Compact Matrix Quantum Group Equivariant Neural Networks

Edward Pearce–Crump¹

Abstract

Group equivariant neural networks have proven effective in modelling a wide range of tasks where the data lives in a classical geometric space and exhibits well-defined group symmetries. However, these networks are not suitable for learning from data that lives in a non-commutative geometry, described formally by non-commutative C^* -algebras, since the C^* -algebra of continuous functions on a compact matrix group is commutative. To address this limitation, we derive the existence of a new type of equivariant neural network, called compact matrix quantum group equivariant neural networks, which encode symmetries that are described by compact matrix quantum groups. We characterise the weight matrices that appear in these neural networks for the easy compact matrix quantum groups, which are defined by set partitions. As a result, we obtain new characterisations of equivariant weight matrices for some compact matrix groups that have not appeared previously in the machine learning literature.

1. Introduction

Spaces in classical (commutative) geometry are given by sets of points with additional structure, and their symmetries are formally described by groups. For example, the compact matrix group SO(3) is the group of all rotations of \mathbb{R}^3 that preserve the origin. A class of neural networks called *group equivariant neural networks* was conceived to take advantage of group symmetries that are inherent in data by encoding them as an inductive bias in their architectures. Many group equivariant neural networks have been designed using representations of groups as their layer spaces, and characterisations of the weight matrices that appear in them have been found for a number of important groups (Zaheer et al., 2017; Ravanbakhsh et al., 2017; Maron et al., 2019a; Finzi et al., 2021; Villar et al., 2021; Pearce-Crump, 2023a;b; 2024; Godfrey et al., 2023). These networks have proven to perform well across a wide range of tasks, including, but not restricted to, particle physics (Bogatskiy et al., 2020), natural language processing (Gordon et al., 2020), the generation of molecules (Satorras et al., 2021), and computer vision (Chatzipantazis et al., 2023).

However, when we consider tasks that involve learning from data that lives in a non-commutative geometry, such as learning properties of random variables in free probability theory (Voiculescu, 1985; Voiculescu et al., 1992), physical phenomena in particle physics (Bhowmick et al., 2014; van Suijlekom, 2014), and properties of exactly solvable models in statistical mechanics (Podlés, 1995), a group equivariant neural network is no longer a valid choice of model. This is because non-commutative geometry, as developed by Connes (1995), is formally described using noncommutative C^* -algebras, which generalize commutative algebras C(X) of continuous functions on compact Hausdorff spaces X, and one can show that the C^* -algebra C(G)of continuous functions on a compact matrix group G is indeed commutative (Appendix B.4). In particular, in noncommutative geometry, the classical notion of a space is replaced by an algebraic structure, thus making traditional group-based symmetries inadequate (Budzyński & Kondracki, 1995; Majid, 2000; Sitarz, 2013). Instead, a more general notion of symmetry is required. Quantum groups, which have been developed by a number of authors (Jimbo, 1985; Drinfeld, 1986; Woronowicz, 1987; 1991; 1998), are an important class of mathematical objects for describing these symmetries.

Given that there exist symmetries that are not captured by groups, this implies that there is a need for a new type of equivariant neural network. In this paper, we focus on quantum groups that were constructed using a topological approach by Woronowicz (1987). Woronowicz generalised compact matrix groups to create a class of quantum groups called *compact matrix quantum groups*. We describe and motivate this construction in Section 3, and provide a comprehensive background on the prerequisite mathematical material, namely, general topology, C^* -algebras, and operator theory, in the Appendix.

The main contributions of our work are as follows. We

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

Proceedings of the 42^{nd} International Conference on Machine Learning, Vancouver, Canada. PMLR 267, 2025. Copyright 2025 by the author(s).

broaden the scope of equivariant neural networks beyond classical symmetries by defining a new class of neural networks that encode, as an inductive bias, symmetries described by compact matrix quantum groups. We call these neural networks compact matrix quantum group equivariant neural networks. In fact, we mathematically derive their existence using Woronowicz's version of Tannaka-Krein duality for compact matrix quantum groups (Woronowicz, 1988), which is introduced in Section 5. After proving that these neural networks exist, we characterise their weight matrices for a large class of compact matrix quantum groups known as *easy*. We focus on easy compact matrix quantum groups not only because they have been studied extensively in the literature (Banica & Speicher, 2009; Banica et al., 2010; Weber, 2013; Freslon & Weber, 2016; Raum & Weber, 2016; Tarrago & Weber, 2016; Weber, 2017; Tarrago & Weber, 2018; Gromada, 2020; Freslon, 2023), but also because the way in which they are defined, namely, by set partitions, makes it possible to obtain a full characterisation of the equivariant weight matrices for these compact matrix quantum groups. We note that our contributions are primarily theoretical in nature: to demonstrate the practical potential of these neural networks, further work is needed to extend the characterisation that we have found for the easy compact matrix quantum groups to the equivariant non-linear layers so that they can be implemented.

We note as a corollory of our work that the new class of neural networks generalises and encompasses all neural networks that are equivariant to compact matrix groups, since we show, in Section 3, that compact matrix groups are, and can be expressed as, compact matrix quantum groups. As a result, we recover characterisations of the weight matrices that have appeared in Ravanbakhsh et al. (2017); Maron et al. (2019a); Pearce-Crump (2023a; 2024); Godfrey et al. (2023), and obtain characterisations of the equivariant weight matrices for three compact matrix groups that have not appeared in the machine learning literature before, namely the hyperoctahedral group H_n , the bistochastic group B_n , and the unitary group U(n).

A word on notation: we use [n] to refer to the set of elements $\{1, \ldots, n\}$, and, for any set A, we denote the set of $n \times n$ matrices with entries in A by $M_n(A)$.

2. Related Work

A few prior related works (Hashimoto et al., 2022; Hataya & Hashimoto, 2023; Hashimoto et al., 2024) use C^* -algebras to design new machine learning architectures, though their approach is motivated by different objectives. While they generalize complex-valued layer spaces and parameters to function-valued ones, our approach extends the concept of group symmetry in group equivariant neural networks to a certain type of quantum group symmetry, thus expand-

ing the range of equivariant models. Consequently, their models need to be trained with a custom gradient descent method, whereas our networks can still be trained with standard backpropagation as both the layer spaces and weight matrices remain complex-valued. Crucially, while both approaches can be used to construct group equivariant neural networks, their models do not encode quantum group symmetry, which is central to our framework. Other works, such as those by Ruhe et al. (2023a;b), use Clifford algebras to develop equivariant neural networks that encode symmetries related to rotations and reflections. However, their approach seems loosely related to ours, in that while both involve noncommutative algebras, their networks capture symmetries in classical geometries whereas ours looks at symmetries in non-commutative geometries. Finally, our approach differs from that of Hoffmann et al. (2020), where they look to learn the algebraic structure from their network. By contrast, we show in Theorem 7.10 that the algebraic structure for easy compact matrix quantum group equivariant neural networks can be derived mathematically from the weight matrices themselves.

3. Non-Commutative Geometry

In this section, we first introduce non-commutative geometry and then motivate the existence of compact matrix quantum groups by studying certain properties of compact matrix groups $G(n) \subseteq GL(n)$.

Two ideas were important in the development of noncommutative geometry. The first was the paradigm shift from thinking in terms of points in a space (geometry) to considering functions on spaces (algebra) as the primary method for solving problems in geometry and topology (Budzyński & Kondracki, 1995; Majid, 2000; Sitarz, 2013). Indeed, whilst one can show (Appendix B.4) that the algebra C(X) of continuous, complex-valued functions on a compact Hausdorff space X is a commutative, unital, C^* algebra, a famous theorem in C^* -algebra theory by Gelfand & Naimark (1943) shows that every commutative, unital, C^* -algebra is, in fact, an algebra of continuous functions C(X) for some compact Hausdorff space X. In this way the geometric notion of a (compact Hausdorff) space X is dualised to the algebra of continuous functions on X.

The second important development was the abstraction of spaces entirely by allowing unital C^* -algebras to be noncommutative (Jimbo, 1985; Drinfeld, 1986; Woronowicz, 1987). In doing so, this created a notion of geometry through these algebras, namely by treating non-commutative C^* -algebras as if they were algebras of continuous functions on (non-existant) *non-commutative spaces*, as an abstract extension of the theorem by Gelfand & Naimark (1943). Consequently, non-commutative C^* -algebras are the foundation of non-commutative geometry, and, under this framework, the notion of a space (of points) is no longer meaningful.

But with these developments, a new notion of symmetry for non-commutative C^* -algebras was needed, since, for any compact matrix group $G(n) \subseteq GL(n)$, one can show (Appendix B.4) that the algebra C(G(n)) is a commutative, unital C^* -algebra.

However, we can study the properties of C(G(n)) to motivate the definition of a compact matrix quantum group, as given by Woronowicz (1987), to describe this new type of symmetry. Firstly, G(n) is not just a compact Hausdorff space, it is a group, and so it comes with a group law $G(n) \times G(n) \to G(n)$ where $(g_1, g_2) \mapsto g_1g_2$. We can express this group law in the language of C^* -algebras as follows: define $\Delta : C(G(n)) \mapsto C(G(n) \times G(n))$ under the mapping $f \mapsto ((g_1, g_2) \mapsto f(g_1g_2))$. Furthermore, since $C(G(n) \times G(n))$ is *-isomorphic to $C(G(n)) \otimes C(G(n))$ under $((g_1, g_2) \mapsto f_1(g_1)f_2(g_2)) \cong f_1 \otimes f_2$, we see that

$$\Delta: C(G(n)) \mapsto C(G(n)) \otimes C(G(n)) \tag{1}$$

is the expression of the group law of G(n) in the language of C^* -algebras.

We generalise this construction by replacing C(G(n)) by a (potentially non-commutative) C^* -algebra A, to define a compact matrix quantum group in the style of Woronowicz (1987).

Definition 3.1. Let A be a unital C^* -algebra, and let $u_{i,j} \in A$, for all $i, j \in [n]$, for some positive integer n. Let u be the $n \times n$ matrix whose (i, j)-entry is $u_{i,j}$, that is, u is an element of the unital C^* -algebra $M_n(A)$. The pair (A, u) is said to be a **compact matrix quantum group** if

- 1. A is the universal C^* -algebra $C^*(u_{i,j}, 1 \le i, j \le n)$,
- 2. u and $u^{\top} = (u_{j,i})$ are invertible matrices, and
- 3. the comultiplication map $\Delta : A \to A \otimes_{\min} A$ defined by

$$\Delta(u_{i,j}) \coloneqq \left(\sum_{k} u_{i,k} \otimes u_{k,j}\right) \tag{2}$$

is a *-homomorphism.

Remark 3.2. In motivating Definition 3.1, we did not specify the tensor product in $C(G(n)) \otimes C(G(n))$ for constructing Δ because C(G(n)) is a nuclear C^* -algebra (Courtney & Gillaspy, 2023, Proposition 10.10), implying that all norms that complete the algebraic tensor product $C(G(n)) \odot C(G(n))$ into a C^* -algebra give the same one. However, in general, the algebraic tensor product of two C^* algebras, $A \odot B$, can be completed in many ways to form a C^* -algebra. One method is to complete with respect to the minimal norm $\|\cdot\|_{\min}$, giving the C^* -algebra $A \otimes_{\min} B$. We review these concepts in Section B of the Appendix. Note that for a compact matrix group $G(n) \subseteq GL(n)$, points 1 and 2 of Definition 3.1 make sense. Indeed, if we define functions $u_{i,j} : G(n) \to \mathbb{C}$ for $i, j \in [n]$ such that $u_{i,j}(g) = g_{i,j}$, then, by the Stone-Weierstrass Theorem (Appendix B.6) and the compactness of G(n), the $u_{i,j}$ generate C(G(n)). Also, the matrices $u := (u_{i,j})$ and u^{\top} are invertible since $u^{-1}(g) = u(g^{-1})$. Furthermore, by defining the comultiplication map Δ for C(G(n)) by

$$\Delta(u_{i,j})(g,h) \coloneqq u_{i,j}(gh) \tag{3}$$

we obtain the following result:

Proposition 3.3 (Appendix: Section C). Compact matrix groups $G(n) \subseteq GL(n)$ form a special class of compact matrix quantum groups.

Moreover, Woronowicz (1987, Theorem 1.5) also showed, using the theorem by Gelfand & Naimark (1943), that if (A, u) is a commutative compact matrix quantum group for some positive integer n, then $A \cong C(G(n))$ for some compact matrix group $G(n) \subseteq GL(n)$. Combining this with Proposition 3.3, we obtain the Fundamental Theorem of Compact Matrix Quantum Groups:

Theorem 3.4 (Fundamental Theorem of Compact Matrix Quantum Groups). Let (A, u) be a compact matrix quantum group for some $n \in \mathbb{N}$. Then A is commutative if and only if $A \cong C(G(n))$ for some compact matrix group $G(n) \subseteq$ GL(n).

Hence, the Fundamental Theorem shows that if (G(n), u) is non-commutative, then we obtain "true" compact matrix quantum groups. Bearing this theorem in mind, convention dictates that we denote A by C(G(n)), even if A is non-commutative, and so we often refer to G(n), or sometimes the pair (G(n), u), as the compact matrix quantum group.

We now provide an example of a "true" compact matrix quantum group to show that they do indeed exist.

Example 3.5. Recall that each element of the symmetric group S_n , considered as a subgroup of GL(n), can be thought of as an $n \times n$ matrix having exactly one entry of 1 in each row and in each column, with all other entries being 0. Writing the elements of $g \in S_n$ as $g_{i,j}$, this condition can be expressed as the following relations:

$$g_{i,j}^2 = g_{i,j} = g_{i,j}^*, \ \sum_{i=1}^n g_{i,j} = \sum_{j=1}^n g_{i,j} = 1$$
 (4)

and

$$g_{i,j}g_{k,l} = g_{k,l}g_{i,j} \tag{5}$$

By the universal property of universal C^* -algebras (Appendix B.24), we have that $C(S_n)$ is the universal C^* -algebra $C^*(E \mid R)$ with generators E given by $\{u_{i,j}, 1 \leq i, j \leq n\}$ and relations R given in (4) and (5) (replacing $g_{i,j}$

with $u_{i,j}$). Wang (1998) defined the symmetric quantum group as a compact matrix quantum group S_n^+ by *liberating* $C(S_n)$ of its commutativity relations to obtain $C(S_n^+)$ as the universal C^* -algebra $C^*(E \mid R)$ where E is the same as for S_n but now the set of relations R is given by (4) only. As a result, the symmetric quantum group S_n^+ is our first example of a compact matrix quantum group that is not a group.

4. Representation Theory of Compact Matrix Quantum Groups

Analogous to compact matrix groups, compact matrix quantum groups have a representation theory. We now motivate the definition of a representation of a compact matrix quantum group. Recall that for a compact matrix group G(n), an *m*-dimensional representation is a choice of *m*dimensional vector space V over \mathbb{C} and group homomorphism $\rho_V : G(n) \to GL(V)$. By choosing a basis of V, we can replace GL(V) by GL(m). The representation ρ_V can itself be represented by functions $v_{i,j} \in C(G(n))$ for $1 \le i, j \le m$ that map $g \in G(n)$ to the (i, j) entry of $\rho_V(g)$, since ρ_V is continuous by the compactness of G(n). We claim the following result.

Lemma 4.1. If Δ is the comultiplication on C(G(n)) defined in (54), then the homomorphism property of the representation satisfies

$$\Delta(v_{i,j})(g,h) = \left(\sum_{k} v_{i,k} \otimes v_{k,j}\right)(g,h)$$
(6)

Proof. Indeed, we have that $\Delta(v_{i,j})(g,h)$ is equal to

$$v_{i,j}(gh) = [\rho_V(gh)]_{i,j} = \sum_k [\rho_V(g)]_{i,k} [\rho_V(h)]_{k,j} = \sum_k v_{i,k}(g) v_{k,j}(h) = \left(\sum_k v_{i,k} \otimes v_{k,j}\right) (g,h) \qquad \Box$$

As a result, we have the following definition.

Definition 4.2. Let (G(n), u) be a compact matrix quantum group. Then an *m*-dimensional representation of G(n) is an element $v \in M_m(C(G(n)))$ such that (2) holds. If a representation v of G(n) has a matrix inverse, we say that v is a **non-degenerate** representation, and if v is unitary, that is, $vv^* = 1 = v^*v$, then we call v a **unitary** representation. *Remark* 4.3. The matrix u given in the definition of a compact matrix quantum group (G(n), u) is an n-dimensional representation called the **fundamental representation** of G(n), and by Definition 3.1, it is a non-degenerate representation of G(n).

