

Harmonics of Learning: Universal Fourier Features Emerge in Invariant Networks

Giovanni Luca Marchetti
KTH Royal Institute of Technology

GLMA@KTH.SE

Christopher J Hillar
Redwood Center for Theoretical Neuroscience

HILLARMATH@GMAIL.COM

Danica Kragic
KTH Royal Institute of Technology

DANI@KTH.SE

Sophia Sanborn
Science Corporation

SOPHIAS@SCIENCE.XYZ

Editors: Shipra Agrawal and Aaron Roth

Abstract

In this work, we formally prove that, under certain conditions, if a neural network is invariant to a finite group then its weights recover the Fourier transform on that group. This provides a mathematical explanation for the emergence of Fourier features – a ubiquitous phenomenon in both biological and artificial learning systems. The results hold even for non-commutative groups, in which case the Fourier transform encodes all the irreducible unitary group representations. Our findings have consequences for the problem of symmetry discovery. Specifically, we demonstrate that the algebraic structure of an unknown group can be recovered from the weights of a network that is at least approximately invariant within certain bounds. Overall, this work contributes to a foundation for an algebraic learning theory of invariant neural network representations.

Keywords: Invariant neural networks, harmonic analysis, group representations




Figure 1: Weights learned by a neural network trained for invariance to planar rotations resemble circular harmonics. Data from [Sanborn et al. \(2023\)](#).

1. Introduction and Related Work

Artificial neural networks trained on natural data exhibit a striking phenomenon: regardless of exact initialization, dataset, or training objective, models trained on the same data domain frequently converge to similar learned representations (Li et al., 2015). For example, the early layer weights of diverse image models tend to converge to Gabor filters and color-contrast detectors (Olah et al., 2020). Remarkably, many of these same features are observed in the visual cortex (Hubel and Wiesel, 1959; Hass and Horwitz, 2013; Willeke et al., 2023), suggesting a form of representational universality that transcends biological and artificial substrates. While such findings are empirically well-established (Räuker et al., 2023), the field lacks theoretical explanations.

Spatially localized versions of canonical 2D Fourier basis functions, such as Gabor filters or wavelets, are perhaps the most frequently observed universal features. They commonly arise in the early layers of vision models – trained with efficient coding (Olshausen and Field, 1997; Bell and Sejnowski, 1997), classification (Olah et al., 2020), temporal coherence (Hurri and Hyvärinen, 2003), and next-step prediction (Fiquet and Simoncelli, 2023) objectives – as well as in the primary visual cortices of diverse mammals – including cats (Hubel and Wiesel, 1962), monkeys (Hubel and Wiesel, 1968), and mice (Dräger, 1975). Non-localized Fourier features have been observed in networks trained to solve tasks that permit cyclic wraparound – for example, modular arithmetic (Nanda et al., 2023), more general group compositions (Chughtai et al., 2023), or invariance to the group of cyclic translations (Sanborn et al., 2023). In the domain of spatial navigation, the so-called *grid cells* of the entorhinal cortex (Moser et al., 2008) display periodic firing patterns at different spatial frequencies as they build a map of space. Their response properties are naturally modeled with the harmonics of the twisted torus (Guanella et al., 2007; Orchard et al., 2013; Gardner et al., 2022). Similar features also emerge in artificial neural networks trained to solve spatial navigation tasks (Banino et al., 2018; Cueva and Wei, 2018; Sorscher et al., 2019; Dorrell et al., 2022; Sorscher et al., 2023; Schaeffer et al., 2023). The ubiquity of these features across diverse learning systems is both striking and unexplained.

In this work, we provide a mathematical explanation for the emergence of Fourier features in learning systems such as neural networks. We argue that the mechanism responsible for this emergence is the downstream *invariance* of the learner to the action of a *group of symmetries* (e.g., planar translations or rotations). Since natural data typically possess symmetries, invariance is a fundamental bias that is injected both implicitly and sometimes explicitly into learning systems (Bronstein et al., 2021; Cohen and Welling, 2016; Kondor and Trivedi, 2018). Motivated by this, we derive theoretical guarantees for the presence of Fourier features in invariant learners that apply to a broad class of machine learning models.

Our results rely on the inextricable link between harmonic analysis and group theory (Folland, 2016). The standard discrete Fourier transform is a special case of more general Fourier transforms on groups, which can be defined by replacing the standard basis of harmonics by irreducible unitary group representations. The latter are equivalent to the familiar definition for cyclic or, more generally, commutative groups, but are more involved for non-commutative ones. In order to accommodate both scenarios, we develop a general theory that applies, in principle, to arbitrary finite groups.

This work represents an attempt to provide mathematical grounding for a general algebraic theory of representation learning, while addressing the *universality hypothesis* for neural networks (Olah et al., 2020; Moschella et al., 2022). A suite of earlier theoretical works (Isely et al., 2010; Hillar and Sommer, 2015; Garfinkle and Hillar, 2019) established such universality for sparse coding

models (Olshausen and Field, 1997), deriving the conditions under which a network will recover the original bases that generate data through sparse linear combinations. In this case, the statistics of the data determine the uniqueness of the representation. Our findings, on the other hand, are purely *algebraic*, since they rely exclusively on the invariance properties of the learner. Given the centrality of invariance to many machine learning tasks, our theory encompasses a broad class of scenarios and neural network architectures, while providing a new perspective on classical neuroscience (Pitts and McCulloch, 1947; Hubel and Wiesel, 1962). As such, it sets a foundation for a learning theory of representations in artificial and biological neural systems, grounded in the mathematics of symmetry.

1.1. Overview of Results

In this section, we provide a non-technical overview of the theoretical results presented in this work. Our main result can be summarized as follows.

Informal Theorem 1 (Theorem 10 and Corollary 11) *If $\varphi(W, x)$ is a parametric function of a certain kind that is invariant in the input variable x to the action of a finite group G , then each component of its weights W coincides with a harmonic of G up to a linear transformation. In particular, when the weights are orthonormal, W coincides with the Fourier transform of G up to linear transformations.*

Here, the term “harmonic” refers to an irreducible unitary representation of G . In particular, one-dimensional unitary representations correspond to homomorphisms with the unit circle $U(1) \subseteq \mathbb{C}$, which is reminiscent of the classical definition via the imaginary exponential. However, *non-commutative* groups can have higher-dimensional irreducible representations, intuitively meaning that harmonics are valued in unitary matrices. In this case, the components of W can be interpreted as *capsules* in the sense of Hinton et al. (2018); i.e., neural units processing matrix-valued signals.

Since harmonic analysis is naturally formalized over the complex numbers, we consider models with complex weights W , which fits into the broader program of complex-valued machine learning (Bassey et al., 2021; Trabelsi et al., 2017; Löwe et al., 2022). We show that the hypothesis on φ in Theorem 1 is satisfied by several machine learning models from the literature. In particular, the theorem applies to the recently-introduced (Bi)Spectral Networks (Sanborn et al., 2023), to single fully-connected layers of McCulloch-Pitts neurons, and, to an extent, to traditional deep networks. As an additional contribution, we generalize Spectral Networks to non-commutative groups.

The group-theoretical Fourier transform encodes the entire group structure of G . Therefore, as a consequence of Theorem 1, the multiplication table of G can be recovered from the weights W of an invariant parametric function φ – a fact empirically demonstrated by Sanborn et al. (2023). This addresses the question of *symmetry discovery* – an established machine learning problem aiming to recover the unknown group of symmetries of data with minimal supervision and prior knowledge (Rao and Ruderman, 1998; Sohl-Dickstein et al., 2010; Desai et al., 2022). Since the multiplication table is a discrete object, it is expected that the invariance constraint on φ can be loosened while still recovering the group correctly. To this end, we prove the following.

Informal Theorem 2 (Theorem 17) *If $\varphi(W, x)$ is “almost invariant” to G according to certain functional bounds and the weights are “almost orthonormal”, then the multiplication table of G can be recovered from W .*

Lastly, we implement a model satisfying the requirements of our theory and demonstrate its symmetry discovery capabilities. To this end, we train it via contrastive learning on an objective encouraging invariance and extract the multiplication table of G from its weights. Our Python implementation is available at a public repository¹.

2. Mathematical Background

We begin by introducing the fundamental concepts from harmonic analysis and group theory used in this paper. For a complete treatment, we refer the reader to [Steinberg \(2012\)](#).

2.1. Groups and Actions

A *group* is an algebraic object whose elements represent abstract symmetries, which can be composed and inverted.

Definition 3 A *group* is a set G equipped with a multiplication map $G \times G \rightarrow G$ denoted by $(g, h) \mapsto gh$, an inversion map $G \rightarrow G$ denoted by $g \mapsto g^{-1}$, and a distinguished identity element $1 \in G$ such that for all $g, h, k \in G$:

$$\begin{array}{lll} \text{Associativity} & \text{Inversion} & \text{Identity} \\ g(hk) = (gh)k & g^{-1}g = gg^{-1} = 1 & g1 = 1g = g \end{array}$$

A map $\rho : G \rightarrow G'$ between groups is called a homomorphism if $\rho(gh) = \rho(g)\rho(h)$ for all $g, h \in G$.

Examples of groups include the permutations of a set and the general linear group $\text{GL}(V)$ of invertible operators over a vector space V , both equipped with the usual composition and inversion of functions. A further example that will be relevant in this work is the *unitary* group $\text{U}(V) \subseteq \text{GL}(V)$ associated to a Hilbert space V , consisting of operators U satisfying $UU^\dagger = I$, where I is the identity matrix.² Groups satisfying $gh = hg$ for all $g, h \in G$ are deemed *commutative*. The idea of a space \mathcal{X} having G as a group of symmetries is abstracted by the notion of *group action*.

