



Homogeneous vector bundles and G-equivariant convolutional neural networks

Downloaded from: <https://research.chalmers.se>, 2025-12-04 23:39 UTC

Citation for the original published paper (version of record):

Aronsson, J. (2022). Homogeneous vector bundles and G-equivariant convolutional neural networks. *Sampling Theory, Signal Processing, and Data Analysis*, 20(2).
<http://dx.doi.org/10.1007/s43670-022-00029-3>

N.B. When citing this work, cite the original published paper.



Homogeneous vector bundles and G -equivariant convolutional neural networks

Jimmy Aronsson¹ 

Received: 11 October 2021 / Accepted: 23 June 2022

© The Author(s) 2022

Abstract

G -equivariant convolutional neural networks (GCNNs) is a geometric deep learning model for data defined on a homogeneous G -space \mathcal{M} . GCNNs are designed to respect the global symmetry in \mathcal{M} , thereby facilitating learning. In this paper, we analyze GCNNs on homogeneous spaces $\mathcal{M} = G/K$ in the case of unimodular Lie groups G and compact subgroups $K \leq G$. We demonstrate that homogeneous vector bundles are the natural setting for GCNNs. We also use reproducing kernel Hilbert spaces (RKHS) to obtain a sufficient criterion for expressing G -equivariant layers as convolutional layers. Finally, stronger results are obtained for some groups via a connection between RKHS and bandwidth.

Keywords Convolutional neural networks · Deep learning · Equivariance · Fiber bundles · Geometry · Symmetry

Mathematics Subject Classification 68T07 · 43A85

1 Introduction

Developments in deep learning have increased dramatically in recent years. Even though multilayer perceptrons [2] and other general-architecture models work well for some tasks, achieving higher levels of performance often requires models that are more tailored to each application, and which incorporate some level of understanding of the data. *Geometric deep learning* [5–7, 13, 34] is the approach of using inherent

This article is part of the topical collection “Recent advances in computational harmonic analysis” edited by Dae Gwan Lee, Ron Levie, Johannes Maly and Hanna Veselovska.

✉ Jimmy Aronsson
jimmyar@chalmers.se

¹ Department of Mathematical Sciences, Chalmers University of Technology, 412 96 Göteborg, Sweden

geometric structure in data, and symmetry derived from geometry, to improve deep learning models.

Convolutional neural networks (CNNs) are among the simplest and most broadly applicable general-architecture models. They have been successfully applied to image classification and segmentation [37, 47, 48], text summarization [38], pose estimation [33], sign language recognition [24], and many other tasks. One of the reasons why CNNs are so successful is that convolutional layers are *translation equivariant*, which means that convolutional layers commute with the translation operator

$$L_y : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2, \quad L_y(x) = x + y, \quad y \in \mathbb{Z}^2. \quad (1.1)$$

In image classification tasks, for instance, \mathbb{Z}^2 represents the pixel lattice and CNNs utilize translation equivariance to identify objects in images regardless of their pixel coordinates. CNNs are examples of geometric deep learning models, as convolutional layers respect the global translation symmetry in \mathbb{Z}^2 .

G-equivariant convolutional neural networks (GCNNs) [10, 12] are generalizations of CNNs. These use *G-equivariant layers* that commute with the translation operator

$$L_g : \mathcal{M} \rightarrow \mathcal{M}, \quad L_g(x) = g \cdot x, \quad g \in G, \quad (1.2)$$

on a homogeneous space \mathcal{M} with global symmetry group G . *G*-equivariance means that GCNNs do not need to learn the global symmetry of \mathcal{M} , it is already built into the network design. GCNNs may therefore focus on learning other relevant features in data and thereby achieve high performance. Consider for example the detection of tumors in digital pathology. Images of tumors can have any orientation, depending on where the tumor is located and its relative position to the camera, and GCNNs with translation equivariant as well as rotation equivariant layers achieve higher accuracy than CNNs [41]. Rotation equivariance is also beneficial in 3D inference problems [44], point cloud recognition [31], and other tasks.

Gauge equivariant neural networks [9, 13, 17, 32] are instead designed to respect *local* symmetries. For example, computations involving vector fields - in meteorology or other areas - require vectors to be expressed in components. This requires choosing a *frame (of reference)* which assigns a basis to each tangent space. Many manifolds do not admit a global frame, however. Computations must instead be performed locally using different local frames for different regions on the manifold. It is then important that numerical results obtained in one frame are compatible with those obtained in others frames, especially on overlapping regions. Computations involving vector fields should thus be equivariant with respect to the choice of local frame. This choice is an internal (*gauge*) degree of freedom; a local symmetry. Gauge equivariant neural networks have also been introduced for problems exhibiting other local symmetries, primarily in lattice gauge theory.

In short, GCNNs preserve global symmetries (such as the translation symmetry in Euclidean space or the rotation symmetry of spheres) while gauge equivariant neural networks preserve local symmetries (internal degrees of freedom). These two types of networks have typically been studied independently of each other, despite sharing many similarities. It would be useful to have a framework that includes both types

of networks, since this would make it possible to analyze gauge equivariant neural networks and GCNNs simultaneously. It would emphasize their similarities as well as their differences, and could serve as a foundation for the study of equivariance in geometric deep learning.

Implementations of GCNNs are based on convolutional layers, but one can define more general G -equivariant layers which could be equally useful in a G -equivariant neural network. It may even be the case that convolutional layers form a small subset of all G -equivariant linear transformations, limiting their expressivity. It is therefore important to investigate the relationship between general G -equivariant layers and the more specific convolutional layers. This relationship is studied in our main result, Theorem 3.22, which we summarize below.

In the following theorem, the inputs to a G -equivariant layer ϕ are called *feature maps*. These are certain functions $f : G \rightarrow V$ into a finite-dimensional vector space V . Examples of feature maps for an ordinary, translation equivariant CNN ($G = \mathbb{Z}^2$) include digital images $f : \mathbb{Z}^2 \rightarrow \mathbb{R}^3$, which map each pixel to an array of RGB values for that pixel. Moreover, vector fields on any homogeneous G -space \mathcal{M} can be viewed as feature maps $f : G \rightarrow \mathbb{R}^{\dim \mathcal{M}}$. See Sect. 3 for details.

Theorem *Let ϕ be a G -equivariant layer and suppose that $\|\phi f(g)\| \leq \|\phi f\|$ for any feature map f and each $g \in G$. Then ϕ is a convolutional layer,*

$$(\phi f)(g) = \int_G \kappa(g^{-1}g') f(g') dg', \quad g \in G, \quad (1.3)$$

with an operator-valued kernel κ .

Our proof of Theorem 3.22 makes use of the fact that feature maps form a Hilbert space with an integral inner product. The required relation $\|\phi f(g)\| \leq \|\phi f\|$ ensures that the evaluation operator $f \mapsto \phi f(g)$ is continuous for each $g \in G$, and the Riesz representation theorem can then be used to construct a convolution kernel κ for the G -equivariant layer ϕ .

Our contributions in this paper are threefold:

- We motivate GCNNs from the point of view of homogeneous vector bundles, and we then demonstrate where GCNNs fit within a general framework that also includes gauge equivariant neural networks.
- Our main result, Theorem 3.22, gives a sufficient criterion for writing a G -equivariant layer as a convolutional layer. It holds for all homogeneous spaces $\mathcal{M} = G/K$ where G is a unimodular Lie group and K is a compact subgroup. We further prove the following corollaries:
 - (1) When G is either discrete abelian or finite, all G -equivariant layers are convolutional layers (Corollary 3.23).
 - (2) Convolutional layers can be computed as integrals over the homogeneous space G/K , despite being defined by integrals over G (Corollary 3.24).
- Our main theorem makes use of reproducing kernel Hilbert spaces (RKHS), i.e., function spaces for which evaluation $f \mapsto f(g)$ is a continuous functional. The relevance of RKHS to the study of GCNNs has not been noted before. In

Appendix A we discuss the deep connection between RKHS and bandlimited functions, which we believe may aid future research on GCNNs.

This work was inspired by a number of papers [9–13, 43]. The papers [9, 12] have been of particular importance, as our work grew from a desire to understand the mathematics of equivariant neural networks in even greater detail.

In the case of compact groups G , the Peter–Weyl theorem and other powerful tools have allowed researchers to study GCNNs using harmonic analysis. Among the most well-known results in this direction is [25, Theorem 1], which uses Fourier analysis to establish that for compact groups G , the layers in a G -equivariant feed-forward neural network must be convolutional layers. The close similarity between [25, Theorem 1] and our main theorem 3.22 is discussed in Sect. 3.4. Finally, others have used the well-known representation theory of $G = SO(3)$ to study rotation equivariant GCNNs for spherical data [15, 16].

The paper is structured as follows. We summarize the relevant machine learning background in Sect. 2.1, give a brief introduction to fiber bundles in Sect. 2.2 and discuss a framework for equivariant neural networks in Sect. 2.3. In Sect. 3.1 we restrict attention to homogeneous spaces and G -equivariance. Sections 3.2–3.3 discuss the relationship between GCNNs, homogeneous vector bundles, and induced representations. This relation is then used to motivate the definition of G -equivariant layers in Sect. 3.4, in which we also discuss convolutional layers and prove our main theorem 3.22. We summarize our work in Sect. 4. Finally, Appendix A discusses the connection between reproducing kernel Hilbert spaces and bandlimited functions.

