
Permutation Equivariant Neural Networks for Symmetric Tensors

Edward Pearce-Crump¹

Abstract

Incorporating permutation equivariance into neural networks has proven to be useful in ensuring that models respect symmetries that exist in data. Symmetric tensors, which naturally appear in statistics, machine learning, and graph theory, are essential for many applications in physics, chemistry, and materials science, amongst others. However, existing research on permutation equivariant models has not explored symmetric tensors as inputs, and most prior work on learning from these tensors has focused on equivariance to Euclidean groups. In this paper, we present two different characterisations of all linear permutation equivariant functions between symmetric power spaces of $\mathbb{R}^n$. We show on two tasks that these functions are highly data efficient compared to standard MLPs and have potential to generalise well to symmetric tensors of different sizes.

1. Introduction

Equivariance has proven to be an effective approach for encoding known structure about a problem directly into a neural network’s architecture, making learning more efficient and interpretable (Cohen & Welling, 2016; 2017; Qi et al., 2017; Ravanbakhsh et al., 2017; Zaheer et al., 2017; Cohen et al., 2018; Esteves et al., 2018; Kondor & Trivedi, 2018; Kondor et al., 2018; Thomas et al., 2018; Weiler et al., 2018; Maron et al., 2019a; Finzi et al., 2021; Satorras et al., 2021; Villar et al., 2021; Yarotsky, 2022; Pearce-Crump, 2023). Equivariant neural networks have been applied successfully in a wide range of problems, including, but not restricted to, particle physics (Bogatskiy et al., 2020; Villar et al., 2021; Gong et al., 2022), biological and medical imaging (Bekkers et al., 2018; Lafarge et al., 2021; Suk et al., 2024), molecular and quantum chemistry (Thomas et al., 2018; Finzi et al., 2020; Fuchs et al., 2020; Miller et al., 2020; Schütt et al.,

2021; Batatia et al., 2022; Batzner et al., 2022; Hoogeboom et al., 2022; Liao & Smidt, 2023), reinforcement learning (van der Pol et al., 2020; Wang et al., 2022a;b;c), computer vision (Marcos et al., 2017; Worrall et al., 2017; Esteves et al., 2019; Deng et al., 2021; Chatzipantazis et al., 2023; Kaba et al., 2023), and modelling compositional structures in natural language (Gordon et al., 2020; Petrache & Trivedi, 2024).

In this paper, we consider linear permutation (S_n) equivariant functions f on symmetric tensors $T \in (\mathbb{R}^n)^{\otimes k}$. In particular, these functions must satisfy

$$f(e_{i_1} \otimes e_{i_2} \otimes \cdots \otimes e_{i_k}) = f(e_{i_{\pi(1)}} \otimes e_{i_{\pi(2)}} \otimes \cdots \otimes e_{i_{\pi(k)}}) \quad (1)$$

for all $\pi \in S_k$, where e_{i_j} are standard basis vectors in $\mathbb{R}^n$, since the coefficients of T , when expressed in this basis, satisfy $T_{i_1, i_2, \dots, i_k} = T_{i_{\pi(1)}, i_{\pi(2)}, \dots, i_{\pi(k)}}$ for all $\pi \in S_k$.

Symmetric tensors and linear permutation equivariant functions on them play a fundamental role in many scientific domains owing to their ability to capture invariant and equivariant relationships in high-dimensional data. In statistics, these functions preserve the defining properties of covariance matrices, which are symmetric and equivariant under permutations of their coordinates (McCullagh, 2018). These functions and tensors are also important in machine learning, where they enable robust parameter estimation in latent variable models (Anandkumar et al., 2014; Jaffe et al., 2018; Goulart et al., 2022). They also appear in clustering algorithms, which must treat all data points equivalently (Khouja, 2022). However, the most notable example is in graph theory, where symmetric adjacency matrices are formed from undirected graphs (Maron et al., 2019a) and high-order adjacency tensors are used to identify affinity relations in hypergraphs (Shashua et al., 2005; Georgii et al., 2011; Ghoshdastidar & Dukkipati, 2017). This framework is essential for many real-world applications: for example, in fluid dynamics simulations, particle interactions must be invariant under exchange (Gao et al., 2022); in materials science, stress-strain relationships must maintain their physical meaning under permutation symmetry (Garanger et al., 2024; Wen et al., 2024); and in neuroscience, brain connectome mapping requires consistent representation of neural pathways (Guha & Guhaniyogi, 2020).

The range and nature of these applications highlight the

¹Department of Mathematics, Imperial College London, United Kingdom. Correspondence to: Edward Pearce-Crump <ep1011@ic.ac.uk>.

importance of achieving a complete and exact characterisation of permutation equivariant linear functions applied to symmetric tensors. However, most studies in the existing machine learning literature on symmetric tensors focus only on equivariance to Euclidean groups, and no prior work on permutation equivariant neural networks has considered the situation where the only inputs are symmetric tensors.

To address this gap, we make the following contributions. We obtain a complete and exact characterisation of all permutation equivariant linear functions between any two symmetric powers of $\mathbb{R}^n$, deriving this result using two different bases. In particular, we show that, by embedding symmetric powers of $\mathbb{R}^n$ into tensor power space, our characterisation includes all linear scalar-valued and vector-valued permutation equivariant functions that are defined on symmetric tensors. As a corollary, we also recover the Deep Sets characterisation (Zaheer et al., 2017). Crucially, we address the implementation challenges that are associated with storing large weight matrices in memory by introducing what we call *map label notation*. This notation makes it possible for us to express the transformation of an input symmetric tensor by a permutation equivariant weight matrix as a series of equations in the input tensor, eliminating the need to store the weight matrices explicitly in memory. In particular, we can then use the same set of equations to compute the transformation between symmetric tensors of fixed orders k and l for any dimension n , eliminating the need to generate separate weight matrices for each choice of n . We validate our approach on two toy problems and show that these functions exhibit high data efficiency and strong potential for generalising well to symmetric tensors of different sizes.

2. Related Work

We describe two important classes of related works.

Equivariant Machine Learning on Symmetric Tensors. Most prior research on learning from symmetric tensors has only considered equivariance to Euclidean groups. Gao et al. (2022) applied tensor contractions to achieve $SO(n)$ -equivariance on high-order symmetric tensors, using their network to simulate fluid systems. Lou & Ganose (2024) and Wen et al. (2024) each introduced $SO(3)$ -equivariant graph neural networks to respectively discover highly anisotropic dielectric crystals and to learn from elasticity tensors, which are fourth-order symmetric tensors that fully characterise a material’s elastic behavior. Heilman et al. (2024) designed an $SE(3)$ -equivariant graph neural network architecture to predict material property tensors directly from crystalline structures. Garanger et al. (2024) constructed a neural network that is equivariant to subgroups of the orthogonal group $O(n)$ in order to learn from symmetric $n \times n$ matrices, and applied it to the constitutive modelling of materials with symmetries.

A few prior works have characterised equivariant functions on symmetric tensors, but none have provided an exact characterisation of permutation equivariant weight matrices for these tensors. Kunisky et al. (2024) used graph moments to characterise $O(n)$ -invariant polynomials on symmetric tensors, while Blum-Smith et al. (2024) introduced a set of $\binom{n+1}{2} + 1$ functions that can universally approximate S_n -invariant functions on almost all real symmetric $n \times n$ matrices.

Permutation Equivariant Neural Networks. These neural networks have been constructed for a number of different types of data. They include permutation equivariant neural networks to learn from set-structured data (Qi et al., 2017; Zaheer et al., 2017; Sverdllov et al., 2024); to learn mappings between different sized sets (Ravanbakhsh et al., 2017); to learn the interaction between sets (Hartford et al., 2018); to learn from graphs and hypergraphs (Maron et al., 2019a; Thiede et al., 2020; Finzi et al., 2021; Morris et al., 2022; Godfrey et al., 2023; Puny et al., 2023; Pearce-Crump, 2024; Pearce-Crump & Knottenbelt, 2024); and to learn from sets of symmetric elements (Maron et al., 2020). Pan & Kondor (2022) studied how to organise the computations that are involved in high-order tensor power networks. Other works have studied the expressive power of permutation equivariant neural networks (Keriven & Peyré, 2019; Maron et al., 2019b; Segol & Lipman, 2020; Yarotsky, 2022).

3. Symmetric Powers of $\mathbb{R}^n$

We present the key definitions and concepts that are needed to study linear permutation equivariant functions on symmetric tensors. We first introduce an action of S_n on a set of k -length indices $S[n]^k$ which will then index a basis of a real vector space, $S^k(\mathbb{R}^n)$, known as the k^{th} symmetric power of $\mathbb{R}^n$. We let $[n] := \{1, \dots, n\}$ throughout.

Define the set $S[n]^k$ to consist of all tuples of the form

$$I = (\underbrace{i_1, \dots, i_1}_{m_1}, \underbrace{i_2, \dots, i_2}_{m_2}, \dots, \underbrace{i_p, \dots, i_p}_{m_p}) \quad (2)$$

such that $p \in [n]$, $1 \leq i_1 < i_2 < \dots < i_p \leq n$, and $\sum m_i = k$. By a “stars and bars” argument, we see that the cardinality of $S[n]^k$ is $\binom{k+n-1}{k}$.

We obtain an action of the symmetric group S_n on $S[n]^k$ by applying the permutation to each element in the tuple and then reordering the result (via a permutation in S_k) to restore ascending numerical order.

Example 3.1. Choosing the element $(1, 1, 1, 2, 3, 3)$ of $S[3]^6$ and applying the permutation (132) in S_3 to it, we first obtain $(3, 3, 3, 1, 2, 2)$ and then reorder it to $(1, 2, 2, 3, 3, 3)$ to give another element of $S[3]^6$.

Consequently, we define the k^{th} symmetric power of $\mathbb{R}^n$, $S^k(\mathbb{R}^n)$, to be the vector space that has a **standard basis**

$\{e_I\}$ indexed by the elements I of $S[n]^k$. Moreover, we can raise the action of S_n on the set $S[n]^k$ to a representation of S_n on $S^k(\mathbb{R}^n)$ in the usual way: we call the representation itself ρ_k^S .

We wish to characterise in full the linear permutation equivariant functions between any two symmetric power spaces of $\mathbb{R}^n$. They are defined as follows.

Definition 3.2. A linear map $f : S^k(\mathbb{R}^n) \rightarrow S^l(\mathbb{R}^n)$ is said to be **permutation equivariant** if, for all $\sigma \in S_n$ and $v \in S^k(\mathbb{R}^n)$

$$f(\rho_k^S(\sigma)[v]) = \rho_l^S(\sigma)[f(v)] \quad (3)$$

We denote the vector space of all such maps by $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$.

Since our goal is to calculate all of the permutation equivariant weight matrices that can appear between any two symmetric power spaces of $\mathbb{R}^n$, it is enough to construct a basis of matrices for $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$, by viewing it as a subspace of $\text{Hom}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ and choosing the standard basis for each symmetric power space, since any weight matrix will be a weighted linear combination of these basis matrices. This characterisation includes the following functions:

Proposition 3.3. *Linear permutation equivariant scalar-valued and vector-valued functions on symmetric tensors in $(\mathbb{R}^n)^{\otimes k}$ are elements of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^0(\mathbb{R}^n))$ and $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^1(\mathbb{R}^n))$, respectively.*

4. Orbit Basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$

We construct our first basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ by associating the basis with the orbits of an action of S_n on the Cartesian product set $S[n]^l \times S[n]^k$ and understanding each orbit as a diagram. We begin with the following result.

Proposition 4.1. *There is a basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ that corresponds bijectively with the orbits coming from the action of S_n on the set $S[n]^l \times S[n]^k$.*

Proof. We see that the vector space $\text{Hom}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ has a standard basis of matrix units

$$\{E_{I,J}\}_{I \in S[n]^l, J \in S[n]^k} \quad (4)$$

where $E_{I,J}$ has a 1 in the (I, J) position and is 0 elsewhere.

Hence, for any standard basis element $e_P \in S^k(\mathbb{R}^n)$, we have that

$$E_{I,J}e_P = \delta_{J,P}e_I \quad (5)$$

and so, for any linear map $f : S^k(\mathbb{R}^n) \rightarrow S^l(\mathbb{R}^n)$, expressing f in the basis of matrix units as

$$f = \sum_{I \in S[n]^l} \sum_{J \in S[n]^k} f_{I,J} E_{I,J} \quad (6)$$

we get that

$$f(e_P) = \sum_{I \in S[n]^l} f_{I,P} e_I \quad (7)$$

Given that $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ is a subspace of $\text{Hom}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$, we have that f is an S_n -equivariant linear map if and only if, for all $\sigma \in S_n$ and standard basis elements $e_J \in S^k(\mathbb{R}^n)$,

$$f(\rho_k^S(\sigma)[e_J]) = \rho_l^S(\sigma)[f(e_J)] \quad (8)$$

(8) holds if and only if

$$\sum_{I \in S[n]^l} f_{I, \sigma(J) \pi_{\sigma(J),k}} e_I = \sum_{I \in S[n]^l} f_{I, J} e_{\sigma(I) \pi_{\sigma(I),l}} \quad (9)$$

for some $\pi_{\sigma(J),k} \in S_k$ and $\pi_{\sigma(I),l} \in S_l$ that restores ascending numerical order in $\sigma(J)$ and $\sigma(I)$ respectively.

(9) now holds if and only if

$$f_{\sigma(I) \pi_{\sigma(I),l}, \sigma(J) \pi_{\sigma(J),k}} = f_{I,J} \quad (10)$$

for all $\sigma \in S_n$, $I \in S[n]^l$ and $J \in S[n]^k$.

Therefore, taking the Cartesian product set $S[n]^l \times S[n]^k$ and writing its elements as $\begin{pmatrix} I \\ J \end{pmatrix}$, we define an action of S_n on $S[n]^l \times S[n]^k$ by

$$\sigma \begin{pmatrix} I \\ J \end{pmatrix} := \begin{pmatrix} \sigma(I) \pi_{\sigma(I),l} \\ \sigma(J) \pi_{\sigma(J),k} \end{pmatrix} \quad (11)$$


Consequently, (10) tells us that the basis elements of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ correspond bijectively with the orbits coming from the action of S_n on $S[n]^l \times S[n]^k$. $\square$

Now, consider an orbit coming from the action of S_n on $S[n]^l \times S[n]^k$. Suppose that $\begin{pmatrix} I \\ J \end{pmatrix}$ is a class representative. We form a diagram in the following way:

We place the values of I in one row and the values of J in a second row directly below it. For each value $x \in [n]$, if x appears in either I or J , then we insert a central green node between the two rows and connect it to every occurrence of x : drawing edges downward from entries in I and upward from entries in J . If x does not appear in either row, then no central green node is added.

Example 4.2. Suppose that $l = 6$, $k = 4$, and $n = 3$. Consider the S_3 orbit of $S[3]^6 \times S[3]^4$ that contains the element $\begin{pmatrix} 111233 \\ 1123 \end{pmatrix}$.

Then, by the construction given above, we obtain






Figure 1. For the $(9, 7)$ -bipartition $\{[3, 1], [2, 2], [2, 1], [1, 3], [1, 0]\}$, we show how to obtain the element in the corresponding S_6 orbit of $S[6]^7 \times S[6]^9$ that comes from labelling its blocks with the 5-length tuple $\{6, 2, 3, 1, 4\}$. By reordering the labelled green nodes of the $(9, 7)$ -orbit bipartition diagram and propagating these values to the ends of the wires, we see that this element is $\begin{pmatrix} 1112236 \\ 122334666 \end{pmatrix}$.

Removing the labels gives the diagram that corresponds to the orbit.

This diagram is, in fact, a diagrammatic representation of a (k, l) -bipartition having at most n blocks that for now we call a (k, l) -orbit bipartition diagram. We define bipartitions below.

Definition 4.3. A (k, l) -bipartition π having some t blocks is the set


$$\{[x_1, y_1], [x_2, y_2], \dots, [x_t, y_t]\} \quad (13)$$

of t pairs $[x_i, y_i]$ such that $x_i, y_i \geq 0$ for all i , not both zero, such that $\sum_{i=1}^t x_i = k$ and $\sum_{i=1}^t y_i = l$. These were first introduced and studied by Major Percy A. MacMahon (1893; 1896; 1899; 1900; 1906; 1908; 1912; 1918). For (k, l) -orbit bipartition diagrams, we call the blocks **spiders** and the black edges **wires**, referring to them as **input wires** or **output wires**, as appropriate.

The number of (k, l) -bipartitions having at most n blocks is the value $p_n(k, l)$. We provide a novel procedure for how to generate all (k, l) -bipartitions having at most n blocks, and hence determine $p_n(k, l)$, in Appendix B.

Note that the (k, l) -bipartition that is obtained from an S_n orbit is independent of the choice of class representative: if we choose $\begin{pmatrix} \sigma(I) \pi_{\sigma(I), l} \\ \sigma(J) \pi_{\sigma(J), k} \end{pmatrix}$ instead and form a diagram using the same procedure given in the red-coloured box above, then the spiders that appear in each of the diagrams (viewed without their labels) will be the same, just potentially in a different order. However, both diagrams give the same (k, l) -bipartition.

Example 4.4. Continuing on from Example 4.2, we see that $\begin{pmatrix} 122333 \\ 1233 \end{pmatrix}$ is in the same orbit as $\begin{pmatrix} 111233 \\ 1123 \end{pmatrix}$, since it is obtained by applying the permutation (132) to $\begin{pmatrix} 111233 \\ 1123 \end{pmatrix}$, then applying permutations in S_6 and S_4 to reorder each component into ascending numerical order. Hence, by applying the procedure given above to $\begin{pmatrix} 122333 \\ 1233 \end{pmatrix}$, we obtain the diagram



which is the same $(4, 6)$ -orbit bipartition diagram as (12), except the spiders are in a different order. However, the underlying $(4, 6)$ -bipartitions are the same.

