phi^{-1} \circ \Phi(\boldsymbol{x}^{(i+1)}, \cdots, \boldsymbol{x}^{(N)}, \cdots, \boldsymbol{x}^{(i-1)})\right) \\ &= \tau\left(\boldsymbol{x}^{(i)}, \Phi^{-1}\left(\Phi\left(\boldsymbol{x}^{(1)}, \cdots, \boldsymbol{x}^{(N)}\right) - \phi(\boldsymbol{x}^{(i)})\right)\right) \\ &= \rho\left(\boldsymbol{x}^{(i)}, \sum_{i=1}^{N} \phi(\boldsymbol{x}^{(i)})\right)\end{aligned}$$

$\square$

F Extension to Complex Numbers

In this section, we formally introduce the nature extension of our Theorem 3.1 to the complex numbers:

Corollary F.1 (Extension to Complex Domain)**.** *For any permutation-invariant function* $f : \mathcal{K}^{N\times D} \to \mathbb{R}$, $\mathcal{K} \subseteq \mathbb{C}$, *there exists continuous functions* $\phi : \mathbb{C}^D \to \mathbb{R}^L$ *and* $\rho : \mathbb{R}^L \to \mathbb{C}$ *such that* $f(\boldsymbol{X}) = \rho\left(\sum_{i\in[N]} \phi(\boldsymbol{x}^{(i)})\right)$ *for every* $j \in [N]$, *where* $L \in [2N(D+1), 4N^5 D^2]$ *when* ϕ *admits LP architecture, and* $L \in [2ND, 8N^2 D^2]$ *when* ϕ *admits LLE architecture* ($\mathcal{K} \in \mathbb{C}_{>0}$)*.*

Proof. We let ϕ first map complex features $\boldsymbol{x}^{(i)} \in \mathbb{C}^D, \forall i \in [N]$ to real features $\widetilde{\boldsymbol{x}}^{(i)} = \begin{bmatrix}\Re(\boldsymbol{x}^{(i)})^\top & \Im(\boldsymbol{x}^{(i)})^\top\end{bmatrix} \in \mathbb{R}^{2D}, \forall i \in [N]$ by divide the real and imaginary parts into separate channels, then utilize either LP or LLE embedding layer to map $\widetilde{\boldsymbol{x}}^{(i)}$ to the latent space. The upper bounds of desired latent space dimension are scaled by 4 for both architectures due to the quadratic dependence on D. Then the same proof of Theorems C.8, C.10, D.1, and D.3 applies. $\square$

G Connection to Unlabled Sensing

Unlabeled sensing [56], also known as linear regression without correspondence [55, 57–59], solves the linear system $\boldsymbol{y} = \boldsymbol{P}\boldsymbol{A}\boldsymbol{x}$ with a given measurement matrix $\boldsymbol{A} \in \mathbb{R}^{M \times N}$ and an unknown permutation $\boldsymbol{P} \in \Pi(M)$. [55, 56] show that as long as $\boldsymbol{A}$ is over-determinant ($M \geq 2N$), such problem is well-posed (i.e., has a unique solution) for almost all cases. Unlabeled sensing shares the similar structure with our LP embedding layer in which a linear layer lifts the feature space to a higher-dimensional ambient space, ensuring the solvability of alignment across each channel. However, our invertibility is defined between the set and embedding spaces, which differs from exact recovery of unknown variables desired in unlabeled sensing [55]. In fact, the well-posedness of unlabeled PCA [54], studying matrix completion with shuffle perturbations, shares the identical definition with our injectivity. But it is noteworthy that the results in [54] are only drawn over *a dense subset of the input space*, while ours are stronger in considering *all possible inputs*. Hence, our theory could potentially bring new insights into the field of unlabeled sensing, which may be of an independent interest.

References

[1] Yann LeCun, Yoshua Bengio, et al. Convolutional networks for images, speech, and time series. *The handbook of brain theory and neural networks*, 3361(10):1995, 1995.

[2] Taco Cohen and Max Welling. Group equivariant convolutional networks. In *International conference on machine learning*, pages 2990–2999. PMLR, 2016.

[3] Michael M Bronstein, Joan Bruna, Yann LeCun, Arthur Szlam, and Pierre Vandergheynst. Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine*, 34(4):18–42, 2017.

[4] Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In *International Conference on Machine Learning*, pages 2747–2755, 2018.

[5] Haggai Maron, Heli Ben-Hamu, Nadav Shamir, and Yaron Lipman. Invariant and equivariant graph networks. In *International Conference on Learning Representations (ICLR)*, 2018.