Definition 4 An *action* by a group G on a set \mathcal{X} is a map $G \times \mathcal{X} \rightarrow \mathcal{X}$ denoted by $(g, x) \mapsto g \cdot x$, satisfying for all $g, h \in G$, $x \in \mathcal{X}$:

$$\begin{array}{ll} \text{Associativity} & \text{Identity} \\ g \cdot (h \cdot x) = (gh) \cdot x & 1 \cdot x = x \end{array}$$

A map $\varphi : \mathcal{X} \rightarrow \mathcal{Z}$ between sets acted upon by G is called *equivariant* if $\varphi(g \cdot x) = g \cdot \varphi(x)$ for all $g \in G, x \in \mathcal{X}$. It is called *invariant* if moreover G acts trivially on \mathcal{Z} or, explicitly, if $\varphi(g \cdot x) = \varphi(x)$.

In general, the following actions can be defined for arbitrary groups: G acts on any set *trivially* by $g \cdot x = x$, and G acts on itself seen as a set via (left) *multiplication* by $g \cdot h = gh$. Further examples are $\text{GL}(V)$ and $\text{U}(V)$ acting on V by evaluating operators.

1. <https://github.com/sophiaas/spectral-universality>

2. Given a linear operator $U : H \rightarrow H'$ between Hilbert spaces, we denote by $U^\dagger : H' \rightarrow H$ its adjoint operator, defined by $\langle Ux, y \rangle = \langle x, U^\dagger y \rangle$ for all $x \in H, y \in H'$.

2.2. Harmonic Analysis on Groups

Harmonic analysis on groups (Folland, 2016) generalizes standard harmonic analysis. We focus here on finite groups for simplicity, which are sufficient for practical applications. This avoids technicalities such as integrability conditions and continuity issues arising for infinite groups. We start by considering commutative groups and cover non-commutative ones in Section 2.3.

Let G be a finite commutative group of order $|G|$. Denote by $\langle G \rangle = \mathbb{C}^G$ the free complex vector space generated by G . Intuitively, an element $x = (x_g)_{g \in G} \in \langle G \rangle$ represents a complex-valued signal over G . The space $\langle G \rangle$ is endowed with the *convolution* product:

$$(x \star y)_g = \sum_{h \in G} x_h y_{h^{-1}g}, \quad (1)$$

and is acted upon by G via $g \cdot x = \delta_g \star x = (x_{g^{-1}h})_{h \in G}$, where δ_g is the canonical basis vector.

Definition 5 *The dual G^\vee of G is the set of homomorphisms $\rho : G \rightarrow \text{U}(\mathbb{C})$, where $\text{U}(\mathbb{C}) \subseteq \mathbb{C}$ is the group of unitary complex numbers equipped with multiplication. It is itself a group when equipped with pointwise composition $(\rho\mu)(g) = \rho(g)\mu(g)$.*

A homomorphism $\rho \in G^\vee$ intuitively represents a *harmonic* over G , generalizing the familiar notion from signal processing. If we endow $\langle G \rangle$ with the canonical scalar product $\langle x, y \rangle = \sum_{g \in G} \bar{x}_g y_g$, then $G^\vee \subseteq \langle G \rangle$ forms an orthogonal basis with all the norms equal to $\sqrt{|G|}$. The linear base-change is, by definition, the Fourier transform over $\langle G \rangle$:

Definition 6 *The Fourier transform is the map $\langle G \rangle \rightarrow \langle G^\vee \rangle$, $x \mapsto \hat{x}$, defined for $\rho \in G^\vee$ as $\hat{x}_\rho = \langle \rho, x \rangle$.*

The Fourier transform is a linear isometry or, equivalently, a unitary operator, up to a multiplicative constant of $|G|$. Moreover, it exchanges the convolution product \star over $\langle G \rangle$ with the Hadamard product \odot over $\langle G^\vee \rangle$. Definition 6 generalizes the usual discrete Fourier transform in the following sense. For an integer $d > 0$, consider the cyclic group C_d with d elements. Concretely, $C_d = \mathbb{Z}/d\mathbb{Z}$ is the group of integers modulo d equipped with addition as composition. The dual G^\vee consists of homomorphisms of the form $\mathbb{Z}/d\mathbb{Z} \ni g \mapsto e^{2\pi\sqrt{-1}gk/d}$ for $k \in \{0, \dots, d-1\}$. Definition 6 specializes then to the familiar Fourier transform.

2.3. Non-Commutative Harmonic Analysis

So far, we have assumed that G is commutative. In this section we briefly discuss the extension of Fourier theory to non-commutative groups. This however requires more elaborate theoretical tools, which we now introduce. To begin, in order to perform harmonic analysis on general groups it is necessary to discuss unitary representations. The latter will play the role of *matrix-valued harmonics*.

Definition 7 *A unitary representation of G is an action by G on a finite-dimensional Hilbert space V via unitary operators or, in other words, a homomorphism $\rho_V : G \rightarrow \text{U}(V)$. A unitary representation is said to be irreducible if V does not contain any non-trivial³ sub-representations.*

We denote by $\text{Irr}(G)$ the set of all irreducible representations of G up to isomorphism. Moreover, for a vector space V we denote by $\text{End}(V)$ the space of its linear operators.

3. The trivial sub-representations of V are 0 and V .

Definition 8 *The Fourier transform is the map $\langle G \rangle \rightarrow \bigoplus_{\rho_V \in \text{Irr}(G)} \text{End}(V)$, $x \mapsto \hat{x}$, defined for $\rho_V \in \text{Irr}(G)$ as:*

$$\hat{x}_{\rho_V} = \sum_{g \in G} \rho_V(g)^\dagger x_g \in \text{End}(V). \quad (2)$$

This generalizes Definition 6 since for a commutative group, ρ_V is irreducible if, and only if, $\dim(V) = 1$. Analogously to the commutative setting, the Fourier transform exchanges the convolution product \star with the point-wise operator composition, which we still denote by \odot . Moreover, the Fourier transform is a unitary operator up to a multiplicative constant of $|G|$ with respect to the normalized Hilbert-Schmidt scalar product on $\text{End}(V)$, given by $\langle A, B \rangle = \dim(V) \text{tr}(A^\dagger B)$. The norm associated to the Hilbert-Schmidt scalar product is the Frobenius norm. The relations between irreducible unitary representations coming from the unitarity of the Fourier transform are known as *Schur orthogonality* relations.

3. Theoretical Results

We now present the primary theoretical contributions of this work. Concretely, we demonstrate that if certain parametric functions are invariant to a group then their weights must almost coincide with harmonics, i.e. irreducible unitary group representations. We start by introducing general algebraic notions and principles, and then proceed to specialize them to machine learning scenarios.

Let G be a finite group, \mathcal{H} be a set, and V_1, \dots, V_k be complex finite-dimensional Hilbert spaces. In what follows, we will consider the space:

$$\mathcal{W} = \langle G \rangle \otimes \bigoplus_i \text{End}(V_i) \simeq \bigoplus_i \text{End}(V_i)^{\oplus G}. \quad (3)$$

\mathcal{W} is a Hilbert space when endowed with the scalar product given by the product of the canonical scalar product over $\langle G \rangle$ and the normalized Hilbert-Schmidt scalar products over $\text{End}(V_i)$. For $W \in \mathcal{W}$, we will denote each of its components as $W_i = (W_i(g))_g \in \text{End}(V_i)^{\oplus G}$. Moreover, we will often interpret elements $W \in \mathcal{W}$ as linear maps $\langle G \rangle \rightarrow \bigoplus_i \text{End}(V_i)$ via $W(x) = \sum_{g \in G} W(g)x_g$ for $x \in \langle G \rangle$, where $W(g) = (W_i(g))_i$. Note that G acts on the left tensor factor of \mathcal{W} while for every i , $U(V_i)$ acts on the right tensor factor of $\langle G \rangle \otimes \text{End}(V_i)$ by composition of operators.

Example 1 *Concretely, given coordinates on $V_i \simeq \mathbb{C}^{d_i}$ for $i = 1, \dots, k$, an element of \mathcal{W} consists of a k -tuple of tensors $W_i \in \mathbb{C}^{d_i \times d_i \times |G|} \simeq \text{End}(V_i)^{\oplus G}$. For example, the dihedral group D_n of isometries of a regular n -gon has $2n$ elements. For n odd, it has 2 irreducible unitary representations of dimension 1, and $(n-1)/2$ of dimension 2, implying that \mathcal{W} consists of a direct sum of copies of \mathbb{C}^{2n} and $\mathbb{C}^{2 \times 2 \times 2n}$.*

Definition 9 *We say that a map $\varphi: \mathcal{W} \rightarrow \mathcal{H}$ has unitary symmetries if for all $W, W' \in \mathcal{W}$ such that $\|W_i\| = \|W'_i\|$ for all i and $\varphi(W) = \varphi(W')$, we have that for every i there exists a unitary operator $U_i \in U(V_i)$ such that $W_i = U_i \cdot W'_i$.*

In the context of machine learning, \mathcal{H} will represent the *hypothesis space*, consisting of functions the model can learn. On the other hand, φ will represent the parametrization of such hypotheses, with its domain \mathcal{W} being the space of weights. Each component $\text{End}(V_i)$ of \mathcal{W} will be responsible for parametrizing a computational unit, i.e. a complex-valued *neuron* in the language of neural networks.