Remark 4.5. Example 4.4 highlights an important point: in drawing a (k, l) -orbit bipartition diagram for a (k, l) -bipartition, the order of the spiders corresponding to the blocks does not matter. Hence each (k, l) -orbit bipartition diagram represents an entire equivalence class of diagrams that correspond to a (k, l) -bipartition.

Now, we also see that all such (k, l) -bipartitions having at most n blocks must appear in this process: representing each (k, l) -bipartition having some $t \in [n]$ blocks (of the form (13)) as a (k, l) -orbit bipartition diagram by connecting, for each $i \in [t]$, x_i wires from below and y_i wires from above to a central green node, we obtain its orbit by following the six steps in the blue-coloured box below.

1. Form all possible t -length tuples with elements in $[n]$, not allowing for repetitions amongst the elements.
2. Then, for each t -length tuple, label the central green nodes in the diagram from left to right with the elements of the tuple.
3. If necessary, reorder the spiders so that they appear from left to right in increasing order of their labels.
4. Then, for each spider, propagate the block label to the end of each wire.
5. This produces the element $\begin{pmatrix} I \\ J \end{pmatrix}$, where the top row's labels are I and the bottom row's labels are J .
6. Doing this for each t -length tuple gives the entire S_n orbit.

Example 4.6. In Figure 1, we give an example of how to obtain, for the $(9, 7)$ -bipartition $\{[3, 1], [2, 2], [2, 1], [1, 3], [1, 0]\}$, the element in the corresponding S_6 orbit of $S[6]^7 \times S[6]^9$ that comes from labelling its blocks with the 5-length tuple $\{6, 2, 3, 1, 4\}$.

In particular, by reordering the labelled green nodes of the corresponding $(9, 7)$ -orbit bipartition diagram and propagating these values to the ends of the wires, we see that this element is $\begin{pmatrix} 1112236 \\ 122334666 \end{pmatrix}$.

Hence, overall, we have shown the following result:

Proposition 4.7. *The orbits that come from the action of S_n on $S[n]^l \times S[n]^k$ correspond bijectively with the (k, l) -bipartitions (or (k, l) -orbit bipartition diagrams) having at most n blocks.*

Combining Proposition 4.1 and Proposition 4.7, we obtain

Theorem 4.8. *There is a basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ that corresponds bijectively with the (k, l) -bipartitions having at most n blocks. Consequently, its dimension is $p_n(k, l)$. Moreover, we call this basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ the **orbit basis**.*

To find the orbit basis element X_π in $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ that corresponds to a (k, l) -bipartition π having at most n blocks, we represent π as a (k, l) -orbit bipartition diagram x_π and follow the same steps 1 through 5 inclusive as before, except we now form the matrix unit $E_{I,J}$ from each element $\begin{pmatrix} I \\ J \end{pmatrix}$ that comes from a t -length tuple. Adding together all of these matrix units (over all of the tuples) gives X_π . Hence we have the following corollary:


Corollary 4.9. *The S_n -equivariant weight matrix from $S^k(\mathbb{R}^n)$ to $S^l(\mathbb{R}^n)$ is $\sum_\pi \lambda_\pi X_\pi$ for weights λ_π , where the sum is over all (k, l) -bipartitions π having at most n blocks.*

Example 4.10. Continuing on from Example 4.6, we see that the $(1112236, 122334666)$ -entry of the basis matrix X_π in $\text{Hom}_{S_6}(S^9(\mathbb{R}^6), S^7(\mathbb{R}^6))$ corresponding to the $(9, 7)$ -bipartition $\pi = \{[3, 1], [2, 2], [2, 1], [1, 3], [1, 0]\}$ is 1.

In the Appendix, we provide a general procedure summarising the results above for calculating the S_n -equivariant weight matrix from $S^k(\mathbb{R}^n)$ to $S^l(\mathbb{R}^n)$ from all of the (k, l) -orbit bipartition diagrams having at most n blocks.

Example 4.11. Suppose that we would like to find the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3) = \mathbb{R}^3$.

We need to consider the $(2, 1)$ -orbit bipartition diagrams having at most $n = 3$ blocks. They are



Each of these four diagrams corresponds to an orbit basis matrix in $\text{Hom}_{S_3}(S^2(\mathbb{R}^3), S^1(\mathbb{R}^3))$ of size 3×6 . Example C.1 gives the calculations of these matrices.

Hence, the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3)$ is of the form

$$\begin{matrix} & \begin{matrix} 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} \lambda_1 & \lambda_3 & \lambda_3 & \lambda_2 & \lambda_4 & \lambda_2 \\ \lambda_2 & \lambda_3 & \lambda_4 & \lambda_1 & \lambda_3 & \lambda_2 \\ \lambda_2 & \lambda_4 & \lambda_3 & \lambda_2 & \lambda_3 & \lambda_1 \end{bmatrix} \end{matrix} \quad (16)$$

for weights $\lambda_1, \lambda_2, \lambda_3, \lambda_4 \in \mathbb{R}$.

5. Diagram Basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$

We are able to obtain a second basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ that we will show is more efficient for performing linear transformations in the following section. To obtain this second basis, we first need to define the following vector space and reformulate the construction given in the previous section as a linear map.

Definition 5.1. We define the formal $\mathbb{R}$ -linear span of the set of all (k, l) -orbit bipartition diagrams to be the vector space $SP_k^l(n)$. We call this vector space the **spherical partition vector space** since it is an adaptation of the spherical partition algebra that was first discovered by Bastías et al. (2024). We also call the generating basis the **orbit basis** of $SP_k^l(n)$.

Definition 5.2. For all non-negative integers k, l and positive integers n , we define a surjective map

$$\Gamma_{k,n}^l : SP_k^l(n) \rightarrow \text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n)) \quad (17)$$

on the orbit basis of $SP_k^l(n)$ as follows, and extend linearly:


$$\Gamma_{k,n}^l(x_\pi) := \begin{cases} X_\pi & \text{if } \pi \text{ has } n \text{ or fewer blocks} \\ 0 & \text{if } \pi \text{ has more than } n \text{ blocks} \end{cases} \quad (18)$$

We now introduce a new diagram d_π for each (k, l) -bipartition π : it is formed in exactly the same way as a (k, l) -orbit bipartition diagram, except the central green nodes are now filled in. We call such a diagram a **(k, l) -bipartition diagram**.

Example 5.3. Let π be the $(9, 7)$ -bipartition

$$\pi := \{[3, 1], [2, 2], [2, 1], [1, 3], [1, 0]\} \quad (19)$$

Then its corresponding $(9, 7)$ -bipartition diagram, d_π , is



As with the orbit bipartition diagrams, the order of the spiders corresponding to the blocks does not matter in drawing the bipartition diagram.

A General Procedure for Calculating the S_n -Equivariant Weight Matrix from $S^k(\mathbb{R}^n)$ to $S^l(\mathbb{R}^n)$ using the Diagram Basis of $SP_k^l(n)$.

Perform the following steps:


1. Calculate all of the (k, l) -bipartitions π that have at most n blocks, and express each (k, l) -bipartition π as a (k, l) -bipartition diagram d_π in $SP_k^l(n)$.
2. Then, if t is the number of blocks in π , form all possible t -length tuples with elements in $[n]$, allowing for repetitions amongst the elements.
3. For each t -length tuple, label the central green nodes in d_π from left to right with the elements of the tuple. Reorder the spiders in d_π such that they are in increasing order from left to right. Fuse together any spiders whose central green nodes are labelled with the same value. Then, for each spider, propagate the block label to the end of each wire. We obtain the matrix unit $E_{I,J}$ from this diagram by setting I to be the labels in the top row and J to be the labels in the bottom row.
4. Since each tuple corresponds to a matrix unit, add together all of the resulting matrix units from all of the tuples to obtain the basis matrix D_π that is associated with d_π . Attach a weight $\lambda_\pi \in \mathbb{R}$ to each matrix D_π . Finally, calculate $\sum \lambda_\pi D_\pi$ to give the overall weight matrix.

We define d_π to be an element of $SP_k^l(n)$ as follows:

$$d_\pi := \sum_{\pi \preceq \theta} x_\theta \quad (21)$$

where $\preceq$ is the partial order defined on all (k, l) -bipartitions as follows: $\pi \preceq \theta$ if every pair in θ is the sum of pairs in π such that each pair in π appears in exactly one sum.

Example 5.4. We see that if $\pi := \{[1, 1], [1, 0]\}$ then, by (21), we have that



$$d_\pi = d_{\theta_1} + d_{\theta_2} \quad (22)$$

Lemma 5.5. The set of (k, l) -bipartition diagrams $\{d_\pi\}$ forms a basis of $SP_k^l(n)$. We call this basis the **diagram basis** of $SP_k^l(n)$.

We can combine the results in this section to define another map

$$\Theta_{k,n}^l : SP_k^l(n) \rightarrow \text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n)) \quad (23)$$

this time on the diagram basis of $SP_k^l(n)$.

Definition 5.6. We define $\Theta_{k,n}^l(d_\pi)$ in the following way. First change the basis in $SP_k^l(n)$ from the diagram basis to the orbit basis using (21), and then apply $\Gamma_{k,n}^l$ to each of the orbit basis diagrams that appear in (21) for d_π . We get that $\Theta_{k,n}^l(d_\pi)$ is defined to be

$$\Gamma_{k,n}^l\left(\sum_{\pi \preceq \theta} x_\theta\right) = \sum_{\pi \preceq \theta} \Gamma_{k,n}^l(x_\theta) = \sum_{\pi \preceq \theta, b(\theta) \leq n} X_\theta \quad (24)$$

where $b(\theta)$ is the number of blocks in θ . Now define

$$D_\pi := \sum_{\pi \preceq \theta, b(\theta) \leq n} X_\theta \quad (25)$$

Clearly, D_π is an element of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$. Hence we have defined $\Theta_{k,n}^l(d_\pi) := D_\pi$. $\Theta_{k,n}^l$ becomes a map on $SP_k^l(n)$ by extending the definition by linearity to the entire diagram basis of $SP_k^l(n)$. We claim the following.

Theorem 5.7. For all non-negative integers k, l and positive integers n , the set $\{D_\pi \mid d_\pi \text{ has at most } n \text{ blocks}\}$ forms a basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$. We call this basis the **diagram basis** of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$. In particular, the map $\Theta_{k,n}^l$ is surjective.

Note that whilst $\Theta_{k,n}^l$ is defined on all diagram basis elements in $SP_k^l(n)$, it is only those that have at most n blocks whose image under $\Theta_{k,n}^l$ creates a basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$. In particular, we have the following corollary.

Corollary 5.8. The S_n -equivariant weight matrix from $S^k(\mathbb{R}^n)$ to $S^l(\mathbb{R}^n)$ is $\sum_\pi \lambda_\pi D_\pi$ for weights λ_π , where the sum is over all (k, l) -bipartitions π having at most n blocks.

We would like to be able to calculate the diagram basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ directly from the diagram basis of $SP_k^l(n)$, without having to go via the orbit basis of $SP_k^l(n)$. We claim the following:

Proposition 5.9. The diagram basis element D_π in $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ that corresponds to a (k, l) -bipartition diagram d_π having some $t \leq n$ blocks can be found as follows: we follow the same steps 1 through 5





Figure 2. After labelling the central green nodes of a $(5, 4)$ -bipartition diagram with the tuple $(1, 2, 1)$ and reordering the spiders into ascending numerical order, we fuse together any spiders that are labelled with the same value.

inclusive in the blue-coloured box as before, except in Step 1 we now allow repetitions amongst the elements in each t -length tuple, and in Step 3, after reordering the spiders, we fuse together any central green nodes that are labelled with the same value. See Figure 2 for an example. Adding together all of the resulting matrix units gives D_π .

In the orange-coloured box, we provide a general procedure for calculating the S_n -equivariant weight matrix from $S^k(\mathbb{R}^n)$ to $S^l(\mathbb{R}^n)$ using the diagram basis of $SP_k^l(n)$.

Example 5.10. We return to Example 4.11 and calculate the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3)$ using the diagram basis of $SP_2^1(3)$ instead.

We now consider the $(2, 1)$ -bipartition diagrams having at most $n = 3$ blocks:



As before, the calculations of the corresponding matrices are given in Example C.2.

Hence, the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3)$ using the diagram basis is of the form

$$\begin{matrix} & \text{1,1} & \text{1,2} & \text{1,3} & \text{2,2} & \text{2,3} & \text{3,3} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} \lambda_{1,2,3,4} & \lambda_{3,4} & \lambda_{3,4} & \lambda_{2,4} & \lambda_4 & \lambda_{2,4} \\ \lambda_{2,4} & \lambda_{3,4} & \lambda_4 & \lambda_{1,2,3,4} & \lambda_{3,4} & \lambda_{2,4} \\ \lambda_{2,4} & \lambda_4 & \lambda_{3,4} & \lambda_{2,4} & \lambda_{3,4} & \lambda_{1,2,3,4} \end{bmatrix} \end{matrix} \quad (27)$$

for weights $\lambda_1, \lambda_2, \lambda_3, \lambda_4 \in \mathbb{R}$, where $\lambda_A := \sum_{i \in A} \lambda_i$.

Remark 5.11. In Example C.3 in the Appendix, we show that we can recover the Deep Sets (Zaheer et al., 2017) characterisation of any S_n -equivariant weight matrix from $\mathbb{R}^n$ to $\mathbb{R}^n$ using the diagram basis of $SP_1^1(n)$.

6. Embedding into Tensor Power Space

There is a useful way of representing any S_n -equivariant weight matrix W from $S^k(\mathbb{R}^n)$ to $S^l(\mathbb{R}^n)$ as a matrix from $(\mathbb{R}^n)^{\otimes k}$ to $(\mathbb{R}^n)^{\otimes l}$, namely by embedding $S^k(\mathbb{R}^n)$ into $(\mathbb{R}^n)^{\otimes k}$ and similarly for $S^l(\mathbb{R}^n)$. We call the latter matrix

the **unrolled** S_n -equivariant weight matrix, and the process of representing the original matrix in this way **unrolling**.

The (I, J) -entry in the unrolled weight matrix is the same as the $(I\pi_{I,l}, J\pi_{J,k})$ -entry in W , where $\pi_{I,l}, \pi_{J,k}$ are any permutations in S_l and S_k , respectively, that restore ascending numerical order in I and J , respectively.

This is because the k^{th} symmetric power of $\mathbb{R}^n$ can be thought of as the quotient space of $(\mathbb{R}^n)^{\otimes k}$ by the ideal generated by elements of the form $x \otimes y - y \otimes x$.

The benefit of unrolling an S_n -equivariant weight matrix to become a map from $(\mathbb{R}^n)^{\otimes k}$ to $(\mathbb{R}^n)^{\otimes l}$ is that we can now operate on symmetric tensors in $(\mathbb{R}^n)^{\otimes k}$ whilst retaining the original properties of the matrix. Going forward, we now consider any weight matrix to be in its unrolled form.

Example 6.1. Continuing on from Example 4.11, we see that the unrolled S_3 -equivariant weight matrix from $(\mathbb{R}^3)^{\otimes 2}$ to $\mathbb{R}^3$ is

$$\begin{matrix} & \text{1,1} & \text{1,2} & \text{1,3} & \text{2,1} & \text{2,2} & \text{2,3} & \text{3,1} & \text{3,2} & \text{3,3} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} \lambda_1 & \lambda_3 & \lambda_3 & \lambda_3 & \lambda_2 & \lambda_4 & \lambda_3 & \lambda_4 & \lambda_2 \\ \lambda_2 & \lambda_3 & \lambda_4 & \lambda_3 & \lambda_1 & \lambda_3 & \lambda_4 & \lambda_3 & \lambda_2 \\ \lambda_2 & \lambda_4 & \lambda_3 & \lambda_4 & \lambda_2 & \lambda_3 & \lambda_3 & \lambda_3 & \lambda_1 \end{bmatrix} \end{matrix} \quad (28)$$

7. Map Label Notation

In practice, when it comes to implementing the transformation of an input tensor T by an unrolled weight matrix W , expressed in tensor form, restrictions on memory make it impractical for the entire weight matrix to be stored in memory. Instead we can use what we term **map label notation** to describe, for each basis matrix D_π that appears in W , the transformation of an input T into its output $D_\pi(T)$. It is here where the diagram basis shines over the orbit basis, in that it is much easier to vectorize the map label for each diagram basis element, since it includes, as a result of (21), all labels that come from all of the possible fusings of the spiders in the bipartition diagram d_π that corresponds to D_π . In fact, instead of calculating a separate weight matrix explicitly for each value of n , when $n \geq l + k$, we will see that a single set of map labels can be used to obtain the transformation for any value of n , and, when $n < l + k$, we only need to retain map labels that come from (k, l) -bipartition diagrams having at most n blocks. Consequently, the only

requirements on memory come from needing to store the input tensor T and the output tensor $W(T)$ in memory.

Definition 7.1. Let d_π be a (k, l) -bipartition diagram. We associate to d_π a set of **map labels**, where each map label consists of an arrow $\leftarrow$, a fixed pattern of indices I on its left hand side, and, on its right hand side, a formal linear combination of tuples of indices that are associated with I .

We immediately use each map label for d_π to obtain the values in $D_\pi(T)$, for unrolled D_π and input tensor $T \in (\mathbb{R}^n)^{\otimes k}$, by replacing, in each map label, I by $D_\pi(T)_I$ and each tuple J that appears on the right hand side by T_J , and treating the result as a map that describes how to find $D_\pi(T)_I$. The resulting **transformation map labels** coming from all of the map labels ultimately determine $D_\pi(T)$.