[6] Alexander Bogatskiy, Brandon Anderson, Jan Offermann, Marwah Roussi, David Miller, and Risi Kondor. Lorentz group equivariant neural network for particle physics. In *International Conference on Machine Learning*, pages 992–1002. PMLR, 2020.

[7] Peihao Wang, Shenghao Yang, Yunyu Liu, Zhangyang Wang, and Pan Li. Equivariant hypergraph diffusion neural operators. In *International Conference on Learning Representations (ICLR)*, 2023.

[8] 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.

[9] Manzil Zaheer, Satwik Kottur, Siamak Ravanbakhsh, Barnabas Poczos, Russ R Salakhutdinov, and Alexander J Smola. Deep sets. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017.

[10] Vinicius Mikuni and Florencia Canelli. Point cloud transformers applied to collider physics. *Machine Learning: Science and Technology*, 2(3):035027, 2021.

[11] Huilin Qu and Loukas Gouskos. Jet tagging via particle clouds. *Physical Review D*, 101(5):056019, 2020.

[12] Hengshuang Zhao, Li Jiang, Jiaya Jia, Philip HS Torr, and Vladlen Koltun. Point transformer. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 16259–16268, 2021.

[13] Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosiorek, Seungjin Choi, and Yee Whye Teh. Set transformer: A framework for attention-based permutation-invariant neural networks. In *International conference on machine learning*, pages 3744–3753. PMLR, 2019.

[14] Yan Zhang, Jonathon Hare, and Adam Prugel-Bennett. Deep set prediction networks. *Advances in Neural Information Processing Systems*, 32, 2019.

[15] Yan Zhang, Jonathon Hare, and Adam Prügel-Bennett. Fspool: Learning set representations with featurewise sort pooling. In *International Conference on Learning Representations*, 2020.

[16] Aditya Grover, Eric Wang, Aaron Zweig, and Stefano Ermon. Stochastic optimization of sorting networks via continuous relaxations. In *International Conference on Learning Representations*, 2020.

[17] 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 *the AAAI Conference on Artificial Intelligence*, volume 33, pages 4602–4609, 2019.

[18] Keyulu Xu, Weihua Hu, Jure Leskovec, and Stefanie Jegelka. How powerful are graph neural networks? In *International Conference on Learning Representations*, 2019.

[19] Gabriele Corso, Luca Cavalleri, Dominique Beaini, Pietro Liò, and Petar Veličković. Principal neighbourhood aggregation for graph nets. *Advances in Neural Information Processing Systems*, 33:13260–13271, 2020.

[20] 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.

[21] Will Hamilton, Zhitao Ying, and Jure Leskovec. Inductive representation learning on large graphs. In *Advances in Neural Information Processing Systems*, 2017.

[22] 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.

[23] Edward Wagstaff, Fabian Fuchs, Martin Engelcke, Ingmar Posner, and Michael A Osborne. On the limitations of representing functions on sets. In *International Conference on Machine Learning*, pages 6487–6494. PMLR, 2019.

[24] Edward Wagstaff, Fabian B Fuchs, Martin Engelcke, Michael A Osborne, and Ingmar Posner. Universal approximation of functions on sets. *Journal of Machine Learning Research*, 23(151):1–56, 2022.

[25] Nimrod Segol and Yaron Lipman. On universal equivariant set networks. In *International Conference on Learning Representations (ICLR)*, 2020.

[26] Aaron Zweig and Joan Bruna. Exponential separations in symmetric neural networks. *arXiv preprint arXiv:2206.01266*, 2022.

[27] George Cybenko. Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, 2(4):303–314, 1989.

[28] Kurt Hornik, Maxwell Stinchcombe, Halbert White, et al. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.

[29] Shupeng Gui, Xiangliang Zhang, Pan Zhong, Shuang Qiu, Mingrui Wu, Jieping Ye, Zhengdao Wang, and Ji Liu. Pine: Universal deep embedding for graph nodes via partial permutation invariant set functions. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 44(2):770–782, 2021.

[30] Nicolas Bourbaki. *Éléments d'histoire des mathématiques*, volume 4. Springer Science & Business Media, 2007.

[31] David Rydh. A minimal set of generators for the ring of multisymmetric functions. In *Annales de l'institut Fourier*, volume 57, pages 1741–1769, 2007.

[32] Adam Santoro, David Raposo, David G Barrett, Mateusz Malinowski, Razvan Pascanu, Peter Battaglia, and Timothy Lillicrap. A simple neural network module for relational reasoning. *Advances in neural information processing systems*, 30, 2017.

[33] R. L. Murphy, B. Srinivasan, V. Rao, and B. Ribeiro. Janossy pooling: Learning deep permutation-invariant functions for variable-size inputs. In *International Conference on Learning Representations (ICLR)*, 2018.