For commutative groups, we simply have $V_i = \mathbb{C} \simeq \text{End}(V_i)$. In general, $\text{End}(V_i)$ can be thought of as parametrizing matrix-valued signals, which, as mentioned in Section 1.1, are computed by neural units sometimes referred to as *capsules* (Hinton et al., 2018; Sabour et al., 2017). Lastly, the components of $\langle G \rangle$ will represent the input space. The fact that the latter consists of scalar signals over G is a simplification of several practical scenarios. However, the results of this section can be generalized, to an extent, to signals over a set acted upon by G – see Section 3.1. The following is an algebraic principle at the core of this work.

Theorem 10 *Suppose that $\varphi: \mathcal{W} \rightarrow \mathcal{H}$ has unitary symmetries and that for some $W \in \mathcal{W}$ the following holds:*

- $\varphi(g \cdot W) = \varphi(W)$ for all $g \in G$.
- W_i , seen as a linear map $\langle G \rangle \rightarrow \text{End}(V_i)$, is surjective for all i .

Then for every i there exist $W'_i \in \text{End}(V_i)$ and an irreducible unitary representation $\rho_i: G \rightarrow \text{U}(V_i)$ such that for all $g \in G$,

$$W_i(g) = W'_i \rho_i(g)^\dagger. \quad (4)$$

We refer to the Appendix for a proof. Note that the surjectivity assumption implies the constraint $\dim(V_i)^2 \leq |G|$ for all i . As a consequence of the result above, the full Fourier transform arises with an additional orthogonality assumption.

Corollary 11 *Suppose that $\varphi: \mathcal{W} \rightarrow \mathcal{H}$ has unitary symmetries and that for some $W \in \mathcal{W}$ the following holds:*

- $\varphi(g \cdot W) = \varphi(W)$ for all $g \in G$.
- W is unitary up to a multiplicative constant, i.e. $W^\dagger W = |G|I$.

Then W is the Fourier transform up to composing each of the components W_i by an operator with Frobenius norm equal to 1.

We refer to the Appendix for a proof. Again, the orthogonality assumption implies that V_1, \dots, V_k are the ambient Hilbert spaces of all the irreducible unitary representations of G up to isomorphism, and in particular $\sum_i \dim(V_i)^2 = |G|$.

We now wish to discuss the other crucial assumption of Theorem 10 requiring that $\varphi(g \cdot W) = \varphi(W)$ for all $g \in G$, which is reminiscent of invariance. However, when \mathcal{H} is a space of functions, we are typically interested in models that are invariant in the input variable rather than the weight variable. Therefore, we introduce the following condition, aimed at reconciling inputs and weights. To this end, suppose that \mathcal{H} is a set of functions $\mathcal{X} \rightarrow \mathcal{Y}$, where \mathcal{X} is a set acted upon by G and \mathcal{Y} is a set. We adhere to the notation $\varphi(W, x) = \varphi(W)(x)$, $x \in \langle G \rangle$, $W \in \mathcal{W}$, for simplicity.

Definition 12 *We say that φ satisfies the adjunction property if $\varphi(W, g \cdot x) = \varphi(g^{-1} \cdot W, x)$ for all $x \in \mathcal{X}, g \in G$.*

The adjunction property implies that if $\varphi(W, x)$ is invariant in x , then $\varphi(g \cdot W, x) = \varphi(W, x)$ for all x, g , recovering the assumption of Theorem 10.

3.1. Extensions

As explained above, the tensor component $\langle G \rangle$ of \mathcal{W} typically represents the input space of a given machine learning model. However, it is often the case that data does not consist of signals over G . This happens, for example, in the case of image data acted upon by the cyclic group via rotations (see Figure 1), since the pixel plane is composed of several copies of G .

One can consider the more general scenario when data consist of complex signals over a finite set \mathcal{S} acted upon by G , therefore replacing $\langle G \rangle = \mathbb{C}^G$ by $\mathbb{C}^{\mathcal{S}}$. Assuming the action over \mathcal{S} is *free*, meaning that $g \cdot s = s$ implies $g = 1$, \mathcal{S} can be decomposed into copies of G deemed *orbits*. Specifically, there is an equivariant isomorphism $\mathcal{S} \simeq G \sqcup \dots \sqcup G$, which in turn induces a linear isomorphism $\mathbb{C}^{\mathcal{S}} \simeq \langle G \rangle^{\oplus p}$, where p is the number of orbits. The results from this section can be extended to this scenario by applying all the arguments to the copies of $\langle G \rangle$ separately, each of which will serve as a domain for its set of irreducible unitary representations.

When the action is not free, it is necessary to take into account *stabilizers*, i.e. $g \in G$ such that $g \cdot s = s$ for some $s \in \mathcal{S}$. Roughly speaking, we expect that the results from this section can be adapted to an extent, obtaining unitary representations “up to stabilizers”. However, the precise meaning of the latter has yet to be clarified, and goes beyond the scope of this work.

3.2. Examples

In this section, we provide examples of machine learning models with unitary symmetries. As anticipated, in the context of machine learning \mathcal{H} and \mathcal{W} represent the hypothesis space and the parameter space, respectively. Indeed, in what follows \mathcal{H} will consist of functions of the form $\langle G \rangle \rightarrow \mathcal{Y}$ for some codomain \mathcal{Y} , and we will adhere to the notation from Definition 12 accordingly. All the models considered in this section satisfy the adjunction property.

3.2.1. SPECTRAL NETWORKS

We start by considering *Spectral Networks* – a class of polynomial machine learning models that inspired this work – and to which our theory applies naturally. These models were introduced by Sanborn et al. (2023) in cubic form and for commutative groups. Here, we generalize them to arbitrary degree and to the non-commutative setting. Spectral Networks are grounded in the algebraic invariant theory of $\langle G \rangle$. We review the latter in the Appendix (Section B), including simple proofs for finite groups of classical results. The overall idea behind Spectral Networks is to approximate the n -order polynomial invariants – deemed spectral invariants – of $\langle G \rangle$ for an *unknown* group G . Specifically, suppose that V_1, \dots, V_k are the ambient Hilbert spaces of the irreducible unitary representations of G . Given a multi-index $\underline{i} = (i_1, \dots, i_n) \in \{1, \dots, k\}^n$, the Spectral Network of order n is defined as the collection of parametric maps $\varphi_{\underline{i}}(W, \cdot) : \langle G \rangle \rightarrow \text{End}(V_{i_1} \otimes \dots \otimes V_{i_n})$:

$$\varphi_{\underline{i}}(W, x) = W_{i_1}(x) \otimes \dots \otimes W_{i_n}(x) \left(W_{i_1}^\dagger \odot \dots \odot W_{i_n}^\dagger \right) (\bar{x}), \quad (5)$$

where $W = \oplus_i W_i \in \mathcal{W} = \langle G \rangle \otimes \oplus_i \text{End}(V_i)$, \odot denotes the G -wise tensor product of operators, and \bar{x} denotes the component-wise conjugate of x . For a commutative G , since $V_i = \mathbb{C}$ for all i , the above expression reduces to $\varphi_{\underline{i}}(W, x) = W_{i_1} x \cdots W_{i_n} x \overline{W_{i_1} \odot \dots \odot W_{i_n} x}$. For $n = 1$ Spectral Networks are called Power-Spectral Networks, and for a commutative G they take the form $\varphi_i(W, x) = |W_i x|^2$. The latter can be simply interpreted as a linear model followed by a function

given by the squared absolute value. Even though the latter is uncommon in machine learning, it has appeared in models of biological neural networks (Adelson and Bergen, 1985).

For simplicity, and without loss of generality, we will consider only the Spectral Networks involving a single unitary representation; that is, we will focus on constant multi-indices $\underline{i} = (i, \dots, i)$ in Equation 5. To this end, let V be a finite-dimensional Hilbert space and \mathcal{H} be the set of functions $\langle G \rangle \rightarrow \text{End}(V^{\otimes n})$ for some $n \in \mathbb{N}$. We also set $\mathcal{W} = \langle G \rangle \otimes \text{End}(V)$. The following is proved in the Appendix.

Proposition 13 *Consider the Spectral Network given by $\varphi(W, x) = W(x)^{\otimes n} W^{\dagger \circ n}(\bar{x})$. Then φ has unitary symmetries.*

3.2.2. MCCULLOCH-PITTS NEURONS AND DEEP NETWORKS

While Spectral Networks provide the most direct application of our theory, in this section we discuss the most common and fundamental neural network primitives in deep learning: the fully-connected McCulloch-Pitts neuron (McCulloch and Pitts, 1943) and the deep neural network. We consider models with complex coefficients and focus on commutative groups, i.e. all the Hilbert spaces V_i from Definition 9 will be equal to \mathbb{C} . A McCulloch-Pitts neuron has the form $\varphi(W, x) = \sigma(Wx)$, where $\sigma: \mathbb{C} \rightarrow \mathcal{Y}$ is a map playing the role of an activation function and $W \in \mathcal{W} = \langle G \rangle$ is the weight vector. For $\sigma(z) = |z|^2$, the McCulloch-Pitts neuron reduces to a commutative Power-Spectral Network, i.e. a Spectral Network with $n = 1$. The hypothesis space \mathcal{H} consists of functions $\langle G \rangle \rightarrow \mathcal{Y}$.