We have provided a full procedure for how to obtain the map labels and the resulting transformation for any (k, l) -bipartition diagram d_π and for any value of n in the green-coloured box. The implementations of Subprocedures I to V, along with illustrative examples for each, are provided in Appendix D. We give definitions of the important terms that appear in the procedure below.

Definition 7.2. Let d_π be a (k, l) -bipartition diagram. A **grouped output** for d_π is any tuple of labels that can be obtained by labelling all the spiders that have output wires, allowing repeats, sorting the spiders in decreasing order of output wires, propagating the labels to the ends of the output wires, and reading off the resulting tuple. A **partially labelled diagram** of d_π is identical to d_π , except all spiders that have output wires are assigned a label, allowing repeats. A **fully labelled diagram** of d_π is identical to d_π , except now all spiders are assigned a label, allowing repeats. A **grouped map label** is obtained from a fully labelled diagram of d_π by propagating the labels to the ends of the wires, placing the tuple on the output wires to the left of an arrow $\leftarrow^g$ and placing the tuple on the input wires to the right of the arrow. A **left grouped map label** is obtained from a grouped map label by unrolling only the right hand side of a grouped map label, and replacing $\leftarrow^g$ with $\leftarrow^{lg}$.

Example 7.3. Continuing on from Example 5.10, and considering once again the $(2, 1)$ -bipartition diagrams having at most $n = 3$ blocks given in (26), we show in Example D.7 that the transformations corresponding to each $(2, 1)$ -bipartition diagram, for $n = 3$, are fully determined by the following transformation map labels — in this case, only one for each $(2, 1)$ -bipartition diagram:

$$D_{\pi_1}(T)_i = T_{i,i} \quad (29)$$

$$D_{\pi_2}(T)_i = \sum_{j=1}^3 T_{j,j} \quad (30)$$




Figure 3. Data efficiency for the synthetic S_{12} -invariant task. The shaded regions depict 95% confidence intervals taken over 3 runs.

$$D_{\pi_3}(T)_i = \sum_{j=1}^3 [T_{i,j} + T_{j,i}] - T_{i,i} \quad (31)$$

$$D_{\pi_4}(T)_i = \sum_{j=1}^3 \sum_{k=1}^3 T_{j,k} \quad (32)$$

where $i \in [3]$.

However, the power of this method lies in the fact that if $n \geq 3$, then we only need to change the value 3 in each sum to n , and if $n = 2$, then we remove π_4 and change the value 3 in each sum to 2.

8. Numerical Experiments

We demonstrate our characterisation on the following toy experiments.

S_{12} -Invariant Task: We evaluate our model on a synthetic S_{12} -invariant task given by the function $f(T) := \sum_{i,j}^{12} T_{i,j,i}$, where T is a 3-order symmetric tensor. We demonstrate the high data efficiency of our model compared with a standard MLP for this task, as shown in Figure 3. We attribute this superior performance to our model’s strong inductive bias.

S_8 -Equivariant Task: We evaluate our model on a synthetic S_8 -equivariant task from $(\mathbb{R}^8)^{\otimes 3}$ to $\mathbb{R}^8$: namely, to extract the diagonal from $8 \times 8 \times 8$ symmetric tensors. We evaluate our model against a standard MLP and a standard S_8 -equivariant model from $(\mathbb{R}^8)^{\otimes 3}$ to $\mathbb{R}^8$. We show the Test

Procedure: Generation of the Transformation Map Labels That Describe the Transformation Corresponding to a (k, l) –Bipartition Diagram of a Symmetric Tensor in $(\mathbb{R}^n)^{\otimes k}$ to a Symmetric Tensor in $(\mathbb{R}^n)^{\otimes l}$.

Input: A (k, l) –bipartition diagram d_π and a value of n .

1. Apply Subprocedure I to obtain all possible grouped outputs for d_π , and for each one, its associated set of partially labelled diagrams, where only spiders that have output wires are labelled in these diagrams.
2. For each grouped output for d_π :
 - (a) Apply Subprocedure II to each partially labelled diagram in the grouped output’s associated set. Take the set union of all fully labelled diagrams that are returned from each call to the subprocedure, giving a set of fully labelled diagrams that are associated with the grouped output.
 - (b) Remove all fully labelled diagrams in the set with more than n spiders. If the set becomes empty, delete the grouped output and proceed to the next iteration of the loop.
 - (c) Apply Subprocedure III to each fully labelled diagram to convert it into a grouped map label.
 - (d) For each grouped map label, apply Subprocedure IV to unroll its right hand side, returning a left grouped map label.
 - (e) Formally add together the right hand side of the resulting left grouped map labels, creating one left grouped map label that is associated with the grouped output.
 - (f) Unroll the left hand side of the left grouped map label according to Subprocedure V to obtain a set of map labels that are associated with the grouped output.
3. The union of the resulting map label sets from all grouped outputs are the map labels that correspond to the (k, l) –bipartition diagram d_π .
4. Finally, for each map label in Step 3, convert each index J on the right hand side to be T_J , and convert the index I on the left hand side to be $D_\pi(T)_I$ to obtain its corresponding transformation map label.

Output: A set of transformation map labels describing the transformation D_π of a symmetric tensor $T \in (\mathbb{R}^n)^{\otimes k}$ to a symmetric tensor $D_\pi(T) \in (\mathbb{R}^n)^{\otimes l}$ for the given value of n , where D_π corresponds to the (k, l) –bipartition diagram d_π .

Mean Squared Error (MSE) for each of these models in Table 1. Furthermore, Table 2 shows that the network appears to generalise well to symmetric tensors having different sizes, even though it was trained on symmetric tensors of a particular size.

Method	Test MSE
MLP	0.6486
PermEquiv	0.0447
SymmPermEquiv	0.0035

Table 1. Diagonal Extraction

n	8	16	32
Test MSE	0.0035	0.0048	0.0088

Table 2. Generalisation

9. Conclusion

We obtained a complete characterisation of permutation equivariant linear functions between symmetric powers of $\mathbb{R}^n$ using two different bases. This characterisation includes all scalar-valued and vector-valued permutation equivariant functions on symmetric tensors in $(\mathbb{R}^n)^{\otimes k}$. We introduced map label notation to address limitations on memory, enabling the efficient transformation of input tensors without explicitly storing large sized weight matrices in memory. We validated our approach through two toy examples and demonstrated the potential for transfer learning to symmetric tensors of different sizes.

Acknowledgements

The author is deeply grateful to Chen Lin for his valuable insights and constructive discussions on designing experiments. This work is supported by the UKRI–EPSRC grant EP/Y028872/1, *Mathematical Foundations of Intelligence: An “Erlangen Programme” for AI*.

Impact Statement

This paper presents work whose goal is to advance the field of Machine Learning. There are many potential societal consequences of our work, none which we feel must be specifically highlighted here.

References

- Anandkumar, A., Ge, R., Hsu, D. J., Kakade, S. M., and Telgarsky, M. Tensor Decompositions for Learning Latent Variable Models. *Journal of Machine Learning Research*, 15:2773–2832, 2014.
- Andrews, G. E. An extension of Carlitz’s bipartition identity. *Proceedings of the American Mathematical Society*, 63 (1):180–184, 1977.
- Andrews, G. E. *The Theory of Partitions*, volume 2. Cambridge University Press, 1998.
- Andrews, G. E. and Paule, P. MacMahon’s Dream. In *Partitions, q -Series, and Modular Forms*, volume 23 of *Developments in Mathematics*, pp. 1–12. Springer, New York, NY, 2012.
- Auluck, F. C. On Partitions of Bipartite Numbers. *Mathematical Proceedings of the Cambridge Philosophical Society*, 49(1):72–83, 1953.
- Bastías, K. O., Martin, P., and Ryom-Hansen, S. On the spherical partition algebra, 2024. arXiv:2402.01890.
- Batatia, I., Kovacs, D. P., Simm, G., Ortner, C., and Csanyi, G. MACE: Higher Order Equivariant Message Passing Neural Networks for Fast and Accurate Force Fields. In *Advances in Neural Information Processing Systems*, volume 35, pp. 11423–11436, 2022.
- Batzner, S., Musaelian, A., Lixin Sun, M. G., Mailoa, J. P., Kornbluth, M., Molinari, N., Smidt, T. E., and Kozinsky, B. E(3)-equivariant graph neural networks for data-efficient and accurate interatomic potentials. *Nature Communications*, 13:2453, 2022.
- Bekkers, E. J., Lafarge, M. W., Veta, M., Eppenhof, K. A. J., Pluim, J. P. W., and Duits, R. Roto-Translation Covariant Convolutional Networks for Medical Image Analysis. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16-20, 2018, Proceedings, Part I*, pp. 440–448. Springer International Publishing, 2018.
- Blum-Smith, B., Huang, N., Cuturi, M., and Villar, S. Learning functions on symmetric matrices and point clouds via lightweight invariant features, 2024. arXiv:2405.08097.
- Bogatskiy, A., Anderson, B., Offermann, J., Roussi, M., Miller, D., and Kondor, R. Lorentz Group Equivariant Neural Network for Particle Physics. In *International Conference on Machine Learning*, pp. 992–1002. PMLR, 2020.
- Carlitz, L. A Problem in Partitions. *Duke Mathematical Journal*, 30(2):203–213, 1963.
- Carlitz, L. and Roselle, D. P. Restricted Bipartite Partitions. *Pacific Journal of Mathematics*, 19(2):221–228, 1966.
- Chatzipantazis, E., Pertigkiozoglou, S., Dobriban, E., and Daniilidis, K. SE(3)-Equivariant Attention Networks for Shape Reconstruction in Function Space. In *The Eleventh International Conference on Learning Representations*, 2023.
- Cohen, T. and Welling, M. Group Equivariant Convolutional Networks. In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2990–2999, New York, USA, 20–22 Jun 2016. PMLR.
- Cohen, T. and Welling, M. Steerable CNNs. In *International Conference on Learning Representations*, 2017.
- Cohen, T., Geiger, M., Köhler, J., and Welling, M. Spherical CNNs. In *International Conference on Learning Representations*, 2018.
- Deng, C., Litany, O., Duan, Y., Poulenard, A., Tagliasacchi, A., and Guibas, L. J. Vector Neurons: A General Framework for SO(3)-Equivariant Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 12200–12209, 2021.
- Esteves, C., Allen-Blanchette, C., Makadia, A., and Daniilidis, K. Learning SO(3) Equivariant Representations with Spherical CNNs. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pp. 52–68, 2018.
- Esteves, C., Xu, Y., Allen-Blanchette, C., and Daniilidis, K. Equivariant Multi-View Networks. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, October 2019.
- Finzi, M., Stanton, S., Izmailov, P., and Wilson, A. G. Generalizing Convolutional Neural Networks for Equivariance to Lie Groups on Arbitrary Continuous Data. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 3165–3176. PMLR, 13–18 Jul 2020.
- Finzi, M., Welling, M., and Wilson, A. G. A Practical Method for Constructing Equivariant Multilayer Perceptrons for Arbitrary Matrix Groups. In *Proceedings of the*

- 38th International Conference on Machine Learning, volume 139 of *Proceedings of Machine Learning Research*, pp. 3318–3328. PMLR, 18–24 Jul 2021.
- Fuchs, F. B., Worrall, D. E., Fischer, V., and Welling, M. SE(3)-transformers: 3D roto-translation equivariant attention networks. In *Proceedings of the 34th International Conference on Neural Information Processing Systems*, 2020.
- Gao, L., Du, Y., Li, H., and Lin, G. RotEqNet: Rotation-equivariant network for fluid systems with symmetric high-order tensors. *Journal of Computational Physics*, 461:111205, 2022.
- Garanger, K., Kraus, J., and Rimoli, J. J. Symmetry-enforcing neural networks with applications to constitutive modeling. *Extreme Mechanics Letters*, 71:102188, 2024.
- Georgii, E., Tsuda, K., and Schölkopf, B. Multi-way set enumeration in weight tensors. *Machine Learning*, 82: 123–155, 2011.
- Ghoshdastidar, D. and Dukkipati, A. Uniform Hypergraph Partitioning: Provable Tensor Methods and Sampling Techniques. *Journal of Machine Learning Research*, 18 (50):1–41, 2017.
- Godfrey, C., Rawson, M. G., Brown, D., and Kvinge, H. Fast computation of permutation equivariant layers with the partition algebra. In *ICLR 2023 Workshop on Physics for Machine Learning*, 2023.
- Gong, S., Meng, Q., Zhang, J., Qu, H., Li, C., Qian, S., Du, W., Ma, Z.-M., and Liu, T.-Y. An Efficient Lorentz Equivariant Graph Neural Network for Jet Tagging. *Journal of High Energy Physics*, 2022(7):1–22, 2022.
- Gordon, J., Lopez-Paz, D., Baroni, M., and Bouchacourt, D. Permutation Equivariant Models for Compositional Generalization in Language. In *International Conference on Learning Representations*, 2020.
- Goulart, J. H. M., Couillet, R., and Comon, P. A Random Matrix Perspective on Random Tensors. *Journal of Machine Learning Research*, 23(264):1–36, 2022.
- Guha, S. and Guhaniyogi, R. Bayesian Generalized Sparse Symmetric Tensor-on-Vector Regression. *Technometrics*, 63(2):160–170, 2020.
- Hartford, J. S., Graham, D. R., Leyton-Brown, K., and Ravanbakhsh, S. Deep Models of Interactions Across Sets. In *Proceedings of the 35th International Conference on Machine Learning*, pp. 1914–1923. PMLR, 2018.
- Heilman, A., Schlesinger, C., and Yan, Q. Equivariant Graph Neural Networks for Prediction of Tensor Material Properties of Crystals, 2024. [arXiv:2406.03563](#).
- Hoogeboom, E., Satorras, V. G., Vignac, C., and Welling, M. Equivariant Diffusion for Molecule Generation in 3D. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162 of *Proceedings of Machine Learning Research*, pp. 8867–8887. PMLR, 17–23 Jul 2022.
- Jaffe, A., Weiss, R., Nadler, B., Carmi, S., and Kluger, Y. Learning Binary Latent Variable Models: A Tensor Eigenpair Approach. In *International Conference on Machine Learning*, pp. 2196–2205. PMLR, 2018.
- Kaba, S.-O., Mondal, A. K., Zhang, Y., Bengio, Y., and Ravanbakhsh, S. Equivariance with Learned Canonicalization Functions. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 15546–15566. PMLR, 23–29 Jul 2023.
- Keriven, N. and Peyré, G. Universal Invariant and Equivariant Graph Neural Networks. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Khouja, R. *Optimization algorithms for the tensor rank approximation problem: application to clustering in machine learning*. PhD thesis, Université Côte d’Azur; Université Libanaise, 2022.
- Kim, J. K. and Hahn, S. G. Partitions of bipartite numbers. *Graphs and Combinatorics*, 13:73–78, 1997.
- Kondor, R. and Trivedi, S. On the Generalization of Equivariance and Convolution in Neural Networks to the Action of Compact Groups. In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pp. 2747–2755. PMLR, 10–15 Jul 2018.
- Kondor, R., Lin, Z., and Trivedi, S. Clebsch–Gordan Nets: a Fully Fourier Space Spherical Convolutional Neural Network. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Kunisky, D., Moore, C., and Wein, A. S. Tensor cumulants for statistical inference on invariant distributions, 2024. [arXiv:2404.18735](#).
- Lafarge, M. W., Bekkers, E. J., Pluim, J. P. W., Duits, R., and Veta, M. Roto-translation equivariant convolutional networks: Application to histopathology image analysis. *Medical Image Analysis*, 68:101849, 2021.
- Landman, B. M., Brown, E. A., and Portier, F. J. Partitions of bi-partite numbers into at most j parts. *Graphs and Combinatorics*, 8:65–73, 1992.