2 Foundations of equivariant neural networks

In this section, we give an introduction to convolutional neural networks (CNNs) and discuss a simple framework for equivariant neural networks. Readers who already know about CNNs and translation equivariance may skip most of this section, except for the last few paragraphs on GCNNs.

2.1 Convolutional neural networks

CNNs were first introduced in 1979 under the name of Neocognitrons, and were used to study visual pattern recognition [20]. In the 1990s, CNNs were successfully applied to problems such as automatic recognition of handwritten digits [29] and face recognition [28]. However, it was arguably not until 2012, when the GPU-based AlexNet CNN outperformed all competition on the ImageNet Large Scale Visual Recognition Challenge [26], that CNNs and other neural networks truly caught the public eye. Industrial work and academic research on deep learning has since soared, and current state-of-the-art deep learning architectures are significantly more powerful and more complex than AlexNet. Yet, convolutional layers remain important components.

In this introduction we focus on *feature maps* that can be represented by functions

$$f : \mathbb{Z}^2 \rightarrow \mathbb{R}^m. \quad (2.1)$$

Digital images, for example, are of this form since each pixel $x \in \mathbb{Z}^2$ is associated to a color array $f(x) \in \mathbb{R}^m$, with $m = 1$ corresponding to grayscale images and $m = 3$ to RGB images. Any 2D ($m = 1$) or 3D ($m > 1$) array with real-valued entries is a feature map of the form (2.1). *Convolutional layers* act on such feature maps by

$$[\kappa \star f](x) = \sum_{y \in \mathbb{Z}^2} \kappa(y - x) f(y), \quad (2.2)$$

given a matrix-valued kernel $\kappa : \mathbb{Z}^2 \rightarrow \text{Hom}(\mathbb{R}^m, \mathbb{R}^n)$ for some $n \in \mathbb{N}$. Observe that the resulting maps $\kappa \star f : \mathbb{Z}^2 \rightarrow \mathbb{R}^n$ are themselves feature maps (2.1) with n channels. Broadly speaking, CNNs consist of convolutional layers (2.2) combined with other transformations, including non-linear activation functions and batch normalization layers. We are primarily interested in convolutional layers and will not go into detail about activation functions or other types of layers. For more extensive descriptions of CNNs, see [1, 21, 45].

Remark 2.1 The name *convolutional layer* is commonly used in the machine learning literature even though (2.2) more closely resembles a cross-correlation. It can be turned into a convolution by replacing the kernel with its involution $\kappa^*(y) = \kappa(-y)$.

CNNs perform very well on image classification and similar machine learning tasks, and are important parts of many state-of-the-art network architectures on such tasks [4, 22, 40, 46]. One reason for their success is *translation equivariance*: Convolutional layers (2.2) commute with the translation operator in the image plane,

$$L_x : \mathbb{Z}^2 \rightarrow \mathbb{Z}^2, \quad L_x(y) = y + x, \quad x \in \mathbb{Z}^2. \quad (2.3)$$

Translation equivariance makes CNNs agnostic to the specific locations of individual pixels, while still taking into account the relative positions of different pixels; images are more easily classified based on relevant features of their subjects, and not based on technical artifacts such as specific pixel coordinates. This observation motivates the introduction of more general convolutional layers that act equivariantly on data points $f : \mathcal{M} \rightarrow V$, where the domain \mathcal{M} is homogeneous with respect to a locally compact group G [10, 12]. Given finite-dimensional vector spaces V, W , convolutional layers are defined as certain vector-valued integrals.

$$\kappa \star f : \mathcal{M} \rightarrow W, \quad (\kappa \star f)(g) = \int_G \kappa(g^{-1}g') f(g') \, dg', \quad (2.4)$$

with operator-valued kernels $\kappa : G \rightarrow \text{Hom}(V, W)$ and integration with respect to a Haar measure on G . For a summary on vector-valued integration on locally compact groups, see [18, Appendix 4]. G -equivariant convolutional neural networks (GCNNs) in their simplest form are sequences

$$\sigma_L \circ \phi_L \circ \cdots \circ \sigma_1 \circ \phi_1, \quad (2.5)$$

of G -equivariant linear transformations (G -equivariant layers) ϕ_ℓ and scalar-valued, non-linear activation functions σ_ℓ . We do not require ϕ_1, \dots, ϕ_L to be convolutional layers (2.4), although only using convolutional layers is certainly an option. One of our main goals in this paper is to better understand the expressivity of convolutional layers, and we do this by comparing them to more general G -equivariant layers. See Definitions 3.14, 3.15 for formal definitions.

2.2 Background on fiber bundles

This paper assumes familiarity with fiber bundles, especially with principal bundles and associated vector bundles. Still, given how central fiber bundles are to this paper, we will present relevant definitions and give a few examples. See [23, 30, 36] for more detailed introductions. This section may be skipped by those who already know about principal bundles and associated vector bundles.

The geometric intuition behind fiber bundles can be summarized as follows: Some manifolds can be formed by taking another, in some sense simpler *base space* \mathcal{M} and attaching a *fiber* F to the base space at each point $x \in \mathcal{M}$. The cylinder $S^1 \times [0, 1]$, for example, can be viewed as the circle $\mathcal{M} = S^1$ together with a thin strip $F = [0, 1]$ that has been glued perpendicularly to the circle at each point $x \in S^1$. The Möbius strip can be visualized in precisely the same way but, in this case, the relative angle between the thin strips and the circle changes from 0 to π as we move around the circle. Locally, in some neighbourhood $U \subset S^1$ around any point $x \in S^1$, the cylinder and the Möbius strip both resemble a rectangle $U \times [0, 1]$, but the two manifolds are not globally equivalent since the Möbius strip is not diffeomorphic to $S^1 \times [0, 1]$.

The point is that both the cylinder and the Möbius strip are examples of smooth fiber bundles over the base space $\mathcal{M} = S^1$ with model fiber $F = [0, 1]$. The former is a *trivial* fiber bundle $\mathcal{M} \times F$, the latter is a non-trivial fiber bundle.

Definition 2.2 [30, p. 268] Let \mathcal{M} and F be smooth manifolds. A (*smooth*) *fiber bundle* over \mathcal{M} with model fiber F is a smooth manifold E together with a surjective smooth map $\pi : E \rightarrow \mathcal{M}$ with the following property: For each $x \in \mathcal{M}$, there is a neighbourhood $U \subseteq \mathcal{M}$ containing x , and a diffeomorphism

$$\xi : \pi^{-1}(U) \rightarrow U \times F, \quad (2.6)$$

called a *local trivialization* of E over U , such that the following diagram commutes:

$$\begin{array}{ccc} \pi^{-1}(U) & \xrightarrow{\xi} & U \times F \\ & \searrow \pi \quad \swarrow \pi_1 & \\ & U & \end{array}$$

That is, $\pi = \pi_1 \circ \xi$ where $\pi_1 : U \times F \rightarrow U$ is the projection onto the first component.

Remark 2.3 Smooth manifolds are assumed to be Hausdorff and second-countable.

The smooth manifolds \mathcal{M} and E are respectively known as the *base space* and the *total space* of the fiber bundle, and the surjective smooth map $\pi : E \rightarrow \mathcal{M}$ is called the *projection*. For each $x \in \mathcal{M}$, the set $E_x = \pi^{-1}(\{x\})$ is diffeomorphic to $\{x\} \times F$ via a local trivialization, and is known as the *fiber at x* . Fiber bundles are typically denoted by their projection $\pi : E \rightarrow \mathcal{M}$, or simply by the total space E if the other components are understood. We will switch freely between these two notations.

One way to use fiber bundles is to deconstruct manifolds into simpler components, as exemplified by the cylinder and the Möbius strip. But fiber bundles are also useful for adding structure to the base space via an appropriate model fiber. *Vector bundles* and *principal bundles* are two prominent examples in which the model fiber is a vector space and a Lie group, respectively.

Definition 2.4 A (real/complex) *vector bundle* is a smooth fiber bundle $\pi : E \rightarrow \mathcal{M}$ for which the model fiber $F = V$ is a (real/complex) vector space and which satisfies the following for each $x \in \mathcal{M}$:

- (i) The fiber $E_x = \pi^{-1}(\{x\})$ is a (real/complex) vector space.
- (ii) If $\xi : \pi^{-1}(U) \rightarrow U \times V$ is a local trivialization on some neighbourhood U of x , then the restriction $\pi : E_x \rightarrow \{x\} \times V$ is a linear isomorphism.

The prototypical example of a vector bundle is the *tangent bundle* $T\mathcal{M}$ associated to any smooth manifold \mathcal{M} . This bundle is discussed in the next example, following a reminder about tangent vectors and tangent spaces.