Proposition 14 *Consider a map $\sigma: \mathbb{C} \rightarrow \mathcal{Y}$ and let $\varphi(W, x) = \sigma(Wx)$. Suppose that $0 \in \mathbb{C}$ is isolated in its fiber of σ , i.e. there exists an open subspace $O \subseteq \mathbb{C}$ such that $\sigma^{-1}(\sigma(0)) \cap O = \{0\}$. Then φ has unitary symmetries.*

We refer to the Appendix for a proof. The above assumption on σ is satisfied by popular activations functions from neuroscience and machine learning, such as the sigmoid and the leaky Rectified Linear Unit (ReLU), applied after taking complex absolute value. Moreover, any non-constant holomorphic map $\sigma: \mathbb{C} \rightarrow \mathbb{C}$ satisfies the assumption, since its fibers are discrete by the Analytic Continuation Theorem.

Example 2 *Consider the logistic regressor with complex weights, corresponding to a McCulloch-Pitts neuron with a sigmoid activation function, i.e., $\sigma(z) = \frac{1}{1+e^{-|z|}} + b$, where $b \in \mathbb{R}$ is the bias term. Since $\sigma^{-1}(\sigma(0)) = \{0\}$, Proposition 14 applies. In particular, if the logistic regressor $\varphi(W, x) = \sigma(Wx)$ is invariant in x to the cyclic group $C_d = \mathbb{Z}/d\mathbb{Z}$, from Theorem 10 it follows that, if $W \neq 0$, then $W(g) = u e^{-2\pi\sqrt{-1}g/d}$ for all $g = 0, \dots, d-1$ and for some nonzero $u \in \mathbb{C}$.*

Next, we discuss the case of classical deep neural networks. We model the latter as $\varphi(W, x) = \chi(\sigma(W(x)))$, where $W \in \langle G \rangle \otimes \mathbb{C}^k$ represents the weights of the first layer containing k neurons, $\sigma: \mathbb{C} \rightarrow \mathbb{R}$ represents the activation function, computed coordinate-wise on \mathbb{C}^k , and $\chi: \mathbb{R}^k \rightarrow \mathbb{R}$ encompasses all the layers after the first one. Therefore, the first layer is assumed to be complex-valued, as required by our theory, while the subsequent layers are allowed to take real values, as typical in practice. Since the weights W only account for the first layer, our main results will guarantee that invariant deep neural networks recover the Fourier transform in that layer, which is consistent with empirical observations (Olah et al., 2020).

Differently from Section 3, for the next result we will restrict \mathcal{W} to the subspace of $\langle G \rangle \otimes \mathbb{C}^k$ consisting of W such that the components W_i are orthonormal. This implies, in particular, the constraint $k \leq |G|$. Note that \mathcal{W} is closed by the actions of G and $U(\mathbb{C})$. The orthonormality condition is anyway necessary in order to recover the full Fourier transform – see Corollary 11.

Proposition 15 *Let $\mathcal{W} = \{W \in \langle G \rangle \otimes \mathbb{C}^k \mid \forall i, j \langle W_i, W_j \rangle = 0\}$, consider maps $\sigma: \mathbb{C} \rightarrow \mathbb{R}$, $\chi: \mathbb{R}^k \rightarrow \mathbb{R}$, and let $\varphi(W, x) = \chi(\sigma(W(x)))$. Suppose that:*

- *There exists an open subspace $O \subseteq \mathbb{R}^k$ containing $\sigma(0)$ where χ is affine with distinct non-vanishing coefficients i.e., $\chi(z) = \sum_i a_i z_i + b$ for $z \in O$ with $0 \neq a_i \neq a_j$ for $i \neq j$.*
- *$\sigma \in C^2(\mathbb{C})$ and $\partial_z \partial_{\bar{z}} \sigma(0) \neq 0$.*

Then φ has unitary symmetries.

We refer to the Appendix for a proof. The hypothesis on χ is motivated by the fact that typical (real-valued) deep networks have piece-wise linear activation functions, such as (leaky) ReLU. As such, they define piece-wise affine maps and therefore are affine when restricted to appropriate open subspaces. Moreover, the hypothesis on the coefficients a_i in Proposition 15 is *generic*, meaning that it defines an open dense subset of $(a_1, \dots, a_k) \in \mathbb{R}^k$.

3.3. Group Recovery

Corollary 11 allows us to recover the group structure of G up to isomorphism from the weights of a map with unitary symmetries. In other words, this enables the recovery of an unknown group in a data-driven manner from the weights of an invariant machine learning model, addressing the problem of symmetry discovery discussed in Section 1.1. The procedure was originally suggested and validated empirically by Sanborn et al. (2023). To this end, assume that φ_i satisfies the requirements of Corollary 11. Moreover, we introduce the additional assumption that $W_i(1) = I \in \text{End}(V_i)$ for all i . If that is the case, W coincides exactly with the Fourier transform by Corollary 11. This implies that the multiplication table of G can be recovered by:

$$gh = \underset{l \in G}{\operatorname{argmin}} \|W(g) \odot W(h) - W(l)\|, \quad (6)$$

where \odot denotes the Hadamard product, i.e. component-wise operator composition. Note that this notation has a different meaning here than in Section 3.2.1. Since the $W(g)$'s are orthogonal, the only possible values for the norms over which the minimum is performed are 0 and $\sqrt{2|G|}$.

The multiplication table of G is a discrete object, while the weights $W \in \mathcal{W}$ can vary continuously. Therefore, it is natural to expect that the invariance condition ($g \cdot W = W$ for all $g \in G$) can be relaxed, while still recovering the multiplication table correctly. In what follows, we analyze relaxations of invariance and give bounds in which the group recovery algorithm holds. We start by introducing a quantity measuring how close a map is to having unitary symmetries. To this end, we assume that \mathcal{H} is a metric space with distance function $\Delta: \mathcal{H} \times \mathcal{H} \rightarrow \mathbb{R}_{\geq 0}$.

Definition 16 *Given a map $\varphi: \mathcal{W} \rightarrow \mathcal{H}$, its unitarity defect is defined for $\delta \in \mathbb{R}_{>0}$ as:*

$$\omega_\varphi(\delta) = \sup_{W, W'} \max_i \inf_{U \in U(V_i)} \|W_i - U \cdot W'_i\|, \quad (7)$$

where the supremum runs over all the $W, W' \in \langle G \rangle \otimes \text{End}(V)$ of the same norm such that $\Delta(\varphi(W), \varphi(W')) \leq \delta$.

Note that φ has unitary symmetries if, and only if, $\omega_\varphi(0) = 0$. We now state our main relaxation result.

Theorem 17 *Suppose that $\varphi: \mathcal{W} \rightarrow \mathcal{H}$ is a map and fix $W \in \mathcal{W}$. Denote $L = \|W^\dagger W - |G|I\|_\infty$, where $\|A\|_\infty = \max_{g,h} |A_{g,h}|$ is the uniform Frobenius norm for $G \times G$ matrices. Suppose that the following holds:*

- For all $g \in G$, $\omega_\varphi(\Delta(\varphi(g \cdot W), \varphi(W))) < \sqrt{\frac{1}{2} - \frac{L}{|G|}} / \left(\sqrt{|G| + L} + 1 \right)$.
- $L \leq \frac{|G|}{2}$.
- $W_i(1) = I \in \text{End}(V_i)$ for all i .

Then Equation 6 holds, i.e. the group recovery algorithm is correct.

We refer to the Appendix for a proof. The assumption on L in the above result is a relaxation of the unitarity assumption in Corollary 11 since $L = 0$ if, and only if, W is unitary up to a multiplicative constant.

We provide an explicit bound for the unitarity defect of the McCulloch-Pitts neurons discussed in Section 3.2. To this end, let \mathcal{H} be the space of continuous functions defined on the unit sphere in $\langle G \rangle$ equipped with the uniform metric (i.e., the L_∞ distance) as Δ .

Proposition 18 *Let $W \in \langle G \rangle$ and consider $\varphi(W, x) = \sigma(Wx)$. Suppose that the activation function $\sigma: \mathbb{C} \rightarrow \mathbb{C}$ is continuous and satisfies the following coercivity condition: there exist constants $C \in \mathbb{R}_{>0}$, $n \in \mathbb{N}$ such that for every $x \in \mathbb{C}$, $|\sigma(0) - \sigma(x)| \geq C|x|^n$. Then the unitarity defect of φ satisfies for $\delta < C$:*

$$\omega_\varphi(\delta) \leq \sqrt{2 \left(1 - \sqrt{1 - \left(\frac{\delta}{C} \right)^{\frac{2}{n}}} \right)}. \quad (8)$$

We refer to the Appendix for a proof. Note that the coercivity condition from above plays the role of the assumption on σ from Proposition 14.

3.3.1. IMPLEMENTATION

We now empirically explore the theory developed in this paper and demonstrate that Spectral Networks are able to recover the group structure in practice. We implement a non-commutative Power-Spectral Network $\varphi_i(W, x) = W_i x W_i^\dagger \bar{x}$ with weights $W \in \mathbb{C}^{d_i \times d_i \times d}$, where $d = |G|$ is the cardinality of the given group and d_1, \dots, d_k are the dimensions of its irreducible unitary representations. As discussed at the beginning of Section 3.3, we force $W_i(1) = I \in \mathbb{C}^{d_i \times d_i}$, where the index 1 is arbitrarily chosen.