- Liao, Y.-L. and Smidt, T. Equiformer: Equivariant Graph Attention Transformer for 3D Atomistic Graphs. In *The Eleventh International Conference on Learning Representations*, 2023.
- Lou, Y. and Ganose, A. M. Discovery of highly anisotropic dielectric crystals with equivariant graph neural networks, 2024. [arXiv:2405.07915](#).
- MacMahon, P. A. XVII. Memoir on the Theory of the Compositions of Numbers. *Philosophical Transactions of the Royal Society of London. (A.)*, 184:835–901, 1893.
- MacMahon, P. A. XVI. Memoir on the Theory of the Partition of Numbers. Part I. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 187:619–673, 1896.
- MacMahon, P. A. Memoir on the Theory of the Partitions of Numbers. Part II. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 192:351–401, 1899.
- MacMahon, P. A. Combinatorial Analysis. The Foundations of a New Theory. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 194:361–386, 1900.
- MacMahon, P. A. Memoir on the Theory of the Partitions of Numbers. Part III. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 205:37–59, 1906.
- MacMahon, P. A. II. Second Memoir on the Compositions of Numbers. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 207(413-426): 65–134, 1908.
- MacMahon, P. A. III. Memoir on the Theory of the Partitions of Numbers. Part V. Partitions in Two-Dimensional Space. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 211(471-483):75–110, 1912.
- MacMahon, P. A. Seventh Memoir on the Partition of Numbers. A Detailed Study of the Enumeration of the Partitions of Multipartite Numbers. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 217: 81–113, 1918.
- Marcos, D., Volpi, M., Komodakis, N., and Tuia, D. Rotation Equivariant Vector Field Networks. In *Proceedings of the IEEE International Conference on Computer Vision*, pp. 5048–5057, 2017.
- Maron, H., Ben-Hamu, H., Shamir, N., and Lipman, Y. Invariant and Equivariant Graph Networks. In *International Conference on Learning Representations*, 2019a.
- Maron, H., Fetaya, E., Segol, N., and Lipman, Y. On the Universality of Invariant Networks. In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pp. 4363–4371. PMLR, 09–15 Jun 2019b.
- Maron, H., Litany, O., Chechik, G., and Fetaya, E. On Learning Sets of Symmetric Elements. In *Proceedings of the 37th International Conference on Machine Learning*, volume 119 of *Proceedings of Machine Learning Research*, pp. 6734–6744. PMLR, 13–18 Jul 2020.
- Mathews, G. B. On the partition of numbers. *Proceedings of the London Mathematical Society*, s1-28(1):486–490, 1896.
- McCullagh, P. *Tensor Methods in Statistics*. Courier Dover Publications, 2018.
- Miller, B. K., Geiger, M., Smidt, T. E., and Noé, F. Relevance of rotationally equivariant convolutions for predicting molecular properties. *arXiv preprint arXiv:2008.08461*, 2020.
- Morris, C., Rattan, G., Kiefer, S., and Ravanbakhsh, S. SpeqNets: Sparsity-aware Permutation-equivariant graph networks. In *Proceedings of the 39th International Conference on Machine Learning*, volume 162, pp. 16017–16042, 17–23 Jul 2022.
- Nanda, V. S. Bipartite partitions. *Mathematical Proceedings of the Cambridge Philosophical Society*, 53(2):273–277, 1957.
- Pan, H. and Kondor, R. Permutation Equivariant Layers for Higher Order Interactions. In *Proceedings of The 25th International Conference on Artificial Intelligence and Statistics*, volume 151 of *Proceedings of Machine Learning Research*, pp. 5987–6001. PMLR, 28–30 Mar 2022.
- Pearce-Crump, E. Brauer’s Group Equivariant Neural Networks. In *Proceedings of the 40th International Conference on Machine Learning*, volume 202 of *Proceedings of Machine Learning Research*, pp. 27461–27482. PMLR, 23–29 Jul 2023.
- Pearce-Crump, E. Connecting Permutation Equivariant Neural Networks and Partition Diagrams. In *Proceedings of the 27th European Conference on Artificial Intelligence*, volume 392, pp. 1511–1518. IOS Press Ebooks, 2024.

- Pearce-Crump, E. and Knottenbelt, W. J. Graph Automorphism Group Equivariant Neural Networks. In *Proceedings of the 41st International Conference on Machine Learning*, volume 235 of *Proceedings of Machine Learning Research*, pp. 40051–40077. PMLR, 21–27 Jul 2024.
- Petrache, M. and Trivedi, S. Position Paper: Generalized grammar rules and structure-based generalization beyond classical equivariance for lexical tasks and transduction, 2024. [arXiv:2402.01629](#).
- Puny, O., Lim, D., Kiani, B., Maron, H., and Lipman, Y. Equivariant Polynomials for Graph Neural Networks. In *International Conference on Machine Learning*, pp. 28191–28222. PMLR, 2023.
- Qi, C. R., Su, H., Mo, K., and Guibas, L. J. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pp. 652–660, 2017.
- Ravanbakhsh, S., Schneider, J., and Póczos, B. Equivariance Through Parameter-Sharing. In *Proceedings of the 34th International Conference on Machine Learning*, volume 70, pp. 2892–2901, 06–11 Aug 2017.
- Satorras, V. G., Hoogeboom, E., and Welling, M. E(n) Equivariant Graph Neural Networks. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9323–9332. PMLR, 18–24 Jul 2021.
- Schütt, K., Unke, O., and Gastegger, M. Equivariant Message Passing for the Prediction of Tensorial Properties and Molecular Spectra. In *Proceedings of the 38th International Conference on Machine Learning*, volume 139 of *Proceedings of Machine Learning Research*, pp. 9377–9388. PMLR, 18–24 Jul 2021.
- Segol, N. and Lipman, Y. On universal equivariant set networks. In *International Conference on Learning Representations*, 2020.
- Shashua, A., Zass, R., and Hazan, T. Multi-way Clustering Using Super-Symmetric Non-Negative Tensor Factorization. In *European Conference on Computer Vision*, pp. 595–608, 2005.
- Stanley, R. P. *Enumerative Combinatorics: Volume 1*. Cambridge University Press, 1997.
- Suk, J., de Haan, P., Lippe, P., Brune, C., and Wolterink, J. M. Mesh neural networks for SE(3)-equivariant hemodynamics estimation on the artery wall. *Computers in Biology and Medicine*, 173, 2024.
- Sverdlov, Y., Springer, I., and Dym, N. Revisiting Multi-Permutation Equivariance through the Lens of Irreducible Representations, 2024. [arXiv:2410.06665](#).
- Thiede, E. H., Hy, T. S., and Kondor, R. The general theory of permutation equivariant neural networks and higher order graph variational encoders, 2020. [arXiv:2004.03990](#).
- Thomas, N., Smidt, T., Kearnes, S., Yang, L., Li, L., Kohlhoff, K., and Riley, P. Tensor field networks: Rotation-and translation-equivariant neural networks for 3d point clouds, 2018. [arXiv:1802.08219](#).
- van der Pol, E., Worrall, D., van Hoof, H., Oliehoek, F., and Welling, M. MDP Homomorphic Networks: Group Symmetries in Reinforcement Learning. In *Advances in Neural Information Processing Systems*, volume 33, pp. 4199–4210, 2020.
- Villar, S., Hogg, D. W., Storey-Fisher, K., Yao, W., and Blum-Smith, B. Scalars are universal: Equivariant machine learning, structured like classical physics. In *Advances in Neural Information Processing Systems*, 2021.
- Wang, D., Jia, M., Zhu, X., Walters, R., and Platt, R. On-Robot Learning with Equivariant Models. In *6th Annual Conference on Robot Learning*, 2022a.
- Wang, D., Walters, R., and Platt, R. SO(2)-Equivariant Reinforcement Learning. In *International Conference on Learning Representations*, 2022b.
- Wang, D., Walters, R., Zhu, X., and Platt, R. Equivariant Q Learning in Spatial Action Spaces. In *Proceedings of the 5th Conference on Robot Learning*, volume 164 of *Proceedings of Machine Learning Research*, pp. 1713–1723. PMLR, 08–11 Nov 2022c.
- Weiler, M. and Cesa, G. General E(2)-Equivariant Steerable CNNs. In *Advances in Neural Information Processing Systems*, volume 32, 2019.
- Weiler, M., Geiger, M., Welling, M., Boomsma, W., and Cohen, T. 3D Steerable CNNs: Learning Rotationally Equivariant Features in Volumetric Data. In *Advances in Neural Information Processing Systems*, volume 31, 2018.
- Wen, M., Horton, M. K., Munro, J. M., Huck, P., and Persson, K. A. An equivariant graph neural network for the elasticity tensors of all seven crystal systems. *Digital Discovery*, 3(5):869–882, 2024.
- Worrall, D. E., Garbin, S. J., Turmukhambetov, D., and Brostow, G. J. Harmonic Networks: Deep Translation and Rotation Equivariance. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2017.

- Wright, E. M. The number of partitions of a large bi-partite number. *Proceedings of the London Mathematical Society*, 3(1):150–160, 1957.
- Wright, E. M. Partitions of large bipartites. *American Journal of Mathematics*, 80(3):643–658, 1958.
- Wright, E. M. Partition of multipartite numbers into k parts. *Journal für die reine und angewandte Mathematik*, 1964 (216):101–112, 1964.
- Yarotsky, D. Universal approximations of invariant maps by neural networks. *Constructive Approximation*, 55(1): 407–474, 2022.
- Zaheer, M., Kottur, S., Ravanbakhsh, S., Poczos, B., Salakhutdinov, R. R., and Smola, A. J. Deep Sets. In *Advances in Neural Information Processing Systems*, volume 30, 2017.

A. Supplementary Proofs and Definitions

Proof of Proposition 3.3. Suppose that f is a linear permutation equivariant function on a symmetric tensor T in $(\mathbb{R}^n)^{\otimes k}$ to $(\mathbb{R}^n)^{\otimes l}$, where l is either 0 or 1.

Then, for all tuples of indices $J = (j_1, \dots, j_k)$ that index the standard basis $e_J \in (\mathbb{R}^n)^{\otimes k}$, we can partition these indices into disjoint sets such that I and J are in the same set if and only if there is a permutation in S_k such that $\sigma(I) = J$.

Moreover, for each such set, f must map each element in the set to the same value, by the definition of f . Hence, for each set, we can select the unique tuple in the set that is in ascending numerical order.

The union of these tuples across all of the sets gives $S[n]^k$, which indexes the standard basis of $S^k(\mathbb{R}^n)$, by construction. Moreover, since l is either 0 or 1, we see that $\mathbb{R} = S^0(\mathbb{R}^n)$ and $\mathbb{R}^n = S^1(\mathbb{R}^n)$, giving the result. $\square$

Proof of Lemma 5.5. Let $B_O := \{x_\pi\}$ be the set of (k, l) -orbit bipartition diagrams and let $B_D := \{d_\pi\}$ be the set of (k, l) -bipartition diagrams.

To see why the set B_D forms a basis of $SP_k^l(n)$, first we form an ordered set of (k, l) -bipartitions by ordering the bipartitions by the number of blocks that they have from smallest to largest, with any arbitrary ordering allowed for a pair of bipartitions that have the same number of blocks. Call this set S_{l+k} .

Then, because the square matrix that maps elements of B_D to linear combinations of the set B_O — whose rows and columns are indexed (in order) by the ordered set S_{l+k} — is unitriangular by (21), it is therefore invertible. Since B_O is the generating basis of $SP_k^l(n)$, we get that B_D also forms a basis of $SP_k^l(n)$, as required. $\square$

Proof of Theorem 5.7. Consider the two sets $B_D := \{D_\pi\}$ and $B_O := \{X_\theta\}$ where each set is over all (k, l) -bipartitions that have at most n blocks. We see that, by the definition of D_π in (25), the transition matrix from B_D to B_O is unitriangular. As B_O forms a basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ by Theorem 4.8, this implies that B_D is also a basis of $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$. Since B_D is a basis, where each element in B_D is the image under $\Theta_{k,n}^l$ of a diagram basis element in $SP_k^l(n)$, and $\Theta_{k,n}^l$ is linear, the surjectivity of $\Theta_{k,n}^l$ is immediate. $\square$

Proof of Proposition 5.9. Suppose that d_π is a diagram basis element in $SP_k^l(n)$. Assume that we have placed the indices I on the top row of d_π and the indices J on the bottom row of d_π .

Consider the following cases.

Case 1: (I, J) does not label the central nodes of d_π consistently.

Then, by (21) and the definition of the partial ordering $\preceq$ on (k, l) -bipartitions, (I, J) does not label the central nodes of any of the x_θ consistently, since such x_θ are formed diagrammatically by fusing together some of the spiders of x_π at the central nodes. Hence the (I, J) -entry of each X_θ that is the image of x_θ under $\Gamma_{k,n}^l$ is 0, and so the (I, J) -entry of D_π is also 0, by (25).

Case 2: (I, J) labels the central nodes of d_π consistently.

Then, by (21), we see that (I, J) labels the central nodes of exactly one x_θ in (21) consistently and *distinctly*. Indeed, this x_θ is obtained by fusing together the central nodes of all of the spiders of the labelled diagram x_π whose labels are the same. Hence the (I, J) -entry for X_θ that is the image of this x_θ under $\Gamma_{k,n}^l$ is 1, and is 0 for the images of all of the other orbit basis diagrams in (21) that are not x_θ . Hence the (I, J) -entry of D_π is 1, by (25). $\square$

B. How to Generate All (k, l) -Bipartition Diagrams Having at Most n Blocks.

In Theorem 4.8 and Theorem 5.7, we constructed two different bases for $\text{Hom}_{S_n}(S^k(\mathbb{R}^n), S^l(\mathbb{R}^n))$ using the set of all (k, l) -bipartitions having at most n blocks. However, we have yet to outline a method for how to generate all of these bipartitions. To the best of our knowledge, there have only been a few studies on (k, l) -bipartitions restricted to a certain number of blocks (Wright, 1964; Landman et al., 1992; Kim & Hahn, 1997). In particular, no-one has given an explicit procedure for generating all (k, l) -bipartitions having at most n blocks: at best, Kim & Hahn (1997) suggest that one

A General Procedure for Calculating the S_n -Equivariant Weight Matrix from $S^k(\mathbb{R}^n)$ to $S^l(\mathbb{R}^n)$ using the Orbit Basis of $SP_k^l(n)$.

Perform the following steps:

1. Calculate all of the (k, l) -bipartitions π that have at most n blocks, and express each (k, l) -bipartition π as a (k, l) -orbit bipartition diagram x_π in $SP_k^l(n)$.
2. Then, if t is the number of blocks in π , form all possible t -length tuples with elements in $[n]$, **not** allowing for repetitions amongst the elements.
3. For each t -length tuple, label the central green nodes in x_π from left to right with the elements of the tuple. Reorder the spiders in x_π such that they are in increasing order from left to right. Then, for each spider, propagate the block label to the end of each wire. We obtain the matrix unit $E_{I,J}$ from this diagram by setting I to be the labels in the top row and J to be the labels in the bottom row.
4. Since each tuple corresponds to a matrix unit, add together all of the resulting matrix units from all of the tuples to obtain the basis matrix X_π that is associated with x_π . Attach a weight $\lambda_\pi \in \mathbb{R}$ to each matrix X_π . Finally, calculate $\sum \lambda_\pi X_\pi$ to give the overall weight matrix.

can obtain them from the set of all $(k + l, 0)$ -bipartitions having at most n blocks, but they do not explain how to do this explicitly.

In the green- and orange-coloured boxes below, we present a procedure for generating these bipartitions, notably without creating any duplicates. While the procedure is described for (k, l) -bipartition diagrams, it applies equally to (k, l) -orbit bipartition diagrams. Only the choice of diagram that is specified in Step 6 of the green-coloured box changes.

To explain why this procedure works, we need to introduce some definitions.

Definition B.1. An **integer partition** λ of m is defined to be a tuple $(\lambda_1, \lambda_2, \dots, \lambda_t)$ of positive integers λ_i such that $\sum_{i=1}^t \lambda_i = m$ and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_t > 0$. We say that λ consists of t **parts**, and we define its length to be t .

We often extend the tuple of any integer partition λ of m that has some $t < n$ parts to have exactly n parts by appending $n - t$ zeros to λ . Hence $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ of non-negative integers λ_i such that $\sum_{i=1}^n \lambda_i = m$ and $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n \geq 0$.

Example B.2. We see that $\lambda = (4, 2, 2)$ is an integer partition of 8 into exactly 3 parts. If $n = 5$, then we write λ as $(4, 2, 2, 0, 0)$ instead.

Integer partitions must be written in non-decreasing order. The following definition relaxes this requirement.

Definition B.3. A **weak composition** μ of m into n parts is an n -length tuple $(\mu_1, \mu_2, \dots, \mu_n)$ of non-negative integers μ_i such that $\sum_{i=1}^n \mu_i = m$.



We say that a weak composition μ is in **partition order** if it is also an integer partition.

Example B.4. $(2, 0, 4, 0, 2)$ is a weak composition of 8. It is not in partition order. However, sorting these entries into non-decreasing order, $(4, 2, 2, 0, 0)$, gives a weak composition of 8 that is in partition order.

Definition B.5. We can represent integer partitions and weak compositions as **Young diagrams**. For each tuple $(a_1, a_2, \dots, a_n)$, we create a diagram consisting of n rows, where row i consists of a_i boxes (by convention, the row number increases downwards).

Note that if a_i is 0 for some i , we leave a row's worth of space for it and move onto the next value in the tuple.

Example B.6. We represent the integer partition $(4, 2, 2, 0, 0)$ and the weak composition $(2, 0, 4, 0, 2)$ as Young diagrams:



(33)

Definition B.7. Let λ be an integer partition of q having exactly n parts, where we allow zero entry parts. Let μ be a weak composition of m into n parts. We say that μ **fits inside** λ if $\lambda_i \geq \mu_i$ for all $i \in [n]$.

We can see this diagrammatically by considering the Young diagrams for each of λ and μ : if we can shade μ in λ without going outside the boundary of λ , then μ fits inside λ , otherwise it does not.

Example B.8. If λ is the integer partition $(6, 4, 4, 2, 0, 0)$ and μ is the weak composition $(2, 3, 0, 2, 0, 0)$, then we see that μ fits inside λ :




However, the weak composition $(3, 5, 3, 2, 0, 1)$ does not fit inside λ , as indicated by the red boxes in the diagram below, which fall outside the boundary of λ :



The key insight behind why our procedure generates all (k, l) -bipartition diagrams having at most n blocks can be stated as follows: every integer partition of $k + l$ having exactly n parts (allowing for zero element parts) can be expressed as a $(k + l, 0)$ -bipartition, and every (k, l) -bipartition having some $t \in [n]$ blocks must come from a integer partition of $k + l$ having exactly t parts. In particular, we can then represent each $(k + l, 0)$ -bipartition as a $(k + l, 0)$ -bipartition diagram (or as a $(k + l, 0)$ -orbit bipartition diagram).

Indeed, if $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_n)$ is an integer partition of $k + l$ having exactly n parts, with zero element parts allowed, then, by first removing all zero element parts to give the integer partition $\lambda = (\lambda_1, \lambda_2, \dots, \lambda_t)$ of $k + l$ for some $t \in [n]$ such that λ_i is positive for all $i \in [t]$, we obtain the $(k + l, 0)$ -bipartition $\lambda = \{[\lambda_1, 0], [\lambda_2, 0], \dots, [\lambda_t, 0]\}$ having precisely t blocks. (Remark: we remove all zero element parts in λ first since the definition of a bipartition does not allow parts $[x_i, y_i]$ where $x_i = 0 = y_i$.) We can then represent the $(k + l, 0)$ -bipartition λ as a one-row $(k + l, 0)$ -bipartition diagram



or with open green circles for a $(k + l, 0)$ -orbit bipartition diagram. We retain the order of the integer partition in the diagram; that is, the number of legs is decreasing from left to right in each spider.