Example 2.5 Let \mathcal{M} be a smooth manifold of dimension $d = \dim \mathcal{M}$ and recall that for each $x \in \mathcal{M}$, there is an associated *tangent space* $T_x\mathcal{M}$, which is a d -dimensional real vector space. Elements of each tangent space are called *tangent vectors* and can be defined as follows. First, consider a coordinate chart $u = (u^1, \dots, u^d) : U \rightarrow \mathbb{R}^d$ for which $x \in U$. The *coordinate basis vectors*

$$\left. \frac{\partial}{\partial u^i} \right|_x \in T_x\mathcal{M}, \quad i = 1, \dots, d, \quad (2.7)$$

act on any smooth function $f : \mathcal{M} \rightarrow \mathbb{R}$ by evaluating the i 'th partial derivative of the smooth function $f \circ u^{-1} : \mathbb{R}^d \rightarrow \mathbb{R}$ at the point $u(x) \in \mathbb{R}^d$:

$$\left. \frac{\partial}{\partial u^i} \right|_x f := \left. \frac{\partial (f \circ u^{-1})}{\partial u^i} \right|_{u(x)}. \quad (2.8)$$

The tangent space $T_x\mathcal{M}$ is defined as the linear span of the coordinate basis vectors, meaning that arbitrary tangent vectors $X_x \in T_x\mathcal{M}$ are linear combinations

$$X_x = \sum_{i=1}^d X_x^i \left. \frac{\partial}{\partial u^i} \right|_x, \quad X_x^1, \dots, X_x^d \in \mathbb{R}. \quad (2.9)$$

Choosing a different coordinate chart would result in new coordinate basis vectors, but their span would remain the same. The tangent space $T_x\mathcal{M}$ is thus well-defined.

The tangent bundle of \mathcal{M} is defined as the disjoint union

$$T\mathcal{M} = \bigsqcup_{x \in \mathcal{M}} T_x\mathcal{M} := \bigcup_{x \in \mathcal{M}} \{(x, X_x) \mid X_x \in T_x\mathcal{M}\}, \quad (2.10)$$

of all tangent spaces on \mathcal{M} , together with a natural topology and a smooth structure that turns $T\mathcal{M}$ into a smooth manifold of dimension $2d$. The projection $\pi : T\mathcal{M} \rightarrow \mathcal{M}$ maps each pair $(x, X_x) \in T\mathcal{M}$ to the point $x \in \mathcal{M}$, implying that the fiber at x can be identified with the tangent space at x :

$$\pi^{-1}(\{x\}) = \{x\} \times T_x\mathcal{M} \simeq T_x\mathcal{M}. \quad (2.11)$$

Moreover, any coordinate chart $u = (u^1, \dots, u^d) : U \rightarrow \mathbb{R}^d$ yields a local trivialization

$$\xi : \pi^{-1}(U) \rightarrow U \times \mathbb{R}^d, \quad (x, X_x) \mapsto (x, (X_x^1, \dots, X_x^d)), \quad (2.12)$$

by expanding $X_x \in T_x\mathcal{M}$ in the coordinate basis for each $x \in U$. The restriction of (2.12) to any $x \in U$ defines a vector space isomorphism between $T_x\mathcal{M}$ and \mathbb{R}^d , and the tangent bundle $T\mathcal{M}$ is thereby a real vector bundle with model fiber \mathbb{R}^d . \square

This example illustrates one benefit of fiber bundles: By attaching a tangent space $T_x\mathcal{M}$ to each point x in the base space \mathcal{M} , thereby forming the tangent bundle $T\mathcal{M}$, we are able to compute directional derivatives of smooth functions on \mathcal{M} . But why bother constructing the bundle $T\mathcal{M}$ instead of just working with the tangent spaces $T_x\mathcal{M}$ individually? Well, thanks to the tangent bundle being a smooth manifold, we can move smoothly between different tangent spaces. This makes it possible to define (*local*) *vector fields* as smooth functions $X : U \rightarrow T\mathcal{M}$ such that $\pi \circ X = \text{Id}_U$, for any open set $U \subseteq \mathcal{M}$. That is, a vector field smoothly attaches a tangent vector $X(x) \in T_x\mathcal{M}$ to the base space at each point $x \in U$. Vector fields are heavily used in pure mathematics as well as in applications, such as electric and magnetic fields in electrodynamics, and wind fields in meteorology. Furthermore, the idea of smoothly assigning an element of the fiber E_x to the base space at each point $x \in U$, extends to arbitrary fiber bundles.

Definition 2.6 Let $\pi : E \rightarrow \mathcal{M}$ be a fiber bundle and let $U \subset \mathcal{M}$ be an open set. A (*local*) *section* is a smooth map $s : U \rightarrow E$ satisfying $\pi \circ s = \text{Id}_U$.

Vector fields are thus sections of the tangent bundle. In Sect. 2.3, we will define *gauges* as sections of principal bundles, and *data points* as certain sections of vector bundles associated to a principal bundle. While this may sound abstract, the idea is that data contains information in the form of a vector at each point of a manifold, just like a digital image $s : \mathbb{Z}^2 \rightarrow \mathbb{R}^3$ contains information in the form of an RGB array $s(x)$ at each point x of the pixel lattice. Modeling data points as sections of vector bundles over \mathcal{M} lets us take the geometry of \mathcal{M} into account, and principal bundles help us utilize relevant symmetries. The symmetries in question are encoded in group actions:

Definition 2.7 Let \mathcal{X} be a smooth manifold and H a Lie group. A smooth function

$$\triangleleft : \mathcal{X} \times H \rightarrow \mathcal{X}, \quad (x, h) \mapsto x \triangleleft h, \quad (2.13)$$

is called a *smooth right H -action* on \mathcal{X} if it is compatible with the group structure:

$$x \triangleleft (hh') = (x \triangleleft h) \triangleleft h', \quad x \in \mathcal{X}, h, h' \in H. \quad (2.14)$$

Note that the identity $e \in H$ stabilizes all points $x \in \mathcal{X}$, in the sense that $x \triangleleft e = x$. A smooth right H -action is called *free* if the identity is the only stabilizer. Furthermore, a smooth right H -action is *transitive* if it can be used to move between any two points. In other words, if for all pairs $x, x' \in \mathcal{X}$, there exists $h \in H$ such that $g' = g \triangleleft h$.

Smooth *left* H -actions $H \times \mathcal{X} \rightarrow \mathcal{X}$, $(h, x) \mapsto h \triangleright x$, are defined analogously.

The Lie group H represents a set of symmetry transformations that can be applied to \mathcal{X} . We are concerned with two different types of symmetries:

- (i) Global symmetries, whereby a Lie group G acts on a smooth manifold \mathcal{M} via a transitive left action $G \times \mathcal{M} \rightarrow \mathcal{M}$ that is typically written as $(g, x) \mapsto g \cdot x$. The symmetry is global because the transitive action can transform any point x to any other point x' by applying a suitable symmetry transformation g . For example, any point on the sphere can reach any other point on the sphere by performing a suitable rotation ($G = SO(3)$, $\mathcal{M} = S^2$).
- (ii) Local symmetries, whereby a Lie group K acts almost independently on each point x of a smooth manifold \mathcal{M} . The action does not move different points $x, x' \in \mathcal{M}$ to each other, but rather represents an internal degree of freedom at each individual point. Local symmetries are modeled by attaching a copy of K to each point x on the base space \mathcal{M} , thereby forming a larger space P for which the local symmetry is more explicit. This is the motivation behind principal bundles.

Definition 2.8 Let K be a Lie group. A smooth fiber bundle $\pi : P \rightarrow \mathcal{M}$ is called a *principal K -bundle* with *structure group* K if there is a free, smooth right K -action

$$P \times K \rightarrow P, \quad (p, k) \mapsto p \triangleleft k, \quad (2.15)$$

with the following properties for each $x \in \mathcal{M}$.

- (i) Let $P_x = \pi^{-1}(\{x\})$ be the fiber at x . Then

$$p \in P_x, k \in K \Rightarrow p \triangleleft k \in P_x. \quad (2.16)$$

That is, the K -action preserves fibers.

- (ii) For each $p \in P_x$, the mapping $k \mapsto p \triangleleft k$ is a diffeomorphism $K \rightarrow P_x$.

Principal bundles are thus natural tools for understanding local symmetries; *gauge* degrees of freedom. In theoretical physics, gauge degrees of freedom are redundancies in the mathematical theory. These redundancies have no physical relevance, they are not observable, but they also cannot be ignored as they are present in the model. For

example, the Yang-Mills equations of motion are underdetermined when modeled as an initial value problem, and cannot be solved without choosing a gauge that fixes the gauge degree of freedom [35].

One important example of a local symmetry, of a gauge degree of freedom, is the freedom to choose bases in tangent spaces.