Following Sanborn et al. (2023), we train the model via *contrastive learning* (Jaiswal et al., 2020). Namely, given a finite dataset \mathcal{D} of pairs (x, y) , where $x, y \in \langle G \rangle \simeq \mathbb{C}^d$ and $x = g \cdot y$ for an unknown $g \in G$, the objective optimized by the model is:

$$\mathcal{L}(W) = \sum_{(x,y) \in \mathcal{D}} \sum_i \|\varphi_i(W, x) - \varphi_i(W, y)\|^2 + \eta \|dI - WW^\dagger\|^2, \quad (9)$$




Figure 2: Multiplication tables inferred by Power-Spectral Networks for various groups. Group elements are labeled with integers. The table is symmetric if, and only if, the group is commutative.

where $\eta > 0$ is a hyper-parameter and $\|\cdot\|$ is the Frobenius norm. The first term in Equation 9 encourages invariance with respect to G while the second one encourages W to be unitary. The model is trained via the Adam optimizer (Kingma and Ba, 2014), which interprets the complex weights as real tensors of doubled dimensionality. The dataset $\mathcal{D} \subseteq \langle G \rangle \oplus \langle G \rangle$ is generated procedurally by first sampling x from a standard Gaussian distribution over $\langle G \rangle \simeq \mathbb{R}^{2d}$, then sampling $g \in G$ uniformly, and finally producing the datapoint $(x, y = g \cdot x) \in \mathcal{D}$. We provide a Python implementation of both the model and the experiments at a public repository (see Section 1.1). The code is available in both the PyTorch (Paszke et al., 2019) and the JAX (Bradbury et al., 2018) frameworks.

Once trained, we evaluate the model by checking whether the multiplication table $M \in \{1, \dots, d\}^{d \times d}$ obtained via the group recovery algorithm described by Equation 6 coincides with the one of G . Since there is no canonical ordering on G , the table is recovered up to a permutation π of $\{1, \dots, d\}$ acting as $(\pi \cdot M)_{i,j} = \pi(M_{\pi^{-1}(i), \pi^{-1}(j)})$. Therefore, we check whether $\pi \cdot M$ coincides with the table of G for all the permutations π . Figure 2 reports the (correct) multiplication tables obtained at convergence for both commutative and non-commutative groups. Specifically, we consider the cyclic group C_6 , the product of cyclic groups $C_2 \times C_2 \times C_2$, and the dihedral group D_3 , which is isomorphic to the symmetric group S_3 .

In order to validate empirically the robustness of the group recovery procedure, we additionally train the model on data corrupted by white noise, i.e. \mathcal{D} consists of pairs $(x, y = g \cdot x + \epsilon)$, where ϵ is sampled from an isotropic Gaussian distribution. We vary the standard deviation of the latter and report in Figure 3 the number of times the multiplication table is recovered correctly across 20 training runs – a metric referred to as “table accuracy”. As can be seen, the group structure is recovered most of the times even with large amounts of noise – up to ~ 0.5 standard deviation. The performance quickly degrades as the noise reaches a critical threshold. This shows empirically that the group recovery procedure is robust to noise, which is in line with the bounds from Theorem 17.

4. Conclusions, Limitations, and Future Work

In this work, we proved that if a machine learning model of a certain kind is invariant to a finite group, then its weights are closely related to the Fourier transform on that group. We discussed how, as a consequence, the algebraic structure of an unknown group can be recovered from a model that is invariant. We established these results for both commutative and non-commutative groups, and




Figure 3: Accuracy for the recovery of the group multiplication tables across 20 training runs as the amount of noise injected into data increases. Since the transition from 1.0 to 0.0 accuracy is sharp, we visualize only the noise regions where the values are non-trivial.

discussed relaxed conditions under which the group recovery procedure holds. Our results represent a first step towards a mathematical explanation of universal features inferred by both biological and artificial neural networks. As such, this work contributes to the line of mathematical results stemming from biology (Sturmfels, 2005).

Due to its open-ended nature, this work is subject to a number of limitations and leaves directions open for future investigation. First, our theory encompasses models with complex-valued weights, which are non-canonical in machine learning. Thus, exploring analogues of the theory over real numbers is an interesting direction that would fit more directly with current practices in the field.

In addition, our theoretical framework encompasses learning models with unitary symmetries. The latter is a technical property satisfied by Spectral Networks and, to an extent, by traditional deep networks. However, it is not clear what other models fit into our framework, and whether the notion is general enough to accommodate other computational primitives fundamental to machine learning, such as attention mechanisms. This is an open question that is worthy of investigation.

Lastly, in this work we focused on groups and their associated harmonics. However, the representations within neural networks or biological systems often resemble imperfect, or *localized*, versions of harmonics, i.e. wavelets, such as Gabors. Since wavelets do not describe group homomorphisms, our theory would need to be extended to accommodate this kind of locality. We suspect that this may be achieved by generalizing the framework to *groupoids* – an algebraic group-like structure that formalizes a locally-defined composition. This, however, goes beyond the scope of our work, and we leave it for future research.

Acknowledgments

This work was supported by the Swedish Research Council, the Knut and Alice Wallenberg Foundation and the European Research Council (ERC-BIRD-884807).

References

- Edward H Adelson and James R Bergen. Spatiotemporal energy models for the perception of motion. *Josa a*, 2(2):284–299, 1985.
- Andrea Banino, Caswell Barry, Benigno Uribe, Charles Blundell, Timothy Lillicrap, Piotr Mirowski, Alexander Pritzel, Martin J Chadwick, Thomas Degris, Joseph Modayil, et al. Vector-based navigation using grid-like representations in artificial agents. *Nature*, 557(7705):429–433, 2018.
- Joshua Basseley, Lijun Qian, and Xianfang Li. A survey of complex-valued neural networks. *arXiv preprint arXiv:2101.12249*, 2021.
- Anthony J Bell and Terrence J Sejnowski. The “independent components” of natural scenes are edge filters. *Vision research*, 37(23):3327–3338, 1997.
- James Bradbury, Roy Frostig, Peter Hawkins, Matthew James Johnson, Chris Leary, Dougal Maclaurin, George Necula, Adam Paszke, Jake VanderPlas, Skye Wanderman-Milne, and Qiao Zhang. JAX: composable transformations of Python+NumPy programs. In *GitHub*, 2018.
- Michael M Bronstein, Joan Bruna, Taco Cohen, and Petar Veličković. Geometric deep learning: Grids, groups, graphs, geodesics, and gauges. *arXiv preprint arXiv:2104.13478*, 2021.
- Bilal Chughtai, Lawrence Chan, and Neel Nanda. A toy model of universality: Reverse engineering how networks learn group operations. *arXiv preprint arXiv:2302.03025*, 2023.
- Taco Cohen and Max Welling. Group equivariant convolutional networks. In Maria Florina Balcan and Kilian Q. Weinberger, editors, *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 2990–2999, 2016.
- Christopher J Cueva and Xue-Xin Wei. Emergence of grid-like representations by training recurrent neural networks to perform spatial localization. *arXiv preprint arXiv:1803.07770*, 2018.
- Krish Desai, Benjamin Nachman, and Jesse Thaler. Symmetry discovery with deep learning. *Physical Review D*, 105(9):096031, 2022.
- William Dorrell, Peter E Latham, Timothy EJ Behrens, and James CR Whittington. Actionable neural representations: Grid cells from minimal constraints. *arXiv preprint arXiv:2209.15563*, 2022.
- Ursula C Dräger. Receptive fields of single cells and topography in mouse visual cortex. *Journal of Comparative Neurology*, 160(3):269–289, 1975.
- Pierre-Étienne H Fiquet and Eero P Simoncelli. Polar prediction of natural videos. *arXiv preprint arXiv:2303.03432*, 2023.
- Gerald B Folland. *A course in abstract harmonic analysis*, volume 29. CRC press, 2016.
- Richard J Gardner, Erik Hermansen, Marius Pachitariu, Yoram Burak, Nils A Baas, Benjamin A Dunn, May-Britt Moser, and Edvard I Moser. Toroidal topology of population activity in grid cells. *Nature*, 602(7895):123–128, 2022.

- Charles J Garfinkle and Christopher J Hillar. On the uniqueness and stability of dictionaries for sparse representation of noisy signals. *IEEE Transactions on Signal Processing*, 67(23):5884–5892, 2019.
- Alexis Guanella, Daniel Kiper, and Paul Verschure. A model of grid cells based on a twisted torus topology. *International journal of neural systems*, 17(04):231–240, 2007.
- Charles A Hass and Gregory D Horwitz. V1 mechanisms underlying chromatic contrast detection. *Journal of Neurophysiology*, 109(10):2483–2494, 2013.
- Christopher J Hillar and Friedrich T Sommer. When can dictionary learning uniquely recover sparse data from subsamples? *IEEE Transactions on Information Theory*, 61(11):6290–6297, 2015.
- Geoffrey E Hinton, Sara Sabour, and Nicholas Frosst. Matrix capsules with em routing. In *International conference on learning representations*, 2018.
- Roger A Horn and Charles R Johnson. *Matrix analysis*. Cambridge university press, 2012.
- Roger A Horn, Roger A Horn, and Charles R Johnson. *Topics in matrix analysis*. Cambridge university press, 1994.
- David H Hubel and Torsten N Wiesel. Receptive fields of single neurones in the cat’s striate cortex. *The Journal of physiology*, 148(3):574, 1959.
- David H Hubel and Torsten N Wiesel. Receptive fields, binocular interaction and functional architecture in the cat’s visual cortex. *The Journal of physiology*, 160(1):106, 1962.
- David H Hubel and Torsten N Wiesel. Receptive fields and functional architecture of monkey striate cortex. *The Journal of physiology*, 195(1):215–243, 1968.
- Jarmo Hurri and Aapo Hyvärinen. Simple-cell-like receptive fields maximize temporal coherence in natural video. *Neural Computation*, 15(3):663–691, 2003.
- Guy Isely, Christopher Hillar, and Fritz Sommer. Deciphering subsampled data: adaptive compressive sampling as a principle of brain communication. *Advances in neural information processing systems*, 23, 2010.
- Ashish Jaiswal, Ashwin Ramesh Babu, Mohammad Zaki Zadeh, Debapriya Banerjee, and Fillia Makedon. A survey on contrastive self-supervised learning. *Technologies*, 9(1):2, 2020.
- Ramakrishna Kakarala. Completeness of bispectrum on compact groups. *arXiv preprint arXiv:0902.0196*, 1, 2009.
- Ramakrishna Kakarala. The bispectrum as a source of phase-sensitive invariants for Fourier descriptors: a group-theoretic approach. *Journal of Mathematical Imaging and Vision*, 44:341–353, 2012.
- Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