Similar to groups, we can also define the tensor product and

complex conjugation of representations, starting with the following definition.

Definition 4.4. Let $v \in M_m(C(G(n)))$ and $w \in M_p(C(G(n)))$. Then the **tensor product** $v \otimes w \in M_{mp}(C(G(n)))$ is simply the Kronecker product of matrices, and the **complex conjugate** is $\bar{v} = (v_{i,j}^*) \in M_m(C(G(n)))$.

Proposition 4.5 (Appendix: Section C). If v, w are representations of a compact matrix quantum group (G(n), u), then both the tensor product and complex conjugate are also representations of G(n).

The following lemma will be useful when it comes to considering the fundamental representation of a compact matrix quantum group.

Lemma 4.6 (Appendix: Section C). Let (G(n), u) be a compact matrix quantum group. Then u^{\top} is unitary if and only if \bar{u} is unitary.

We can also define the analogous concept of equivariance for compact matrix quantum groups.

Definition 4.7. Let $v \in M_m(C(G(n)))$ and $w \in M_p(C(G(n)))$ be representations of a compact matrix quantum group G(n). A map $\phi : \mathbb{C}^m \to \mathbb{C}^p$ is said to be G(n)-equivariant (also known as an intertwiner) if $\phi v = w\phi$. The set of all such *linear* maps is denoted by $\operatorname{Hom}_{G(n)}(v, w)$, and it is, in fact, a vector space.

Remark 4.8. It is worth noting that the expression $\phi v = w\phi$ given in Definition 4.7 has a different interpretation relative to its sister definition for group equivariance. The representation $v \in M_m(C(G(n)))$ coacts on the space \mathbb{C}^m in the following sense: if $x \in \mathbb{C}^m$, then $vx \in \mathbb{C}^m \otimes C(G(n))$. Hence v can be thought of as a map $\mathbb{C}^m \to \mathbb{C}^m \otimes C(G(n))$ or as an element of $M_m(\mathbb{C}) \otimes C(G(n))$. In fact, any representation on a vector space V can be described by a map $V \to V \otimes C(G(n))$. Consequently, $\phi \in \text{Hom}_{G(n)}(v, w)$ if and only if it satisfies the commutative square:

$$\begin{array}{ccc} \mathbb{C}^{m} \xrightarrow{v} \mathbb{C}^{m} \otimes C(G(n)) \\ \phi & & & \downarrow \phi \\ \mathbb{C}^{p} \xrightarrow{w} \mathbb{C}^{p} \otimes C(G(n)) \end{array}$$

$$(7)$$

where the right-hand arrow maps $x \otimes f \mapsto \phi x \otimes f$.

Definition 4.9. We say that two representations $v \in M_m(C(G(n)))$ and $w \in M_m(C(G(n)))$ are **equivalent** if there is an invertible linear map $\phi : \mathbb{C}^m \to \mathbb{C}^m$ that intertwines v and w.

One can show (Gromada, 2020, Proposition 2.2.3) that every non-degenerate representation is equivalent to a unitary representation. Hence, going forward, we always choose the fundamental representation u of a compact matrix quantum group and its transpose to be unitary. By Lemma 4.6, we therefore have that both u and \bar{u} are unitary.

5. Tannaka–Krein Duality for Two-Coloured Representation Categories

Woronowicz's version of Tannaka–Krein duality is an important result in the theory of compact matrix quantum groups. Informally, the duality says that we can construct a compact matrix quantum group just by knowing its fundamental representation category. We will use this duality to derive compact matrix quantum group equivariant neural networks in Section 6. We begin with some definitions.

Definition 5.1. Consider the **two-coloured set** $\{\circ, \bullet\}$ consisting of a white point and a black point. For any nonnegative integer k, we can construct a **word** w of length k as a string of k colours from $\{\circ, \bullet\}$. If k = 0, we define \emptyset to be the empty word. The **concatenation** of two words, w_1, w_2 , is the concatenation of their two strings and is written as $w_1 \cdot w_2$. We define a **homomorphism on words**, $w \mapsto \overline{w}$, by first defining it on the individual colours by $\overline{\circ} := \bullet, \overline{\bullet} := \circ$, and then applying it element wise to the word w. If w is a word, then its **involution** w^* is the word read backwards together with its colours inverted.

We can use words of colours to create tensor products of the fundamental representation of a compact matrix quantum group as follows.

Definition 5.3. Let (G(n), u) be a compact matrix quantum group. Let $u^{\circ} \coloneqq u$ and $u^{\bullet} \coloneqq \overline{u}$. Then, for any word w formed from the two-coloured set $\{\circ, \bullet\}$, we define $u^{\otimes w}$ to be the corresponding tensor product of representations u° and u^{\bullet} .

Moreover, if w_k and w_l are words of lengths k, l respectively, then we define $\operatorname{FundRep}_{G(n)}(w_k, w_l)$ to be $\operatorname{Hom}_{G(n)}(u^{\otimes w_k}, u^{\otimes w_l})$, that is, the vector space

$$\{\phi: (\mathbb{C}^n)^{\otimes k} \to (\mathbb{C}^n)^{\otimes l} \mid \phi u^{\otimes w_k} = u^{\otimes w_l} \phi\}$$
(8)

Example 5.4. If w is the word $\circ \bullet \circ$, then $u^{\otimes w} = u \otimes \overline{u} \otimes \overline{u} \otimes u$.

We can also construct so-called representation categories using words of colours as follows.

Definition 5.5. A two-coloured representation category C is a collection $C(w_k, w_l)$, where w_k, w_l are two words constructed from $\{\circ, \bullet\}$ of lengths k, l respectively, that are subspaces of the set of all linear maps $(\mathbb{C}^n)^{\otimes k} \to (\mathbb{C}^n)^{\otimes l}$ and which satisfy the following axioms.

- 1. If $\phi_1 \in \mathcal{C}(w_k, w_l)$, $\phi_2 \in \mathcal{C}(w_q, w_m)$, then $\phi_1 \otimes \phi_2 \in \mathcal{C}(w_k \cdot w_q, w_l \cdot w_m)$.
- 2. If $\phi_1 \in \mathcal{C}(w_k, w_l)$, $\phi_2 \in \mathcal{C}(w_l, w_m)$, then $\phi_2 \circ \phi_1 \in \mathcal{C}(w_k, w_m)$.

- 3. If $\phi \in \mathcal{C}(w_k, w_l)$, then $\phi^* \in \mathcal{C}(w_l, w_k)$.
- 4. For every word w (having some length k), we have $1_n^{\otimes k} \in \mathcal{C}(w, w)$, and
- 5. The map $R : 1 \mapsto \sum_{i=1}^{n} e_i \otimes e_i$ is in both $\mathcal{C}(\emptyset, \circ \bullet)$ and $\mathcal{C}(\emptyset, \bullet \circ)$.

Definition 5.6. A one-coloured representation category is a two-coloured representation category where $u^{\circ} = u^{\bullet}$. Hence, the words correspond precisely to natural numbers because, for any natural number k, there is only one word of length k in this case.

The following two theorems describe the relationship between two-coloured representation categories and compact matrix quantum groups.

Theorem 5.7 (Appendix: Section C). If (G(n), u) is a compact matrix quantum group, then FundRep_{G(n)} is a two-coloured representation category.

Theorem 5.8 (Woronowicz–Tannaka–Krein Duality, (Woronowicz, 1988)). Let C be a two-coloured representation category. Then there exists a unique compact matrix quantum group (G(n), u) such that $\operatorname{FundRep}_{G(n)} = C$.

Remark 5.9 (Appendix C.1). There is a second part to Theorem 5.8, which states how to construct the unique compact matrix quantum group (G(n), u) from C. We provide this part in the Appendix for brevity.

6. Compact Matrix Quantum Group Equivariant Neural Networks

In this section we present the major contribution of this paper: we derive a new class of neural networks which we term *compact matrix quantum group equivariant neural networks* using Woronowicz's formulation of Tannaka–Krein duality. In the following definition, when referring to a word of length k_l , we write w_l instead of w_{k_l} to suppress multiple subscripts.

Definition 6.1. Let (G(n), u) be a compact matrix quantum group with fundamental representation $u \in M_n(C(G(n)))$. A **compact matrix quantum group equivariant neural network** for G(n) is a function f_{NN} that is a composition of some L layer functions

$$f_{NN} \coloneqq f_L \circ \ldots \circ f_l \circ \ldots \circ f_1 \tag{9}$$

such that the l^{th} layer function is equivariant to G(n) and is a function between representation spaces of G(n) under tensor products and complex conjugations of u

$$f_l: ((\mathbb{C}^n)^{\otimes k_{l-1}}, u^{\otimes w_{l-1}}) \to ((\mathbb{C}^n)^{\otimes k_l}, u^{\otimes w_l})$$
(10)

that is itself a composition

$$f_l \coloneqq \sigma_l \circ \phi_l \tag{11}$$

of a learnable, linear function in FundRep_{G(n)} (w_{l-1}, w_l)

$$\phi_l: ((\mathbb{C}^n)^{\otimes k_{l-1}}, u^{\otimes w_{l-1}}) \to ((\mathbb{C}^n)^{\otimes k_l}, u^{\otimes w_l})$$
(12)

that is equivariant to G(n), together with a fixed, non-linear activation function

$$\sigma_l: ((\mathbb{C}^n)^{\otimes k_{l-1}}, u^{\otimes w_l}) \to ((\mathbb{C}^n)^{\otimes k_l}, u^{\otimes w_l})$$
(13)

that is also equivariant to G(n). By choosing the standard basis of \mathbb{C}^n , the learnable, linear functions ϕ_l that are equivariant to G(n) become **matrices** that have **weights** that can be **learned**. As usual, we call these matrices (**equivariant**) weight matrices.

Remark 6.2. A compact matrix quantum group equivariant neural network is well-defined since the fundamental representation category $\operatorname{FundRep}_{G(n)}$ of a compact matrix quantum group (G(n), u) is a two-coloured representation category by Theorem 5.7.

We immediately obtain the following corollary for compact matrix groups $G(n) \subseteq GL(n)$.

Corollary 6.3. Let f_{NN} be a compact matrix quantum group equivariant neural network for a compact matrix **group** (G(n), u). If all of the words that are used in the network only consist of the white point \circ , then f_{NN} is, in fact, a compact matrix group equivariant neural network for G(n). Moreover, every compact matrix group equivariant neural network is a compact matrix quantum group equivariant neural network.

7. Easy Compact Matrix Quantum Groups

Having just derived the existence of a compact matrix quantum group equivariant neural network, we show, in this section, that we can characterise the weight matrices that appear in these neural networks for the so-called "easy" compact matrix quantum groups. A compact matrix quantum group is said to be **easy** if its two-coloured representation category is determined under a functor by a so-called two-coloured category of partitions. We first explore two-coloured categories of partitions, then define the functor, and finally construct both the easy compact matrix quantum group and characterise its equivariant weight matrices.

Our first goal is to define a two-coloured category of partitions. We need a few preliminary definitions.

Definition 7.1. For all non-negative integers l and k, a set **partition** π of [l + k] is a partition of the set [l + k] into a disjoint union of subsets, each of which is called a **block**. A **(one-coloured)** (k, l)-**partition diagram** d_{π} represents a set partition π as a diagram, where we place l white nodes on the top row, labelled left to right by $1, \ldots, l$, where we place k white nodes on the bottom row, labelled left to right by $l + 1, \ldots, l + k$, and the edges between the vertices

correspond to the connected components of the set partition π .

Definition 7.2. A (two-coloured) (w_k, w_l) -partition diagram is a (k, l)-partition diagram together with twocoloured words w_k, w_l of lengths k, l respectively such that the vertices on the top row have the same colours as the word w_l , from left to right, and the vertices on the bottom row have the same colours as the word w_k , from left to right. We use the same notation d_{π} to refer to the two-coloured version, since w_k and w_l will be known in making reference to the diagram. We define the **two-coloured partition vector space** $P_{w_k}^{w_l}(n)$ to be the \mathbb{C} -linear span of the set of all (w_k, w_l) -partition diagrams.

Example 7.3. Suppose that π is the set partition of [5+6] given by

$$\{1, 6 \mid 2, 3 \mid 4, 8 \mid 5, 9 \mid 7, 10 \mid 11\}$$
(14)

Then



is the two-coloured (w_k, w_l) -partition diagram d_{π} where w_k is the word $\circ \bullet \circ \bullet \circ \bullet \circ$ and w_l is the word $\bullet \circ \bullet \bullet \circ$.

We can now define a category for two-coloured partition diagrams, as follows:

Definition 7.4. $\mathcal{P}(n)$ is the category whose objects are the two-coloured words and, for any pair of objects w_k and w_l , the morphism space $\operatorname{Hom}_{\mathcal{P}(n)}(w_k, w_l)$ is defined to be $P_{w_k}^{w_l}(n)$.

We provide definitions of the vertical, horizontal and involution operations on morphisms in Section C of the Appendix. The unit object is the empty word \emptyset .

One can show the following result.

Proposition 7.5 (Appendix: Section C). $\mathcal{P}(n)$ is a strict \mathbb{C} -linear monoidal category.

We achieve our first goal with the following definition.

Definition 7.6. A two-coloured category of partitions $\mathcal{K}(n)$ is any subcategory of $\mathcal{P}(n)$ such that

- if w_k, w_l are objects in K(n), then Hom_{K(n)}(w_k, w_l) is a subspace of Hom_{P(n)}(w_k, w_l),
- 2. the identity partition diagram is an element of $\operatorname{Hom}_{\mathcal{K}(n)}(\circ, \circ)$,
- 3. the top-row pair partition diagram corresponding to the set partition {1, 2} of {1, 2} superimposed, on the one hand, with the word ○●, and on the other, with the word ●○, is in Hom_{K(n)}(Ø, ○●) and Hom_{K(n)}(Ø, ●○), respectively, and

4. the morphism spaces are closed under the vertical, horizontal and involution operations.

By Proposition 7.5, it is immediate that a **two-coloured** category of partitions $\mathcal{K}(n)$ is a strict \mathbb{C} -linear monoidal category.

We now introduce the critical functor for our purposes: we show that this functor takes two-coloured (w_k, w_l) -partition diagrams d_{π} that live in a two-coloured category of partitions $\mathcal{K}(n)$ to linear maps $\phi_{\pi} : (\mathbb{C}^n)^{\otimes k} \to (\mathbb{C}^n)^{\otimes l}$ that live in a two-coloured representation category. By choosing the standard basis for \mathbb{C}^n , we obtain $n^l \times n^k$ matrices instead, which are denoted by D_{π} in what follows.

Definition 7.7. Suppose that d_{π} is a (w_k, w_l) -partition diagram, where w_k, w_l are two-coloured words of lengths k, l respectively. We define D_{π} as follows.

Associate the indices i_1, i_2, \ldots, i_l with the vertices in the top row of d_{π} and j_1, j_2, \ldots, j_k with the vertices in the bottom row of d_{π} . Then, if $S_{\pi}((I, J))$ is defined to be the set of tuples $(I, J) \in [n]^{l+k}$ such that

if x, y are in the same block of π , then $i_x = i_y$ (16)

(where we have momentarily replaced the elements of J by $i_{l+m} \coloneqq j_m$ for all $m \in [k]$), we define

$$D_{\pi} \coloneqq \sum_{I \in [n]^l, J \in [n]^k} \delta_{\pi, (I, J)} E_{I, J}$$
(17)

where

$$\delta_{\pi,(I,J)} \coloneqq \begin{cases} 1 & \text{if } (I,J) \in S_{\pi}((I,J)) \\ 0 & \text{otherwise} \end{cases}$$
(18)

and $E_{I,J}$ is the $n^l \times n^k$ matrix having a 1 in the (I, J) position and is 0 elsewhere. We extend the definition of the map $d_{\pi} \mapsto D_{\pi}$ linearly to $P_{w_k}^{w_l}(n)$.

Example 7.8. If d_{π} is the two-coloured $(\circ \bullet \circ \circ \circ \circ, \bullet \circ \bullet \circ)$ -partition diagram given in Example 7.3, then, if n = 3, for example, we see from



that the (1, 2, 2, 1, 3 | 1, 2, 1, 3, 2, 2)-entry of D_{π} is 1, whereas, we see that, from



the (1, 1, 2, 1, 3 | 2, 2, 1, 3, 1, 2)-entry of D_{π} is 0.

Consequently, we obtain the following key result, which makes it possible for us to construct a two-coloured representation category from any two-coloured category of partitions.