Example 2.9 Elements of the frame bundle $F\mathcal{M}$ over a smooth manifold \mathcal{M} are (in a precise sense) bases in the tangent spaces $T_x\mathcal{M}$. Given a local coordinate chart $u = (u^1, \dots, u^d) : U \rightarrow \mathbb{R}^d$, for example, the coordinate basis

$$p = \left(\frac{\partial}{\partial u^1} \Big|_x, \dots, \frac{\partial}{\partial u^d} \Big|_x \right), \quad (2.17)$$

is an element of the frame bundle $F\mathcal{M}$ for each $x \in U$. The projection $\pi : F\mathcal{M} \rightarrow \mathcal{M}$ sends each basis $p = (p_1, \dots, p_d)$ in $T_x\mathcal{M}$ to the point $\pi(p) = x$, which in particular means that the fiber $F\mathcal{M}_x = \pi^{-1}(\{x\})$ is the set of all bases in $T_x\mathcal{M}$.

We may define a smooth right $GL(d, \mathbb{R})$ -action

$$\triangleleft : F\mathcal{M} \times GL(d, \mathbb{R}) \rightarrow F\mathcal{M}, \quad (p, k) \mapsto p \triangleleft k, \quad (2.18)$$

by essentially applying a change-of-basis matrix $k \in GL(d, \mathbb{R})$ to the elements of any basis $p \in F\mathcal{M}_x$. A more precise definition of this action would require us to utilize the linear isomorphism between $T_x\mathcal{M}$ and \mathbb{R}^d and the induced Lie group isomorphism between $GL(T_x\mathcal{M})$ and $GL(d, \mathbb{R})$, but let us simplify the notation by avoiding that step. The action then satisfies

$$(p \triangleleft k)_i = k^{-1} p_i, \quad i = 1, \dots, d, \quad (2.19)$$

which defines the basis vectors $(p \triangleleft k)_i \in T_x\mathcal{M}$ in terms of the basis vectors $p_i \in T_x\mathcal{M}$. Note that the action 2.18 preserves the fibers $F\mathcal{M}_x$ because it maps between different bases in the tangent space $T_x\mathcal{M}$ for each $x \in \mathcal{M}$. Moreover, any basis $p \in F\mathcal{M}_x$ can be mapped to any other basis $\tilde{p} \in F\mathcal{M}_x$ by a unique change-of-basis matrix, so the mapping

$$K \rightarrow F\mathcal{M}_x, \quad k \mapsto p \triangleleft k, \quad (2.20)$$

is a bijection for each $p \in F\mathcal{M}_x$ and each $x \in \mathcal{M}$. In fact, the smooth structure on $F\mathcal{M}$ is such that (2.20) is a diffeomorphism. The frame bundle $F\mathcal{M}$ is therefore a principal bundle over \mathcal{M} with structure group $GL(d, \mathbb{R})$. Local sections $\omega : U \rightarrow F\mathcal{M}$ are called *local frames (of reference)* and are useful e.g. in theoretical physics. \square

The tangent bundle $T\mathcal{M}$ is *associated* to the frame bundle $F\mathcal{M}$, in the sense that tangent vectors $X \in T\mathcal{M}$ can be expanded in a basis $p \in F\mathcal{M}$. This is an example of a more general phenomenon whereby a vector bundle can be associated to a principal bundle. Let us take a closer look at this example before formally defining *associated vector bundles*.

Example 2.10 Fix a tangent vector $X \in T_x \mathcal{M}$ and a basis $p = (p_1, \dots, p_d) \in F\mathcal{M}_x$, for some $x \in \mathcal{M}$. If we collect the components of $X = \sum_i X^i p_i$ in an array

$$X(p) = [X^1 \dots X^d]^T \in \mathbb{R}^d, \quad (2.21)$$

then the following tuple is a basis-dependent decomposition of X :

$$(p, X(p)) = (p_1, \dots, p_d, X^1, \dots, X^d) \in F\mathcal{M} \times \mathbb{R}^d. \quad (2.22)$$

Performing a change of basis $p \mapsto p \triangleleft k$ by (2.19), for some $k \in GL(d, \mathbb{R})$, transforms this decomposition by

$$(p, X(p)) \mapsto (p \triangleleft k, X(p \triangleleft k)) = (p \triangleleft k, kX(p)). \quad (2.23)$$

This transformation does not change the tangent vector X itself, it only changes the decomposition of X in terms of basis vectors and components. For this reason, if we define an equivalence relation on $F\mathcal{M} \times \mathbb{R}^d$ by

$$(p, X(p)) \sim (p \triangleleft k, X(p \triangleleft k)), \quad k \in K, \quad (2.24)$$

and condense all decompositions $(p, X(p))$, for all bases p , into their equivalence class $[p, X(p)]$, then the result is a basis-independent description of the tangent vector X . In fact, the quotient space $(F\mathcal{M} \times \mathbb{R}^d)/\sim$ is a smooth vector bundle that is isomorphic to the tangent bundle $T\mathcal{M}$. \square

This example hints towards the following construction of vector bundles associated to an arbitrary principal bundle P . Associated vector bundles are useful because they make local symmetries in vector bundles more explicit. Just compare the two objects $X \in T\mathcal{M}$ and $[p, X(p)] \in (F\mathcal{M} \times \mathbb{R}^d)/\sim$, both of which represent a tangent vector. The local symmetry, that tangent vectors do not depend on the choice of basis p , is much more visible in the object $[p, X(p)]$ than in X .

Definition 2.11 Let $\pi : P \rightarrow \mathcal{M}$ be a principal bundle with structure group K and let $\rho : K \rightarrow GL(V_\rho)$ be a representation on a finite-dimensional vector space V_ρ . Define an equivalence relation on $P \times V_\rho$ by

$$(p, v) \sim (p \triangleleft k, \rho(k)^{-1}v), \quad p \in P, v \in V_\rho, k \in K. \quad (2.25)$$

The quotient space $P \times_\rho V_\rho = (P \times V_\rho)/\sim$, together with the mapping

$$\pi_\rho : P \times_\rho V_\rho \rightarrow \mathcal{M}, \quad \pi_\rho([p, v]) = \pi(p), \quad (2.26)$$

is called an *associated vector bundle* over \mathcal{M} .

Lemma 2.12 [23, Sect. 10.7] *The associated bundle $\pi_\rho : P \times_\rho V_\rho \rightarrow \mathcal{M}$ is a smooth vector bundle.*

2.3 Gauge theory and the equivariant framework

In this section we describe a mathematical framework for equivariant neural networks. This framework is based on gauge theoretic concepts, which makes it suitable for describing gauge equivariant neural networks, but it is equally suitable for GCNNs. In fact, the framework unifies existing works on both types of networks [9, 12].

Gauge theory originated in physics as a way to model local symmetry. In quantum electrodynamics, for example, the electron wave function can be locally phase shifted, $\psi \mapsto e^{i\alpha}\psi$, with no physically observable consequence, so quantum electrodynamics is said to possess a $U(1)$ gauge symmetry. Mathematicians have later adopted gauge theory in order to study other types of local symmetries. The introduction of gauge equivariant deep learning models has been suggested by deep learning practitioners and physicists alike. For example, [9] investigates the structure of gauge equivariant layers used for vector fields, tensor fields, and more general fields. Physicists have introduced gauge equivariant neural networks for applications in, e.g., lattice gauge theory [3, 17, 32].

Definition 2.13 Let $\pi : P \rightarrow \mathcal{M}$ be a principal K -bundle and let $U \subseteq \mathcal{M}$ be open.

- (i) A (local) gauge is a section $\omega : U \rightarrow P$.
- (ii) A gauge transformation is an automorphism $\chi : P \rightarrow P$ that preserves fibers ($\pi \circ \chi = \pi$) and that is gauge equivariant, in the sense that

$$\chi(p \triangleleft k) = \chi(p) \triangleleft k, \quad p \in P, k \in K. \quad (2.27)$$

Example 2.14 Local frames are local gauges $\omega : U \rightarrow F\mathcal{M}$ of the frame bundle $P = F\mathcal{M}$. Since each element $p \in F\mathcal{M}$ is a basis in the tangent space $T_{\pi(p)}\mathcal{M}$, and the structure group of the frame bundle is $K = GL(d, \mathbb{R})$, gauge transformations are maps $\chi : F\mathcal{M} \rightarrow F\mathcal{M}$ that are equivariant to changes of basis $p \mapsto p \triangleleft k$. \square

Equivariant neural networks use the language of principal and associated bundles. In the remainder of this subsection, let $E_\rho = P \times_\rho V_\rho$ and $E_\sigma = P \times_\sigma V_\sigma$ be associated bundles, given a principal bundle $\pi : P \rightarrow \mathcal{M}$ over a smooth manifold \mathcal{M} . Further let $\Gamma_c(E_\rho)$ and $\Gamma_c(E_\sigma)$ be the vector spaces of compactly supported continuous sections of E_ρ and E_σ , respectively.

Definition 2.15 A data point is a section $s \in \Gamma_c(E_\rho)$.

Remark 2.16 We restrict attention to complex representations (ρ, V_ρ) to simplify the mathematical theory. Our decision to focus on compactly supported sections was also made for mathematical reasons: G -equivariant layers are defined in Sect. 3.4 in terms of an induced representation, which lives on the completion of $\Gamma_c(E_\rho)$ with respect to an inner product.