- Risi Kondor and Shubhendu Trivedi. On the generalization of equivariance and convolution in neural networks to the action of compact groups. In Jennifer Dy and Andreas Krause, editors, *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2747–2755, 2018.
- Yixuan Li, Jason Yosinski, Jeff Clune, Hod Lipson, and John Hopcroft. Convergent learning: Do different neural networks learn the same representations? *arXiv preprint arXiv:1511.07543*, 2015.
- Sindy Löwe, Phillip Lippe, Maja Rudolph, and Max Welling. Complex-valued autoencoders for object discovery. *arXiv preprint arXiv:2204.02075*, 2022.
- Warren S McCulloch and Walter Pitts. A logical calculus of the ideas immanent in nervous activity. *The bulletin of mathematical biophysics*, 5:115–133, 1943.
- Luca Moschella, Valentino Maiorca, Marco Fumero, Antonio Norelli, Francesco Locatello, and Emanuele Rodola. Relative representations enable zero-shot latent space communication. *arXiv preprint arXiv:2209.15430*, 2022.
- Edvard I Moser, Emilio Kropff, and May-Britt Moser. Place cells, grid cells, and the brain’s spatial representation system. *Annu. Rev. Neurosci.*, 31:69–89, 2008.
- Neel Nanda, Lawrence Chan, Tom Liberum, Jess Smith, and Jacob Steinhardt. Progress measures for grokking via mechanistic interpretability. *arXiv preprint arXiv:2301.05217*, 2023.
- Chris Olah, Nick Cammarata, Ludwig Schubert, Gabriel Goh, Michael Petrov, and Shan Carter. An overview of early vision in inceptionV1. *Distill*, 5(4):e00024–002, 2020.
- Bruno A Olshausen and David J Field. Sparse coding with an overcomplete basis set: A strategy employed by V1? *Vision research*, 37(23):3311–3325, 1997.
- Jeff Orchard, Hao Yang, and Xiang Ji. Does the entorhinal cortex use the Fourier transform? *Frontiers in computational neuroscience*, 7:179, 2013.
- Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- Walter Pitts and Warren S McCulloch. How we know universals: The perception of auditory and visual forms. *The Bulletin of mathematical biophysics*, 9:127–147, 1947.
- Rajesh Rao and Daniel Ruderman. Learning Lie groups for invariant visual perception. *Advances in neural information processing systems*, 11, 1998.
- Tilman Räuher, Anson Ho, Stephen Casper, and Dylan Hadfield-Menell. Toward transparent ai: A survey on interpreting the inner structures of deep neural networks. In *2023 IEEE Conference on Secure and Trustworthy Machine Learning (SaTML)*, pages 464–483. IEEE, 2023.
- Sara Sabour, Nicholas Frosst, and Geoffrey E Hinton. Dynamic routing between capsules. *Advances in neural information processing systems*, 30, 2017.

- Sophia Sanborn, Christian Shewmake, Bruno Olshausen, and Christopher Hillar. Bispectral neural networks. *International Conference on Learning Representations (ICLR)*, 2023.
- Rylan Schaeffer, Mikail Khona, Tzuhsuan Ma, Cristóbal Eyzaguirre, Sanmi Koyejo, and Ila Rani Fiete. Self-supervised learning of representations for space generates multi-modular grid cells. *arXiv preprint arXiv:2311.02316*, 2023.
- Fethi Smach, Cedric Lemaître, Jean-Paul Gauthier, Johel Miteran, and Mohamed Atri. Generalized Fourier descriptors with applications to objects recognition in SVM context. *Journal of mathematical imaging and Vision*, 30:43–71, 2008.
- Jascha Sohl-Dickstein, Ching Ming Wang, and Bruno A Olshausen. An unsupervised algorithm for learning Lie group transformations. *arXiv preprint arXiv:1001.1027*, 2010.
- Ben Sorscher, Gabriel Mel, Surya Ganguli, and Samuel Ocko. A unified theory for the origin of grid cells through the lens of pattern formation. *Advances in neural information processing systems*, 32, 2019.
- Ben Sorscher, Gabriel C Mel, Samuel A Ocko, Lisa M Giocomo, and Surya Ganguli. A unified theory for the computational and mechanistic origins of grid cells. *Neuron*, 111(1):121–137, 2023.
- Benjamin Steinberg. *Representation theory of finite groups: an introductory approach*. Springer, 2012.
- Bernd Sturmfels. Can biology lead to new theorems. *Annual report of the Clay Mathematics Institute*, 1468:13–26, 2005.
- Bernd Sturmfels. *Algorithms in invariant theory*. Springer Science & Business Media, 2008.
- Chiheb Trabelsi, Olexa Bilaniuk, Ying Zhang, Dmitriy Serdyuk, Sandeep Subramanian, João Felipe Santos, Soroush Mehri, Negar Rostamzadeh, Yoshua Bengio, and Christopher J Pal. Deep complex networks (2017). *arXiv preprint arXiv:1705.09792*, 2017.
- Konstantin F Willeke, Kelli Restivo, Katrin Franke, Arne F Nix, Santiago A Cadena, Tori Shinn, Cate Nealley, Gabby Rodriguez, Saumil Patel, Alexander S Ecker, et al. Deep learning-driven characterization of single cell tuning in primate visual area V4 unveils topological organization. *bioRxiv*, pages 2023–05, 2023.

Appendix A. Proofs of Theoretical Results

A.1. Proof of Theorem 10

Proof Since φ has unitary symmetries and $\|g \cdot W_i\| = \|W_i\|$ for all i , it follows that for every $g \in G$ and every i there exists $\rho_i(g) \in \mathbf{U}(V_i)$ such that:

$$g \cdot W_i = \rho_i(g) \cdot W_i. \quad (10)$$

In particular, by considering the component with index $1 \in G$ on both sides of Equation 10, we see that $W_i(g^{-1}) = W_i(1) \rho_i(g)$. We wish to show that ρ_i is a homomorphism. To this end, for all $g, h \in G$ it holds that:

$$\rho_i(gh) \cdot W_i = (gh) \cdot W_i = g \cdot (\rho_i(h) \cdot W_i) = \rho_i(h) \cdot (g \cdot W_i) = (\rho_i(g)\rho_i(h)) \cdot W_i. \quad (11)$$

By the surjectivity hypothesis, the set $\{W_i(g)\}_{g \in G}$ generates $\text{End}(V_i)$ as a vector space. Therefore, Equation 11 implies that $\rho_i(gh) = \rho_i(g)\rho_i(h)$, as desired. Lastly, note that ρ_i is irreducible by the surjectivity assumption. Indeed, a non-trivial linear subspace of V_i fixed by $\rho_i(g)$ for all g would be sent by $W_i(g)$ into a fixed proper subspace due to Equation 10. This contradicts the surjectivity of W_i . ■

A.2. Proof of Corollary 11

Proof Firstly, note that the unitarity assumption implies the surjectivity assumption from Theorem 10. Therefore, it follows that for every i , there exists an irreducible unitary representation $\rho_i: G \rightarrow \mathbf{U}(V_i)$ such that $W_i(g) = W_i(1) \rho_i(g)^\dagger$. We wish to show that if $i \neq j$ then ρ_i and ρ_j are non-isomorphic representations. If not, the orthogonality assumption implies that:

$$0 = \sum_{g \in G} \overline{W_i(g)} \otimes W_j(g) = \sum_{g \in G} \left(\overline{W_i(1)} \otimes W_j(1) \right) \left(\rho_i(g)^\top \otimes \rho_j(g)^\dagger \right) = \quad (12)$$

$$= |G| \overline{W_i(1)} \otimes W_j(1), \quad (13)$$

where \top denotes the transpose and where the last identity follows from the Schur orthogonality relations. But then $W_i(1)$ or $W_j(1)$ is vanishing, which contradicts the unitarity assumption.

Lastly, in order to compute the Frobenius norm of $W_i(1)$, note that:

$$|G| \|W_i(1)\|^2 = \sum_{g \in G} \|W_i(g) \rho_i(g)\|^2 = \sum_{g \in G} \|W_i(g)\|^2 = |G|, \quad (14)$$

from which $\|W_i(1)\| = 1$ follows. ■

A.3. Proof of Proposition 13

In order to prove this proposition, we will need some technical facts from matrix algebra. We start by showing the following uniqueness result.