Theorem 7.9. For any two-coloured category of partitions $\mathcal{K}(n)$, we can form a category $\mathcal{C}_{\mathcal{K}}(n)$ whose objects are the same as those in $\mathcal{K}(n)$, and, for any pair of objects w_k and w_l , the morphism space $\operatorname{Hom}_{\mathcal{C}_{\mathcal{K}}(n)}(w_k, w_l)$ is the image of $\operatorname{Hom}_{\mathcal{K}(n)}(w_k, w_l)$ under $d_{\pi} \mapsto D_{\pi}$, that is $\operatorname{Hom}_{\mathcal{C}_{\mathcal{K}}(n)}(w_k, w_l)$ is defined to be

$$\{D_{\pi} \mid d_{\pi} \in \operatorname{Hom}_{\mathcal{K}(n)}(w_k, w_l)\}$$
(21)

Then the map $d_{\pi} \mapsto D_{\pi}$ defines a strict \mathbb{C} -linear monoidal functor

$$\Theta: \mathcal{K}(n) \to \mathcal{C}_{\mathcal{K}}(n) \tag{22}$$

such that $C_{\mathcal{K}}(n)$ is a two-coloured representation category. [Appendix: Section C].

Hence we can use Woronowicz-Tannaka-Krein duality given in Theorem 5.8 to not only construct the easy compact matrix quantum group but also to characterise the weight matrices that appear in the neural network that was derived in Definition 6.1, as follows.

Theorem 7.10. A two-coloured category of partitions $\mathcal{K}(n)$ determines a unique compact matrix quantum group (G(n), u). We call these compact matrix quantum groups **easy**, and they can be constructed as the universal C^* -algebra $C^*(E \mid R)$, where $E = \{u_{i,j}, i, j \in [n]\}$ and $R = \{D_{\pi}u^{\otimes w_k} = u^{\otimes w_l}D_{\pi}$ for all $d_{\pi} \in \mathcal{K}(n)(w_k, w_l)\}$. The weight matrix W for the learnable linear function (12) in the neural network that is equivariant to G(n) is given by

$$W = \sum_{\pi \mid d_{\pi} \in \mathcal{K}(n)(w_{l-1}, w_l)} w_{\pi} D_{\pi}$$
(23)

8. Examples

We noted in the Introduction that easy compact matrix quantum groups have been well studied in the literature and many classes of them have been characterised in full. We introduce a few examples, characterise their weight matrices, and refer the reader to the referenced works for more details.

Note in the following that it is enough to provide the spanning set of matrices (21) for an easy compact matrix quantum group, since, by (23), the weight matrix is a weighted linear combination of these matrices.

Firstly, we are able to characterise the equivariant weight matrices for some compact matrix groups that have not appeared previously in the machine learning literature by combining Theorem 3.4 and Corollary 6.3. Namely, if we

Procedure 1: How to Calculate the Weight Matrix of an Equivariant Linear Layer Function from $((\mathbb{C}^n)^{\otimes k}, u^{\otimes w_k}) \rightarrow ((\mathbb{C}^n)^{\otimes l}, u^{\otimes w_l})$ for an Easy Compact Matrix Quantum Group (G(n), u).

Assume that $\mathcal{K}(n)$ is the two-coloured category of partitions that defines the easy compact matrix quantum group (G(n), u) as in Theorem 7.10.

Perform the following steps:

- 1. Calculate all of the two-coloured (w_k, w_l) -partition diagrams d_{π} that live in $\mathcal{K}(n)$. These diagrams form a basis of the morphism space $\operatorname{Hom}_{\mathcal{K}(n)}(w_k, w_l)$.
- 2. Apply the function $d_{\pi} \mapsto D_{\pi}$ to each diagram to obtain its associated spanning set matrix D_{π} .
- 3. Attach a weight $w_{\pi} \in \mathbb{C}$ to each matrix D_{π} .
- 4. Finally, calculate $\sum w_{\pi} D_{\pi}$ to give the overall weight matrix.

consider one-coloured partition categories that contain the swap partition diagram



then, by Theorem 7.10, we obtain easy compact matrix groups (as the C^* -algebra is commutative) and spanning sets of matrices (21), which allow us to characterise their weight matrices by (23). The groups themselves were characterised by Banica & Speicher (2009). We highlight two examples here that are new, and provide more examples in the Appendix that recover characterisations that were known to the machine learning community before.

Example 8.1. The first is the **hyperoctahedral group** H_n , the symmetry group of the hypercube. Banica & Speicher (2009) showed that the spanning set of matrices (21) is the image under $d_{\pi} \mapsto D_{\pi}$ of all (k, l)-partition diagrams whose blocks have even size.

Example 8.2. The second is the **bistochastic group** B_n , which is the group of orthogonal matrices having sum 1 in each row and column. Again, Banica & Speicher (2009) showed that the spanning set of matrices (21) is the image under $d_{\pi} \mapsto D_{\pi}$ of all (k, l)-partition diagrams whose blocks have size one or two.

Secondly, in the case where the partition categories are onecoloured but they do not contain the swap partition diagram (24), we obtain true compact matrix quantum groups (as the C^* -algebra is non-commutative). Again we highlight two examples.

Example 8.3. Wang (1998) discovered the **symmetric quantum group** $C(S_n^+)$, which we introduced in Section 3. Banica & Speicher (2009) showed that the spanning set of matrices (21) for $C(S_n^+)$ is the image under $d_{\pi} \mapsto D_{\pi}$ of all non-

crossing (k, l)-partition diagrams (i.e. no two connected components cross).

Example 8.4. (Wang, 1995) originally discovered the **or-thogonal quantum group** $C(O(n)^+)$ as a universal C^* -algebra. Banica & Speicher (2009) showed that the spanning set of matrices (21) for $C(O(n)^+)$ is the image under $d_{\pi} \mapsto D_{\pi}$ of all non-crossing (k, l)-partition diagrams whose blocks come in pairs, and also recovered $C(O(n)^+)$ itself in this way (using the result of Theorem 7.10).

Finally, for two-coloured partition categories, the most important examples are the **unitary group** U(n) and the **unitary quantum group** $U(n)^+$. These depend on the twocoloured words w_k and w_l that are chosen, but Tarrago & Weber (2016; 2018) showed that, for U(n), the spanning set of matrices is the image under $d_{\pi} \mapsto D_{\pi}$ of all (w_k, w_l) -partition diagrams whose blocks come in pairs such that if two vertices of a block are in the same row, then they have different colours, otherwise they have the same colours, and for $U(n)^+$ (defined by Theorem 7.10), it is the same except all crossing partition diagrams are first removed.

9. Conclusion

We have derived the existence of compact matrix quantum group equivariant neural networks for learning from data that lives in a non-commutative geometry and has symmetries that are described by compact matrix quantum groups. We have characterised the linear layers of these neural networks for the easy compact matrix quantum groups. As future work, we suggest that it would be good to extend this characterisation to the equivariant non-linear layers, which would enable us to demonstrate in practice what these neural networks promise in theory.

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

- Banica, T. and Speicher, R. Liberation of Orthogonal Lie Groups. *Advances in Mathematics*, 222(4):1461–1501, 2009. ISSN 0001-8708. URL https://doi.org/ 10.1016/j.aim.2009.06.009.
- Banica, T., Curran, S., and Speicher, R. Classification results for easy quantum groups. *Pacific Journal of Mathematics*, 247(1):1–26, 2010.
- Barcelo, H. and Ram, A. Combinatorial Representation Theory, 1997. arXiv:math/9707221.
- Benkart, G. and Halverson, T. Partition algebras $P_k(n)$ with 2k > n and the fundamental theorems of invariant theory for the symmetric group S_n . J. London Math. Soc, 99(2): 194–224, 2019a. URL https://doi.org/10.111 2/jlms.12175.
- Benkart, G. and Halverson, T. Partition Algebras and the Invariant Theory of the Symmetric Group. In Barcelo, H., Karaali, G., and Orellana, R. (eds.), *Recent Trends in Algebraic Combinatorics*, volume 16 of *Association for Women in Mathematics Series*, pp. 1–41. Springer, 2019b. URL https://doi.org/10.1007/978-3-030 -05141-9.
- Bhowmick, J., Goswami, D., Skalski, A., and Soltan, P. M. Quantum gauge symmetries in noncommutative geometry. *Journal of Noncommutative Geometry*, 8(2):433–471, 2014.
- Blackadar, B. Operator Algebras, volume 122 of Encyclopaedia of Mathematical Sciences. Springer-Verlag, Berlin, 2006. URL https://doi.org/10.1007/ 3-540-28517-2. Theory of C*-algebras and von Neumann algebras, Operator Algebras and Non-commutative Geometry, III.
- 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.
- Bourbaki, N. General Topology: Chapters 1–4. Springer Book Archive. Springer-Verlag Berlin Heidelberg, 1995. URL https://doi.org/10.1007/978-3-6 42-61701-0. Original French edition published by MASSON, Paris, 1971.

- Brauer, R. On Algebras Which Are Connected with the Semisimple Continuous Groups. *Ann. Math.*, 38:857– 872, 1937. ISSN 0003486X. URL https://doi.or g/10.2307/1968843.
- Budzyński, R. and Kondracki, W. Quantum Spaces: Notes and Comments on a Lecture by S. L. Woronowicz. *Banach Center Publications*, 34(1):37–45, 1995.
- Caspers, M. Quantum symmetry. Snapshots of modern mathematics from Oberwolfach, 0, 2020. URL 10.147 60/SNAP-2020-009-EN.
- Ceccherini-Silberstein, T., Scarabotti, F., and Tolli, F. *Representation Theory of the Symmetric Groups*. Cambridge University Press, 2010.
- Chatzipantazis, E., Pertigkiozoglou, S., Dobriban, E., and Daniilidis, K. \$\mathrm{SE}(3)\$-equivariant attention networks for shape reconstruction in function space. In *The Eleventh International Conference on Learning Representations*, 2023. URL https://openreview.n et/forum?id=RDy3IbvjMqT.
- Cohen, T. and Welling, M. Group Equivariant Convolutional Networks. In Balcan, M. F. and Weinberger, K. Q. (eds.), *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pp. 2990–2999, New York, New York, USA, 20–22 Jun 2016. PMLR. URL https://proc eedings.mlr.press/v48/cohenc16.html.
- Connes, A. *Noncommutative Geometry*. Elsevier Science, January 1995.

Courtney, K. and Gillaspy, E. Notes on C*-Algebras, 2023.

- Drinfeld, V. G. Quantum groups. Zap. Nauchn. Semin., 155: 18–49, 1986. doi: 10.1007/BF01247086.
- 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. URL https://proceedings.mlr.press/v139/f inzi21a.html.
- Freslon, A. On the Partition Approach to Schur-Weyl Duality and Free Quantum Groups. *Transformation Groups*, 22:707–751, 2017. URL https://doi.org/10.1 007/s00031-016-9410-9.
- Freslon, A. Compact Matrix Quantum Groups and Their Combinatorics. Cambridge University Press, 2023. URL https://doi.org/10.1017/9781009345705.

- Freslon, A. and Weber, M. On the representation theory of partition (easy) quantum groups. *Journal für die reine und angewandte Mathematik (Crelles Journal)*, 2016(720): 155–197, 2016. URL https://doi.org/10.151 5/crelle-2014-0049.
- Gelfand, I. and Naimark, M. On the imbedding of normed rings into the ring of operators in Hilbert space. *Matematiceskij sbornik*, 54(2):197–217, 1943. URL http: //eudml.org/doc/65219.
- 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. URL https://openre view.net/forum?id=VXwts-IZFi.
- Goodman, R. and Wallach, N. R. Symmetry, Representations and Invariants. Springer, 2009.
- 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. URL https://op enreview.net/forum?id=SylVNerFvr.
- Gromada, D. Compact Matrix Quantum Groups and Their Representation Categories. PhD thesis, Universität des Saarlandes, 2020.
- Halverson, T. and Ram, A. Partition algebras. *European J. Combin.*, 26(6):869–921, 2005. ISSN 0195-6698. URL https://doi.org/10.1016/j.ejc.2004.06. 005.
- Halverson, T. and Ram, A. Gems from the Work of Georgia Benkart. *Notices of the American Mathematical Society*, 69(3):375–384, March 2022. URL https://doi.or g/10.1090/noti2447.
- Hashimoto, Y., Wang, Z., and Matsui, T. C*-algebra Net: A New Approach Generalizing Neural Network Parameters to C*-algebra. In Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., and Sabato, S. (eds.), Proceedings of the 39th International Conference on Machine Learning, volume 162 of Proceedings of Machine Learning Research, pp. 8523–8534. PMLR, 17–23 Jul 2022. URL https://proceedings.mlr.press/v1 62/hashimoto22a.html.
- Hashimoto, Y., Ikeda, M., and Kadri, H. Position: C*-Algebraic Machine Learning – Moving in a New Direction. In Proceedings of the 41st International Conference on Machine Learning, volume 235 of Proceedings of Machine Learning Research, pp. 17667–17679. PMLR, 21–27 Jul 2024. URL https://proceedings.ml r.press/v235/hashimoto24a.html.

- Hataya, R. and Hashimoto, Y. Noncommutative C^* -algebra Net: Learning Neural Networks with Powerful Product Structure in C^* -algebra, 2023. arXiv:2302.01191.
- Hessam, H., Khalkhali, M., Pagliaroli, N., and Verhoeven, L. S. From noncommutative geometry to random matrix theory. *Journal of Physics A: Mathematical and Theoretical*, 55(41):413002, 2022. doi: 10.1088/1751-8121/ac8f c5.
- Hoffmann, J., Schmitt, S., Osindero, S., Simonyan, K., and Elsen, E. Algebranets, 2020. arXiv:2006.07360.
- Jimbo, M. A q-Difference Analogue of U(g) and the Yang-Baxter Equation. *Letters in Mathematical Physics*, 10: 63–69, July 1985. doi: 10.1007/BF00704588.
- Jones, V. The Potts model and the symmetric group. In Araki, H., Kawahigashi, Y., and Kosaki, H. (eds.), *Subfactors: Proceedings of the Taniguchi Symposium on Operator Algebras (Kyuzeso, 1993)*, pp. 259–267. World Scientific, 1994.
- 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. Curran Associates, Inc., 2018. URL https://proceedings.neurips.cc/paper /2018/file/a3fc981af450752046be17918 5ebc8b5-Paper.pdf.
- Lehrer, G. I. and Zhang, R. B. The Brauer category and invariant theory. *Journal of the European Mathematical Society*, 17:2311–2351, 2012. URL https://ems.pr ess/content/serial-article-files/728 7.
- Leinster, T. *Basic Category Theory*. Cambridge University Press, 2014.
- Lim, L.-H. and Nelson, B. J. What is an equivariant neural network?, 2022. arXiv:2205.07362.
- Lin, H. An Introduction to the Classification of Amenable *C*-Algebras*. World Scientific, 2001.
- Maaßen, L. Representation Categories of Compact Matrix Quantum Groups. Ph.D. thesis, RWTH Aachen University, Aachen, Germany, 2021.
- Majid, S. Non-Commutative Geometry and Quantum Groups. *Philosophical Transactions: Mathematical, Physical and Engineering Sciences*, 358(1765):89–109, January 2000.
- Maron, H., Ben-Hamu, H., Shamir, N., and Lipman, Y. Invariant and Equivariant Graph Networks. In *International Conference on Learning Representations*, 2019a. URL

https://openreview.net/forum?id=Syx7
2jC9tm.

- Maron, H., Fetaya, E., Segol, N., and Lipman, Y. On the Universality of Invariant Networks. In Chaudhuri, K. and Salakhutdinov, R. (eds.), *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. URL https://proc eedings.mlr.press/v97/maron19a.html.
- Maron, H., Litany, O., Chechik, G., and Fetaya, E. On Learning Sets of Symmetric Elements. In III, H. D. and Singh, A. (eds.), *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. URL https://proceedi ngs.mlr.press/v119/maron20a.html.

Munkres, J. R. Topology. Pearson, 2018.

- 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 2023a. URL https://proceedings.mlr.press/v202/pearce-crump23a.html.
- Pearce-Crump, E. How Jellyfish Characterise Alternating Group Equivariant Neural Networks. In Proceedings of the 40th International Conference on Machine Learning, volume 202 of Proceedings of Machine Learning Research, pp. 27483–27495. PMLR, 23–29 Jul 2023b. URL https://proceedings.mlr.press/v2 02/pearce-crump23b.html.
- 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. A Diagrammatic Approach to Improve Computational Efficiency in Group Equivariant Neural Networks, 2024. arXiv:2412.10837.
- Podlés, P. Symmetries of quantum spaces. Subgroups and quotient spaces of quantum SU(2) and SO(3) groups. *Communications in Mathematical Physics*, 170:1–20, 1995. doi: 10.1007/BF02099436.
- 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.