Definition 2.17 A feature map $f : P \rightarrow V_\rho$ is a compactly supported continuous function that satisfies the transformation property

$$f(p \triangleleft k) = \rho(k)^{-1} f(p), \quad (2.28)$$

for all $p \in P, k \in K$. The vector space of such feature maps is denoted $C_c(P; \rho)$

Data points and feature maps are, in a sense, dual to each other: Each data point in $\Gamma_c(E_\rho)$ is of the form

$$s_f(x) = [p, f(p)], \quad (2.29)$$

for a feature map $f \in C_c(P; \rho)$, where $p \in P_x$ is any element of the fiber at $x \in \mathcal{M}$. In order to see that (2.29) does not depend on the choice of p , recall from Definition 2.8 that the mapping $K \rightarrow P_x$ given by $k \mapsto p \triangleright k$ is a diffeomorphism, for each $p \in P$. For this reason, given another element $p' \in P_x$, there exists a unique $k \in K$ such that $p' = p \triangleleft k$ and

$$[p', f(p')] = [p \triangleleft k, f(p \triangleleft k)] = [p \triangleleft k, \rho(k^{-1})f(p)] = [p, f(p)]. \quad (2.30)$$

That is, the equivalence class $[p, f(p)]$ only depends on the basepoint x .

Lemma 2.18 [23, Sect. 10.12] *The linear map $C_c(P; \rho) \rightarrow \Gamma_c(E_\rho)$, $f \mapsto s_f$ is a vector space isomorphism.*

We are almost ready to define general and gauge equivariant layers. Before doing so, however, we must say how gauge transformations $\chi : P \rightarrow P$ act on data points. Let $\theta_\chi : P \rightarrow K$ be the uniquely defined map satisfying $\chi(p) = p \triangleleft \theta_\chi(p)$ for all $p \in P$, and define the following action on the associated bundle E_ρ :

$$\chi \cdot [p, v] = [\chi(p), v] = [p, \rho(\theta_\chi(p))v], \quad [p, v] \in E_\rho. \quad (2.31)$$

The corresponding action on data points is given by

$$(\chi \cdot s_f)(x) = [p, \rho(\theta_\chi(p))f(p)] = s_{\rho(\theta_\chi)f}(x), \quad s_f \in \Gamma_c(E_\rho). \quad (2.32)$$

We distinguish between general layers and more specific gauge equivariant layers, as G -equivariant layers in GCNNs will only be a special case of the former.

Definition 2.19 Let $E_\rho = P \times_\rho V_\rho$ and $E_\sigma = P \times_\sigma V_\sigma$ be associated bundles.

- (i) A (linear) layer is a linear map $\Phi : \Gamma_c(E_\rho) \rightarrow \Gamma_c(E_\sigma)$.
- (ii) A layer Φ is gauge equivariant if, for all gauge transformations $\chi : P \rightarrow P$,

$$\Phi \circ \chi = \chi \circ \Phi. \quad (2.33)$$

In equivariant neural networks, data points are sent through a sequence of layers, which are mixed with non-linear activation functions. Again, we focus on individual layers in this paper, and leave the analysis of equivariant activation functions and multi-layer networks to future work. The fiber bundle-theoretic concepts discussed in this part describe two kinds of equivariant neural networks:

- (i) Gauge equivariant neural networks, which respect local gauge symmetry and whose layers are gauge equivariant.

- (ii) GCNNs, which respect global translation symmetry in homogeneous G -spaces \mathcal{M} , and whose layers are G -equivariant (Definition 3.14).

Remark 2.20 Note that the definition of a layer in [9] deviates from ours by considering sections supported on a single coordinate chart. That work further investigates the structure of layers under additional assumptions.

It follows from Lemma 2.18 that a (gauge equivariant) layer $\Phi : \Gamma_c(E_\rho) \rightarrow \Gamma_c(E_\sigma)$ induces a unique linear map $\phi : C_c(P; \rho) \rightarrow C_c(P; \sigma)$ such that $\Phi s_f = s_{\phi f}$. Writing data points as $s_f = [\cdot, f]$ allows us to also express this relation as $\Phi[\cdot, f] = [\cdot, \phi f]$. We think of Φ and ϕ as two sides of the same coin, and use the name (gauge equivariant) layer for both maps.

$$\begin{array}{ccc} \Gamma_c(E_\rho) & \xrightarrow{\Phi} & \Gamma_c(E_\sigma) \\ \uparrow \downarrow & & \uparrow \downarrow \\ C_c(P; \rho) & \xrightarrow{\phi} & C_c(P; \sigma) \end{array}$$

Example 2.21 Let $T : V_\rho \rightarrow V_\sigma$ be a linear transformation and consider the layer

$$\phi : C_c(P; \rho) \rightarrow C_c(P; \sigma), \quad (\phi f)(p) = T(f(p)), \quad (2.34)$$

for $p \in P$, $f \in C_c(P; \rho)$. Since f and ϕf are feature maps and thereby satisfy (2.28), the linear transformation T must satisfy

$$\sigma(k)T(f(p)) = T(f(p \triangleleft k^{-1})) = T(\rho(k)f(p)), \quad (2.35)$$

for all $k \in K$, $p \in P$, $f \in C_c(P; \rho)$. This can be seen to imply that $\sigma \circ T = T \circ \rho$, so T intertwines the representations ρ and σ . Another way to arrive at this conclusion is to analyze when the corresponding layer

$$\Phi : \Gamma_c(E_\rho) \rightarrow \Gamma_c(E_\sigma), \quad \Phi s_f = [\cdot, \phi f], \quad (2.36)$$

is well-defined.

Now consider a gauge transformation $\chi : P \rightarrow P$ and its induced map $\theta_\chi : P \rightarrow K$. Because T is an intertwiner,

$$(\Phi \circ \chi)s_f = s_{\phi\rho(\theta_\chi)f} = s_{\sigma(\theta_\chi)\phi f} = (\chi \circ \Phi)s_f, \quad (2.37)$$

hence the layer Φ is automatically gauge equivariant. \square

As this example illustrates, gauge equivariance is tightly connected to intertwining properties of ϕ . Rearranging (2.37) gives the following result.

Lemma 2.22 *A general layer $\Phi : \Gamma_c(E_\rho) \rightarrow \Gamma_c(E_\sigma)$ is gauge equivariant iff*

$$\phi \circ \rho(\theta_\chi) f = \sigma(\theta_\chi) \circ \phi f, \quad (2.38)$$

for all gauge transformations $\chi : P \rightarrow P$ and all feature maps $f \in C_c(P; \rho)$.

This concludes our discussion of gauge theory and of equivariant neural networks. The framework for the latter is evidently very general, consisting of layers and non-linear activation functions between data points. There are advantages of working at this level of generality: Ordinary (non-equivariant) neural networks have a multitude of different types of layers, many of them linear. Equivariant analogues of such layers are likely to satisfy either Definition 2.19(ii) (gauge equivariance) or Definition 3.14 (G -equivariance; compatibility with the global symmetry when \mathcal{M} is homogeneous space), depending on the relevant type of equivariance. Results that can be proven in this general framework will thus hold for many different equivariant neural networks. One example is Theorem 3.22 below, which characterizes the structure of abstract G -equivariant layers in any GCNN.

3 G -equivariant convolutional neural networks

Recall that GCNNs generalize ordinary CNNs to use G -equivariant layers and data points defined on homogeneous G -spaces \mathcal{M} . Let us give a recap on homogeneous spaces and global symmetry, before discussing homogeneous vector bundles, sections, and induced representations. We will demonstrate that GCNNs and G -equivariant layers (originally defined in [12]) are most naturally understood from the perspective of homogeneous vector bundles. We will then apply reproducing kernel Hilbert spaces to understand which G -equivariant layers are expressible as convolutional layers.

3.1 Homogeneous spaces

Definition 3.1 Let G be a Lie group. A smooth manifold \mathcal{M} is called a *homogeneous G -space* if there exists a smooth, transitive left G -action

$$G \times \mathcal{M} \rightarrow \mathcal{M}, \quad (g, x) \mapsto g \cdot x. \quad (3.1)$$

Since the action (3.1) is transitive, we can arbitrarily fix an element $o \in \mathcal{M}$ that we call the *origin*, and express any other point $x \in \mathcal{M}$ as $x = g \cdot o$ for some $g \in G$. This group element is typically not unique, but observe that

$$g \cdot o = g' \cdot o \iff g^{-1}g' \in H_o, \quad (3.2)$$

where $H_o = \{g \in G \mid g \cdot o = o\}$ is the isotropy group of the origin o . In other words, there is a one-to-one correspondence between points $x \in \mathcal{M}$ and cosets $gH_o \in G/H_o$.

Proposition 3.2 [30, Theorem 21.18] *Let \mathcal{M} be a homogeneous G -space and fix an origin $o \in \mathcal{M}$. The isotropy group H_o is a closed subgroup of G , and the map*

$$F_o : G/H_o \rightarrow \mathcal{M}, \quad gH_o \mapsto g \cdot o, \quad (3.3)$$

is a diffeomorphism.