Lemma 19 *Let $d' \geq d$ and $A, B \in \mathbb{C}^{d' \times d}$. If $AA^\dagger = BB^\dagger$, then there exists a unitary matrix $U \in \mathbb{C}^{d \times d}$ such that $A = BU$.*

Proof From the polar decomposition for matrices (Horn et al., 1994, Theorem 3.1.9), we know that:

$$A = PV, B = QW, \quad (15)$$

where $P, Q \in \mathbb{C}^{d' \times d'}$ are Hermitian positive semidefinite and $V, W \in \mathbb{C}^{d' \times d}$ have orthonormal rows. Also, $P^2 = AA^\dagger = BB^\dagger = Q^2$ from which it follows that $P = Q$ by uniqueness of square roots of positive semidefinite Hermitian matrices (Horn and Johnson, 2012, Theorem 7.2.6). In particular, we have:

$$A = PV = QV = QWW^\dagger V = BU, \quad (16)$$

with $U = W^\dagger V$ unitary. ■

Remark 20 We note that if A and B are real matrices, then the conclusion holds with U being a real orthogonal matrix.

Next, we show that positive semidefinite Hermitian matrices possess unique tensor roots.

Lemma 21 Let $A, B \in \mathbb{C}^{d \times d}$ be Hermitian and positive semidefinite. If $A^{\otimes n} = B^{\otimes n}$ for some $n > 0$, then $A = B$.

Proof From the Spectral Theorem we know that:

$$A = UDU^\dagger, \quad (17)$$

where U is unitary and D is diagonal. Since $A^{\otimes n} = B^{\otimes n}$, it follows that $D^{\otimes n} = C^{\otimes n}$, where $C = U^\dagger B U$. Note that the (point-wise) Hadamard product of matrices is a submatrix of the tensor (Kronecker) product. In particular, the off-diagonal entries of C must vanish. On the diagonal we have $(D_{i,i})^n = (C_{i,i})^n$ for every i , and therefore $D_{i,i} = C_{i,i}$ since they are non-negative. This shows that $C = D$, implying that $B = UDU^\dagger = A$. ■

Putting together the above lemmas, we obtain the following.

Lemma 22 Let $A_1, \dots, A_k, B_1, \dots, B_k \in \mathbb{C}^{d \times d}$. Suppose that for some $n > 0$, for all $x \in \mathbb{C}^k$:

$$\left(\sum_i x_i A_i \right)^{\otimes n} \left(\sum_i \bar{x}_i A_i^{\dagger \otimes n} \right) = \left(\sum_i x_i B_i \right)^{\otimes n} \left(\sum_i \bar{x}_i B_i^{\dagger \otimes n} \right). \quad (18)$$

Then there exists a unitary matrix $U \in \mathbb{C}^{d \times d}$ such that $A_i = B_i U$ for every i .

Proof By multilinearity of the tensor product we see that for all $x \in \mathbb{C}^k$:

$$\sum_{i_1, \dots, i_{n+1}} x_{i_1} \cdots x_{i_n} \bar{x}_{i_{n+1}} (A_{i_1} \otimes \cdots \otimes A_{i_n}) A_{i_{n+1}}^{\dagger \otimes n} = \quad (19)$$

$$= \sum_{i_1, \dots, i_{n+1}} x_{i_1} \cdots x_{i_n} \bar{x}_{i_{n+1}} (A_{i_1} A_{i_{n+1}}^\dagger) \otimes \cdots \otimes (A_{i_n} A_{i_{n+1}}^\dagger) = \quad (20)$$

$$= \sum_{i_1, \dots, i_{n+1}} x_{i_1} \cdots x_{i_n} \bar{x}_{i_{n+1}} (B_{i_1} B_{i_{n+1}}^\dagger) \otimes \cdots \otimes (B_{i_n} B_{i_{n+1}}^\dagger). \quad (21)$$

Since a polynomial vanishes as function if and only if it is the zero polynomial, it follows that $(A_{i_1}A_{i_{n+1}}^\dagger) \otimes \cdots \otimes (A_{i_n}A_{i_{n+1}}^\dagger) = (B_{i_1}B_{i_{n+1}}^\dagger) \otimes \cdots \otimes (B_{i_n}B_{i_{n+1}}^\dagger)$ for all i_1, \dots, i_{n+1} , and in particular $(A_iA_j^\dagger)^{\otimes n} = (B_iB_j^\dagger)^{\otimes n}$ for all i, j . From Lemma 21 we conclude that $A_iA_j = B_iB_j$ for all i, j , which can be rephrased as $AA^\dagger = BB^\dagger$, where A, B are the $(dk) \times d$ matrices obtained by concatenating the A_i 's and B_i 's, respectively. From Lemma 19 we conclude that $A = BU$ for a unitary $d \times d$ matrix U , as desired. ■

Proposition 13 now follows immediately from Lemma 22 by setting $A_i = W(g_i)$ and $B_i = W'(g_i)$ for $g_i \in G$ and $W, W' \in \langle G \rangle \otimes \text{End}(V)$ of the same norm.

A.4. Proof of Proposition 14

Proof Pick $W, W' \in \langle G \rangle$ of the same norm such that $\varphi(W, x) = \varphi(W', x)$ for all $x \in \langle G \rangle$. Given the open set $O \subseteq \mathbb{C}$ from the hypothesis on σ , consider $O' = \{x \in \langle G \rangle \mid Wx, W'x \in O\}$, which is open and non-empty since $0 \in O'$. For $x \in O'$, $Wx = 0$ implies $\varphi(W, x) = \varphi(W', x) = 0$, from which it follows that $W'x = 0$ by definition of O . Therefore, W and W' share the same orthogonal complement, implying that $W' = \rho W$ for some $\rho \in \mathbb{C}$. Since W and W' have the same norm, we conclude that $\rho \in \text{U}(\mathbb{C})$. ■

A.5. Proof of Proposition 15

Proof Pick $W, W' \in \mathcal{W}$ such that $\|W_i\| = \|W'_i\|$ for all i and $\varphi(W, x) = \varphi(W', x)$ for all $x \in \langle G \rangle$. Given the open set $O \subseteq \mathbb{R}^k$ from the hypothesis on χ , consider the set $O' = \{x \in \langle G \rangle \mid \sigma(W(x)), \sigma(W'(x)) \in O\}$, which is open and non-empty since $0 \in O'$. For $x \in O'$, an immediate calculation shows that:

$$\partial_x \partial_{\bar{x}} \varphi(W, x) = \sum_i a_i \partial_z \partial_{\bar{z}} \sigma(W_i x) W_i \otimes \bar{W}_i, \quad (22)$$

with the a_i 's being distinct, and similarly for $\varphi(W', x)$. By specializing to $x = 0$, we deduce that:

$$\cancel{\partial_z \partial_{\bar{z}} \sigma(0)} \sum_i a_i W_i \otimes \bar{W}_i = \cancel{\partial_z \partial_{\bar{z}} \sigma(0)} \sum_i a_i W'_i \otimes \bar{W}'_i. \quad (23)$$

Since both the sets $\{W_i\}_i$ and $\{W'_i\}_i$ are orthonormal by hypothesis on \mathcal{W} , both sides of Equation 23 define a spectral decomposition, i.e. a decomposition into projections over orthonormal vectors. From the hypothesis on χ it follows that the eigenvalues $a_i \partial_z \partial_{\bar{z}} \sigma(W_i x)$ of the operators in Equation 23 are distinct and non-vanishing, implying that the eigenvectors coincide up to a multiplicative scalar, i.e. $W_i = \rho_i W'_i$ for some $\rho_i \in \text{U}(\mathbb{C})$, as desired. ■

A.6. Proof of Theorem 17

Proof Firstly, the definition of L implies the inequalities $|G| - L \leq \|W(g)\|^2 \leq |G| + L$ and $|\langle W(g), W(h) \rangle| \leq L$ for all $g, h \in G$. In particular, we have:

$$\|W(g) - W(h)\|^2 = \|W(g)\|^2 + \|W(h)\|^2 - 2\text{Re}(\langle W(g), W(h) \rangle) \geq 2|G| - 4L. \quad (24)$$

By hypothesis, for every i and $g \in G$ there exists $\rho_i(g) \in \mathbf{U}(V_i)$ such that:

$$\|g^{-1} \cdot W_i - \rho_i(g) \cdot W_i\| < \frac{\sqrt{\frac{1}{2} - \frac{L}{|G|}}}{\sqrt{|G| + L + 1}}. \quad (25)$$

In particular, $\|W_i(gh) - \rho_i(g)W_i(h)\|$ is bounded by the same quantity for all $g, h \in G$. Therefore, via the triangle inequality we see that:

$$\|W_i(g)W_i(h) - W_i(gh)\| \leq \quad (26)$$

$$\leq \|W_i(h)W_i(g) - \rho_i(g)W_i(h)\| + \|W_i(gh) - \rho_i(g)W_i(h)\| = \quad (27)$$

$$= \underbrace{\|W_i(h)\|}_{\leq \sqrt{|G|+L}} \|W_i(g) - \rho_i(g)W_i(1)\| + \|W_i(gh) - \rho_i(g)W_i(h)\| < \quad (28)$$

$$\leq \sqrt{|G|+L} < \sqrt{\frac{1}{2} - \frac{L}{|G|}}. \quad (29)$$

Equation 26 implies that:

$$\|W(g) \odot W(h) - W(gh)\| < \sqrt{|G|} \sqrt{\frac{1}{2} - \frac{L}{|G|}} = \frac{\sqrt{2|G| - 4L}}{2}. \quad (30)$$

Since $\|W(p) - W(q)\| \geq \sqrt{2|G| - 4L}$ for all $p \neq q \in G$, $W(g) \odot W(h)$ is closer to $W(gh)$ than to any other $W(q)$ for $q \in G$, which immediately implies the desired result. \blacksquare

A.7. Proof of Proposition 18

In order to prove this proposition, we will need the following technical fact from linear algebra.