[34] Akiyoshi Sannai, Yuuki Takai, and Matthieu Cordonnier. Universal approximations of permutation invariant/equivariant functions by deep neural networks. *arXiv preprint arXiv:1903.01939*, 2019.

[35] Branko Ćurgus and Vania Mascioni. Roots and polynomials as homeomorphic spaces. *Expositiones Mathematicae*, 24(1):81–95, 2006.

[36] Andrew R Barron. Universal approximation bounds for superpositions of a sigmoidal function. *IEEE Transactions on Information theory*, 39(3):930–945, 1993.

[37] Djork-Arné Clevert, Thomas Unterthiner, and Sepp Hochreiter. Fast and accurate deep network learning by exponential linear units (elus). *arXiv preprint arXiv:1511.07289*, 2015.

[38] Fernando Q Gouvêa. Was cantor surprised? *The American Mathematical Monthly*, 118(3):198–209, 2011.

[39] Dmitry Yarotsky. Error bounds for approximations with deep relu networks. *Neural Networks*, 94:103–114, 2017.

[40] Shiyu Liang and R Srikant. Why deep neural networks for function approximation? In *International Conference on Learning Representations*, 2017.

[41] Joe Kileel, Matthew Trager, and Joan Bruna. On the expressive power of deep polynomial neural networks. *Advances in neural information processing systems*, 32, 2019.

[42] Nadav Cohen, Or Sharir, and Amnon Shashua. On the expressive power of deep learning: A tensor analysis. In *Conference on learning theory*, pages 698–728. PMLR, 2016.

[43] Maithra Raghu, Ben Poole, Jon Kleinberg, Surya Ganguli, and Jascha Sohl-Dickstein. On the expressive power of deep neural networks. In *international conference on machine learning*, pages 2847–2854. PMLR, 2017.

[44] Dmitry Yarotsky. Universal approximations of invariant maps by neural networks. *Constructive Approximation*, 55(1):407–474, 2022.

[45] Ding-Xuan Zhou. Universality of deep convolutional neural networks. *Applied and computational harmonic analysis*, 48(2):787–794, 2020.

[46] Justin Gilmer, Samuel S Schoenholz, Patrick F Riley, Oriol Vinyals, and George E Dahl. Neural message passing for quantum chemistry. In *International Conference on Machine Learning (ICML)*, 2017.

[47] Thomas N Kipf and Max Welling. Semi-supervised classification with graph convolutional networks. In *International Conference on Learning Representations (ICLR)*, 2017.

[48] Nicolas Keriven and Gabriel Peyré. Universal invariant and equivariant graph neural networks. In *Advances in Neural Information Processing Systems*, pages 7090–7099, 2019.

[49] Zhengdao Chen, Soledad Villar, Lei Chen, and Joan Bruna. On the equivalence between graph isomorphism testing and function approximation with gnns. In *Advances in Neural Information Processing Systems*, pages 15868–15876, 2019.

[50] Zhengdao Chen, Lei Chen, Soledad Villar, and Joan Bruna. Can graph neural networks count substructures? volume 33, 2020.

[51] Waïss Azizian and Marc Lelarge. Expressive power of invariant and equivariant graph neural networks. In *ICLR 2021-International Conference on Learning Representations*, 2021.

[52] Andreas Loukas. What graph neural networks cannot learn: depth vs width. In *International Conference on Learning Representations*, 2020.

[53] Manolis C Tsakiris. Low-rank matrix completion theory via plücker coordinates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2023.

[54] Yunzhen Yao, Liangzu Peng, and Manolis Tsakiris. Unlabeled principal component analysis. *Advances in Neural Information Processing Systems*, 34:30452–30464, 2021.

[55] Manolis Tsakiris and Liangzu Peng. Homomorphic sensing. In *International Conference on Machine Learning*, pages 6335–6344. PMLR, 2019.

[56] Jayakrishnan Unnikrishnan, Saeid Haghighatshoar, and Martin Vetterli. Unlabeled sensing with random linear measurements. *IEEE Transactions on Information Theory*, 64(5):3237–3253, 2018.

[57] Daniel J Hsu, Kevin Shi, and Xiaorui Sun. Linear regression without correspondence. *Advances in Neural Information Processing Systems*, 30, 2017.

[58] Manolis C Tsakiris, Liangzu Peng, Aldo Conca, Laurent Kneip, Yuanming Shi, and Hayoung Choi. An algebraic-geometric approach for linear regression without correspondences. *IEEE Transactions on Information Theory*, 66(8):5130–5144, 2020.

[59] Liangzu Peng and Manolis C Tsakiris. Linear regression without correspondences via concave minimization. *IEEE Signal Processing Letters*, 27:1580–1584, 2020.