Homogeneous spaces are globally symmetric in the sense that any element $o \in \mathcal{M}$ may be chosen as origin. Given another choice of origin $o' \in \mathcal{M}$, the quotient spaces $G/H_{o'} \simeq G/H_o$ are diffeomorphically related by a translation in G - more precisely, by the composition $F_{o'}^{-1} \circ F_o$. For example, Euclidean space $\mathcal{M} = \mathbb{R}^d$ is homogeneous with respect to translations, allowing any point to be considered as origin. Similarly, the rotationally symmetric sphere $\mathcal{M} = S^2$ does not have a unique north pole.

Recall from Sect. 2.1 that convolutional neural networks utilize the homogeneity of the pixel lattice. We argued that translation equivariance, that is, the property that convolutional layers commute with the translation operator in \mathbb{Z}^2 , help CNNs learn relevant features in images. We want to extend this property to other types of data defined on more general homogeneous spaces. The first step in this direction is the following definition of G -equivariant functions on \mathcal{M} .

Definition 3.3 Let G be a Lie group and let \mathcal{M} and \mathcal{N} be homogeneous G -spaces. We say that a function $f : \mathcal{M} \rightarrow \mathcal{N}$ is G -equivariant if

$$f(g \cdot x) = g \cdot f(x), \quad (3.4)$$

for all $g \in G$ and all $x \in \mathcal{M}$.

The diffeomorphism F_o in (3.3) is G -equivariant since, for all $g \in G$,

$$F_o(gH_o) = g \cdot o = g \cdot F_o(H_o). \quad (3.5)$$

We end this part with the following proposition, which is instrumental in relating homogeneous vector bundles to the equivariance framework in Sect. 2.3.

Proposition 3.4 [39, Sect. 7.5] *Let G be a Lie group and let $H \leq G$ be a closed subgroup. Then the quotient map*

$$q : G \rightarrow G/H, \quad g \mapsto gH, \quad (3.6)$$

defines a smooth principal H -bundle over the homogeneous G -space $\mathcal{M} = G/H$.

3.2 Homogeneous vector bundles

Vector bundles may inherit global symmetry from a homogeneous base space; the transitive action $(g, x) \mapsto g \cdot x$ may induce linear maps $E_x \mapsto E_{gx}$ between fibers. Such bundles are naturally called *homogeneous* and, because this symmetry is also

encoded in its sections (data points), we will show that homogeneous vector bundles is the natural setting for studying GCNNs.

From this point on, we restrict attention to homogeneous spaces $\mathcal{M} = G/K$ where G is a unimodular Lie group and $K \leq G$ is a compact subgroup. Elements of the homogeneous space are interchangably denoted as $x \in \mathcal{M}$ or $gK \in G/K$. The origin $o = q(e) = eK \in G/K$ is the coset containing the identity element $e \in G$.

Examples of unimodular Lie groups include all compact or abelian Lie groups, all finite groups and discrete groups, the Euclidean groups, and many others [18, 19].

Definition 3.5 [42, 5.2.1] Let \mathcal{M} be a homogeneous G -space and let $\pi : E \rightarrow \mathcal{M}$ be a smooth vector bundle with fibers E_x . We say that E is *homogeneous* if there is a smooth left G -action $G \times E \rightarrow E$ satisfying

$$g \cdot E_x = E_{g \cdot x}, \quad (3.7)$$

and such that the induced map $L_{g,x} : E_x \rightarrow E_{g \cdot x}$ is linear, for all $g \in G, x \in \mathcal{M}$.

Example 3.6 The frame bundle $F\mathcal{M}$ is a homogeneous vector bundle whenever \mathcal{M} is a homogeneous space, and the same is true of any associated bundle $F\mathcal{M} \times_\rho V_\rho$. In particular, the tangent bundle $T\mathcal{M}$ is a homogeneous vector bundle. \square

Example 3.7 For any finite-dimensional K -representation (ρ, V_ρ) , the associated bundle $E_\rho = G \times_\rho V_\rho$ is a homogeneous vector bundle with respect to the left action

$$g \cdot [g', v] = [gg', v]. \quad (3.8)$$

\square

All homogeneous vector bundles E are of the form $G \times_\rho V_\rho$, up to isomorphism. To understand why, consider the fiber E_o at the origin $o = q(e) \in G/K$ and observe that the restriction of (3.7) to E_o and elements $k \in K$ yields invertible linear maps

$$L_k : E_o \rightarrow E_{k \cdot o} = E_o. \quad (3.9)$$

The defining properties of group actions ensure that $\rho(k) = L_k$ is a finite-dimensional K -representation on E_o . Moreover, because the linear maps $L_{g,x}$ are isomorphisms, any element v' of any fiber E_x can be obtained as the image $v' = L_{g,o}(v) =: L_g(v)$ for some choices of $g \in q^{-1}(\{x\})$ and $v \in E_o$. The mapping

$$\xi : G \times E_o \rightarrow E, \quad (g, v) \mapsto L_g(v), \quad (3.10)$$

is thus surjective. It is not injective, though, since the relation

$$L_g = L_g \circ L_k \circ L_{k^{-1}} = L_{gk} \circ \rho(k^{-1}), \quad (3.11)$$

implies that $\xi(g, v) = L_g(v) = L_{gk}(\rho(k^{-1})v) = \xi(gk, \rho(k^{-1})v)$ for $k \in K$. However, the same argument shows that ξ is made injective by passing to the quotient $G \times_\rho E_o$.

Lemma 3.8 [42,5.2.3] *The map*

$$G \times_{\rho} E_o \rightarrow E, \quad [g, v] \mapsto L_g(v), \quad (3.12)$$

is an isomorphism of homogeneous vector bundles.

We now have two perspectives on bundles $G \times_{\rho} V_{\rho}$: As bundles associated to the principal bundle $P = G$, and as homogeneous vector bundles (up to isomorphism). The former perspective offers a connection to the framework in Sect. 2.3, whereas the latter motivates Definition 3.14 of G -equivariant layers in Sect. 3.4.

3.3 Induced representations

Let us show the relationship between homogeneous vector bundles and induced representations, which will be an essential ingredient in the definition of G -equivariant layers. To this end, let (ρ, V_{ρ}) be a finite-dimensional unitary K -representation and consider the homogeneous vector bundle $E_{\rho} = G \times_{\rho} V_{\rho}$.

Lemma 3.9 [42,5.2.7] *The unitary structure*

$$[g, v], [g, w]_x := \langle v, w \rangle_{\rho}, \quad (x = q(g)), \quad (3.13)$$

defines a complete inner product on each fiber E_x , making E_{ρ} into a Hilbert bundle with $L_{g,x}$ unitary. This unitary structure is unique in that, if we identify V_{ρ} with E_o in the canonical manner, then the inner product on V_{ρ} so induced agrees with $\langle \cdot, \cdot \rangle_{\rho}$.

We also need the following measure on G/K :

Theorem 3.10 Quotient Integral Formula [14, Sect. 1.5] *There is a unique G -invariant, nonzero Radon measure dx on G/K such that the following quotient integral formula holds for every $f \in C_c(G)$:*

$$\int_G f(g) dg = \int_{G/K} \int_K f(xk) dk dx. \quad (3.14)$$

Using these two ingredients, we make $\Gamma_c(E_{\rho})$ into a pre-Hilbert space with respect to the inner product

$$\langle s, s' \rangle_{L^2(E_{\rho})} := \int_{G/K} \langle s(x), s'(x) \rangle_x dx, \quad s, s' \in \Gamma_c(E_{\rho}), \quad (3.15)$$

and we denote its completion $L^2(E_{\rho})$. Similarly, $C_c(G; \rho)$ is a pre-Hilbert space with respect to the inner product

$$\langle f, f' \rangle_{L^2(G; \rho)} = \int_G \langle f(g), f'(g) \rangle_{\rho} dg, \quad f, f' \in C_c(G; \rho), \quad (3.16)$$

the completion of which is denoted $L^2(G; \rho)$.

Definition 3.11 The G -representations

$$\mathrm{ind}_K^G \rho(g) : L^2(E_\rho) \rightarrow L^2(E_\rho), \quad (\mathrm{ind}_K^G \rho(g)s)(x) = g \cdot s(g^{-1}x), \quad (3.17)$$

$$\mathrm{Ind}_K^G \rho(g) : L^2(G; \rho) \rightarrow L^2(G; \rho), \quad (\mathrm{Ind}_K^G \rho(g)f)(g') = f(g^{-1}g'). \quad (3.18)$$

are called *induced representations*, or *representations induced by ρ* .