To see why every (k, l) -bipartition having t blocks must come from a integer partition of $k + l$ having exactly t parts, let $\pi = \{[x_1, y_1], [x_2, y_2], \dots, [x_t, y_t]\}$ be a (k, l) -bipartition having t blocks. We form a $(k + l, 0)$ -bipartition by subtracting, for each part i , y_i from the second element of the bracket and then adding it to the first element of the bracket. This gives the $(k + l, 0)$ -bipartition $\{[x_1 + y_1, 0], [x_2 + y_2, 0], \dots, [x_t + y_t, 0]\}$. Without loss of generality, we can assume that the parts in this $(k + l, 0)$ -bipartition have been reordered in decreasing size order from left to right. Hence, this $(k + l, 0)$ -bipartition is precisely the integer partition $(x_1 + y_1, x_2 + y_2, \dots, x_t + y_t)$ of $k + l$ into exactly t parts.

Hence, to generate all (k, l) -bipartitions having at most n blocks, we first need to generate all integer partitions of $k + l$ having exactly n parts, with zero element parts allowed, and express them as $(k + l, 0)$ -bipartitions (or as bipartition diagrams). We then need to work out how to create all possible (k, l) -bipartitions (or diagrams) from each such $(k + l, 0)$ -bipartition (diagram), which is the reverse procedure of the process described in the previous paragraph.

We see that, for a given $(k + l, 0)$ -bipartition diagram having t blocks, where $t \in [n]$, to obtain a (k, l) -bipartition diagram, we need to turn l legs up across all of the t spiders. This is the same as creating a weak composition μ of l into exactly n

parts, where we set the last $n - t$ parts in μ to 0. Hence, calculating all weak compositions of l into exactly n parts will be enough to determine all of the (k, l) -bipartition diagrams from a given $(k + l, 0)$ -bipartition diagram having at most n blocks.

In effect, this produces a matrix where the integer partitions λ of $k + l$ into exactly n parts, with zero element parts allowed, index the rows, and the weak compositions μ of l into exactly n parts index the columns. We then need to see, for each (λ, μ) pair, whether the elements of μ are a well-fitting combination for turning up the legs of λ . We call such a pair **valid**. A simple way to determine if each (λ, μ) pair is valid is to see whether μ fits inside λ . If it does, then the pair is valid, otherwise it is not. Then, for each valid pair (λ, μ) , we can create a (k, l) -bipartition diagram by turning up μ_i legs in each spider i of λ , viewed as a $(k + l, 0)$ -bipartition diagram.

Hence, if $\mu = (\mu_1, \mu_2, \dots, \mu_t, 0, \dots, 0)$ is a weak composition of l into exactly n parts such that (λ, μ) is a valid pair, where λ is the $(k + l, 0)$ -bipartition diagram given in (36), then, for each spider i in (36), we can turn up μ_i legs to obtain the (k, l) -bipartition diagram

(37)

At this stage, this would be enough to generate all (k, l) -bipartition diagrams having at most n blocks. However, this procedure in its current state would create potentially multiple duplicate (k, l) -bipartition diagrams from any $(k + l, 0)$ -bipartition diagram where spiders have the same number of legs; that is, this procedure will overgenerate the (k, l) -bipartition diagrams having at most n blocks that we need. Indeed, if block i and block j have the same number of legs, that is, $\lambda_i = \lambda_j$, then, if μ^1 is the weak composition with entry μ_i in position i and entry μ_j in position j such that (λ, μ^1) is valid, and μ^2 is the weak composition with entry μ_j in position i and entry μ_i in position j such that (λ, μ^2) is valid, with the entries of μ^1 and μ^2 being equal at all other indices in the tuples, we see that using μ^1 to turn up the legs of λ gives the same (k, l) -bipartition diagram as using μ^2 to turn up the legs of λ , even though the spiders might be in a different order.

Hence, we can improve this procedure so that only the exact number of (k, l) -bipartition diagrams having at most n blocks are generated, in the following way: after the matrix that is indexed by (λ, μ) pairings is created, perform what we have termed the (λ, μ) Duplication Test to eliminate all pairings that fail the test first before checking the rest for validity.

The (λ, μ) Duplication Test does the following: to avoid duplicating (k, l) -bipartition diagrams from a given $(k + l, 0)$ -bipartition diagram λ , for spiders whose legs have the same size, we only need to turn up the legs on these spiders in decreasing order, since any other order will produce duplicate diagrams. That is, in the subtuple of a weak composition formed by taking those elements of the weak composition at the same indices as the blocks of the spiders, we only need to consider the subtuple that is in partition order. Hence we can set any (λ, μ) pairing that fails this Duplication Test to be invalid.

We can also speed up the process of determining whether a (λ, μ) pair is valid or not as follows. Suppose that λ has non-zero entries only in its first t parts. If μ contains a non-zero entry in any index from $t + 1$ to n , then this pair is automatically invalid. Hence, for a given λ having non-zero entries only in its first t parts, we only need to consider the validity of (λ, μ) pairings where μ is of the form $\mu = (\mu_1, \mu_2, \dots, \mu_t, 0, \dots, 0)$, with the μ_i being non-negative integers. We perform this before the (λ, μ) Duplication Test.

There is one further optimisation that we can make to the entire procedure. As it stands, we chose to turn l legs up in a $(k + l, 0)$ -bipartition diagram and found all possible ways of doing this by considering weak combinations of l into exactly n parts. In fact, we could choose instead to not turn up k legs in total, which would be the same as considering all weak combinations of k into exactly n parts. As the number of weak combinations of a number m into n parts increases as m increases, by considering the sizes of k and l , we can reduce the number of weak combinations that we need to consider for each $(k + l, 0)$ -bipartition diagram by choosing the smaller value between k and l .

Example B.9. Suppose that we would like to find all $(3, 2)$ -bipartition diagrams having at most 3 blocks. We follow the procedure that is given in the green- and orange-coloured boxes below.

Procedure: How to Generate All (k, l) -Bipartition Diagrams Having at Most n Blocks.

1. Calculate all integer partitions λ of $k + l$ into at most n parts, and express each integer partition as an n -length tuple whose elements are in decreasing order, allowing zero element parts.
2. Choose the smaller value between k and l . Let $m = \min(k, l)$. Calculate all weak compositions μ of m into exactly n parts.
3. Create a matrix where the integer partitions λ index the rows and the weak compositions μ index the columns.
4. Put 0 in any (λ, μ) -entry where μ has a non-zero entry in any index that is greater than the number of parts in λ .
5. Now consider only those λ that have parts that are equal. For each such λ , put 0 in each remaining (λ, μ) -entry that fails the (λ, μ) Duplication Test.
6. For each remaining (λ, μ) -entry, determine whether μ fits inside λ , viewing each as a Young diagram. If it does, put 1 in that entry, and call the pair (λ, μ) **valid**. Otherwise, put 0 in that entry.
7. For each valid (λ, μ) pair, remove all zero element parts of λ and express it as a $(k + l, 0)$ -bipartition diagram. Hence, each block i in λ is a spider that corresponds to the pair $[\lambda_i, 0]$, viewing λ as a $(k + l, 0)$ -bipartition. Then:
 - If $m = l$, then, for each block i in λ , turn μ_i legs up to the top row.
 - Otherwise ($m = k$), for each block i in λ , turn $\lambda_i - \mu_i$ legs up to the top row.

Note that this procedure also works for generating (k, l) -orbit bipartition diagrams; simply express each λ in Step 6 as a (k, l) -orbit bipartition diagram instead.

Procedure: The (λ, μ) Duplication Test.

1. Suppose that λ has t non-zero parts, where $t \in [n]$. Partition $[t]$ into blocks, where two elements i, j are in the same block if and only if $\lambda_i = \lambda_j$.
2. For each block in the partition whose size is greater than 1, create a subtuple from μ by only considering those elements in μ that are indexed by the elements of the block.
3. Consider each subtuple in turn. If a subtuple is not in partition order, then we say that (λ, μ) has **failed the duplication test**.
4. Otherwise, we say that (λ, μ) has **passed the duplication test**.

Step 1: We calculate all integer partitions λ of $3 + 2 = 5$ into at most $n = 3$ parts, and express each as a 3-length tuple whose elements are in decreasing order, allowing zero element parts.

They are: $(5, 0, 0)$, $(4, 1, 0)$, $(3, 2, 0)$, $(3, 1, 1)$, and $(2, 2, 1)$.

Step 2: As $l = 2 < 3 = k$, we calculate all weak compositions μ of $m = \min(3, 2) = 2$ into exactly $n = 3$ parts.

They are: $(2, 0, 0)$, $(0, 2, 0)$, $(0, 0, 2)$, $(1, 1, 0)$, $(1, 0, 1)$, and $(0, 1, 1)$.

Steps 3, 4 and 5: We create a matrix where the integer partitions λ index the rows and the weak compositions μ index the columns. We put 0 in any (λ, μ) -entry where μ has a non-zero entry in any index that is greater than the number of parts in λ . We then look to find any remaining (λ, μ) -entries that fail the Duplication Test. For this test, we only need to consider those λ that have repeating parts: they are $(3, 1, 1)$ and $(2, 2, 1)$.

Consider $\lambda = (3, 1, 1)$ and $\mu = (1, 0, 1)$. We form the partition $\{1 \mid 2, 3\}$ of $[3]$ from λ , and form the subtuple $(0, 1)$ of μ from the only block in the partition of size greater than 1. As this subtuple is not in partition order, we see that this pair (λ, μ) has failed the Duplication Test.

Likewise, consider $\lambda = (2, 2, 1)$ and $\mu = (0, 1, 1)$. We form the partition $\{1, 2 \mid 3\}$ of $[3]$ from λ , and form the subtuple $(0, 1)$ of μ from the only block in the partition of size greater than 1. As this subtuple is not in partition order, we see that this pair (λ, μ) has failed the Duplication Test.

For these two pairs, we put 0 in the (λ, μ) -entry of the matrix. All other pairs pass the Duplication Test.

Step 6: For each remaining (λ, μ) -entry, we determine whether μ fits inside λ , viewing each as a Young diagram.

At the end of this step, we see that the matrix contains the following entries:

$$\begin{array}{c}
 \begin{array}{cccccc}
 (2,0,0) & (0,2,0) & (0,0,2) & (1,1,0) & (1,0,1) & (0,1,1)
 \end{array} \\
 \begin{array}{c}
 (5,0,0) \\
 (4,1,0) \\
 (3,2,0) \\
 (3,1,1) \\
 (2,2,1)
 \end{array}
 \begin{bmatrix}
 1 & 0 & 0 & 0 & 0 & 0 \\
 1 & 0 & 0 & 1 & 0 & 0 \\
 1 & 1 & 0 & 1 & 0 & 0 \\
 1 & 0 & 0 & 1 & 0 & 1 \\
 1 & 1 & 0 & 1 & 1 & 0
 \end{bmatrix}
 \end{array} \tag{38}$$

At this point, we can see that $p_3(3, 2) = 13$.

Step 7: We form the $(3, 2)$ -bipartition diagram that is associated with each valid (λ, μ) pair.

They are:

$$\begin{array}{c}
 \begin{array}{cccccc}
 (2,0,0) & (0,2,0) & (0,0,2) & (1,1,0) & (1,0,1) & (0,1,1)
 \end{array} \\
 \begin{array}{c}
 (5,0,0) \\
 (4,1,0) \\
 (3,2,0) \\
 (3,1,1) \\
 (2,2,1)
 \end{array}
 \begin{bmatrix}
 \begin{array}{c} \diagup \quad \diagdown \\ \bullet \\ \diagdown \quad \diagup \end{array} & 0 & 0 & 0 & 0 & 0 \\
 \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & 0 & \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & 0 \\
 \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & 0 \\
 \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & 0 & \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} \\
 \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \quad \bullet \quad \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & 0 & \begin{array}{c} \diagup \quad \diagdown \\ \bullet \quad \bullet \quad \bullet \quad \bullet \quad \bullet \\ \diagdown \quad \diagup \end{array} & 0 & 0
 \end{bmatrix}
 \end{array} \tag{39}$$

To finish this section, we note that there exists a more generic partition function, $p(k, l)$, which counts the number of possible

(k, l) -bipartitions without any restriction on the number of blocks. In particular, the sequence $p(k, k)$ (i.e $k = l$) is the OEIS sequence A002774. Bipartitions without any restriction on the number of blocks have been studied in many prior works (MacMahon, 1893; 1896; Mathews, 1896; MacMahon, 1899; 1918; Auluck, 1953; Nanda, 1957; Wright, 1957; 1958; Carlitz, 1963; Carlitz & Roselle, 1966; Andrews, 1977; 1998). It is fair to say that most of these authors, with the exception of MacMahon and Mathews, primarily focused on determining either generating functions or asymptotic formulas for $p(k, l)$.

We see from our analysis that $p(k, l)$ is equal to $p_{k+l}(k, l)$, the number of (k, l) -bipartitions having at most $k + l$ parts. We give the value of $p(k, l)$ up to $k = l = 5$ in Table 3. Note that $p(k, l)$ is symmetric in k and l , and that $p(k, 0) = p(0, k)$ is precisely the number of integer partitions of k (into at most k parts).

Table 3. The number of (k, l) -bipartitions: $p(k, l)$


$k \backslash l$	0	1	2	3	4	5
0	1	1	2	3	5	7
1	1	2	4	7	12	19
2	2	4	9	16	29	47
3	3	7	16	31	57	97
4	5	12	29	57	109	189
5	7	19	47	97	189	339

(40)

C. Supplementary Weight Matrix Calculations

Example C.1 (Calculations for Example 4.11). Recall that we would like to find the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3) = \mathbb{R}^3$ using the orbit basis of $SP_2^1(3)$.


We need to consider the $(2, 1)$ -orbit bipartition diagrams having at most $n = 3$ blocks. They are



For example, for the second diagram which has two blocks, we see that, since $n = 3$, there are only six possible 2-length tuples having distinct elements in $[3]$:

$$(1, 2); (1, 3); (2, 1); (2, 3); (3, 1); (3, 2) \quad (42)$$

Assigning these tuples to the central green nodes of the diagram and reordering the spiders where appropriate gives



By propagating the central values to the ends of the wires in each diagram, then reading the I tuple off the top row and the J tuple off the bottom row to form the matrix unit $E_{I,J}$ for each diagram, and then adding these matrix units together gives the following orbit basis matrix X_π corresponding to the second diagram in (41):

$$\begin{matrix} & 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \end{bmatrix} \end{matrix} \quad (44)$$

Hence, in full, from left to right, the four diagrams in (41) correspond to the following basis matrices X_π of size 3×6 :

$$\begin{matrix} & 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \end{matrix} ; \begin{matrix} & 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 0 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 0 \end{bmatrix} \end{matrix} \quad (45)$$

$$\begin{matrix} & 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 0 \end{bmatrix} \end{matrix} ; \begin{matrix} & 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \end{bmatrix} \end{matrix} \quad (46)$$


Hence, the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3)$ using the orbit basis of $SP_2^1(3)$ is of the form

$$\begin{matrix} & 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} \lambda_1 & \lambda_3 & \lambda_3 & \lambda_2 & \lambda_4 & \lambda_2 \\ \lambda_2 & \lambda_3 & \lambda_4 & \lambda_1 & \lambda_3 & \lambda_2 \\ \lambda_2 & \lambda_4 & \lambda_3 & \lambda_2 & \lambda_3 & \lambda_1 \end{bmatrix} \end{matrix} \quad (47)$$

for weights $\lambda_1, \lambda_2, \lambda_3, \lambda_4 \in \mathbb{R}$.

Example C.2 (Calculations for Example 5.10). Recall that we would like to find the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3) = \mathbb{R}^3$ using the diagram basis of $SP_2^1(3)$ instead.


We need to consider the $(2, 1)$ -bipartition diagrams having at most $n = 3$ blocks. They are




For example, for the third diagram which has two blocks, we see that, since $n = 3$, there are now nine possible 2-length tuples whose elements are in $[3]$, since the elements in each tuple no longer need to be distinct:

$$(1, 1); (1, 2); (1, 3); (2, 1); (2, 2); (2, 3); (3, 1); (3, 2); (3, 3) \quad (49)$$

Assigning these tuples to the central green nodes of the diagram, reordering the spiders where appropriate, and fusing together any spiders whose central green nodes are labelled with the same value gives



and



By propagating the central values to the ends of the wires in each diagram, then reading the I tuple off the top row and the J tuple off the bottom row to form the matrix unit $E_{I,J}$ for each diagram, and then adding these matrix units together gives the following diagram basis matrix D_π corresponding to the third diagram in (48):

$$\begin{matrix} & \begin{matrix} 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} \end{matrix} \quad (52)$$

Hence, in full, from left to right, the four diagrams in (41) correspond to the following basis matrices D_π of size 3×6 :

$$\begin{matrix} & \begin{matrix} 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \end{matrix} ; \begin{matrix} & \begin{matrix} 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 & 0 & 1 \end{bmatrix} \end{matrix} \quad (53)$$

$$\begin{matrix} & \begin{matrix} 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 1 & 1 & 0 \\ 0 & 0 & 1 & 0 & 1 & 1 \end{bmatrix} \end{matrix} ; \begin{matrix} & \begin{matrix} 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \end{matrix} \quad (54)$$

Hence, the S_3 -equivariant weight matrix from $S^2(\mathbb{R}^3)$ to $S^1(\mathbb{R}^3)$ using the diagram basis of $SP_2^1(3)$ is of the form

$$\begin{matrix} & \begin{matrix} 1,1 & 1,2 & 1,3 & 2,2 & 2,3 & 3,3 \end{matrix} \\ \begin{matrix} 1 \\ 2 \\ 3 \end{matrix} & \begin{bmatrix} \lambda_{1,2,3,4} & \lambda_{3,4} & \lambda_{3,4} & \lambda_{2,4} & \lambda_4 & \lambda_{2,4} \\ \lambda_{2,4} & \lambda_{3,4} & \lambda_4 & \lambda_{1,2,3,4} & \lambda_{3,4} & \lambda_{2,4} \\ \lambda_{2,4} & \lambda_4 & \lambda_{3,4} & \lambda_{2,4} & \lambda_{3,4} & \lambda_{1,2,3,4} \end{bmatrix} \end{matrix} \quad (55)$$

for weights $\lambda_1, \lambda_2, \lambda_3, \lambda_4 \in \mathbb{R}$, where $\lambda_A := \sum_{i \in A} \lambda_i$.