Lemma 23 *Let H be a finite-dimensional complex Hilbert space, $v, w \in H$ normal and $\varepsilon \in \mathbb{R}$ such that $0 < \varepsilon < 1$. Suppose that for every normal x orthogonal to w , it holds that $|\langle x, v \rangle| \leq \varepsilon$. Then there exists $\rho \in \mathbf{U}(\mathbb{C})$ such that:*

$$\|v - \rho w\| \leq \sqrt{2 \left(1 - \sqrt{1 - \varepsilon^2}\right)}. \quad (31)$$

Proof Consider an orthogonal decomposition $v = \langle w_1, v \rangle w_1 + \langle w_2, v \rangle w_2$, where $w_1 \in w^\perp$ is normal and $w_2 = \rho w$ for some $\rho \in \mathbf{U}(\mathbb{C})$ such that $\langle w_2, v \rangle \in \mathbb{R}_{\geq 0}$. It follows that:

$$1 = \|v\|^2 = |\langle w_1, v \rangle|^2 + \langle w_2, v \rangle^2. \quad (32)$$

The hypothesis implies then that $\langle w_2, v \rangle \geq \sqrt{1 - \varepsilon^2}$. Therefore, we have:

$$\|v - w_2\|^2 = 2 - 2\langle w_2, v \rangle \leq 2 \left(1 - \sqrt{1 - \varepsilon^2}\right), \quad (33)$$

as desired. \blacksquare

Remark 24 Note that the right-hand side of Equation 31 is bounded by the concise quantity $\sqrt{2\varepsilon}$.

We are now ready to prove Proposition 18.

Proof Consider $\delta \in \mathbb{R}_{>0}$ and $W, W' \in \langle G \rangle$ of the same norm such that $\Delta(\varphi(W), \varphi(W')) \leq \delta$. The latter and the hypotheses together imply that if $x \in \langle G \rangle$ is normal such that $W'x = 0$, then:

$$C |Wx|^n \leq |\sigma(0) - \sigma(Wx)| \leq \delta. \quad (34)$$

By Lemma 23, there exists $\rho \in U(\mathbb{C})$ such that:

$$\|W - \rho W'\| \leq \sqrt{2 \left(1 - \sqrt{1 - \left(\frac{\delta}{C} \right)^{\frac{2}{n}}} \right)}, \quad (35)$$

from which the claim follows. ■

Appendix B. Spectral Invariants

In this section, we overview the theory of *invariants* over $\langle G \rangle$, i.e. (polynomial) maps $\langle G \rangle \rightarrow \mathbb{C}$ that are invariant with respect to the action by G . To this end, we recall the following notion for a commutative group G .

Definition 25 Fix $n > 0$ and $\underline{\rho} = (\rho_1, \dots, \rho_n) \in (G^\vee)^n$. The spectrum of order n associated to $\underline{\rho}$ is defined for $x \in \langle G \rangle$ as:

$$\beta_{\underline{\rho}}(x) = \hat{x}_{\rho_1} \cdots \hat{x}_{\rho_n} \overline{\hat{x}_{\rho_1 \cdots \rho_n}}. \quad (36)$$

The spectra of order n are invariant polynomials of degree $n + 1$ containing one conjugate variable. The presence of the latter is necessary for invariance. For $n = 1, 2$ they are alternatively referred to as *power spectra* and *bispectra* respectively. Note that the power spectra reduce simply to $\beta_{\rho}(x) = |\hat{x}_{\rho}|^2$, $\rho \in G^\vee$, and constitute a standard tool in signal processing. Bispectra, together with higher-order spectra, were first introduced by Kakarala (2012). It is immediate to see that spectra generate all the polynomial invariants of $\langle G \rangle$ – see also (Sturmfels, 2008, Theorem 2.1.4).

Proposition 26 The space of polynomial invariants of degree $n + 1$ over $\langle G \rangle$ (with one conjugate variable) is generated as a complex vector space by the spectra of order n .

Proof This follows from interpreting the invariance condition via the Fourier transform. Namely, given $\underline{\rho} \in (G^\vee)^{n+1}$ consider the monomial over $\langle G^\vee \rangle$ defined by $\hat{x}_{\rho_1} \cdots \hat{x}_{\rho_n} \overline{\hat{x}_{\rho_{n+1}}}$. Since $g \cdot \hat{x} = (\overline{\rho(g)} \hat{x}_\lambda)_{\lambda \in G^\vee}$, the monomial is invariant if and only if $\rho_1(g) \cdots \rho_n(g) \overline{\rho_{n+1}(g)} = 1$ for all $g \in G$, i.e. $\rho_{n+1} = \rho_1 \cdots \rho_n$. Since monomials linearly generate polynomials, the claim follows. ■

A remarkable aspect of spectra of even order is the fact that they jointly determine real (generic) elements of $\langle G \rangle$ up to the action by G – a property known as *completeness*. This was first shown by Smach et al. (2008). For convenience, we report below a simple proof for finite commutative groups.

Proposition 27 Fix n even. Suppose that $x, y \in \mathbb{R}^G \subseteq \langle G \rangle$ are such that $\hat{x}_\rho, \hat{y}_\rho \neq 0$ for all $\rho \in G^\vee$. If $\beta_{\underline{\rho}}(x) = \beta_{\underline{\rho}}(y)$ for all $\underline{\rho} \in (G^\vee)^n$ then $x = g \cdot y$ for some $g \in G$.

Proof By setting $\underline{\rho} = (1, \dots, 1)$ we see that $\beta_{\underline{\rho}}(x) = \hat{x}_1^{n+1} = \beta_{\underline{\rho}}(y) = \hat{y}_1^{n+1} \in \mathbb{R} \setminus \{0\}$ and therefore $\hat{x}_1 = \hat{y}_1$ since n is even. For $\rho \in G^\vee$, by setting $\underline{\rho} = (\rho, \bar{\rho}, 1, \dots, 1)$, we see that $\hat{x}_\rho \hat{x}_{\bar{\rho}} \hat{x}_1^{n-1} = |\hat{x}_\rho|^2 \hat{x}_1^{n-1} = |\hat{y}_\rho|^2 \hat{y}_1^{n-1}$ and therefore $|\hat{x}_\rho| = |\hat{y}_\rho|$. Note that here we relied on the fact that x and y are real. This implies that the following map $\eta : G^\vee \rightarrow \mathbb{C}$ takes values in $U(\mathbb{C})$:

$$\eta(\rho) = \frac{\hat{x}_\rho}{\hat{y}_\rho}. \quad (37)$$

Now, $\eta(1) = 1$ since $\hat{x}_1 = \hat{y}_1$. By setting $\rho = (\rho, \mu, \bar{\rho}\bar{\mu}, 1, \dots, 1)$ we see that $\eta(\rho)\eta(\mu) = \eta(\bar{\rho}\bar{\mu})$ for all $\rho, \mu \in G^\vee$ and therefore $\eta \in (G^\vee)^\vee$. Since the Fourier transform sends $G^\vee \subseteq \langle G \rangle$ to $(G^\vee)^\vee \subseteq \langle G^\vee \rangle$, there exists $g \in G$ such that $\eta(\rho) = \overline{\rho(g)}$. This means that $\hat{x}_\rho = \overline{\rho(g)} \hat{y}_\rho$, which implies $x = g \cdot y$ by the equivariance properties of the Fourier transform. \blacksquare

Spectral invariants can be defined in the non-commutative case, but give rise to subtleties. First, we replace the group structure of G^\vee with the tensor product \otimes of unitary representations. However, $\text{Irr}(G)$ is not closed with respect to \otimes . This is circumvented by considering *Clebsch-Gordan coefficients*, i.e. irreducible unitary sub-representations of tensor products. This leads to the following definition of operator-valued spectra.

Definition 28 Fix $n > 0$ and $\underline{\rho} = (\rho_{V_1}, \dots, \rho_{V_n}) \in \text{Irr}(G)^n$. The spectrum of order n associated to $\underline{\rho}$ is defined for $x \in \langle G \rangle$ as:

$$\beta_{\underline{\rho}}(x) = \hat{x}_{\rho_{V_1}} \otimes \dots \otimes \hat{x}_{\rho_{V_n}} \left(\hat{x}_{\rho_{T_1}}^\dagger \oplus \dots \oplus \hat{x}_{\rho_{T_k}}^\dagger \right) \in \text{End}(V_1 \otimes \dots \otimes V_n), \quad (38)$$

where the direct sum runs over the k irreducible unitary representations appearing in an orthogonal decomposition $V_1 \otimes \dots \otimes V_n = T_1 \oplus \dots \oplus T_k$.

The completeness of spectra of order $n \geq 2$ (Proposition 27) extends to the non-commutative case (Kakarala, 2009). The proof we provide of Proposition 27 translates almost identically to non-commutative groups, assuming $x, y \in \mathbb{R}^G$ have non-singular Fourier coefficients, i.e., $\hat{x}_{\rho_V}, \hat{y}_{\rho_V} \in \text{GL}(V) \subseteq \text{End}(V)$. The only substantial difference in the proof is that, in order to conclude that $\eta(\rho) = \rho(g)^\dagger$ for some $g \in G$, it is necessary to invoke Tannaka-Krein duality.