Both $\mathrm{ind}_K^G \rho$ and $\mathrm{Ind}_K^G \rho$ are unitary [42, 5.3.2] and may be identified:

Lemma 3.12 [42, 5.3.4] *The induced representations $\mathrm{ind}_K^G \rho$ and $\mathrm{Ind}_K^G \rho$ are unitarily equivalent.*

Proof First observe that the isomorphism $C_c(G; \rho) \rightarrow \Gamma_c(E_\rho)$, $f \mapsto s_f$ is unitary, which follows by combining the quotient integral formula (3.14), the unitarity of ρ , and the compactness of K : For all $f, f' \in C_c(G; \rho)$, the map $g \mapsto \langle f(g), f'(g) \rangle_\rho$ lies in $C_c(G)$ and so

$$\begin{aligned} \langle f, f' \rangle_{L^2(G; \rho)} &= \int_{G/K} \int_K \langle f(xk), f'(xk) \rangle_\rho \, dk \, dx \\ &= \int_{G/K} \langle f(x), f'(x) \rangle_\rho \, dx \\ &= \int_{G/K} \langle [x, f(x)], [x, f'(x)] \rangle_x \, dx = \langle s_f, s_{f'} \rangle_{L^2(E_\rho)}. \end{aligned} \quad (3.19)$$

The same map $f \mapsto s_f$ satisfies

$$(\mathrm{ind}_K^G(\rho)s_f)(x) = g \cdot s_f(g^{-1}x) = [x, f(g^{-1}x)] = s_{\mathrm{Ind}_K^G(\rho)f}(x), \quad (3.20)$$

so it extends to a unitary isomorphism $L^2(G; \rho) \rightarrow L^2(E_\rho)$ intertwining the induced representations. \square

To gain a better understanding of the induced representations, consider the Bochner space $L^2(G, V)$, the space of square-integrable functions $f : G \rightarrow V$ that take values in a finite-dimensional Hilbert space V . It is itself a Hilbert space with inner product

$$\langle f, f' \rangle_{L^2(G, V)} = \int_G \langle f(g), f'(g) \rangle_V \, dg'. \quad (3.21)$$

The induced representation $(\mathrm{Ind}_K^G \rho, L^2(G; \rho))$ is nothing but the restriction of the left regular representation Λ on $L^2(G, V_\rho)$ to a closed, invariant subspace. Furthermore, Λ is intimately related to the left regular representation λ on $L^2(G)$, as the following lemma shows. The proof of this lemma is a short calculation.

Lemma 3.13 *Let V be a finite-dimensional Hilbert space and equip $L^2(G) \otimes V$ with the tensor product inner product. Then the natural unitary isomorphism*

$$\begin{aligned} A : L^2(G) \otimes V &\rightarrow L^2(G, V) \\ f \otimes v &\mapsto fv \end{aligned} \quad (3.22)$$

intertwines $\lambda \otimes \text{Id}_V$ with Λ .

This lemma also shows that, if we choose an orthonormal basis $e_1, \dots, e_{\dim V} \in V$, elements of $L^2(G, V)$ are simply linear combinations $f = \sum_i f^i e_i$ with component functions $f^i \in L^2(G)$. We use this fact in some calculations of vector-valued integrals, and the component functions will also be important in Section A.

3.4 G -equivariant and convolutional layers

Given a homogeneous G -space \mathcal{M} , we observed that vector bundles $\pi : E \rightarrow \mathcal{M}$ may inherit the global symmetry of \mathcal{M} . We took a closer look at such homogeneous vector bundles and found that they are isomorphic to associated bundles $G \times_\rho V_\rho$, and therefore fit within the equivariance framework of Section 2.3. We also saw how the global symmetry of \mathcal{M} is encoded in data points and feature maps via induced representations, and we want G -equivariant layers to preserve this global symmetry.

Consider homogeneous vector bundles $E_\rho = G \times_\rho V_\rho$ and $E_\sigma = G \times_\sigma V_\sigma$, and recall Definition 2.19 of layers as general linear maps $\Phi : \Gamma_c(E_\rho) \rightarrow \Gamma_c(E_\sigma)$. We are mainly interested in *bounded* layers from an application point of view, and we can make this restriction now that the domain and codomain are normed spaces. Furthermore, any bounded linear map $\Gamma_c(E_\rho) \rightarrow \Gamma_c(E_\sigma)$ can be uniquely extended to a bounded linear map

$$\Phi : L^2(E_\rho) \rightarrow L^2(E_\sigma), \quad (3.23)$$

and we assume this extension has already been made.

Definition 3.14 A bounded linear map $\Phi : L^2(E_\rho) \rightarrow L^2(E_\sigma)$ is a G -equivariant layer if it intertwines the induced representations:

$$\Phi \circ \text{ind}_K^G \rho = \text{ind}_K^G \sigma \circ \Phi. \quad (3.24)$$

That is, G -equivariant layers are elements $\Phi \in \text{Hom}_G(L^2(E_\rho), L^2(E_\sigma))$.

Lemma 3.12 states that the induced representations $\text{ind}_K^G \rho$ and $\text{Ind}_K^G \rho$ are unitarily equivalent. It follows that any G -equivariant layer $\Phi : L^2(E_\rho) \rightarrow L^2(E_\sigma)$ induces a bounded linear map $\phi : L^2(G; \rho) \rightarrow L^2(G; \sigma)$ satisfying

$$\phi \circ \text{Ind}_K^G \rho = \text{Ind}_K^G \sigma \circ \phi, \quad (3.25)$$

and vice versa. Even though Φ acts on data points and thereby has a more geometric interpretation than ϕ , convolutional layers will be special cases of ϕ . This is because

integrals of vector-valued feature maps $f \in L^2(G; \rho)$ are easier to define than integrals of bundle-valued data points $s \in L^2(E_\rho)$. For this reason, let us extend the notion of G -equivariant layers to include both Φ and ϕ , just like we did for (gauge equivariant) layers in Sect. 2.3. Again, we view layers Φ and ϕ as two sides of the same coin.

Aside from minor technical differences, Definition 3.14 coincides with the definition of equivariant maps in [12]. We have thus motivated and defined G -equivariant layers, as used in GCNNs, almost directly from the definition of homogeneous vector bundles and a desire for layers to respect the homogeneity. For this reason, we argue that homogeneous vector bundles form the natural setting for GCNNs.

Let us now define convolutional layers.

Definition 3.15 A convolutional layer $L^2(G; \rho) \rightarrow L^2(G; \sigma)$ is a bounded operator of the form

$$[\kappa \star f](g) = \int_G \kappa(g^{-1}g')f(g') \, dg', \quad f \in L^2(G; \rho), \quad (3.26)$$

for each $g \in G$. Here, $\kappa : G \rightarrow \text{Hom}(V_\rho, V_\sigma)$ is an operator-valued kernel.

Of course, not any function $\kappa : G \rightarrow \text{Hom}(V_\rho, V_\sigma)$ can be chosen as the kernel of a convolutional layer. The kernel must ensure both that (3.26) is bounded and that $\kappa \star f \in L^2(G; \sigma)$ for each $f \in L^2(G; \rho)$. We give a sufficient condition for boundedness in Lemma 3.17 and the other requirement has been studied in detail in [12, 27].

The next result is an almost immediate consequence of the Fubini-Tonelli theorem.

Proposition 3.16 The adjoint of (3.26) is the integral operator

$$[f * \kappa^*](g) = \int_G \kappa^*(g'^{-1}g)f(g') \, dg', \quad f \in L^2(G; \sigma), \quad (3.27)$$

where κ^* is the pointwise adjoint of κ . That is, $(\kappa \star \cdot)^* = \cdot * \kappa^*$.

The next lemma investigates the boundedness of (3.26) and (3.27) in terms of the kernel matrix elements $\kappa_{ij} : G \rightarrow \mathbb{C}$, for any pair of orthonormal bases in V_ρ and V_σ .

Lemma 3.17 The operators (3.26)–(3.27) are bounded if $\kappa_{ij} \in L^1(G)$ for all i, j .

Proof We need only prove that (3.27) is bounded, its adjoint (3.26) will be bounded as well. Choose orthonormal bases $e_1, \dots, e_{\dim V_\rho} \in V_\rho$ and $\tilde{e}_1, \dots, \tilde{e}_{\dim V_\sigma} \in V_\sigma$ and observe that, because $L^2(G; \sigma) \subset L^2(G, V_\sigma)$, Lemma 3.13 enables the decomposition of $f \in L^2(G; \sigma)$ into component functions $f^i \in L^2(G)$:

$$f = \sum_{i=1}^{\dim V_\sigma} f^i \tilde{e}_i. \quad (3.28)$$

The kernel κ is similarly decomposed into matrix elements $\kappa_{ij} = \langle \tilde{e}_j, \kappa e_i \rangle_\sigma$, which satisfy $\kappa_{ji}^* = \overline{\kappa_{ij}}$. The integral (3.27) now takes the form

$$[f * \kappa^*](g) = \sum_{j=1}^{\dim V_\rho} \left(\int_G \sum_{i=1}^{\dim V_\sigma} \kappa_{ji}^*(g'^{-1}g) f^i(g') dg' \right) e_j, \quad (3.29)$$

hence

$$\|f * \kappa^*\|_{L^2(G; \rho)}^2 \leq \sum_{i,j} \int_G \left| \int_G \kappa_{ji}^*(g'^{-1}g) f^i(g') dg' \right|^2 dg = \sum_{i,j} \|f^i * \kappa_{ji}^*\|_2^2. \quad (3.30)$$

Applying Young's convolution inequality to each term yields the desired bound

$$\|f * \kappa^*\|_{L^2(G; \rho)} \leq \sum_{i,j} \|\kappa_{ij}\|_1^2 \|f^i\|_2^2 \leq M \sum_i \|f^i\|_2^2 = M \|f\|_{L^2(G; \sigma)}^2, \quad (3.31)$$

where $M = \sum_{i,j} \|\kappa_{ij}\|_1^2 < \infty$ if $\kappa_{ij} \in L^1(G)$ for all i, j . \square

We are interested in convolutional layers partly because they are concrete examples of G -equivariant layers, which we show next.