Example C.3 (Deep Sets S_n -Equivariant Weight Matrix from $\mathbb{R}^n$ to $\mathbb{R}^n$). To find the Deep Sets S_n -equivariant weight matrix from $\mathbb{R}^n$ to $\mathbb{R}^n$, we need to find the diagram basis of $\text{Hom}_{S_n}(S^1(\mathbb{R}^n), S^1(\mathbb{R}^n))$ from the diagram basis of $SP_1^1(n)$, since $S^1(\mathbb{R}^n) = \mathbb{R}^n$.

We need to consider all $(1, 1)$ -bipartition diagrams having at most n blocks. Assuming that $n > 1$, they are



For the first diagram, the only valid 1-length tuples are the elements of $[n]$. Hence, the matrix that corresponds to this $(1, 1)$ -bipartition diagram is the $n \times n$ identity matrix. For the second diagram, all possible 2-length tuples with elements from $[n]$ are valid, since the elements in each tuple do not need to be distinct. Hence, the matrix that corresponds to this $(1, 1)$ -bipartition diagram is the $n \times n$ all ones matrix.

Hence, the (i, j) -entry of the S_n -equivariant weight matrix from $\mathbb{R}^n$ to $\mathbb{R}^n$ is $\lambda_1 + \lambda_2$ if $i = j$, and is λ_2 if $i \neq j$, for weights $\lambda_1, \lambda_2 \in \mathbb{R}$. This is precisely the Deep Sets characterisation given in [Zaheer et al. \(2017\)](#).

D. Map Label Implementation and Examples

At the end of the paper, we present complete implementations of the five subprocedures (I to V) that appear in the procedure for generating the transformation map labels that correspond to a (k, l) -bipartition diagram for a given value of n . We provide illustrative examples of each of the subprocedures below.

D.1. Subprocedure I

Example D.1. Consider the $(3, 2)$ -bipartition diagram d_π given by



We follow Subprocedure I to obtain all possible grouped outputs for d_π and, for each one, its associated set of partially labelled diagrams. (In this case, the resulting diagrams will also be fully labelled, that is, there would be no need to apply Subprocedure II to the outcome).

Step 1: We sort the spiders in decreasing order of output wires, breaking ties by sorting them in decreasing order of input wires. This gives



Step 2: We partition the spiders into groupings where spiders are in the same grouping if and only if they have the same number of output wires. Labelling each spider with O_x , where x denotes the number of output wires that the spider has, we see that, for the diagram given in (58), the partition π is $\{O_1, O_1\}$, which consists of only one grouping, $\pi_B = \{O_1, O_1\}$.

Step 3: For π_B , we form a new partition of the spiders such that spiders are in the same grouping in the new partition $\pi_{B,I}$ if and only if they have the same number of input wires. Labelling each spider in π_B with I_x , where x denotes the number of input wires that the spider has, we see that $\pi_{B,I}$ is $\{I_2 \mid I_1\}$.


Step 4: We calculate all integer partitions of t , where t is the number of spiders in d_π that have output wires. Here $t = 2$, since both spiders in (58) have output wires. The integer partitions of 2 are $1 + 1$ and 2 only.

Step 5: Consider each integer partition λ in turn.

a): For $\lambda = 1 + 1$, we form the multiset $\{i, j\}$ and compute all partitions of this multiset that correspond to the original partition structure of π , modulo isomorphism. There is only one non-isomorphic partition of the multiset, namely $L = \{i, j\}$ itself, which has only one grouping, $L_B = \{i, j\}$.

We now compute all partitions of the grouping L_B that correspond to the partition structure of $\pi_{B,I}$. This gives $\{i \mid j\}$ and $\{j \mid i\}$, which we convert into ordered tuples, that is, $T_{L,B} = \{(i, j), (j, i)\}$.

Since L had only one grouping, we see that $T_{1+1,L} = \{(i, j), (j, i)\}$. Hence the set $D_{1+1,L}$ consists only of the following two diagrams



where, in each diagram, we are unable to merge any spiders. Hence $R_{1+1,L} = \emptyset$, and so $P_{1+1,L} = D_{1+1,L}$.

Moreover, since there was only one non-isomorphic partition of the multiset $\{i, j\}$, we see that $P_{1+1} = P_{1+1,L}$.


b): For $\lambda = 2$, we form the multiset $\{i, i\}$ and compute all partitions of this multiset that correspond to the original partition structure of π , modulo isomorphism. There is only one non-isomorphic partition of the multiset, namely $L = \{i, i\}$ itself, which has only one grouping, $L_B = \{i, i\}$.

We now compute all partitions of the grouping L_B that correspond to the partition structure of $\pi_{B,I}$. This gives $\{i \mid i\}$ only, which we convert into an ordered tuple, that is, $T_{L,B} = \{(i, i)\}$.

Since L had only one grouping, we see that $T_{2,L} = \{(i, i)\}$. Hence the set $D_{2,L}$ consists only of the following diagram



We merge spiders that are labelled with the same label to form:


(61)


It is clear that, despite the new grouping, no new partially labelled diagrams can appear from (61), and so we get that $R_{2,L} = \emptyset$. Hence $P_{2,L} = D_{2,L}$.

Moreover, since there was only one non-isomorphic partition of the multiset $\{i, i\}$, we see that $P_2 = P_{2,L}$.

Step 6: We now take the set union of P_λ over all integer partitions λ of $t = 2$, and call the new set P . Hence P has three elements, namely the two diagrams given in (59) and the one diagram given in (61). By reordering the labelled spiders in lexicographical order, and partitioning P into subsets such that diagrams are in the same subset if and only if their tuples of output labels are the same, we obtain the following two sets as output of Subprocedure I for the $(3, 2)$ -bipartition diagram d_π :

$$(i, j) = \left\{ \begin{array}{c} \text{diagram 1} \end{array} \right\} \quad \text{and} \quad (i, i) = \left\{ \begin{array}{c} \text{diagram 2} \end{array} \right\} \quad (62)$$

Example D.2. Consider instead the $(1, 4)$ -bipartition diagram d_π given by


(63)

We follow Subprocedure I to obtain all possible grouped outputs for d_π and, for each one, its associated set of partially labelled diagrams. (Once again, the resulting diagrams will also be fully labelled, that is, there would be no need to apply Subprocedure II to the outcome).

Step 1: We sort the spiders in decreasing order of output wires, breaking ties by sorting them in decreasing order of input wires. We see that the spiders in (63) are already sorted.

Step 2: We partition the spiders into groupings where spiders are in the same grouping if and only if they have the same number of output wires. Labelling each spider with O_x , where x denotes the number of output wires that the spider has, we see that, for the diagram given in (63), the partition π is $\{O_2 \mid O_1, O_1\}$, which consists of two groupings, $\pi_{B_1} = \{O_2\}$ and $\pi_{B_2} = \{O_1, O_1\}$.


Step 3: For each grouping π_{B_i} , we form a new partition of the spiders such that spiders are in the same grouping in the new partition $\pi_{B_i, I}$ if and only if they have the same number of input wires. Labelling each spider in π_{B_i} with I_x , where x denotes the number of input wires that the spider has, we see that $\pi_{B_1, I} = \{I_1\}$ and $\pi_{B_2, I} = \{I_0, I_0\}$.

Step 4: We calculate all integer partitions of t , where t is the number of spiders in d_π that have output wires. Here $t = 3$, since all three spiders in (63) have output wires. The integer partitions of 3 are $1 + 1 + 1$, $2 + 1$ and 3 .

Step 5: Consider each integer partition λ in turn.

a): For $\lambda = 1 + 1 + 1$, we form the multiset $\{i, j, k\}$ and compute all partitions of this multiset that correspond to the original partition structure of π , modulo isomorphism. Despite there being three partitions, namely $\{i \mid j, k\}$, $\{j \mid i, k\}$, and $\{k \mid i, j\}$, all three are equivalent under a relabelling of the labels. Hence, there is only one non-isomorphic partition of the multiset, namely $L = \{i \mid j, k\}$, which has two groupings: $L_{B_1} = \{i\}$ and $L_{B_2} = \{j, k\}$.

For each grouping L_{B_i} , we now compute all partitions of the grouping L_{B_i} that correspond to the partition structure of $\pi_{B_i, I}$. For L_{B_1} , this gives $\{i\}$ only, and for L_{B_2} , this gives $\{j, k\}$ only. Converting these into ordered tuples, we get that $T_{L, B_1} = \{(i)\}$ and $T_{L, B_2} = \{(j, k)\}$. Hence $T_{1+1+1, L}$ is $\{(i, j, k)\}$. Consequently, the set $D_{1+1+1, L}$ consists only of the diagram


(64)

Since we are unable to merge any spiders, we have that $R_{1+1+1,L} = \emptyset$, and so $P_{1+1+1,L} = D_{1+1+1,L}$.

Moreover, since there was only one non-isomorphic partition of the multiset $\{i, j, k\}$, we see that $P_{1+1+1} = P_{1+1+1,L}$.

b): For $\lambda = 2 + 1$, we form the multiset $\{i, i, j\}$ and compute all partitions of this multiset that correspond to the original partition structure of π , modulo isomorphism. There are two non-isomorphic partitions of the multiset, namely $L^1 = \{i \mid i, j\}$ and $L^2 = \{j \mid i, i\}$. L^1 has two groupings, $L_{B_1}^1 = \{i\}$ and $L_{B_2}^1 = \{i, j\}$, and L^2 has two groupings, $L_{B_1}^2 = \{j\}$ and $L_{B_2}^2 = \{i, i\}$.

i): Consider L^1 . For each grouping $L_{B_i}^1$, we now compute all partitions of the grouping $L_{B_i}^1$ that correspond to the partition structure of $\pi_{B_i,I}$. For $L_{B_1}^1$, this gives $\{i\}$ only, and for $L_{B_2}^1$, this gives $\{i, j\}$ only. Converting these into ordered tuples, we get that $T_{L^1,B_1} = \{(i)\}$ and $T_{L^1,B_2} = \{(i, j)\}$. Hence T_{2+1,L^1} is $\{(i, i, j)\}$. Consequently, the set D_{2+1,L^1} consists only of the diagram



Merging spiders that are labelled with the same label gives



It is clear that, despite the new grouping, no new partially labelled diagrams can appear from (66), since there is only one spider with three output wires and one spider with one output wire. Hence, we get that $R_{2+1,L^1} = \emptyset$, and so $P_{2+1,L^1} = D_{2+1,L^1}$.

ii): Now consider L^2 . For each grouping $L_{B_i}^2$, we now compute all partitions of the grouping $L_{B_i}^2$ that correspond to the partition structure of $\pi_{B_i,I}$. For $L_{B_1}^2$, this gives $\{j\}$ only, and for $L_{B_2}^2$, this gives $\{i, i\}$ only. Converting these into ordered tuples, we get that $T_{L^2,B_1} = \{(j)\}$ and $T_{L^2,B_2} = \{(i, i)\}$. Hence T_{2+1,L^2} is $\{(j, i, i)\}$. Consequently, the set D_{2+1,L^2} consists only of the diagram



Merging spiders that are labelled with the same label gives



However, this time, with the new grouping, we will form new partially labelled diagrams from (68). Let the newly merged bipartition diagram in (68) be d_θ . We retain but ignore its labels for now. We see that d_θ is already sorted in decreasing order on output wires, with ties broken by input wires, and so we partition the spiders into groupings based on their output wires. Hence $\theta = \{O_2, O_2\}$. Moreover, for the single individual grouping in θ , θ_B , we form a new partition of θ_B such that spiders are in the same grouping if and only if they have the same number of input wires. Hence we have that $\theta_{B,I} = \{I_1 \mid I_0\}$. Removing the labels from θ , we obtain a partition of labels $\tilde{L}$ into individual groupings $\tilde{L}_B$. Hence $\tilde{L}_B = \{j, i\}$. We now compute all partitions of $\tilde{L}_B$ that correspond to the partition structure of $\theta_{B,I}$. This gives $\{j \mid i\}$ and $\{i \mid j\}$, which we convert into ordered tuples. Hence $T_{\tilde{L},B} = \{(j, i), (i, j)\}$, and since there is only one grouping in $\tilde{L}$, we have that $T_{1+1,\tilde{L}} = \{(j, i), (i, j)\}$. Creating a copy of (68) without its labels for each element l of $T_{1+1,\tilde{L}}$, we label it with l . This gives back (68) and a new diagram




both of which we add to R_{2+1,L^2} .

Taking the set union of D_{2+1,L^2} and R_{2+1,L^2} gives P_{2+1,L^2} , which in this case is equal to R_{2+1,L^2} .


As there were two non-isomorphic partitions of the multiset $\{i, i, j\}$, we see that $P_{2+1} = P_{2+1,L^1} \cup P_{2+1,L^2}$.

c): For $\lambda = 3$, we form the multiset $\{i, i, i\}$ and compute all partitions of this multiset that correspond to the original partition structure of π , modulo isomorphism. Hence, there is only one non-isomorphic partition of the multiset, namely $L = \{i \mid i, i\}$, which has two groupings: $L_{B_1} = \{i\}$ and $L_{B_2} = \{i, i\}$.

For each grouping L_{B_i} , we now compute all partitions of the grouping L_{B_i} that correspond to the partition structure of $\pi_{B_i,I}$. For L_{B_1} , this gives $\{i\}$ only, and for L_{B_2} , this gives $\{i, i\}$ only. Converting these into ordered tuples, we get that $T_{L,B_1} = \{(i)\}$ and $T_{L,B_2} = \{(i, i)\}$. Hence $T_{3,L}$ is $\{(i, i, i)\}$. Consequently, the set $D_{3,L}$ consists only of the diagram


(70)

Merging spiders that are labelled with the same label gives


(71)

It is clear that, despite the new grouping, no new partially labelled diagrams can appear from (71), since there is only one spider with four output wires. Hence, we get that $R_{3,L} = \emptyset$, and so $P_{3,L} = D_{3,L}$.

Moreover, since there was only one non-isomorphic partition of the multiset $\{i, i, i\}$, we see that $P_3 = P_{3,L}$.

Step 6: We now take the set union of P_λ over all integer partitions λ of $t = 3$, and call the new set P . Hence P has five elements, namely the diagrams given in (64), (66), (68), (69), and (71). By reordering the labelled spiders in lexicographical order, and partitioning P into subsets such that diagrams are in the same subset if and only if their tuples of output labels are the same, we obtain the following four sets as output of Subprocedure I for the $(1, 4)$ -bipartition diagram d_π :

$$(i, i, j, k) = \left\{ \begin{array}{c} \text{Diagram with 4 output wires: first two are 'i' (blue), last two are 'j' (blue) and 'k' (blue)} \end{array} \right\} \quad (72)$$


$$(i, i, i, j) = \left\{ \begin{array}{c} \text{Diagram with 4 output wires: first three are 'i' (blue), last is 'j' (blue)} \end{array} \right\} \quad (73)$$

$$(i, i, j, j) = \left\{ \begin{array}{c} \text{Diagram with 4 output wires: first two are 'i' (blue), last two are 'j' (blue) and 'j' (blue)} \end{array} \right\} \quad (74)$$

$$(i, i, i, i) = \left\{ \begin{array}{c} \text{Diagram with 4 output wires: all four are 'i' (blue)} \end{array} \right\} \quad (75)$$

D.2. Subprocedure II

Example D.3. Consider the partially labelled diagram for the grouped output (i, i, j) of the $(4, 3)$ -bipartition diagram d_π


(76)

We follow Subprocedure II to obtain all possible fully labelled diagrams for this particular partially labelled diagram.

Step 1: We see that the partially labelled diagram has $t = 2$ labelled spiders, and also $s = 3$ spiders with no label. Hence L_X , the set of fixed labels, is $\{i, j\}$.

Step 2: We see that the unlabelled spiders in (76) are already sorted in decreasing order of input wires.

Step 3: We partition the sorted unlabelled spiders into groupings where spiders are in the same grouping if and only if they have the same number of input wires. We see that $\pi = \{I_2 \mid I_1, I_1\}$, which has two individual groupings $\pi_{B_1} = \{I_2\}$ and $\pi_{B_2} = \{I_1, I_1\}$.

Step 4: We consider all possible ways of labelling the unlabelled spiders, where k of them are fixed and the remaining $s - k$ of them are free, for $k = 0, 1, 2, 3$.

Consider each value of k in turn:

i) $k = 0$: 0 fixed labels and 3 free labels.


First, we create the multiset $\{F, F, F\}$, where F denotes a free symbol. We compute all partitions of this multiset that correspond to the original partition structure of π . In this case, there is only one such partition, and so $\Pi_{X,F} = \{\{F \mid F, F\}\}$.