- Raum, S. and Weber, M. The combinatorics of an algebraic class of easy quantum groups. *Infin. Dimens. Anal. Quantum Probab. Relat. Top.*, 17(3), 2014.
- Raum, S. and Weber, M. Easy quantum groups and quantum subgroups of a semi-direct product quantum group. J. Noncommut. Geom., 9:1261–1293, 2015.
- Raum, S. and Weber, M. The Full Classification of Orthogonal Easy Quantum Groups. *Commun. Math. Phys.*, 341: 751–779, 2016. URL https://doi.org/10.100 7/s00220-015-2537-z.
- Ravanbakhsh, S. Universal Equivariant Multilayer Perceptrons. In Proceedings of the 37th International Conference on Machine Learning, ICML 2020, 13-18 July 2020, Virtual Event, volume 119 of Proceedings of Machine Learning Research, pp. 7996–8006. PMLR, 2020. URL http://proceedings.mlr.press/v119/rav anbakhsh20a.html.
- 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.
- Riehl, E. *Category Theory in Context*. Dover Publications, 2017.
- Rørdam, M., Larsen, F., and Laustsen, N. J. An Introduction to K-Theory for C*-Algebras, volume 49 of London Mathematical Society Student Texts. Cambridge University Press, Cambridge, 2000.
- Ruhe, D., Brandstetter, J., and Forré, P. Clifford Group Equivariant Neural Networks. In Advances in Neural Information Processing Systems, volume 36, pp. 62922– 62990. Curran Associates, Inc., 2023a. URL https: //proceedings.neurips.cc/paper_files /paper/2023/file/c6e0125e14ea3d1a3de 3c33fd2d49fc4-Paper-Conference.pdf.
- Ruhe, D., Gupta, J. K., De Keninck, S., Welling, M., and Brandstetter, J. Geometric Clifford Algebra Networks. In Proceedings of the 40th International Conference on Machine Learning, volume 202 of Proceedings of Machine Learning Research, pp. 29306–29337. PMLR, 23–29 Jul 2023b. URL https://proceedings.mlr.pres s/v202/ruhe23a.html.
- Sagan, B. E. The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions. Springer, 2000.
- 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. URL https://proceedings.mlr.press/v139/s atorras21a.html.

- Segal, E. Group Representation Theory. Course Notes for Imperial College London, 2014.
- Sitarz, A. A Friendly Overview of Noncommutative Geometry. *Acta Physica Polonica*. *B*, 44(12), 2013.
- Speicher, R. and Weber, M. ISEM24 C*-Algebras and Dynamics Lecture Notes. Course Notes for Universität des Saarlandes, 2020.
- Takesaki, M. Theory of Operator Algebras. I, volume 124 of Encyclopaedia of Mathematical Sciences. Springer-Verlag, Berlin, 2002.
- Tarrago, P. and Weber, M. Unitary Easy Quantum Groups: The Free Case and the Group Case. *International Mathematics Research Notices*, 2017(18):5710–5750, 08 2016. ISSN 1073-7928. URL https://doi.org/10.1 093/imrn/rnw185.
- Tarrago, P. and Weber, M. The classification of tensor categories of two-colored noncrossing partitions. *Journal* of Combinatorial Theory, Series A, 154:464–506, 2018. ISSN 0097-3165. URL https://doi.org/10.1 016/j.jcta.2017.09.003.
- van Suijlekom, W. D. *Noncommutative Geometry and Particle Physics*. Springer Netherlands, 2014. doi: 10.1007/978-94-017-9162-5.
- 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. URL https://openreview.net/forum?id= ba27-RzNaIv.
- Voiculescu, D. Symmetries of some reduced free product C*-algebras. In Operator Algebras and Their Connections with Topology and Ergodic Theory (Busteni, 1983), volume 1132 of Lecture Notes in Mathematics, pp. 556– 588. Springer, Berlin, 1985.
- Voiculescu, D. V., Dykema, K. J., and Nica, A. Free Random Variables. American Mathematical Society, 1992.
- Wang, S. Free products of compact quantum groups. Comm. Math. Phys., 167(3):671–692, 1995.
- Wang, S. Quantum symmetry groups of finite spaces. Comm. Math. Phys., 195(1):195–211, 1998.
- Weber, M. On the classification of easy quantum groups. Advances in Mathematics, 245:500–533, 2013. ISSN 0001-8708. URL https://doi.org/10.1016/j. aim.2013.06.019.

- Weber, M. Introduction to compact (matrix) quantum groups and Banica-Speicher (easy) quantum groups. *Proc Math Sci*, 127:881–933, 2017. URL https://doi.org/ 10.1007/s12044-017-0362-3.
- Weber, M. Quantum symmetry. Snapshots of modern mathematics from Oberwolfach, 5, 2020. URL 10.14760/SNAP-2020-005-EN.
- Woronowicz, S. L. Compact Matrix Pseudogroups. Communications in Mathematical Physics, 111:613–665, 1987.
- Woronowicz, S. L. Tannaka-Krein duality for compact matrix pseudogroups. Twisted SU(N) groups. *Inventiones* mathematicae, 93(1):35–76, 1988.
- Woronowicz, S. L. A remark on compact matrix quantum groups. Letters in Mathematical Physics, 21:35–39, 1991.
- Woronowicz, S. L. Compact quantum groups. symétries quantiques, 1998.
- 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. Curran Associates, Inc., 2017. URL https: //proceedings.neurips.cc/paper/2017/ file/f22e4747dalaa27e363d86d40ff442f e-Paper.pdf.

A. General Topology

We introduce some concepts from general topology — also known as point set topology — that are used in the theory of C^* -algebras (Appendix: Section B). General topology provides an abstract mathematical framework for studying the properties of a wide range of mathematical structures. The mathematical structures in general topology are the so-called topological spaces, namely, sets containing a family of subsets that satisfy some nice properties, together with the so-called continuous functions, which are the structure-preserving maps between topological spaces. The family of subsets that is chosen for a set is called a topology, and a set can have different topologies assigned to it. In this way the notion of a continuous function depends on the topology that is assigned to the domain and the codomain of the function. Our treatment is only an overview of the main ideas and results: we will mostly state these without proof and so we refer the reader to any standard reference, such as Bourbaki (1995) or Munkres (2018), for more details.

We define the first key concept of general topology. A **topology** on a set X is a family of subsets τ_X of X such that

- 1. Both the empty set and X are elements of τ_X ;
- 2. Any union of elements of τ_X is an element of τ_X ;
- 3. Any intersection of finitely many elements of τ_X is an element of τ_X .

We call a set X together with a topology τ_X on X a **topological space**, written as a pair (X, τ_X) . When the topology of a topological space (X, τ_X) is understood, we often just refer to the topological space as X. We sometimes call the elements of X **points**. We call the elements of τ_X the **open sets** of (X, τ_X) . A subset of (X, τ_X) is said to be **closed** if its complement is an open set. Note that by changing the topology that is associated with a set, the definition of the open sets changes.

We said that the maps of interest between topological spaces are the continuous functions, which are defined as follows. A function $f: (X, \tau_X) \to (Y, \tau_Y)$ between topological spaces is said to be **continuous** if f is a function on the underlying sets such that if U is an open set of Y, then its preimage, $f^{-1}(U)$, is an open set of X. One can show using this definition that f is continuous if and only if the preimage of a closed set in Y is a closed set in X. The continuous functions are important because they preserve certain properties of topological spaces, many of which will be defined in what follows.

One useful notion that helps to define topologies on sets is that of a **base**, or sometimes **basis**. A base *B* for a topological space (X, τ_X) is a collection of open sets in τ_X such that every open set can be written as a union of elements in *B*. We say that *B* generates the topology τ_X .

An important class of topological spaces are **metric spaces**, which come with a notion of distance between any pair of points in the set. A metric space is a set X together with a function $d_X : X \times X \to \mathbb{R}$, called a **metric**, such that the following axioms hold:

1. $d_X(x,y) \ge 0$ for all $x, y \in X$ and $d_X(x,y) = 0$ if and only if x = y;

2.
$$d_X(x,y) = d_X(y,x)$$
 for all $x, y \in X$:

3. $d_X(x,y) \leq d_X(x,z) + d_X(z,y)$ for all $x, y, z \in X$.

A metric space is often denoted by the pair (X, d_X) or simply by X when the metric d_X is understood. Some of the most important sets in a metric space are the **open balls**. For any point x in a metric space (X, d_X) and for any real number $\epsilon > 0$, the open ball of radius ϵ around x, written $B(x, \epsilon)$, is defined to be the set $\{y \in X \mid d_X(x, y) < \epsilon\}$. When a metric space is given a topology that is induced by the metric, known as the **metric topology**, then the metric space together with the metric topology becomes a topological space. The metric topology is generated by the base that is given by all of the open balls that are defined by the metric.

A particularly important class of metric spaces are the **normed vector spaces**, which are commonly referred to simply as **normed spaces**. To define them, we first need to define the concept of a **norm** on a vector space. A norm on a complex vector space X is a function $\|\cdot\| : X \to \mathbb{R}$ such that

1. $||x|| \ge 0$ for all $x \in X$ and if ||x|| = 0 then x = 0;

- 2. $\|\lambda x\| \leq |\lambda| \|y\|$ for all $x \in X$ and $\lambda \in \mathbb{C}$;
- 3. $||x + y|| \le ||x|| + ||y||$ for all $x, y \in X$;

If X is an algebra, then a norm on X must also satisfy the submultiplicativity axiom, namely

4. $||xy|| \le ||x|| ||y||$ for all $x, y \in X$

A normed space is a complex vector space (or algebra) X that is equipped with a norm $\|\cdot\| : X \to \mathbb{R}$. Normed spaces X become metric spaces under the metric that is induced by the norm, namely

$$d_X(x,y) \coloneqq \|x - y\| \text{ for all } x, y \in X$$
(25)

In particular, one can show that d_X satisfies the axioms of a metric on X.

One nice property for a topological space to have is the **Hausdorff property**. The Hausdorff property defines a notion of separation between points in a topological space that will be particularly useful when we consider limits of sequences below. To define the Hausdorff property, we first need the definition of a **neighbourhood**. A neighbourhood of a point p in a topological space (X, τ_X) is a subset V of X such that V contains an open set U that contains p, that is, $p \in U \subseteq V \subseteq X$. Note that V does not necessarily need to be an open set itself. We say that a topological space (X, τ_X) is **Hausdorff** if, for any two distinct points $x, y \in X$, there exists a neighbourhood V of x and a neighbourhood W of y such that V and W are disjoint. In particular, one can show that every metric space (X, d_X) is Hausdorff.

We now look at the **closure** of a set in a topological space (X, τ_X) . Every subset V of a topological space (X, τ_X) has a closure which is denoted by \overline{V} : it is the intersection of all closed sets in (X, τ_X) that contain V. This definition is not particularly helpful, so to make it more useful we need the following definition of a **limit point**. For a subset V of a topological space (X, τ_X) , a limit point of V is a point $x \in X$ such that every neighbourhood of x contains a point of V that is not x. One can then show that the closure of V is the union of V and the set containing all limit points of V. We say that a subset V of a topological space (X, τ_X) is **dense** in X if its closure is X, that is, $\overline{V} = X$. Moreover, one can show that the closure of a closed set is the set itself. Hence if V is dense in X and V is closed, then V = X. All of these definitions apply immediately in the case where X is a metric space.

We continue now with metric spaces and consider sequences of points in a metric space. A sequence (x_n) in a metric space (X, d_X) is said to be **Cauchy** if for all $\epsilon > 0$, there exists a positive integer N such that for all $m, n \ge N$, $d_X(x_m, x_n) < \epsilon$. A sequence (x_n) in a metric space (X, d_X) is said to **converge** to a limit $x \in X$ if for all $\epsilon > 0$, there exists a positive integer N such that for all $n \ge N$, $d_X(x_x, x) < \epsilon$. In particular, one can show that if a sequence converges in a metric space, then the limit must be unique. This is a direct consequence of a metric space being Hausdorff. Moreover, continuous functions preserve the convergence of sequences, in that if $f : (X, d_X) \to (Y, d_Y)$ is a continuous function between metric spaces and (x_n) is a sequence that converges to $x \in X$, then the sequence $(f(x_n))$ in Y converges to $f(x) \in Y$. Particularly nice metric spaces are those that are said to be complete: a metric space (X, d_X) is **complete** if every Cauchy sequence in (X, d_X) converges. In particular, a normed space that is complete with the metric induced by the norm is known as a **Banach algebra**, which is the starting point for the material that appears in Section B of the Appendix. One further useful result is that a subset of a complete metric space is complete if and only if it is closed.

A notable point is that closed sets have an alternative characterisation when X is a metric space. Firstly, note that if (x_n) is a sequence that lives in a subset V of X that converges, then its limit x is not necessarily in V (but it is in X). However, one can show that a subset V of a metric space (X, d_X) is **closed** if and only if every sequence (x_n) in V converges to a limit in V.

We said that complete metric spaces are particularly nice: they allow us to solve problems iteratively using Cauchy sequences, since we know that they have a limit. However, even if a metric space is incomplete, one can create a complete metric space from it in a process known as **completion**. We first need the following definition. A map $i : (X, d_X) \to (Y, d_Y)$ is said to be an **isometry** if it satisfies

$$d_Y(i(x_1), i(x_2)) = d_X(x_1, x_2) \tag{26}$$

for all $x_1, x_2 \in X$. A surjective isometry between metric spaces is called an **isomorphism of metric spaces**, or simply an **isomorphism**. Hence, we have that a metric space $(\tilde{X}, d_{\tilde{X}})$ is called the completion of a metric space (X, d_X) if the following conditions are satisfied:

- there is an isometry $i: X \to \tilde{X}$
- the image i(X) is dense in \tilde{X} , and
- the space $(\tilde{X}, d_{\tilde{X}})$ is complete.

In fact, one can show that every metric space has a unique completion, up to isomorphism. In particular, every incomplete metric space can be completed.

One of the most important concepts in general topology is **compactness**, which is a notion of finiteness for topological spaces. In general topology, compact topological spaces play an analogous role to finite dimensional vector spaces in the theory of vector spaces or to finite sets in the theory of sets. The motivation for compactness of a topological space X comes from wanting to understand the conditions that X needs to satisfy such that any continuous map $f : X \to \mathbb{C}$ is bounded. Recall that a function $f : X \to \mathbb{C}$ is **bounded** if there exists a non-negative real number M such that $|f(x)| \le M$ for all $x \in X$.

One way in which a continuous function $f: X \to \mathbb{C}$ will be bounded is if it is bounded on a finite number of subsets of X whose union is X itself. This provides sufficient motivation for the following definitions. If (X, τ_X) is a topological space, then a **cover** of X is a family $(U_i)_{i \in I}$ of subsets of X such that $\bigcup_{i \in I} U_i = X$. The cover is **finite** if the indexing set I is finite, and the cover is **open** if U_i is open for each $i \in I$. If $(U_i)_{i \in I}$ is a cover of X, then we say that $(U_j)_{j \in J}$ is a subcover of $(U_i)_{i \in I}$ if $J \subseteq I$ and $(U_j)_{j \in J}$ is itself a cover of X. Consequently, a topological space (X, τ_X) is **compact** if every open cover of X has a finite subcover. Hence, every continuous function $f: X \to \mathbb{C}$ on a compact topological space X is bounded. We study C(X), the space of continuous functions $f: X \to \mathbb{C}$, in the case where X is a compact Hausdorff space in Section B. It is also worth noting that every finite topological space is compact since any open cover must already be finite.

One can also show that every closed subspace of a compact space is compact, where the subspace is given the **subspace** topology: if V is a subspace of a topological space (X, τ_X) , then the subspace topology τ_V on V is defined by $\{U \cap V \mid U \in \tau_X\}$. Although it is not true that every compact subspace of a compact space is closed, we do have that every compact subspace of a Hausdorff space is closed. Moreover, if (X, τ_X) is compact and $f : (X, \tau_X) \to (Y, \tau_Y)$ is a continuous function, then f(X) is compact. And if (X, τ_X) is compact and $f : X \to \mathbb{C}$ is a continuous function, then f(X) is bounded. For metric spaces (X, d_X) , one can show that compactness is equivalent to sequential compactness, which means that every sequence in X has a convergent subsequence. Furthermore, if the metric space is a Euclidean space, say \mathbb{R}^n or \mathbb{C}^n , then one can show that a subspace of a Euclidean space is compact if and only if that subspace is closed and bounded.