Next, we calculate all possible size 0 multisets choosing labels from L_X , allowing repeats, and all possible size 3 multisets choosing free labels, introducing them as needed, and also allowing repeats. This gives $\{k, l, m\}$, $\{k, k, l\}$ and $\{k, k, k\}$. We consider each of these in turn:

- we compute all partitions of $\{k, l, m\}$ that correspond to the original partition structure of the only partition in $\Pi_{X,F}$, namely $\{F \mid F, F\}$, modulo isomorphism. Despite there being three partitions, namely $\{k \mid l, m\}$, $\{l \mid k, m\}$, and $\{m \mid k, l\}$, all three are equivalent under a relabelling of the free labels. Hence there is only one non-isomorphic partition of the multiset, namely $\{k \mid l, m\}$ only. Creating a copy of (76), we label the spiders in order from left to right to create the fully labelled diagram



- we compute all partitions of $\{k, k, l\}$ that correspond to the original partition structure of the only partition in $\Pi_{X,F}$, namely $\{F \mid F, F\}$, modulo isomorphism. This gives $\{k \mid k, l\}$ and $\{l \mid k, k\}$. As before, we label the spiders in order from left to right to create the following two fully labelled diagrams



- we compute all partitions of $\{k, k, k\}$ that correspond to the original partition structure of the only partition in $\Pi_{X,F}$, namely $\{F \mid F, F\}$, modulo isomorphism. This gives $\{k \mid k, k\}$ only. As before, we label the spiders in order from left to right to create the following two fully labelled diagrams



ii) $k = 1$: 1 fixed label and 2 free labels.

First, we create the multiset $\{X, F, F\}$, where X denotes a fixed symbol and F denotes a free symbol. We compute all partitions of this multiset that correspond to the original partition structure of π . There are two such partitions, and so $\Pi_{X,F} = \{\{X \mid F, F\}, \{F \mid X, F\}\}$.

Next, we calculate all possible size 1 multisets choosing labels from L_X , allowing repeats, and all possible size 1 multisets choosing free labels, introducing them as needed, and also allowing repeats. Taking their Cartesian product gives all possible ways of choosing 1 fixed label and 1 free label. This gives $\{i, k, l\}$, $\{i, k, k\}$, $\{j, k, l\}$, and $\{j, k, k\}$. We consider each of these in turn:

- for each partition $\pi_{X,F} \in \Pi_{X,F}$, we compute all partitions of $\{i, k, l\}$ that correspond to the original partition structure of $\pi_{X,F}$, modulo isomorphism. For $\{X \mid F, F\}$, this gives $\{i \mid k, l\}$ only. For $\{F \mid X, F\}$, this initially gives $\{k \mid i, l\}$ and $\{l \mid i, k\}$, but these two are equivalent under a relabelling of the free labels, hence there is only $\{k \mid i, l\}$.

For each partition, we create a copy of (76) and label the spiders in order from left to right to create the two fully labelled diagrams

(80)

A similar process occurs for $\{j, k, l\}$, giving

(81)

- for each partition $\pi_{X,F} \in \Pi_{X,F}$, we compute all partitions of $\{i, k, k\}$ that correspond to the original partition structure of $\pi_{X,F}$, modulo isomorphism. For $\{X \mid F, F\}$, this gives $\{i \mid k, k\}$ only. For $\{F \mid X, F\}$, this gives $\{k \mid i, k\}$ only.

As before, we create a copy of (76) and label the spiders in order from left to right to create the two fully labelled diagrams

(82)

A similar process occurs for $\{j, k, k\}$, giving

(83)

iii) $k = 2$: 2 fixed labels and 1 free label.

First, we create the multiset $\{X, X, F\}$, where X denotes a fixed symbol and F denotes a free symbol. We compute all partitions of this multiset that correspond to the original partition structure of π . There are two such partitions, and so $\Pi_{X,F} = \{\{X \mid X, F\}, \{F \mid X, X\}\}$.

Next, we calculate all possible size 2 multisets choosing labels from L_X , allowing repeats, and all possible size 1 multisets choosing free labels, introducing them as needed, and also allowing repeats. Taking their Cartesian product gives all possible ways of choosing 2 fixed labels and 1 free label. This gives $\{i, j, k\}$, $\{i, i, k\}$, and $\{j, j, k\}$. We consider each of these in turn:

- for each partition $\pi_{X,F} \in \Pi_{X,F}$, we compute all partitions of $\{i, j, k\}$ that correspond to the original partition structure of $\pi_{X,F}$, modulo isomorphism. For $\{X \mid X, F\}$, this gives $\{i \mid j, k\}$ and $\{j \mid i, k\}$ (which are not isomorphic). For $\{F \mid X, X\}$, this gives $\{k \mid i, j\}$ only.


As usual, we create a copy of (76) and label the spiders in order from left to right to create the three fully labelled diagrams

(84)


and

(85)

- for each partition $\pi_{X,F} \in \Pi_{X,F}$, we compute all partitions of $\{i, i, k\}$ that correspond to the original partition structure of $\pi_{X,F}$, modulo isomorphism. For $\{X \mid X, F\}$, this gives $\{i \mid i, k\}$ only. For $\{F \mid X, X\}$, this gives $\{k \mid i, i\}$ only. As usual, we create a copy of (76) and label the spiders in order from left to right to create the two fully labelled diagrams


(86)

- a similar process to $\{i, i, k\}$ occurs for $\{j, j, k\}$, hence we obtain



(87)

iv) $k = 3$: 3 fixed labels and 0 free labels.


Despite there only being two fixed labels in L_X , we can still label the $s = 3$ unlabelled spiders in (76) with them. First, we create the multiset $\{X, X, X\}$, where X denotes a fixed symbol. We compute all partitions of this multiset that correspond to the original partition structure of π . There is only one such partition, and so $\Pi_{X,F} = \{\{X \mid X, X\}\}$.

Next, we calculate all possible size 3 multisets choosing labels from L_X , allowing repeats, and all possible size 0 multisets choosing free labels, introducing them as needed, and also allowing repeats. Taking their Cartesian product gives all possible ways of choosing 3 fixed labels and 0 free labels. This gives $\{i, i, j\}$, $\{i, i, i\}$, $\{j, j, i\}$, and $\{j, j, j\}$. We consider each of these in turn:


- we compute all partitions of $\{i, i, j\}$ that correspond to the original partition structure of the only partition in $\Pi_{X,F}$, namely $\{X \mid X, X\}$, modulo isomorphism. This gives $\{i \mid i, j\}$ and $\{j \mid i, i\}$. As before, we label the spiders in order from left to right to create the following two fully labelled diagrams


(88)

A similar process occurs for $\{j, j, i\}$, giving


(89)

- we compute all partitions of $\{i, i, i\}$ that correspond to the original partition structure of the only partition in $\Pi_{X,F}$, namely $\{X \mid X, X\}$, modulo isomorphism. This gives $\{i \mid i, i\}$ only. A similar process for $\{j, j, j\}$ gives $\{j \mid j, j\}$ only. Hence we obtain the following two fully labelled diagrams


(90)

Step 5: To finish, we merge all of the spiders in each of the diagrams, remove duplicates, and form a set as a result. This is the set of fully labelled diagrams corresponding to the partially labelled diagram (76) for the grouped output (i, i, j) of the $(4, 3)$ -bipartition diagram d_π .

To ensure brevity, we will not reproduce the final set of merged spiders here. However, we can see that only three duplicates are removed: the right hand diagram in (86) is removed since it is the same as the left hand diagram in (82); the right hand diagram in (87) is removed since it is the same as the left hand diagram in (83); and the right hand diagram in (89) is removed since it is the same as the right hand diagram in (88).

D.3. Subprocedures III, IV, and V

Example D.4. In Example D.1, we considered the $(3, 2)$ -bipartition diagram (57) and, by applying Subprocedure I, obtained two sets of partially labelled diagrams that were given in (62), one for each possible grouped output for (57). In fact, we noted that the diagrams in each set were actually fully labelled, hence we could skip applying Subprocedure II to each of them and proceed to applying Subprocedure III to each fully labelled diagram.

Following Subprocedure III, we see that the grouped map labels that correspond to the diagrams in the set (i, j) of (62) are

$$i, j \xleftarrow{g} i, i, j \quad (91)$$

and

$$i, j \xleftarrow{g} j, j, i \quad (92)$$

where $i \neq j$, and the grouped map label that corresponds to the diagram in the set (i, i) of (62) is

$$i, i \xleftarrow{g} i, i, i \quad (93)$$

We now apply Subprocedure IV to each of these grouped map labels — that is, we unroll the right hand side of each grouped map label to obtain its left grouped map label:

$$i, j \xleftarrow{lg} i, i, j + i, j, i + j, i, i \quad (94)$$

$$i, j \xleftarrow{lg} j, j, i + j, i, j + i, j, j \quad (95)$$

$$i, i \xleftarrow{lg} i, i, i \quad (96)$$

For each grouped output, we can add together the right hand side of every left grouped map label where this grouped output appears on the left hand side to obtain a single left grouped map label for each grouped output:

$$i, j \xleftarrow{lg} i, i, j + i, j, i + j, i, i + j, j, i + j, i, j + i, j, j \quad (97)$$

and

$$i, i \xleftarrow{lg} i, i, i \quad (98)$$

Hence, by applying Subprocedure V and converting the resulting map labels to transformation map labels, we see that, for an input symmetric tensor $T \in (\mathbb{R}^n)^{\otimes 3}$, where $n \geq 2$, the output symmetric tensor $D_\pi(T) \in (\mathbb{R}^n)^{\otimes 2}$, where d_π is given in (57), is given by

$$D_\pi(T)_{i,j} \leftarrow T_{i,i,j} + T_{i,j,i} + T_{j,i,i} + T_{j,j,i} + T_{j,i,j} + T_{i,j,j} \quad (99)$$

for all $i \neq j \in [n]$, and


$$D_\pi(T)_{i,i} \leftarrow T_{i,i,i} \quad (100)$$

for all $i \in [n]$.

As a trivial case, note that if $n = 1$, then we would not consider the fully labelled diagrams in the set (i, j) of (62), and so the only transformation map label would be

$$D_\pi(T)_{1,1} \leftarrow T_{1,1,1} \quad (101)$$

Example D.5. We can compare Example D.4 directly with the $(3, 4)$ -bipartition diagram d_π


(102)

Subprocedure I will give two sets:

$$(i, i, j, j) = \left\{ \begin{array}{c} \text{Diagram 1} \\ \text{Diagram 2} \\ \text{Diagram 3} \\ \text{Diagram 4} \end{array} \right\} \quad \text{and} \quad (i, i, i, i) = \left\{ \begin{array}{c} \text{Diagram 5} \end{array} \right\} \quad (103)$$

Once again we can skip over Subprocedure II since each diagram in each set is fully labelled. Applying Subprocedure III and then Subprocedure IV, we obtain the following left grouped map labels

$$i, i, j, j \xleftarrow{lg} i, j, j + j, i, j + j, j, i + i, i, j + i, j, i + j, i, i \quad (104)$$

where $i \neq j$, and

$$i, i, i, i \xleftarrow{lg} i, i, i \quad (105)$$

We now apply Subprocedure V to obtain a set of map labels for each left grouped map label: (104) becomes

$$i, i, j, j \leftarrow i, j, j + j, i, j + j, j, i + i, i, j + i, j, i + j, i, i \quad (106)$$

$$i, j, i, j \leftarrow i, j, j + j, i, j + j, j, i + i, i, j + i, j, i + j, i, i \quad (107)$$

$$i, j, j, i \leftarrow i, j, j + j, i, j + j, j, i + i, i, j + i, j, i + j, i, i \quad (108)$$

and (105) becomes

$$i, i, i, i \leftarrow i, i, i \quad (109)$$

Hence, by converting the resulting map labels to transformation map labels, we see that, for an input symmetric tensor $T \in (\mathbb{R}^n)^{\otimes 3}$, where $n \geq 2$, the output symmetric tensor $D_\pi(T) \in (\mathbb{R}^n)^{\otimes 4}$, where d_π is given in (102), is given by

$$D_\pi(T)_{i,i,j,j} \leftarrow T_{i,j,j} + T_{j,i,j} + T_{j,j,i} + T_{i,i,j} + T_{i,j,i} + T_{j,i,i} \quad (110)$$

$$D_\pi(T)_{i,j,i,j} \leftarrow T_{i,j,j} + T_{j,i,j} + T_{j,j,i} + T_{i,i,j} + T_{i,j,i} + T_{j,i,i} \quad (111)$$

$$D_\pi(T)_{i,j,j,i} \leftarrow T_{i,j,j} + T_{j,i,j} + T_{j,j,i} + T_{i,i,j} + T_{i,j,i} + T_{j,i,i} \quad (112)$$

for all $i \neq j \in [n]$, and

$$D_\pi(T)_{i,i,i,i} \leftarrow T_{i,i,i} \quad (113)$$

for all $i \in [n]$.

Once again, as a trivial case, note that if $n = 1$, then we would not consider the fully labelled diagrams in the set (i, i, j, j) of (103), and so the only transformation map label would be

$$D_\pi(T)_{1,1,1,1} \leftarrow T_{1,1,1} \quad (114)$$

Example D.6. We obtain the left grouped map labels for the $(3, 1)$ -bipartition diagram d_π



as follows.

We apply Subprocedure I to (115) to obtain


$$(i) = \left\{ \begin{array}{c} \text{Diagram with three green dots on the left and one on the right. The top dot on the left is labeled } i \text{ (blue). The middle and bottom dots on the left are connected by a V-shaped line. The bottom dot on the left is connected to the bottom dot on the right by a vertical line.} \end{array} \right\} \quad (116)$$

as the only possible grouped output for (115) with its associated set of partially labelled diagrams.

We then apply Subprocedure II to the only partially labelled diagram in (i) , giving five fully labelled diagrams for this grouped output:



and


(118)

Following Subprocedure III, we see that the grouped map labels that correspond to these fully labelled diagrams are

$$i \stackrel{g}{\leftarrow} \sum_{\substack{j, k \\ j \neq k \\ j \neq i \\ k \neq i}} j, j, k \quad i \stackrel{g}{\leftarrow} \sum_{\substack{j \\ j \neq i}} j, j, j \quad i \stackrel{g}{\leftarrow} \sum_{\substack{j \\ j \neq i}} i, i, j \quad (119)$$

and

$$i \stackrel{g}{\leftarrow} \sum_{\substack{j \\ j \neq i}} i, j, j \quad i \stackrel{g}{\leftarrow} i, i, i \quad (120)$$

We now apply Subprocedure IV to each of these grouped map labels — unrolling the right hand side of each grouped map label, we obtain its left grouped map label:

$$i \stackrel{lg}{\leftarrow} \sum_{\substack{j, k \\ j \neq k \\ j \neq i \\ k \neq i}} [j, j, k + j, k, j + k, j, j] \quad i \stackrel{lg}{\leftarrow} \sum_{\substack{j \\ j \neq i}} j, j, j \quad i \stackrel{lg}{\leftarrow} \sum_{\substack{j \\ j \neq i}} [i, i, j + i, j, i + j, i, i] \quad (121)$$

and

$$i \stackrel{lg}{\leftarrow} \sum_{\substack{j \\ j \neq i}} [i, j, j + j, i, j + j, j, i] \quad i \stackrel{lg}{\leftarrow} i, i, i \quad (122)$$

Given that there is only one grouped output, we can add together the right hand side of every left grouped map label above to obtain a single left grouped map label for this grouped output:

$$i \stackrel{lg}{\leftarrow} \sum_{\substack{j, k \\ j \neq k \\ j \neq i \\ k \neq i}} [j, j, k + j, k, j + k, j, j] + \sum_{\substack{j \\ j \neq i}} j, j, j + \sum_{\substack{j \\ j \neq i}} [i, i, j + i, j, i + j, i, i] + \sum_{\substack{j \\ j \neq i}} [i, j, j + j, i, j + j, j, i] + i, i, i \quad (123)$$

Tidying the right hand side of (123) gives

$$i \stackrel{lg}{\leftarrow} \sum_{j, k=1}^n [j, j, k + j, k, j + k, j, j] - 2 \sum_{j=1}^n j, j, j \quad (124)$$

Hence, by applying Subprocedure V and converting the resulting map label to a transformation map label, we see that, for an input symmetric tensor $T \in (\mathbb{R}^n)^{\otimes 3}$, where $n \geq 3$, the output (symmetric) tensor $D_\pi(T) \in \mathbb{R}^n$, where d_π is given in (115), is given by


$$D_\pi(T)_i \leftarrow \sum_{j, k=1}^n [T_{j,j,k} + T_{j,k,j} + T_{k,j,j}] - 2 \sum_{j=1}^n T_{j,j,j} \quad (125)$$

for all $i \in [n]$.

Finally, in a similar way to the preceding examples, it is possible to obtain the transformation map label in the trivial cases $n = 1$ and $n = 2$.

D.4. Map Label Calculations for Example 7.3.

Example D.7 (Calculations for Example 7.3). We calculate the transformation map labels for each of the four $(2,1)$ -bipartition diagrams that were given in (26) for $n = 3$. We reproduce them below, for ease:


(126)

In the following, we refer to these diagrams as d_{π_1} , d_{π_2} , d_{π_3} , d_{π_4} , respectively. We follow the procedure that was given in the main paper to obtain the transformation map labels.

For d_{π_1} , it is clear that the only grouped output, and its associated set of partially labelled diagrams, is

$$(i) = \left\{ \begin{array}{c} | \\ \bullet \\ \diagup \quad \diagdown \end{array} \right\} \quad (127)$$

Since there is only one diagram in the set, and it is fully labelled, we skip over Subprocedure II to apply Subprocedure III, obtaining the grouped map label

$$i \xleftarrow{g} i, i \quad (128)$$

Applying Subprocedure IV to unroll its right hand side does nothing, hence we obtain

$$i \xleftarrow{lg} i, i \quad (129)$$

Finally, there is nothing to unroll on its left hand side either, hence we obtain the map label

$$i \leftarrow i, i \quad (130)$$

which we can convert into the transformation map label

$$D_{\pi_1}(T)_i \leftarrow T_{i,i} \quad (131)$$

for $i \in [3]$. This is (29) in the main paper.

For d_{π_2} , once again there is only one grouped output, whose associated set of partially labelled diagrams is

$$(i) = \left\{ \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ \diagup \quad \diagdown \end{array} \right\} \quad (132)$$

As there is only one diagram in the set, we now apply Subprocedure II to it to obtain two fully labelled diagrams that are associated with the grouped output

$$(i) = \left\{ \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ \diagup \quad \diagdown \end{array}, \begin{array}{c} | \\ \bullet \\ \diagup \quad \diagdown \end{array} \right\} \quad (133)$$

From these fully labelled diagrams, we obtain the following grouped map labels

$$i \xleftarrow{g} \sum_{\substack{j=1 \\ j \neq i}}^3 j, j \quad \text{and} \quad i \xleftarrow{g} i, i \quad (134)$$

Applying Subprocedure IV to unroll the right hand side for each grouped map label does nothing, hence we obtain

$$i \xleftarrow{lg} \sum_{\substack{j=1 \\ j \neq i}}^3 j, j \quad \text{and} \quad i \xleftarrow{lg} i, i \quad (135)$$

Adding together the right hand side of the resulting left grouped map labels gives

$$i \stackrel{lg}{\leftarrow} \sum_{j=1}^3 j, j \quad (136)$$

Finally, there is nothing to unroll on its left hand side either, hence we obtain the map label

$$i \leftarrow \sum_{j=1}^3 j, j \quad (137)$$

which we can convert into the transformation map label

$$D_{\pi_2}(T)_i \leftarrow \sum_{j=1}^3 T_{j,j} \quad (138)$$

for $i \in [3]$. This is (30) in the main paper.

For d_{π_3} , once again there is only one grouped output, whose associated set of partially labelled diagrams is

$$(i) = \left\{ \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array} \right\} \quad (139)$$

As there is only one diagram in the set, we now apply Subprocedure II to it to obtain two fully labelled diagrams that are associated with the grouped output

$$(i) = \left\{ \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array}, \begin{array}{c} | \\ \bullet \\ \diagdown \end{array} \right\} \quad (140)$$

From these, we obtain the following grouped map labels

$$i \stackrel{g}{\leftarrow} \sum_{\substack{j=1 \\ j \neq i}}^3 i, j \quad \text{and} \quad i \stackrel{g}{\leftarrow} i, i \quad (141)$$

Applying Subprocedure IV to unroll the right hand side for each grouped map label gives

$$i \stackrel{lg}{\leftarrow} \sum_{\substack{j=1 \\ j \neq i}}^3 [i, j + j, i] \quad \text{and} \quad i \stackrel{lg}{\leftarrow} i, i \quad (142)$$

Adding together the right hand side of the resulting left grouped map labels and tidying gives

$$i \stackrel{lg}{\leftarrow} \sum_{j=1}^3 [i, j + j, i] - i, i \quad (143)$$

Finally, there is nothing to unroll on its left hand side either, hence we obtain the map label

$$i \leftarrow \sum_{j=1}^3 [i, j + j, i] - i, i \quad (144)$$

which we can convert into the transformation map label

$$D_{\pi_3}(T)_i \leftarrow \sum_{j=1}^3 [T_{i,j} + T_{j,i}] - T_{i,i} \quad (145)$$

for $i \in [3]$. This is (31) in the main paper.

Finally, for d_{π_4} , once again there is only one grouped output, whose associated set of partially labelled diagrams is

$$(i) = \left\{ \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array} \right\} \quad (146)$$

As there is only one diagram in the set, we now apply Subprocedure II to it to obtain four fully labelled diagrams that are associated with the grouped output

$$(i) = \left\{ \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array}, \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array}, \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array}, \begin{array}{c} | \\ \bullet \\ | \end{array} \quad \begin{array}{c} | \\ \bullet \\ | \end{array} \right\} \quad (147)$$

From these, we obtain the following grouped map labels

$$i \stackrel{g}{\leftarrow} \sum_{\substack{j,k \\ j \neq k \\ j \neq i \\ k \neq i}}^3 j, k \quad i \stackrel{g}{\leftarrow} \sum_{\substack{j=1 \\ j \neq i}}^3 j, j \quad i \stackrel{g}{\leftarrow} \sum_{\substack{j=1 \\ j \neq i}}^3 i, j \quad i \stackrel{g}{\leftarrow} i, i \quad (148)$$

Applying Subprocedure IV to unroll the right hand side for each grouped map label gives

$$i \stackrel{lg}{\leftarrow} \sum_{\substack{j,k \\ j \neq k \\ j \neq i \\ k \neq i}}^3 j, k \quad i \stackrel{lg}{\leftarrow} \sum_{\substack{j=1 \\ j \neq i}}^3 j, j \quad i \stackrel{lg}{\leftarrow} \sum_{\substack{j=1 \\ j \neq i}}^3 [i, j + j, i] \quad i \stackrel{lg}{\leftarrow} i, i \quad (149)$$

Adding together the right hand side of the resulting left grouped map labels and tidying gives

$$i \stackrel{lg}{\leftarrow} \sum_{j,k=1}^3 j, k \quad (150)$$

Finally, there is nothing to unroll on its left hand side either, hence we obtain the map label

$$i \leftarrow \sum_{j,k=1}^3 j, k \quad (151)$$

which we can convert into the transformation map label

$$D_{\pi_4}(T)_i \leftarrow \sum_{j,k=1}^3 T_{j,k} \quad (152)$$

for $i \in [3]$. This is (32) in the main paper.

E. Implementation Details

We describe the architecture and training details that were used for each of the two tasks. In implementing¹ the linear permutation equivariant functions, we used batch vectorised operations for each of the unrolled basis transformations that came from the corresponding transformation map labels.

S_{12} -Invariant Task: We randomly generated a synthetic data set consisting of 5000 symmetric tensors, split into 90% training and 10% test. We trained two models which had the following specifications:

¹The code is available at <https://github.com/epearcecrump/symmetrictensors>.

- SymmPermEquiv, the S_{12} -invariant model on symmetric tensors in $(\mathbb{R}^{12})^{\otimes 3}$, consisted of a single, trainable, linear permutation equivariant function from symmetric tensors in $(\mathbb{R}^{12})^{\otimes 3}$ to $\mathbb{R}$. The total number of parameters was 3 and the training time was 2.26 seconds.
- MLP consisted of a single standard linear layer from $\mathbb{R}^{12 \times 12 \times 12}$ to $\mathbb{R}$. The total number of parameters was 1728 and the training time was 1.46 seconds.

Both models were optimised with stochastic gradient descent with a learning rate of 0.0001. We trained both models for 50 epochs with a batch size of 50.

S_8 -Equivariant Task: We randomly generated a synthetic data set consisting of 10000 symmetric tensors, split into 90% training and 10% test. We trained three models which had the following specifications:

- SymmPermEquiv, the S_8 -equivariant model on symmetric tensors in $(\mathbb{R}^8)^{\otimes 3}$, consisted of a single, trainable, linear permutation equivariant function from symmetric tensors in $(\mathbb{R}^8)^{\otimes 3}$ to symmetric tensors in $\mathbb{R}^8$. The total number of parameters was 7 and the training time was 12.01 seconds.
- PermEquiv, the standard S_8 -equivariant model (Maron et al., 2019a; Pearce-Crump, 2024; Godfrey et al., 2023), consisted of a single, trainable, linear permutation equivariant function from $(\mathbb{R}^8)^{\otimes 3}$ to $\mathbb{R}^8$. The total number of parameters was 15 and the training time was 12.17 seconds.
- MLP consisted of a single standard linear layer from $\mathbb{R}^{8 \times 8 \times 8}$ to $\mathbb{R}^8$. The total number of parameters was 4096 and the training time was 4.38 seconds.

Both models were optimised with stochastic gradient descent with a learning rate of 0.0001. We trained both models for 50 epochs with a batch size of 50. In Table 1, we used a batch size of 50 to obtain the test mean squared error (MSE). In Table 2, we generated a test set of 1000 tensors, one for each of $n = 8, 16, 32$, using the parameters of the model that was trained in the task itself for each of these models.

Comparison between Map Label and Weight Matrix Implementations: We are grateful to the reviewers for the following suggestion: they recommended that we should provide empirical evidence to justify the time and space requirements of our map label notation method compared with the standard implementation that uses unrolled weight matrices.

For the S_{12} -invariant task, the training time of SymmPermEquiv using the map label implementation was 2.26 seconds, compared to 127.55 seconds using a weight matrix implementation, resulting in a $60\times$ speedup.

For the S_8 -equivariant task, the training time of SymmPermEquiv using the map label implementation was 12.01 seconds, compared to 2451.45 seconds using a weight matrix implementation, resulting in a $200\times$ speedup.

These results demonstrate that what we have said in theory is true in practice, namely that the standard weight matrix approach becomes unfeasible as n, k and l become larger, owing to constraints on storing the weight matrices in memory, particularly in training the networks.

Subprocedure I: How to Obtain All Possible Grouped Outputs of a (k, l) -Bipartition Diagram.

Input: A (k, l) -bipartition diagram d_π .

1. Sort the spiders in decreasing order of output wires. For spiders that have the same number of output wires, break ties by sorting them in decreasing order of input wires. For now, ignore in the following all spiders that do not have output wires.
2. Partition the spiders having output wires into groupings where spiders are in the same grouping if and only if they have the same number of output wires. Call this partition π and the individual groupings $\{\pi_B\}$.
3. For each grouping π_B , form a new partition of the spiders in π_B such that spiders are in the same grouping in the new partition if and only if they have the same number of input wires. Let $\pi_{B,I}$ denote the new partition that corresponds to the grouping π_B .
4. Suppose that the number of spiders in d_π that have output wires is t . Calculate all integer partitions of t .
5. Then, for each integer partition λ of t :
 - (a) Assign a distinct label $(i, j, k, \dots)$ to each part of λ such that the label appears as many times as the size of the part. Treating the labels as a multiset, compute all partitions of this multiset that correspond to the original partition structure of π , modulo isomorphism: that is, identifying partitions that are equivalent under a relabelling of the labels. Denote the non-isomorphic partitions of the multiset of labels by $\{L\}$, and for each L , denote its individual groupings by $\{L_B\}$.
 - (b) For each non-isomorphic partition L :
 - i. For each grouping L_B , compute all partitions of L_B that correspond to the partition structure of $\pi_{B,I}$. Convert each partition of L_B into an ordered tuple of labels by replacing curly brackets with curved brackets and the vertical bar that splits groupings with a comma. Let $T_{L,B}$ denote the resulting set of tuples corresponding to L_B .
 - ii. Form the Cartesian product of the sets $T_{L,B}$, maintaining the order of the individual groupings π_B in π . Denote the Cartesian product set of ordered labels by $T_{\lambda,L}$.
 - iii. For each tuple $l \in T_{\lambda,L}$, create a copy of the sorted diagram of d_π and label the output spiders with l , resulting in a set $D_{\lambda,L}$ of partially labelled diagrams that are indexed by $T_{\lambda,L}$.
 - iv. Create an empty set $R_{\lambda,L}$. For each partially labelled diagram in $D_{\lambda,L}$, merge spiders that are labelled with the same label. If no merging took place, continue, otherwise:
 - A. Reconsider the newly merged bipartition diagram d_θ (retaining but ignoring its labels for now) and apply Steps 1, 2, and 3 to d_θ . This gives a partition θ whose individual groupings are $\{\theta_B\}$, and for each grouping θ_B , it is further partitioned into $\theta_{B,I}$.
 - B. For each grouping θ_B , form a set from the labels of the grouping and remove the labels from the diagram. This gives a partition of labels $\tilde{L}$ into individual groupings $\{\tilde{L}_B\}$.
 - C. Apply Step 5b to this partition of labels $\tilde{L}$ for $\theta_{B,I}$ up to and including Step 5(b)ii.
 - D. Then, for each tuple l in the newly created $T_{\tilde{\lambda},\tilde{L}}$, create a copy of the sorted diagram of d_θ with its labels removed and label the output spiders with l . Add the resulting diagram for each tuple l to $R_{\lambda,L}$.
 - v. Take the set union of $D_{\lambda,L}$ and $R_{\lambda,L}$. This gives a full set of partially labelled diagrams for each non-isomorphic partition of the multiset of labels L . Denote this set by $P_{\lambda,L}$.
 - (c) Take the set union of the $P_{\lambda,L}$ over all non-isomorphic partitions of the multiset of labels L . This gives a full set of partially labelled diagrams for each integer partition λ of t . Denote this set by P_λ .
6. Take the set union of P_λ over all integer partitions λ of t . Denote this set by P . For each diagram in P , reorder the labelled spiders in lexicographical order. Propagate the labels to the ends of the output wires. Partition the set P into subsets such that diagrams are in the same subset if and only if their tuples of output indices are the same. Label each subset with its shared tuple of output labels; each label represents a grouped output for d_π .

Output: All possible grouped outputs for d_π and, for each one, its associated set of partially labelled diagrams.

Subprocedure II: How to Determine All Fully Labelled Diagrams for a Partially Labelled Diagram of a Grouped Output.

Input: A partially labelled diagram for a grouped output of a (k, l) -bipartition diagram d_π .

1. Suppose that the partially labelled diagram has t labelled spiders, each with a separate label, and s spiders with no label. Call the labels in the partially labelled diagram fixed. Let L_X denote the set of fixed labels. Any other labels that are created in the following are said to be free.
2. Leaving the labelled spiders as they are, sort the unlabelled spiders in decreasing order of input wires.
3. Partition the sorted unlabelled spiders into groupings where spiders are in the same grouping if and only if they have the same number of input wires. Call this partition π and the individual groupings $\{\pi_B\}$.
4. Consider all possible ways of labelling the unlabelled spiders, where k of them are fixed and the remaining $s - k$ of them are free, for $k = 0, 1, \dots, s$. Then, for each value of k :
 - (a) Create a multiset of k X symbols (to denote fixed) and $s - k$ F symbols (to denote free). Compute all partitions of this multiset that correspond to the original partition structure of π . Denote this set by $\Pi_{X,F}$.
 - (b) Now calculate all possible k -size multisets choosing labels from L_X , allowing repeats, and all possible $(s - k)$ -size multisets choosing free labels (by introducing them as needed), also allowing repeats. Taking their Cartesian product gives all possible ways of choosing k fixed indices and $s - k$ free indices.
 - (c) For each pair of k fixed indices and $s - k$ free indices in this set:
 - i. For each partition $\pi_{X,F} \in \Pi_{X,F}$, compute all partitions of this pair of k fixed indices and $s - k$ free indices that correspond to the original partition structure of $\pi_{X,F}$, modulo isomorphism: that is, identifying partitions that are equivalent under a relabelling of the free indices only.
 - ii. For each non-isomorphic partition, create a copy of the sorted partially labelled diagram from Step 2 and label the output spiders in order from left to right to create a fully labelled diagram.
 - (d) Take the set union of the set of fully labelled diagrams for all pairs of k fixed indices and $s - k$ free indices.
5. Finally, take the set union of the set of the fully labelled diagrams for all values of k . Merge all of the spiders in each fully labelled diagram, and remove any duplicates. If it is desired, propagate the labels to the ends of the wires.

Output: A set of fully labelled diagrams for the partially labelled diagram of a grouped output.

Subprocedure III: How to Convert a Fully Labelled Diagram into a Grouped Map Label.

Input: A fully labelled diagram of a (k, l) -bipartition diagram d_π .

Convert the fully labelled diagram into a grouped map label as follows:

1. Create the arrow $\xleftarrow{g}$.
2. On its left, write the tuple that is formed from propagating the labels of the fully labelled diagram to the output wires. In particular, distinct labels that appear on the left hand side of the arrow represent distinct values.
3. On its right, write the tuple that is formed from propagating the labels of the fully labelled diagram to the input wires, summing over any labels that do not appear on the left hand side, and not allowing any summation labels to equal each other nor any label that appears on the left hand side.

Output: A grouped map label.

Subprocedure IV: How to Convert the Grouped Map Label for a Fully Labelled Diagram into a Left Grouped Map Label.

Input: A grouped map label for a fully labelled diagram of a (k, l) -bipartition diagram d_π .

1. Consider the right hand side of the grouped map label. The summation indices correspond to free indices in Subprocedure II, and all other indices correspond to fixed indices in Subprocedure II. (Each fixed index should also appear on the left hand side of the grouped map label.)
2. Replace the tuple of indices by a sum of all permutations of the tuple, modulo isomorphism under a relabelling of the free indices only. Also replace $\leftarrow^g$ with $\leftarrow^{lg}$.

Output: A left grouped map label for a fully labelled diagram of a (k, l) -bipartition diagram d_π .

Subprocedure V: How to Unroll the Left Hand Side of a Left Grouped Map Label.

Input: A left grouped map label.

Create a set of map labels for the left grouped map label by unrolling the left hand side as follows:

1. Calculate all permutations of the left hand side of the left grouped map label modulo isomorphism: any two permutations are said to be equivalent if the first can be bijectively relabelled to become the second, using the multiset of letters of letters formed from the left hand side of the left grouped map label.
2. For each resulting permutation, create a map label by letting it be the left hand side of the arrow $\leftarrow$. To its right is exactly the right hand side of the left grouped map label, unchanged.

Output: A set of map labels corresponding to the left grouped map label.