Some of the most important compact topological spaces that we consider are **compact groups**. Firstly, a **topological group** is a topological space (G, τ_G) that has a group structure defined on it such that the group operations of multiplication and inversion

$$(g,h) \mapsto gh, \quad g \mapsto g^{-1}$$
 (27)

are both continuous functions. A topological group is called a compact group if the underlying topological space is compact. Note that we always choose our topological groups to be Hausdorff, and so this property is always assumed when we refer to a topological group. Note also that, in the theory of topological groups, a subgroup of a topological group is always assumed to be closed. In particular, a closed subgroup of a compact group is a compact group.

In the main paper, we study compact matrix groups. The most notable example of a compact matrix group is the **orthogonal group** O(n) of $n \times n$ matrices M such that $M \mapsto M^{\top}M = I$. To show that O(n) is compact, it is enough to show that it is closed and bounded, since O(n) is a subset of the Euclidean space $\mathbb{R}^{n \times n}$. Given that one can show that singleton sets in a metric space are closed and that polynomials are continuous functions, O(n) is closed since it is the preimage of the singleton set $\{I\}$ under the continuous map $f : \mathbb{R}^{n \times n} \to \mathbb{R}^{n \times n}$ that is given by $M \mapsto M^{\top}M$. O(n) is bounded since the Euclidean norm of an orthogonal $n \times n$ matrix is \sqrt{n} . Moreover, any finite group — in particular, the **symmetric group** S_n — is compact.

B. An Overview of C*-Algebras and Operator Theory

We provide an overview of the key results that appear in the theory of C^* -algebras, which is an area of mathematics that lies at the intersection of algebra and topology. The theory has its roots in the development of quantum mechanics in the 1920s, where the non-commutativity of bounded, linear operators on a Hilbert space is used to model physical observables. Gelfand & Naimark (1943) developed an abstract definition of C^* -algebras that did not require the use of bounded, linear operators on a Hilbert space. Forty-five years later, Woronowicz (1987; 1988) used C^* -algebras to develop compact matrix quantum groups. Since a proper treatment of the theory of C^* -algebras would amount to several hundreds of pages, we focus solely on the key concepts that also appear in the theory of compact matrix quantum groups given in Section 3; notably, tensor products of C^* -algebras and the construction of the universal C^* -algebra from a set of generators and relations on those generators. When we quote a result without proof, we provide a reference to the literature where its proof can be found. We also rely on the background material on general topology that appeared in Section A throughout. For a more comprehensive treatment on C^* -algebras, we recommend any of the following references: Blackadar (2006); Courtney & Gillaspy (2023); Gromada (2020); Lin (2001); Maaßen (2021); Speicher & Weber (2020) and Rørdam et al. (2000).

Definition B.1. We build up to the definition of a C^* -algebra with the following definitions.

- An algebra A over C is a complex vector space together with a bilinear, associative multiplication A × A → A satisfying λ(xy) = (λx)y = x(λy) for x, y ∈ A and λ ∈ C.
- An algebra A is **unital** if it contains an element $1 \in A$, called the **unit**, such that 1x = x1 = x for all $x \in A$.
- A subset $B \subseteq A$ is called a **subalgebra** if B is an algebra with respect to the same operations.
- A subset $I \subseteq A$ is called an **ideal** if I is a subalgebra of A that is invariant with respect to both left and right multiplication by elements of A.
- A normed algebra is an algebra A that has a norm ||·|| : A → C such that it is submultiplicative: ||xy|| ≤ ||x|| ||y|| for all x, y ∈ A.
- A **Banach algebra** is a normed algebra A which is complete with respect to the metric induced by the norm $\|\cdot\|$.
- An involution on an algebra A is an antilinear map $*: A \to A$ such that $(x^*)^* = x$ and $(xy)^* = y^*x^*$ for all $x, y \in A$.
- A *-algebra is an algebra A that has an involution. Similarly, we define *-subalgebra and *-ideal.
- A Banach *-algebra is a Banach algebra A that has an involution.

Consequently, we have that

Definition B.2. A C^* -algebra is a Banach *-algebra A that satisfies the C^* -identity: $||x^*x|| = ||x||^2$ for all $x \in A$.

Note that a closed *-subalgebra of a C^* -algebra is also a C^* -algebra.

Definition B.3. We say that an element a of a C^* -algebra A is **self-adjoint** if $a^* = a$, and, if A is unital, we say that a is **unitary** if $a^*a = aa^* = 1$.

Example B.4. We give some examples of C^* -algebras.

- 1. The complex numbers \mathbb{C} form a unital C^* -algebra.
- 2. Let \mathcal{H} be a complex Hilbert space with inner product denoted by $\langle \cdot, \cdot \rangle$. The collection of bounded linear operators $\mathcal{B}(\mathcal{H})$ is a unital C^* -algebra, where the norm is the operator norm

$$|T|| \coloneqq \sup\{ \|Tv\|_{\mathcal{H}} \mid v \in \mathcal{H}, \|v\|_{\mathcal{H}} \le 1 \}$$

$$(28)$$

for all $T \in \mathcal{B}(\mathcal{H})$. Note that multiplication is given by the composition of operators, and that the involution is the adjoint, defined with respect to the inner product by $\langle T^*v, w \rangle = \langle v, Tw \rangle$ for all $T \in \mathcal{B}(\mathcal{H})$ and $v, w \in \mathcal{H}$. Also, we see that $\mathcal{B}(\mathcal{H})$ is non-commutative.

3. For any positive integer n, let $M_n(\mathbb{C})$ denote the set of $n \times n$ matrices with elements in \mathbb{C} . Then $M_n(\mathbb{C})$ is a unital C^* -algebra, where the norm is the operator norm

$$\|M\| \coloneqq \sup\{\|Mv\|_2 \mid v \in \mathbb{C}^n, \|v\|_2 \le 1\}$$
(29)

where $\|\cdot\|_2$ is the Euclidean norm on \mathbb{C}^n , and the involution is given by the conjugate transpose. This is a special case of the previous example, where we have chosen $\mathcal{H} = \mathbb{C}^n$ and picked a basis to represent linear transformations as matrices.

4. Let X be a compact Hausdorff space. Then the space of continuous functions on X,

$$C(X) \coloneqq \{f : X \to \mathbb{C} \mid f \text{ is continuous}\}$$
(30)

is a unital C^* -algebra where the norm is the supremum norm

$$\|f\|_{\infty} \coloneqq \sup\{|f(x)| \mid x \in X\}$$
(31)

such that $(f+g)(x) \coloneqq f(x) + g(x)$, $(\lambda f)(x) \coloneqq \lambda f(x)$, $(fg)(x) \coloneqq f(x)g(x)$, $1(x) \coloneqq 1$ and $f^*(x) \coloneqq \overline{f(x)}$ for all $f, g \in C(X), x \in X$ and $\lambda \in \mathbb{C}$. Note, in particular, that C(X) is commutative. The supremum norm is indeed a norm since the compactness of X implies that every continuous function $f: X \to \mathbb{C}$ is bounded, and so $||f||_{\infty} < \infty$ for all $f \in C(X)$. The norm is also submultiplicative because

$$|fg(x)| = |f(x)g(x)| = |f(x)||g(x)| \le ||f||_{\infty} ||g||_{\infty}$$
(32)

which implies that $||fg||_{\infty} \leq ||f||_{\infty} ||g||_{\infty}$. Similar reasoning on the fact that $|f(x)\overline{f(x)}| = |f(x)|^2$ shows that the supremum norm satisfies the C^* -identity. Finally, the supremum norm is complete because one can show that any Cauchy sequence of continuous functions on X converges uniformly on X such that the limit is also a continuous function on X.

An important theorem for the C^* -algebra of continuous functions on a compact Hausdorff space X is the Stone-Weierstrass theorem. Its proof can be found in Speicher & Weber (2020, Theorem 3.3). First we need the following definition.

Definition B.5. Let A be an algebra of functions from some set X to the complex numbers \mathbb{C} . We say that A separates points in X if, for all $x \neq y \in X$, there is a function $f \in A$ such that $f(x) \neq f(y)$.

Consequently, we have that

Theorem B.6 (Stone-Weierstrass). Let X be a compact Hausdorff space. If A is a unital *-subalgebra of C(X) such that A separates points in X, then A is dense in C(X). In particular, if A is closed, then A = C(X).

We also have homomorphisms on C^* -algebras:

Definition B.7. If A, B are algebras, then a linear map $\phi : A \to B$ such that $\phi(xy) = \phi(x)\phi(y)$ for all $x, y \in A$ is called an **algebra homomorphism**.

If A, B are *-algebras, then if $\phi : A \to B$ is an algebra homomorphism that preserves the involution operation, that is, $\phi(x^*) = \phi(x)^*$ for all $x \in A$, then ϕ is called a *-homomorphism.

If A, B are *-algebras, then a bijective *-homomorphism $\phi : A \to B$ is called a *-isomorphism.

We use the properties of $\mathcal{B}(\mathcal{H})$ to obtain representations of C^* -algebras:

Definition B.8. A representation of an algebra A is a choice of complex Hilbert space \mathcal{H} and an algebra homomorphism $\pi : A \to \mathcal{B}(\mathcal{H})$.

A *-representation of a *-algebra A is a representation of A that preserves the involution operation.

A representation is said to be **faithful** if it is injective.

There is another important theorem relating to the unital C^* -algebra of continuous functions on a compact Hausdorff space X. This is the Gelfand–Naimark Theorem that we referenced in Section 3, and since it is the first of two theorems by Gelfand and Naimark that appears in this section, we will call it Gelfand–Naimark I. Its proof can be found in Blackadar (2006, Theorem II.2.2.4).

Theorem B.9 (Gelfand–Naimark I). Let A be a commutative, unital C^* -algebra. Then A is *-isomorphic to C(X), where

$$X \coloneqq \{\phi : A \to \mathbb{C} \mid \phi \text{ is a non-zero }^*\text{-homomorphism}\}$$
(33)

is a compact Hausdorff space. The *-isomorphism is given by $f : A \to C(X)$ where $f(a) : X \to \mathbb{C}$ for all $a \in A$ is such that $f(a)[\phi] = \phi(a)$ for all $\phi \in X$.

We now look at tensor products of C^* -algebras. First, we have the following definition.

Definition B.10. Let A, B be *-algebras. The **algebraic tensor product**, written $A \odot B$, is defined to be the algebraic tensor product of the vector spaces A and B together with a multiplication operation that is given by

$$(a_1 \odot b_1)(a_2 \odot b_2) = a_1 a_2 \odot b_1 b_2 \tag{34}$$

and an involution operation that is given by

$$(a \odot b)^* = a^* \odot b^* \tag{35}$$

In particular, $A \odot B$ is a *-algebra.

To define a tensor product on C^* -algebras A and B such that the result is also a C^* -algebra, we need to be able to define a norm $\|\cdot\|$ that satisfies the C^* -identity on $A \odot B$, and then complete $(A \odot B, \|\cdot\|)$ to obtain a C^* -algebra. Any norm that satisfies the C^* -identity is said to be a C^* -norm. Although it can be shown, as a consequence of the Gelfand–Naimark–Segal Theorem, that C^* -norms on algebraic tensor products of C^* -algebras always exist — see the definition of the spatial norm below — one of the main difficulties behind this construction is that in many cases there can be more than one C^* -norm that can be associated with $A \odot B$ (Courtney & Gillaspy, 2023), and so in those cases there are many possible definitions of a tensor product on C^* -algebras A and B such that the result is a C^* -algebra. Hence, in what follows, we choose only to outline the main ideas that lead to the construction of a particular tensor product called the **minimal tensor product** since this tensor product will be used in the definition of a compact matrix quantum group.

We start with another theorem by Gelfand–Naimark, which is technically a corollary of the Gelfand–Naimark–Segal Theorem. Its proof can be found in Speicher & Weber (2020, Theorem 5.19).

Theorem B.11 (Gelfand–Naimark II). Every C^* -algebra A has a faithful *-representation $\pi : A \to \mathcal{B}(\mathcal{H})$ on some Hilbert space \mathcal{H} .

Hence, for C^* -algebras A and B, we have faithful *-representations $\pi_A : A \to \mathcal{B}(\mathcal{H}_A)$ and $\pi_B : B \to \mathcal{B}(\mathcal{H}_B)$. One can show (Courtney & Gillaspy, 2023, Corollary 11.15) that this induces a faithful *-representation $\pi_A \odot \pi_B : A \odot B \to \mathcal{B}(\mathcal{H}_A \otimes \mathcal{H}_B)$ such that $(\pi_A \odot \pi_B)(a \odot b) = \pi_A(a) \otimes \pi_B(b)$ for all $a \in A$ and $b \in B$.

Hence we can define a norm $\|\cdot\|_*$ on $A \odot B$, known as the **spatial norm**, that is given by

$$\|x\|_* = \|(\pi_A \odot \pi_B)(x)\|_{\mathcal{B}(\mathcal{H}_A \otimes \mathcal{H}_B)} \tag{36}$$

which is a C^* -norm by the fact that the norm on $\mathcal{B}(\mathcal{H}_A \otimes \mathcal{H}_B)$ is a C^* -norm together with the injectivity of $\pi_A \odot \pi_B$.

Takesaki (2002) showed that the spatial norm is independent of the faithful *-representation, and that it is the minimal C^* -norm on $A \odot B$; that is, for all C^* -norms γ on $A \odot B$, we have that

$$\|x\|_* \le \gamma(x) \tag{37}$$

for all $x \in A \odot B$. Hence we call $||x||_*$ the **minimum norm**, and write it as $||x||_{\min}$.

Consequently, the completion of $A \odot B$, for C^* -algebras A and B, with respect to the minimal norm $||x||_{\min}$, is called the **minimal tensor product** and is denoted by $A \otimes_{\min} B$.

One can also show (Courtney & Gillaspy, 2023, Corollary 11.29) that for a pair of *-homomorphisms $\phi_1 : A_1 \to B_1$ and $\phi_2 : A_2 \to B_2$ on *-algebras A_1, A_2, B_1, B_2 , the algebraic tensor product $\phi_1 \odot \phi_2 : A_1 \odot A_2 \to B_1 \odot B_2$ extends to a *-homomorphism $\phi_1 \otimes_{\min} \phi_2 : A_1 \otimes_{\min} A_2 \to B_1 \otimes_{\min} B_2$.

We now focus on $M_n(A)$, the set of $n \times n$ matrices with entries in a C^{*}-algebra A, that is

$$M_n(A) \coloneqq \{(a_{i,j}) \mid a_{i,j} \in A, 1 \le i, j \le n\}$$

$$(38)$$

This comes with a natural involution that is given by $(a_{i,j})^* = (a_{i,i}^*)$ for all $(a_{i,j}) \in M_n(A)$.

We use the following definition to create a C^* -norm on $M_n(A)$.

Definition B.12. Suppose that $\phi : A \to B$ is a linear map between *-algebras A and B. Then, for all positive integers n, we define the linear map $\phi^{(n)} : M_n(A) \to M_n(B)$ to be $\phi^{(n)}((a_{i,j})) = (\phi(a_{i,j}))$. The map $\phi^{(n)}$ is called a **matrix amplification** of ϕ .

We create the C^* -norm on $M_n(A)$ as follows. By Gelfand–Naimark II, we have a faithful *-representation $\pi : A \to \mathcal{B}(\mathcal{H})$ for some Hilbert space \mathcal{H} . This induces a faithful *-representation $\pi^{(n)} : M_n(A) \to \mathcal{B}(\mathcal{H}^n)$ since one can show Rørdam et al. (2000, Section 6.1) that $M_n(\mathcal{B}(\mathcal{H}))$ is *-isomorphic to $\mathcal{B}(\mathcal{H}^n)$. Then we can define a function $\|\cdot\| : M_n(A) \to \mathbb{C}$ by

$$\|(a_{i,j})\| \coloneqq \left\|\pi^{(n)}((a_{i,j}))\right\|_{\mathcal{B}(\mathcal{H}^n)}$$
(39)

which is a C^* -norm since $\pi^{(n)}$ is injective and the norm on $\mathcal{B}(\mathcal{H}^n)$ is a C^* -norm.

One can show the following inequality, whose proof is outlined in Rørdam et al. (2000, Exercise 1.13). Lemma B.13. Let $(T_{i,j}) \in M_n(\mathcal{B}(H))$. Then

$$\|T_{i,j}\|_{\mathcal{B}(\mathcal{H})} \le \|(T_{i,j})\|_{M_n(\mathcal{B}(\mathcal{H}))} \le \sum_{i,j} \|T_{i,j}\|_{\mathcal{B}(\mathcal{H})}$$

$$\tag{40}$$

Lemma B.13 implies that the image of $M_n(A)$ under $\pi^{(n)}$, which is *-isomorphic to $M_n(A)$ since $\pi^{(n)}$ is injective, is a closed *-subalgebra of the C*-algebra $\mathcal{B}(\mathcal{H}^n)$. Hence $M_n(A)$ is itself a C*-algebra. Moreover, the norm on $M_n(A)$ defined in (39) is unique because a *-algebra admits at most one norm that makes it into a C*-algebra Lin (2001, Corollary 1.2.8).

Expressing this result as a theorem, we have that

Theorem B.14. If A is a C^* -algebra, then $M_n(A)$ is also a C^* -algebra whose norm is defined in (39).

One can also show the following result.

Lemma B.15. Suppose that A is a C^* -algebra and let n be any positive integer. Define a map $\phi : M_n(A) \to M_n(\mathbb{C}) \odot A$ by

$$\phi((a_{i,j})) = \sum_{i,j=1}^{n} E_{i,j} \odot a_{i,j}$$
(41)

where $E_{i,j}$ is the matrix unit with a 1 in the (i, j)-entry, and 0 otherwise. Then ϕ is a *-isomorphism.

We use the *-isomorphism between $M_n(A)$ and $M_n(\mathbb{C}) \odot A$ to obtain the following isomorphism.

Proposition B.16. Suppose that A is a C^* -algebra and let n be any positive integer. Then there is a unique C^* -norm on the algebraic tensor product $M_n(\mathbb{C}) \odot A$. Hence we write $M_n(\mathbb{C}) \otimes A$, and have that

$$M_n(\mathbb{C}) \otimes A = M_n(\mathbb{C}) \odot A \cong M_n(A) \tag{42}$$

Proof. The *-isomorphism together with the C*-norm given on $M_n(A)$ defines the existence of a C*-norm on $M_n(\mathbb{C}) \odot A$, namely

$$\|(\lambda_{i,j}) \odot a\| = \|(\lambda_{i,j}a)\| \tag{43}$$

 $M_n(\mathbb{C}) \odot A$ is now *-isomorphic to a closed *-subalgebra of the C*-algebra $M_n(A)$, namely $M_n(A)$ itself, hence it is a C*-algebra. The result is now immediate because a *-algebra admits at most one norm making it into a C*-algebra. \Box

This means that we can use the *-isomorphism ϕ to identify elements of $M_n(A)$ with elements of $M_n(\mathbb{C}) \otimes A$, namely

$$(a_{i,j}) \leftrightarrow \sum_{i,j=1}^{n} E_{i,j} \otimes a_{i,j}$$
(44)

We now turn to the concept of **universal** C^* -algebras. We know that we can create C^* -algebras by starting with a *-algebra, finding a C^* -norm for it, and then completing the pair. We would like to be able to create C^* -algebras using sets of generators and relations. We follow the presentation that is given in Speicher & Weber (2020, Chapter 6).

We begin by creating a *-algebra as follows.

Definition B.17. Let $E = \{x_i \mid i \in I\}$ be a set of symbols over an index set I. Let $E^* = \{x_i^* \mid i \in I\}$ be another set of symbols that is disjoint from E. Then

- A non-commutative monomial in E ∪ E* is a word of non-commuting symbols x^{α1}_{i1} ··· x^{αm}_{im} for some positive integer m such that i₁,..., i_m ∈ I and α_j ∈ {±1} for all i ∈ [m].
- A non-commutative polynomial in $E \cup E^*$ is a formal linear combination of non-commutative monomials in $E \cup E^*$ with coefficients in \mathbb{C} .
- We can define a **concatenation operation** on non-commutative monomials in $E \cup E^*$:

$$(x_{i_1}^{\alpha_1}\cdots x_{i_m}^{\alpha_m})\cdot (x_{j_1}^{\beta_1}\cdots x_{j_m}^{\beta_n}) = x_{i_1}^{\alpha_1}\cdots x_{i_m}^{\alpha_m}x_{j_1}^{\beta_1}\cdots x_{j_m}^{\beta_n}$$
(45)

for any non-commutative monomials $x_{i_1}^{\alpha_1} \cdots x_{i_m}^{\alpha_m}$ and $x_{j_1}^{\beta_1} \cdots x_{j_m}^{\beta_n}$ in $E \cup E^*$. This operation is extended linearly to to non-commutative polynomials in $E \cup E^*$.

• We can also define an **involution operation** on non-commutative monomials in $E \cup E^*$:

$$(x_{i_1}^{\alpha_1} \cdots x_{i_m}^{\alpha_m})^* = x_{i_1}^{\alpha_1^*} \cdots x_{i_m}^{\alpha_m^*}$$
(46)

where α_j^* is 1 if $\alpha_j = *$ and is * if $\alpha_j = 1$ for all $j \in [m]$. This operation is extended antilinearly to non-commutative polynomials in $E \cup E^*$.

Then, we have the following definition.

Definition B.18. Let E, E^* be as before. The **free *-algebra on the generator set** E, P(E), is the set of non-commutative polynomials in $E \cup E^*$ where addition and scalar multiplication are the standard operations, and multiplication and involution are given by the concatenation and involution operations given in (45) and (46), respectively.

The *-algebra P(E) is free in the sense that there are no relations amongst the polynomials in P(E).

For any subset R of polynomials in P(E), we can create the *-ideal I(R) that is generated by R. Considering P(E) and I(R) as algebras, we can certainly form the quotient algebra P(E)/I(R). Now, since P(E) and I(R) are *-algebras, the involution operation on P(E)/I(R) given by

$$(a+I)^* \coloneqq a^* + I \tag{47}$$

for all $a \in P(E)$ is well-defined, and so in this case P(E)/I(R) is also a (quotient) *-algebra. Hence, we have that

Definition B.19. The universal *-algebra with generators E and relations R, A(E | R), is the quotient *-algebra P(E)/I(R).

We now aim to create a C^* -norm on a *-algebra, although this *-algebra isn't quite $A(E \mid R)$. For this, we first need the definition of a C^* -seminorm.

Definition B.20. A C^* -seminorm on a *-algebra A is a map $\gamma : A \to \mathbb{R}$ such that

1. γ is a seminorm, that is

- $\gamma(x) \ge 0$ for all $x \in A$
- $\gamma(\lambda x) = |\lambda|\gamma(x)$ for all $x \in A$ and $\lambda \in \mathbb{C}$
- $\gamma(x+y) \leq \gamma(x) + \gamma(y)$ for all $x, y \in A$
- 2. γ is submultiplicative, that is, $\gamma(xy) \leq \gamma(x)\gamma(y)$ for all $x, y \in A$, and
- 3. γ satisfies the C*-identity, that is $\gamma(x^*x) = \gamma(x)^2$.

Note that a norm on an algebra A is a seminorm $\gamma : A \to \mathbb{R}$ such that if $\gamma(x) = 0$ then x = 0. Hence a C^{*}-norm is a C^{*}-seminorm satisfying this additional property.

We make the following definition.

Definition B.21. We define a function $\|\cdot\|_{A(E|R)} : A(E \mid R) \to \mathbb{R}$ by

$$\|x\|_{A(E|R)} \coloneqq \sup\{\gamma(x) \mid \gamma \text{ is a } C^* \text{-seminorm on } A(E \mid R)\}$$
(48)

It is possible that $||x||_{A(E|R)}$ is infinity for some $x \in A(E | R)$, so we only allow ourselves to consider those E and R such that $||x||_{A(E|R)} < \infty$ for all A(E | R). In this case, one can show that $||x||_{A(E|R)}$ is a C^* -seminorm: for example, for all $x, y \in A(E | R)$ and any C^* -seminorm γ on A(E | R), we have that

$$\gamma(xy) \le \gamma(x)\gamma(y) \le \|x\|_{A(E|R)} \|y\|_{A(E|R)}$$
(49)

and so we see that $\|\cdot\|_{A(E|R)}$ is submultiplicative. Also, we see that

$$I_0 \coloneqq \{ x \in A(E \mid R) \mid \|x\|_{A(E|R)} = 0 \}$$
(50)

is a *-ideal. Hence, $A(E \mid R)/I_0$ is a quotient *-algebra having a C^* -norm $||x||_{A(E|R)/I_0}$ that is obtained from the C^* -seminorm on $A(E \mid R)$ under the quotient map $A(E \mid R) \to A(E \mid R)/I_0$. Hence, we have that

Definition B.22. The universal C^* -algebra with generators E and relations R, $C^*(E | R)$, is the completion of $A(E | R)/I_0$ with respect to $||x||_{A(E|R)/I_0}$. It is indeed a C^* -algebra, by construction.

The following lemma is very useful for showing the existence of the universal C^* -algebra $C^*(E \mid R)$.

Lemma B.23. Let $E = \{x_i \mid i \in I\}$ (with E^* understood) be a set of generators and let $R \subseteq P(E)$ be a set of relations. If there is a constant C > 0 such that the value of γ on any generator is less than or equal to C and all C^* -seminorms γ on $A(E \mid R)$, then $||x|| < \infty$ for all $x \in A(E \mid R)$. Hence by the above construction $C^*(E \mid R)$ exists.

Proof. Let γ be any C^* -seminorm on $A(E \mid R)$. Then, since any monomial in $E \cup E^*$ has a positive integer length, we can pick an arbitrary monomial in $E \cup E^*$ having some arbitrary length m. Then, by the submultiplicative property, we have that the value of γ on this monomial is less than or equal to C^m , and so by taking the supremum over all C^* -seminorms we see that any element of $A(E \mid R)$ has a norm that is bounded, as required. \Box

One can show Speicher & Weber (2020, Proposition 6.7) that every universal C^* -algebra $C^*(E \mid R)$ satisfies the following so-called universal property.

Proposition B.24 (Universal Property). Let $E = \{x_i \mid i \in I\}$ be a set of generators and let $R \subseteq P(E)$ be a set of relations such that the universal C^* -algebra $C^*(E \mid R)$ exists. Suppose that A is a C^* -algebra having a subset $E' = \{y_i \mid i \in I\}$ that satisfies the relations R, that is, all polynomials in R become zero when we replace each x_i by y_i for all $i \in I$. Then there exists a unique *-homomorphism $\phi : C^*(E \mid R) \to A$ such that $x_i \mapsto y_i$ for all $i \in I$.

Example B.25. For $n \ge 2$, we have a *-isomorphism of C^* -algebras

$$M_n(\mathbb{C}) \cong C^*(u_{i,j}, 1 \le i, j \le n \mid u_{i,j}^* = u_{j,i}, u_{i,k}u_{l,j} = \delta_{k,l}u_{i,j})$$
(51)

Proof. Let $C^*(E \mid R) \coloneqq C^*(u_{i,j}, 1 \le i, j \le n \mid u_{i,j}^* = u_{j,i}, u_{i,k}u_{l,j} = \delta_{k,l}u_{i,j})$. We first need to show that $C^*(E \mid R)$ exists, that is, that $||x|| < \infty$ for all $x \in A(E \mid R)$. This is immediate from Lemma B.23 if we can show that there is some constant C such that $\gamma(u_{i,j}) \le C$ for all $1 \le i, j \le n$ and for all C^* -seminorms γ on $A(E \mid R)$.

Let γ be a C^* -seminorm on $A(E \mid R)$. By the C^* -identity and the relations R we have that

$$\gamma(u_{j,j})^2 = \gamma(u_{j,j}^* u_{j,j}) = \gamma(u_{j,j} u_{j,j}) = \gamma(\delta_{j,j} u_{j,j}) = \gamma(u_{j,j})$$
(52)

Hence $\gamma(u_{j,j})$ is either 0 or 1 for all $j \in [n]$ and so, again by the C^{*}-identity and the relations R, we have that

$$\gamma(u_{i,j})^2 = \gamma(u_{i,j}^* u_{i,j}) = \gamma(u_{j,i} u_{i,j}) = \gamma(\delta_{i,i} u_{j,j}) = \gamma(u_{j,j})$$
(53)

Hence $\gamma(u_{i,j})$ is bounded above by 1 for all $1 \le i, j \le n$ and all C*-seminorms γ on $A(E \mid R)$, as required.

We now prove the *-isomorphism $M_n(\mathbb{C}) \cong C^*(E \mid R)$. The matrix units $E_{i,j} \in M_n(\mathbb{C})$ satisfy the relations R, since $E_{i,j}^* = E_{i,j}^\top = E_{j,i}$ and $E_{i,k}E_{l,j} = \delta_{k,l}E_{i,j}$, hence, by the universal property, there exists a surjective *-homomorphism $\phi : C^*(E \mid R) \to M_n(\mathbb{C})$ such that $u_{i,j} \mapsto E_{i,j}$ for all $i, j \in [n]$. Looking at the relations R we see that the monomials in $C^*(E \mid R)$ can be at most the elements $u_{i,j}$, but since ϕ is surjective the monomials are exactly the elements $u_{i,j}$. Hence, as a vector space, $C^*(E \mid R)$ has dimension n^2 and so ϕ is also injective. Hence ϕ is a *-isomorphism, as required.

C. Missing Proofs and Supplementary Content

We include proofs of results that appeared in the main paper, and provide the full statement of Woronowicz–Tannaka–Krein Duality (Woronowicz, 1988), the first part of which appeared in Theorem 5.8.

Proof of Proposition 3.3. Let $G(n) \subseteq GL(n)$ be a compact matrix group. Consider C(G(n)), the C^* -algebra of continuous functions $G(n) \to \mathbb{C}$, which is commutative. Firstly, we can define functions $u_{i,j} : G(n) \to \mathbb{C}$ for $i, j \in [n]$ such that $u_{i,j}(g) = g_{i,j}$. By the Stone-Weierstrass Theorem (Theorem B.6) and the compactness of G(n), the $u_{i,j}$ generate C(G(n)) and the matrices $u := (u_{i,j})$ and u^{\top} are invertible since $u^{-1}(g) = u(g^{-1})$.

Moreover, we have that the comultiplication map Δ for C(G(n)) is given by

$$\Delta(u_{i,j})(g,h) = u_{i,j}(gh) \tag{54}$$

since

$$\left(\sum_{k} u_{i,k} \otimes u_{k,j}\right)(g,h) = \sum_{k} u_{i,k}(g)u_{k,j}(h) = \sum_{k} g_{i,k}h_{k,j} = (gh)_{i,j} = u_{i,j}(gh)$$
(55)

that is, Δ coincides with matrix multiplication in G(n).

The definition of Δ that is given in (54) is a *-homomorphism because

$$\Delta(u_{i,j}u_{k,l})(g,h) = u_{i,j}u_{k,l}(gh) \tag{56}$$

$$=u_{i,j}(gh)u_{k,l}(gh) \tag{57}$$

$$=\Delta(u_{i,j})(g,h)\Delta(u_{k,l})(g,h)$$
(58)

$$=\Delta(u_{i,j})\Delta(u_{k,l})(g,h)$$
(59)

and

$$\Delta(u_{i,j}^*)(g,h) = u_{i,j}^*(gh) = \overline{u_{i,j}(gh)} = \overline{\Delta(u_{i,j})(g,h)} = \Delta(u_{i,j})^*(g,h)$$
(60)

Note that in (57) we have used that C(G(n)) has a pointwise multiplication, and likewise, in (59), we have used that $C(G(n)) \times C(G(n))$ has a pointwise multiplication.

Proof of Proposition 4.5. We have that $v \otimes w$ is a representation of G(n) because

$$\Delta(v_{i,j}w_{k,l}) = \Delta(v_{i,j})\Delta(w_{k,l}) = \left(\sum_{x} v_{i,x} \otimes v_{x,j}\right) \left(\sum_{y} w_{k,y} \otimes w_{y,l}\right) = \sum_{x,y} v_{i,x}w_{k,y} \otimes v_{x,j}w_{y,l}$$
(61)

and \bar{v} is a representation of G(n) because

$$\Delta(v_{i,j}^*) = \Delta(v_{i,j})^* = \left(\sum_k v_{i,k} \otimes v_{k,j}\right)^* = \sum_k v_{i,k}^* \otimes v_{k,j}^*$$
(62)

where in both (61) and (62) we have used that Δ is a *-homomorphism.

Proof of Lemma 4.6. We have that

$$\bar{u}^* = (u_{i,j}^*)^* = ((u_{j,i}^*)^*) = (u_{j,i}) = u^\top$$
(63)

and also

$$(u^{\top})^* = (u_{j,i})^* = (u_{i,j}^*) = \bar{u}$$
(64)

Hence

$$u^{\top}(u^{\top})^* = 1 = (u^{\top})^* u^{\top} \iff \bar{u}\bar{u}^* = 1 = \bar{u}^*\bar{u}$$
 (65)

as required.

Proof of Theorem 5.7. It is easy to show that $\operatorname{FundRep}_{G(n)}$ satisfies the first four axioms of Definition 5.5. For example, the second axiom says that if $\phi_1 \in \operatorname{FundRep}_{G(n)}(w_k, w_l)$ and $\phi_2 \in \operatorname{FundRep}_{G(n)}(w_l, w_m)$, then we have that

$$(\phi_2 \circ \phi_1) u^{\otimes w_k} = \phi_2 u^{\otimes w_l} \phi_1 = u^{\otimes w_m} (\phi_2 \circ \phi_1)$$
(66)

and so $\phi_2 \circ \phi_1 \in \operatorname{FundRep}_{G(n)}(w_k, w_m)$.

For the fifth axiom, we need to show that the operator $R : 1 \mapsto \sum_{i=1}^{n} e_i \otimes e_i$ is in both $\operatorname{FundRep}_{G(n)}(\emptyset, \bullet \bullet)$ and $\operatorname{FundRep}_{G(n)}(\emptyset, \bullet \circ)$.

For the first case, we have that

$$R \in \operatorname{FundRep}_{G(n)}(\emptyset, \circ \bullet) \iff R = (u \otimes \bar{u})R \tag{67}$$

By applying both sides to 1, we have that (67) holds if and only if

I

$$\sum_{i=1}^{n} e_i \otimes e_i = \sum_{i=1}^{n} u e_i \otimes \bar{u} e_i = \sum_{i=1}^{n} \sum_{j,k=1}^{n} e_j \otimes e_k \otimes u_{j,i} u_{k,i}^* = \sum_{j,k=1}^{n} e_j \otimes e_k \otimes (uu^*)_{j,k}$$
(68)

Hence, by equating coefficients, we see that $R \in \text{FundRep}_{G(n)}(\emptyset, \circ \bullet)$ if and only if $uu^* = 1$. For the second case, a similar calculation shows that $R \in \text{FundRep}_{G(n)}(\emptyset, \bullet \circ)$ if and only if $\bar{u}\bar{u}^* = 1$. Since u and \bar{u} are both unitary, we obtain the result.

As promised in Remark 5.9, we provide the full statement of Woronowicz–Tannaka–Krein Duality (Woronowicz, 1988) here, the first part of which already appeared in Theorem 5.8.

Theorem C.1. [Woronowicz–Tannaka–Krein Duality] Let C be a two-coloured representation category. Then there exists a unique compact matrix quantum group (G(n), u) such that $\operatorname{FundRep}_{G(n)} = C$. If $E = \{x_{i,j} \mid 1 \le i, j \le n\}$ is a set of generators for the free *-algebra P(E), then the *-ideal $I_{G(n)} \subseteq P(E)$ that determines G(n) is given by

$$span\left\{ \left[\phi x^{\otimes w_k} - x^{\otimes w_l}\phi\right]_{I,J} \mid \phi \in \mathcal{C}(w_k, w_l), I \in [n]^l, J \in [n]^k \right\}$$

$$\tag{69}$$

In Definition 7.4, we also said that we would provide definitions of the following three \mathbb{C} -(bi)linear operations on twocoloured set partition diagrams

composition: • :
$$P_{w_l}^{w_m}(n) \times P_{w_k}^{w_l}(n) \to P_{w_k}^{w_m}(n)$$
 (70)

tensor product:
$$\otimes : P_{w_k}^{w_l}(n) \times P_{w_a}^{w_m}(n) \to P_{w_k \cdot w_a}^{w_l \cdot w_m}(n)$$
 (71)

involution:
$$*: P_{w_k}^{w_l}(n) \to P_{w_l}^{w_k}(n)$$
 (72)

We define these operations as follows.

Definition C.2 (Composition). Let $d_{\pi_1} \in P_{w_k}^{w_l}(n)$ and $d_{\pi_2} \in P_{w_l}^{w_m}(n)$. First, we concatenate the diagrams, written $d_{\pi_2} \circ d_{\pi_1}$, by putting d_{π_1} below d_{π_2} , concatenating the edges in the middle row of vertices, and then removing all connected components that lie entirely in the middle row of the concatenated diagrams. Let $c(d_{\pi_2}, d_{\pi_1})$ be the number of connected components that are removed from the middle row in $d_{\pi_2} \circ d_{\pi_1}$. Then the composition is defined, using infix notation, as

$$d_{\pi_2} \bullet d_{\pi_1} \coloneqq n^{c(d_{\pi_2}, d_{\pi_1})} (d_{\pi_2} \circ d_{\pi_1}) \tag{73}$$

Example C.3. If d_{π_2} is the two-coloured $(\circ \circ \bullet \circ \circ \circ, \circ \bullet \circ)$ -partition diagram

and d_{π_1} is the two-coloured (• • • • • • • • • • • • • • •)-partition diagram



then we have that $d_{\pi_2} \circ d_{\pi_1}$ is the $(\bullet \circ \bullet \circ \circ \circ, \circ \bullet \bullet \circ)$ -partition diagram

and so $d_{\pi_2} \bullet d_{\pi_1}$ is the diagram (76) multiplied by *n*, since one connected component was removed from the middle row of $d_{\pi_2} \circ d_{\pi_1}$.

Definition C.4 (Tensor Product). Let $d_{\pi_1} \in P_{w_k}^{w_l}(n)$ and $d_{\pi_2} \in P_{w_q}^{w_m}(n)$. Then $d_{\pi_1} \otimes d_{\pi_2}$ is defined to be the $(w_k \cdot w_q, w_l \cdot w_m)$ -partition diagram obtained by horizontally placing d_{π_1} to the left of d_{π_2} without any overlapping of vertices.

Example C.5. If d_{π_1} is the two-coloured $(\circ \bullet, \bullet \circ \bullet)$ -partition diagram

and d_{π_2} is the two-coloured $(\circ \bullet \circ \bullet, \circ \circ \bullet)$ -partition diagram

then we have that $d_{\pi_1} \otimes d_{\pi_2}$ is the $(\circ \bullet \circ \bullet \circ \bullet, \bullet \circ \bullet \circ \circ \circ)$ -partition diagram

Definition C.6 (Involution). Let $d_{\pi} \in P_{w_k}^{w_l}(n)$. Then $d_{\pi}^* \in P_{w_l}^{w_k}(n)$ is defined by reflecting d_{π} in the horizontal axis.

Example C.7. Recalling the two-coloured $(\circ \bullet \circ \circ)$ -partition diagram d_{π} from Example 7.3 (reprinted below for ease)

we have that d_{π}^* is the $(\bullet \circ \bullet \circ \circ, \circ \bullet \circ \circ \circ)$ -partition diagram

With this, we can now provide the following proof of Proposition 7.5.

Proof of Proposition 7.5. $\mathcal{P}(n)$ is a strict monoidal category because the bifunctor (horizontal operation) on objects reduces to the concatenation operation on words, which is associative, and the bifunctor on morphisms is the tensor product of linear combinations of two-coloured partition diagrams that is given in Definition C.4, which is associative by definition.

 $\mathcal{P}(n)$ is \mathbb{C} -linear because the morphism space between any two objects is by definition a vector space, and the composition of morphisms is \mathbb{C} -bilinear by definition. For the same reason, the bifunctor is also \mathbb{C} -bilinear.

Proof of Theorem 7.9. We prove this theorem in a series of steps. We first consider the image $C_{\mathcal{P}}(n)$ of the two-coloured partition category $\mathcal{P}(n)$ under $d_{\pi} \mapsto D_{\pi}$.

Step 1: We need to show that $C_{\mathcal{P}}(n)$ is actually a category.

The objects and the morphisms have been defined in the statement of the theorem. For every triple of objects, the composition rule is given by matrix multiplication, which is associative. Furthermore, the identity morphism for every object w_k of length k is the $n^k \times n^k$ identity matrix. Note also that the unit object is \emptyset , the empty word.

Step 2: We need to show that $\Theta : \mathcal{P}(n) \to \mathcal{C}_{\mathcal{P}}(n)$ is a functor.

- Θ maps the objects of $\mathcal{P}(n)$ to the objects of $\mathcal{C}_{\mathcal{P}}(n)$ since Θ maps each word w_k to itself, by definition.
- It is enough to show that Θ(g)Θ(f) = Θ(g f) on arbitrary basis elements of arbitrary morphism spaces where the codomain of f is the domain of g because the morphism spaces are vector spaces.

Let $f = d_{\pi_1}$ be a (w_k, w_l) -partition diagram, and let $g = d_{\pi_2}$ be a (w_l, w_m) -partition diagram. Then, by Definition C.2, we have that

$$d_{\pi_2} \bullet d_{\pi_1} = n^{c(d_{\pi_2}, d_{\pi_1})} d_{\pi_3} \tag{82}$$

where d_{π_3} be the composition $d_{\pi_2} \circ d_{\pi_1}$ expressed as a (w_k, w_m) -partition diagram. Then

$$\Theta(g \bullet f) = \Theta(d_{\pi_2} \bullet d_{\pi_1}) = n^{c(d_{\pi_2}, d_{\pi_1})} D_{\pi_3}$$
(83)

We also have that

$$\Theta(g)\Theta(f) = D_{\pi_2} D_{\pi_1} = \left(\sum_{I \in [n]^m, K \in [n]^l} \delta_{\pi_2, (I,K)} E_{I,K}\right) \left(\sum_{L \in [n]^l, J \in [n]^k} \delta_{\pi_1, (L,J)} E_{L,J}\right)$$
(84)

$$= \sum_{I \in [n]^m, K \in [n]^l, J \in [n]^k} \delta_{\pi_2, (I,K)} \delta_{\pi_1, (K,J)} E_{I,K} E_{K,J}$$
(85)

For fixed I, J, consider

$$\sum_{K \in [n]^l} \delta_{\pi_2, (I,K)} \delta_{\pi_1, (K,J)} \tag{86}$$

This is equal to

$$n^{c(d_{\pi_2},d_{\pi_1})}\delta_{\pi_3,(I,J)}$$
(87)

Indeed, for fixed $I, J, \delta_{\pi_3,(I,J)}$ is 1 if and only if both $\delta_{\pi_2,(I,K)}$ and $\delta_{\pi_1,(K,J)}$ are 1 for some $K \in [n]^l$ since d_{π_3} is the composition $d_{\pi_2} \circ d_{\pi_1}$. The number of such K is determined only by the vertices that appear in connected components that are removed from the middle row of $d_{\pi_2} \circ d_{\pi_1}$, since, for fixed I, J, only these vertices can be freely chosen if we want both $\delta_{\pi_2,(I,K)}$ and $\delta_{\pi_1,(K,J)}$ to be 1. However, since the entries in each connected component must be the same for both $\delta_{\pi_2,(I,K)}$ and $\delta_{\pi_1,(K,J)}$ to be 1, this implies that the number of $K \in [n]^l$ such that both $\delta_{\pi_2,(I,K)}$ and $\delta_{\pi_1,(K,J)}$ are 1 is $n^{c(d_{\pi_2},d_{\pi_1})}$.

Hence (85) becomes

$$\sum_{\in [n]^m, J \in [n]^k} n^{c(d_{\pi_2}, d_{\pi_1})} \delta_{\pi_3, (I, J)} E_{I, J} = n^{c(d_{\pi_2}, d_{\pi_1})} D_{\pi_3}$$
(88)

and so, by (83), we have that $\Theta(g)\Theta(f) = \Theta(g \bullet f)$, as required.

T

• For each object w_k in $\mathcal{P}(n)$, the identity morphism 1_{w_k} is the two-coloured (w_k, w_k) -partition diagram d_{π} where π is the set partition

$$\{1, k+1 \mid 2, k+2, |\cdots \mid k, 2k\}$$
(89)

of [2k], and its image under $d_{\pi} \mapsto D_{\pi}$ is the $n^k \times n^k$ identity matrix, as required.

Step 3: We need to show that $C_{\mathcal{P}}(n)$ is a strict *C*-linear monoidal category.

We first consider the tensor product of arbitrary basis elements of arbitrary morphism spaces. Let d_{π} be a (w_k, w_l) -partition diagram, and let d_{τ} be a (w_m, w_q) -partition diagram. Then, defining d_{ω} to be the $(w_k \cdot w_q, w_l \cdot w_m)$ -partition diagram, where ω is the set partition $\pi \cup \tau$ of [l + m + k + q], we see that

$$D_{\pi} \otimes D_{\tau} = \left(\sum_{I \in [n]^l, J \in [n]^k} \delta_{\pi, (I,J)} E_{I,J}\right) \otimes \left(\sum_{X \in [n]^m, Y \in [n]^q} \delta_{\tau, (X,Y)} E_{X,Y}\right)$$
(90)

$$= \sum_{(I,X)\in[n]^{l+m}, (J,Y)\in[n]^{k+q}} \delta_{\omega,(I,X),(J,Y)} E_{(I,X),(J,Y)}$$
(91)

$$D_{\omega}$$
 (92)

since $S_{\pi}((I,J)) \cup S_{\tau}((X,Y)) = S_{\omega}((I,X),(J,Y)).$

_

The bifunctor on objects of $C_{\mathcal{P}}(n)$ reduces to the concatenation operation on words, which is associative. Also, since the bifunctor on morphisms in $\mathcal{C}_{\mathcal{P}}(n)$ is the tensor product of linear combinations of images of two-coloured partition diagrams, it is associative because the tensor product of images of two-coloured partition diagrams is associative, both by the above calculation and because taking unions of set partitions is associative. Hence, $\mathcal{C}_{\mathcal{P}}(n)$ is a strict monoidal category.

 $C_{\mathcal{P}}(n)$ is \mathbb{C} -linear because the morphism space between any two objects is by definition a vector space, and the composition of morphisms is \mathbb{C} -bilinear by definition. For the same reason, the bifunctor is also \mathbb{C} -bilinear.

Step 4: We need to show that $\Theta : \mathcal{P}(n) \to \mathcal{C}_{\mathcal{P}}(n)$ is a strict \mathcal{C} -linear monoidal functor.

- 1. Θ preserves the tensor product on objects, since $\Theta(w_k) = w_k$ for any object in $\mathcal{P}(n)$, and the tensor product in both categories is given by the concatenation of words.
- 2. It is enough to show that $\Theta(f) \otimes \Theta(g) = \Theta(f \otimes g)$ on arbitrary basis elements of arbitrary morphism spaces as the morphism spaces are vector spaces. This result is now immediate from the calculation given in (90).
- 3. It is clear that Θ sends the unit object \emptyset in $\mathcal{P}(n)$ to \emptyset , which is the unit object in $\mathcal{C}_{\mathcal{P}}(n)$.
- 4. This is immediate because the map $d_{\pi} \mapsto D_{\pi}$ is \mathbb{C} -linear by Definition 7.7.

Step 5: We need to show that $\mathcal{C}_{\mathcal{P}}(n)$ is a two-coloured representation category.

The tensor product, concatenation and identity axioms have already been shown in previous steps. For the involution, if D_{π} is an element of $\operatorname{Hom}_{\mathcal{C}_{\mathcal{P}}(n)}(w_k, w_l)$ and is the image of a (w_k, w_l) -partition diagram d_{π} under Θ , then

$$(D_{\pi})^{*} = \left(\sum_{I \in [n]^{l}, J \in [n]^{k}} \delta_{\pi, (I, J)} E_{I, J}\right)^{*} = \sum_{J \in [n]^{k}, I \in [n]^{l}} \delta_{\pi^{*}, (J, I)} E_{J, I} = D_{\pi^{*}}$$
(93)

As D_{π^*} is the image of the (w_l, w_k) -partition diagram d_{π^*} under Θ , this implies that $(D_{\pi})^*$ is in $\operatorname{Hom}_{\mathcal{C}_{\mathcal{P}}(n)}(w_l, w_k)$, as required.

Finally, we see that the map $R: 1 \mapsto \sum_{i=1}^{n} e_i \otimes e_i$ is in both $\operatorname{Hom}_{\mathcal{C}_{\mathcal{P}}(n)}(\emptyset, \bullet \bullet)$ and $\operatorname{Hom}_{\mathcal{C}_{\mathcal{P}}(n)}(\emptyset, \bullet \circ)$ since it is the image under $d_{\pi} \mapsto D_{\pi}$ of the top-row pair partition diagram corresponding to the set partition $\{1, 2\}$ of $\{1, 2\}$ superimposed, on the one hand, with the word $w_2 := \bullet \bullet$, and on the other hand, with the word $w_2 := \bullet \circ$.

Step 6: We need to show that $d_{\pi} \mapsto D_{\pi}$ defines a strict \mathbb{C} -linear monoidal functor $\Theta : \mathcal{K}(n) \to \mathcal{C}_{\mathcal{K}}(n)$ such that $\mathcal{C}_{\mathcal{K}}(n)$ is a two-coloured representation category.

This is now immediate by Definition 7.6, the definition of a two-coloured category of partitions, and since the functor on $\mathcal{P}(n)$ is strict \mathbb{C} -linear monoidal whose image is a two-coloured representation category. In particular, the morphism spaces for $\mathcal{K}(n)$ are closed under composition, tensor product and involution, and so these are inherited by their image under Θ . Also, the map $R: 1 \mapsto \sum_{i=1}^{n} e_i \otimes e_i$ is in $\operatorname{Hom}_{\mathcal{C}_{\mathcal{K}}(n)}(\emptyset, \circ \bullet)$ and $\operatorname{Hom}_{\mathcal{C}_{\mathcal{K}}(n)}(\emptyset, \circ \circ)$ since the top row pair partition diagram corresponding to the set partition $\{1, 2\}$ of $\{1, 2\}$ superimposed either with $\circ \bullet$ or with $\bullet \circ$ is in both $\operatorname{Hom}_{\mathcal{K}(n)}(\emptyset, \circ \bullet)$ and $\operatorname{Hom}_{\mathcal{K}(n)}(\emptyset, \bullet \circ)$.

D. Additional Examples

We provide additional examples and content for the examples that were given in Section 8. We begin with one-coloured partition categories, which have been fully characterised by Raum & Weber (2016).

D.1. One-Coloured Partition Categories

One-coloured partition categories can be divided into four cases: group, non-crossing, half-liberated, and the rest. The two most important cases are the group and non-crossing ones: we refer the reader either to Banica et al. (2010) or to Gromada (2020) for results relating to the two other cases. The corresponding compact matrix quantum groups G(n) are known in the literature as **easy orthogonal compact matrix quantum groups**, or sometimes **Banica–Speicher quantum groups**, because $S_n \subseteq G(n) \subseteq O(n)^+$.

Group: Banica & Speicher (2009) provided a full characterisation: there are only six compact matrix groups such that $S_n \subseteq G(n) \subseteq O(n)^+$. Consequently, by Corollary 6.3, we obtain characterisations of six compact matrix group equivariant neural networks, of which two – the symmetric group S_n (Ravanbakhsh et al., 2017; Maron et al., 2019a; Pearce-Crump, 2024; Godfrey et al., 2023) and the orthogonal group O(n) (Pearce-Crump, 2023a) – were known previously to the machine learning community. Said differently, we have found characterisations of the weight matrices that appear in four new compact matrix group equivariant neural networks, two of which we discussed in Section 8. We state the result of Banica & Speicher (2009) below.

Theorem D.1 (Banica & Speicher (2009) [Theorem 2.8]). There are only six easy orthogonal compact matrix groups such that $S_n \subseteq G(n) \subseteq O(n)^+$ – the groups and the spanning sets that determine the weight matrices of Theorem 7.10 are given below. In what follows, when we refer to the image of a (k, l)-partition diagram, we mean the image under the map $d_{\pi} \mapsto D_{\pi}$ given in Definition 7.7.

- The symmetric group S_n : the image of all (k, l)-partition diagrams.
- The orthogonal group O(n): the image of all (k, l)-partition diagrams whose blocks come in pairs.
- The hyperoctahedral group H_n (the symmetry group of the hypercube): the image of all (k, l)-partition diagrams having blocks of even size.
- The bistochastic group B_n (the group of orthogonal matrices having sum 1 in each row and in each column): the image of all (k, l)-partition diagrams having blocks of size one or two.
- The modified symmetric group $S'_n \coloneqq \mathbb{Z}_2 \times S_n$: the image of all (k, l)-partition diagrams where the number of blocks of odd size is even.
- The modified bistochastic group B'_n := Z₂ × B_n: the image of all (k, l)-partition diagrams having an even number of blocks of size one and any number of blocks of size two.

Remark D.2. We know from Godfrey et al. (2023) that we can improve upon the spanning set for the symmetric group S_n given in Theorem D.1 by removing all (k, l)-partition diagrams that have more than n blocks from the set of all (k, l)-partition diagrams before taking their image. The resulting set of matrices is the diagram basis for S_n .

Non-Crossing: In order to state the classifications for the non-crossing case, we first need the following definition.

Definition D.3. A set partition diagram d_{π} corresponding to a set partition π of [l + k] is said to be **crossing** if there exist four integers $1 \le x_1 < x_2 < x_3 < x_4 \le l + k$ satisfying:

- 1. x_1 and x_3 are in the same block
- 2. x_2 and x_4 are in the same block, and
- 3. x_1 and x_2 are not in the same block.

Otherwise, d_{π} is said to be **non-crossing**.

The six easy orthogonal compact matrix groups that were found by Banica & Speicher (2009) were **liberated** by the same authors to create six compact matrix quantum groups – which are true quantum groups – that fall into the non-crossing case. Liberation means that crossings are not allowed in any of the set partition diagrams. Weber (2013) found a seventh using this method. We call these compact matrix quantum groups **free** since their categories of partitions $\mathcal{K}(n)$ only contain non-crossing set partition diagrams. We have the following result.



Figure 1. The fifteen (2, 2)-partition diagrams.

Theorem D.4. There are only seven free easy orthogonal compact matrix quantum groups. We list the quantum groups and the spanning sets that determine the weight matrices of Theorem 7.10 below.

The first six are the symmetric quantum group S_n^+ , the orthogonal quantum group $O(n)^+$, the hyperoctahedral quantum group H_n^+ , the bistochastic quantum group B_n^+ , the modified symmetric quantum group $S_n'^+$, and the modified bistochastic quantum group $B_n'^+$, whose spanning sets are determined by the same sets of set partition diagrams that were given in Theorem D.1 for their "sister" groups, but with all crossing set partition diagrams removed.

The seventh is the freely modified bistochastic quantum group $B_n^{\#+}$. To obtain its corresponding spanning set, we consider the same set of set partition diagrams as for $B_n'^+$, but now temporarily label each vertex with a colour from $\{\circ, \bullet\}$ in an alternating fashion, starting with \bullet in the top right corner and moving counterclockwise, ultimately finishing in the bottom right corner. We only retain those set partition diagrams where the blocks of size two pair one \circ with one \bullet , and take its image under $d_{\pi} \mapsto D_{\pi}$, removing the colours in the process.

In fact, we obtain a stronger classification for the non-crossing case due to the following theorem.

Theorem D.5 (Banica & Speicher (2009) [Theorem 3.8]). The spanning sets for the seven compact matrix quantum groups in the non-crossing case are bases when $n \ge 4$.

Remark D.6. By Theorem 7.10, note that each easy compact matrix quantum group (G(n), u) that appears in Theorems D.1 and D.4 is *defined by* the two-coloured category of partitions $\mathcal{K}(n)$ that is associated with it. That is, Banica & Speicher (2009) actually showed that there are only six one-coloured category of partitions that contain the swap partition diagram, that is, the (2, 2)-partition diagram



and together with Weber (2013) they showed that there are only seven one-coloured category of partitions that are noncrossing. With these restrictions on the possible one-coloured category of partitions in place they could immediately apply Woronowicz–Tannaka–Krein duality to construct the easy compact matrix quantum groups from each of these one-coloured category of partitions. In fact, we have, again by Theorem 7.10, that each compact matrix quantum group is precisely the universal C^* -algebra

$$C(G(n)) = C^*(u_{i,j}, 1 \le i, j \le n \mid D_\pi u^{\otimes w_k} = u^{\otimes w_l} D_\pi \text{ for all } d_\pi \in \mathcal{K}(n)(w_k, w_l))$$

$$(95)$$

Example D.7. We calculate the weight matrices that appear in the compact matrix quantum group equivariant neural networks for all of the easy orthogonal compact matrix quantum groups that appear in the group and non-crossing case when n = k = l = 2. Recall from Theorems D.1 and D.4 that the weight matrices are formed as a weighted linear combination of the image under $d_{\pi} \mapsto D_{\pi}$ of some subset of all of the possible (2, 2)-partition diagrams.

There are B(4) = 15 (2, 2)-partition diagrams in total, where B(4) is the fourth Bell number. They are shown in Figure 1. We number the diagrams from i = 1 to 15, going from left to right and then top to bottom, assigning a weight w_i to each diagram. **Procedure 2:** How to Calculate the (I, J)-entry of each Equivariant Spanning Set Matrix D_{π} from $((\mathbb{C}^n)^{\otimes k}, u^{\otimes w_k}) \to ((\mathbb{C}^n)^{\otimes l}, u^{\otimes w_l})$ for an Easy Compact Matrix Quantum Group (G(n), u).

We assume that d_{π} is a two-coloured (w_k, w_l) -partition diagram in the two-coloured category of partitions $\mathcal{K}(n)$ that uniquely determines (G(n), u). We perform the following steps:

- 1. Place the indices I on the top row of d_{π} and the indices J on the bottom row of d_{π} .
- 2. If all of the vertices in each block in d_{π} have been overlaid with the same number, then the (I, J) entry of D_{π} is 1, otherwise it is 0.

Figure 2 shows the images of the fifteen (2, 2)-partition diagrams under the map $d_{\pi} \mapsto D_{\pi}$. We apply the characterisations of Theorems D.1 and D.4 to obtain the following results.

For the group case, we have that

- The Symmetric Group S_n : The weight matrix is a weighted linear combination of the images of all fifteen (2, 2)-partition diagrams.
- The Orthogonal Group O(n): The weight matrix is a weighted linear combination of the images of diagrams six, seven and eight.
- The Hyperoctahedral Group H_n : The weight matrix is (coincidentally) the same as for the orthogonal group in this case, since the (2, 2)-partition diagrams with blocks of even size are the same as the (2, 2)-partition diagrams whose blocks come in pairs.
- The Bistochastic Group B_n : The weight matrix is a weighted linear combination of the images of diagrams six through fifteen, inclusive.
- The Modified Symmetric Group S'_n : The weight matrix is (coincidentally) the same as for the symmetric group, since the number of blocks of odd size in any (2, 2)-partition diagram is either 0, 2 or 4.
- The Modified Bistochastic Group B'_n : The weight matrix is (coincidentally) the same as for the bistochastic group, since the number of blocks of odd size in any (2, 2)-partition diagram is either 0, 2 or 4.

and for the non-crossing case, we have that

- The Symmetric Quantum Group S_n^+ : The weight matrix is a weighted linear combination of the images of all diagrams except for seven, which is the only crossing diagram.
- The Orthogonal Quantum Group $O(n)^+$: The weight matrix is a weighted linear combination of the images of diagrams six and eight.
- The Hyperoctahedral Quantum Group H_n^+ : The weight matrix is the same as for $O(n)^+$.
- The Bistochastic Quantum Group B_n^+ : The weight matrix is a weighted linear combination of the images of diagrams six and eight through fifteen, inclusive.
- The Modified Symmetric Quantum Group S'_n^+ : The weight matrix is the same as for S_n^+ .
- The Modified Bistochastic Quantum Group B'_n : The weight matrix is the same as for B_n^+ .
- The Freely Modified Bistochastic Quantum Group $B_n^{\#+}$: We colour in the diagrams that appear in the weight matrix for B_n^+ , and retain only the diagrams where the blocks of size two pair one \circ with one \bullet . Consequently, we remove diagrams eleven and fourteen from the set, and so the weight matrix is a weighted linear combination of the images of diagrams six, eight, nine, ten, twelve, thirteen and fifteen.

d_{π}	D_{π}	d_{π}	D_{π}
	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$\begin{bmatrix} 1 & 2 \\ 0 & \cdots & 0 \\ 0 & 0 \end{bmatrix}$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$\begin{bmatrix} 1 & 2 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$ \begin{array}{c} 1 & 2 \\ $	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
$\begin{bmatrix} 1 & 2 \\ 0 & \\ 3 & 4 \end{bmatrix}$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$\begin{bmatrix} 1 & 2 \\ 0 & 0 \end{bmatrix}$ $\begin{bmatrix} 0 & -0 \\ 3 & 4 \end{bmatrix}$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$\begin{bmatrix} 1 & 2 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
$ \begin{array}{c} 1 & 2 \\ $	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
$ \begin{array}{c} 1 \\ 2 \\ 3 \\ 4 \end{array} $	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$	$ \begin{array}{cccccccccccccccccccccccccccccccccccc$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$
$\begin{bmatrix} 1 & & 2 \\ 0 & & 0 \\ 0 & & 0 \end{bmatrix}$	$\begin{array}{c ccccccccccccccccccccccccccccccccccc$		

Figure 2. For n = 2, we display the images under the map $d_{\pi} \mapsto D_{\pi}$ for each of the (2, 2)-partition diagrams.

D.2. Two-Coloured Partition Categories

Two-coloured partition categories are much richer than one-coloured partition categories because the generators $u_{i,j}$ of the corresponding compact matrix quantum groups are no longer self-adjoint. They were first investigated by Freslon & Weber (2016). Unfortunately, at the time of writing, a full characterisation of these categories is unknown. However, Tarrago & Weber (2016; 2018) classified all two-coloured partition categories in the group and non-crossing case. In particular, there are seven series of two-coloured partition categories in the group case and twelve series of two-coloured partition categories in the group case and twelve series of two-coloured partition categories in the non-crossing case. The corresponding compact matrix quantum groups G(n) are known as **easy unitary compact matrix quantum groups**, since they satisfy $S_n \subseteq G(n) \subseteq U(n)^+$. We focus on the characterisation for only the two most important series: the unitary group U(n) and the unitary quantum group $U(n)^+$, referring the reader to Tarrago & Weber (2016, Theorem 5.3) for the characterisations of the other series.

Tarrago & Weber (2016; 2018) showed that the two-coloured category of partitions for the unitary group U(n) has, for fixed words w_k and w_l , a morphism space that is spanned by all two-coloured (w_k, w_l) -partition diagrams whose blocks come in pairs such that if two vertices of a block are in the same row, then they have different colours, otherwise they have the same colours. Tarrago & Weber (2016; 2018) also showed that the two-coloured category of partitions for the unitary quantum group $U(n)^+$ has, for words w_k and w_l , a morphism space that is spanned by all *non-crossing* two-coloured (w_k, w_l) -partition diagrams satisfying the same conditions as those for the unitary group. As before, we obtain the $n^l \times n^k$ weight matrices of Theorem 7.10 for these compact matrix quantum groups by taking the images of the two-coloured partition diagrams that span the respective morphism spaces under the map $d_{\pi} \mapsto D_{\pi}$ given in Definition 7.7.

This introduces the rather interesting point for the unitary group U(n). If we pick two words w_k and w_l , then the $n^l \times n^k$ weight matrix corresponding to these words that appears in a compact matrix quantum group equivariant neural network for the unitary group U(n) is the \mathbb{C} -linear span of the image of all (w_k, w_l) -partition diagrams that live in the two-coloured partition category for U(n). However, if we consider compact matrix group equivariant neural networks for U(n) instead, as the words can only be formed from white points, then not only must we have that k equals l, by the characterisation of the two-coloured partition category for U(n), but also that the $n^k \times n^k$ weight matrix is the image of all permutations in the symmetric group S_k , expressed as (k, k)-partition diagrams. We see that this is the classic version of Schur–Weyl duality. *Example* D.8. We calculate the weight matrices that appear in compact matrix quantum group equivariant neural networks for the unitary group U(n) and unitary quantum group $U(n)^+$ when n = k = l = 2.

The easiest way to do this is to consider the one-coloured partitions that come in pairs and then superimpose colours onto each row of vertices to see which two-coloured partitions match the classification result.

The one-coloured (2,2)-partition diagrams that come in pairs are



We label the diagrams as 1 to 3 from left to right. Note that the matrices obtained under $d_{\pi} \mapsto D_{\pi}$ are the same no matter the colours on the vertices, hence we can reuse the matrices for diagrams six, seven and eight from Figure 2 in order to calculate the weight matrices.

We now consider pairs of words. As there are only 4 words of length 2, this implies that there are 16 such pairs of words. We obtain the following result for the weight matrices for U(2), where for each cell, we take weighted linear combinations of the matrices D_{π} for each set partition diagram in the cell.

Similarly, for $U(2)^+$, the weight matrices are weighted linear combinations of the images of the following set partition

diagrams, where we have removed the second diagram from (97).

Note that the top left hand cell of (97) provides a classification of the weight matrix appearing in the compact matrix group equivariant neural network for U(2) when n = k = l = 2.