Proposition 3.18 *Convolutional layers are G -equivariant layers.*

Proof Convolutional layers $\kappa \star : L^2(G; \rho) \rightarrow L^2(G; \sigma)$ are bounded linear operators by definition, so the only thing we need to prove is that $\kappa \star$ intertwines the induced representations. This follows immediately from left-invariance of the Haar measure: For each $f \in L^2(G; \rho)$ and all $g, h \in G$,

$$\begin{aligned} [\kappa \star (\text{Ind}_K^G \rho(g)f)](h) &= \int_G \kappa(h^{-1}g') f(g^{-1}g') dg' \quad (g' \mapsto gg') \\ &= \int_G \kappa((g^{-1}h)^{-1}g') f(g') dg' = [\kappa \star f](g^{-1}h), \end{aligned} \quad (3.32)$$

hence $[\kappa \star \text{Ind}_K^G \rho(g)f] = \text{Ind}_K^G \sigma(g)[\kappa \star f]$. \square

Example 3.19 Let us show where ordinary CNNs fit in the present context. CNNs represent the case $G = \mathbb{Z}^2$ when $K = \{0\}$ is the trivial subgroup. The homogeneous space is $G/K = \mathbb{Z}^2/\{0\} = \mathbb{Z}^2$ and the quotient map $q : G \rightarrow G/K$ is thus the identity map on \mathbb{Z}^2 . Its inverse, the identity map $\omega : G/K \rightarrow G$, is a globally defined gauge that eliminates the need for gauge equivariance, as we may choose to work exclusively in this one gauge. This is just a reflection of the fact that

$$G = \mathbb{Z}^2 = \mathbb{Z}^2/\{0\} \times \{0\} = G/K \times K, \quad (3.33)$$

is (obviously) trivial as a principal bundle. Its associated bundles $E_\rho = \mathbb{Z}^2 \times_\rho V_\rho$ are also trivial: partly because the finite-dimensional K -representation σ must be trivial, and partly because each equivalence class $[g, v]$ only contains a single representative. These reasons are, of course, due to the triviality of K .

This is not to say that the equivariant framework of Sect. 2.3 is uninteresting when dealing with CNNs, or with GCNNs for other homogeneous spaces $\mathcal{M} = G/K$ with K trivial. We saw in Sects. 3.2–3.3 how the homogeneity gives rise to induced representations, which encode the global symmetry in both data points and feature maps. This is a useful perspective to have, and G -equivariant layers are interesting even when the bundles are trivial.

Triviality of the associated bundles, $E_\rho \simeq \mathbb{Z}^2 \times \mathbb{C}^m$ where $m = \dim V_\sigma$,¹ implies that data points and feature maps are general square-integrable functions,

$$L^2(E_\rho) \simeq L^2(\mathbb{Z}^2; \rho) \simeq L^2(\mathbb{Z}^2, \mathbb{C}^m), \quad (3.34)$$

and are thereby extensions of compactly supported functions $f : \mathbb{Z}^2 \rightarrow \mathbb{C}^m$. This ties well into the discussion in Sect. 2.1. Convolutional layers (3.26) reduce to bounded linear operators $L^2(\mathbb{Z}^2, \mathbb{C}^m) \rightarrow L^2(\mathbb{Z}^2, \mathbb{C}^n)$ and take the form

$$(\kappa \star f)(x) = \sum_{y \in \mathbb{Z}^2} \kappa(y - x) f(y), \quad (3.35)$$

as the Haar measure on \mathbb{Z}^2 is the counting measure. The kernel $\kappa : \mathbb{Z}^2 \rightarrow \text{Hom}(\mathbb{C}^m, \mathbb{C}^n)$ is finitely supported in practice, so boundedness of (3.35) is ensured by Lemma 3.17.

Interestingly, all \mathbb{Z}^2 -equivariant layers are convolutional layers; there are no other types of \mathbb{Z}^2 -equivariant layers than (3.35). This follows from Theorem 3.22 and is proven in Corollary 3.23 below. \square

For more general groups G , it is no longer true that all G -equivariant layers are convolutional layers; we give an example of this fact in Example 3.21. Implementations of GCNNs, however, are usually based on convolutional layers, or on analogous layers in the Fourier domain. What consequences does the restriction to convolutional layers have for the expressivity of GCNNs? Can we tell whether a given G -equivariant layer is expressible as a convolutional layer? The answer to this last question, it turns out, requires the following notion of reproducing kernel Hilbert spaces.

Definition 3.20 Let G be a group, let V be a finite-dimensional normed vector space, and let \mathcal{H} be a Hilbert space of functions $G \rightarrow V$. Then \mathcal{H} is a *reproducing kernel Hilbert space (RKHS)* if the evaluation operator

$$\mathcal{E}_g : \mathcal{H} \rightarrow V, \quad f \mapsto f(g), \quad (3.36)$$

is bounded for all $g \in G$. Moreover, by *left-invariant RKH subspace* $\mathcal{H} \subseteq L^2(G, V)$ we mean a closed subspace that is both a RKHS and an invariant subspace for the left regular representation Λ on $L^2(G, V)$.

¹ Recall that we focus on complex vector bundles, hence the use of \mathbb{C}^m instead of \mathbb{R}^m .

The name RKHS is due to the existence of a kernel-type function that reproduces all elements of \mathcal{H} . To see how, choose an orthonormal basis $e_1, \dots, e_{\dim V} \in V$ and write elements $v \in V$ as linear combinations $v = \sum_i v^i e_i$. The projection $P_i(v) = v^i$ onto the i 'th component is always continuous, so the composition $\mathcal{E}_{g,i} := P_i \circ \mathcal{E}_g$ is a continuous linear functional

$$\mathcal{E}_{g,i} : \mathcal{H} \rightarrow \mathbb{C}, \quad f \mapsto f^i(g), \quad (3.37)$$

for all $g \in G$ and $i = 1, \dots, \dim V$. By the Riesz representation theorem, there are elements $\varphi_{g,i} \in \mathcal{H}$ such that $f^i(g) = \mathcal{E}_{g,i}(f) = \langle f, \varphi_{g,i} \rangle$, hence

$$f(g) = \sum_{i=1}^{\dim V} \langle f, \varphi_{g,i} \rangle e_i. \quad (3.38)$$

Now, if $\mathcal{H} \subseteq L^2(G, V)$ is a left-invariant RKH subspace, expanding both $f = \sum_j f^j e_j$ and $\varphi_{g,i} = \sum_j \varphi_{g,i}^j e_j$ in the orthonormal basis in V yields the formula

$$f(g) = \sum_i \left\langle \sum_j f^j e_j, \sum_k \varphi_{g,i}^k e_k \right\rangle e_i \quad (3.39)$$

$$= \sum_i \left(\sum_j \int_G \overline{\varphi_{g,i}^j(g')} f^j(g') dg' \right) e_i \quad (3.40)$$

$$= \int_G \left(\sum_{i,j} \overline{\varphi_{g,i}^j(g')} f^j(g') e_i \right) dg' = \int_G \varphi_g^*(g') f(g') dg', \quad (3.41)$$

where φ_g^* is the conjugate transpose of the matrix $(\varphi_g)_i^j = \varphi_{g,i}^j$. By left-invariance,

$$f(g) = (\Lambda(g^{-1})f)(e) = \int_G \varphi_e^*(g^{-1}g') f(g') dg', \quad (3.42)$$

and so $f \in \mathcal{H}$ is *reproduced* by the operator-valued kernel $\varphi_e^* : G \rightarrow \text{Hom}(V)$. We rename φ_e^* as κ to emphasize the similarity between (3.42) and convolutional layers. The reproducing kernel κ is unique and thus independent of the choice of basis in V , which follows from uniqueness in the Riesz representation theorem.

It is now clear why left-invariant RKH subspaces of $L^2(G, V)$ are relevant when discussing convolutional layers, as the latter are given by integral operators similar to (3.42). In order to show that an abstract G -equivariant layer $\phi : L^2(G; \rho) \rightarrow L^2(G; \sigma)$ can be written as a convolutional layer, it is almost necessary for it to act in a RKHS:

Example 3.21 The identity map $\phi : L^2(G; \sigma) \rightarrow L^2(G; \sigma)$ is clearly a G -equivariant layer regardless of G, K, σ , but it is only a convolutional layer if $L^2(G; \sigma)$ is a RKHS.

This is because when ϕ is the identity, (3.26) becomes the reproducing property

$$f(g) = \int_G \kappa(g^{-1}g') f(g') \, dg', \quad f \in L^2(G; \sigma). \quad (3.43)$$

It follows that not every G -equivariant layer is a convolutional layer, because $L^2(G; \sigma)$ is not always a RKHS. When σ is the trivial representation, for instance, $L^2(G; \sigma)$ reduces to $L^2(G)$ which is not a RKHS when G is nondiscrete [19, Theorem 2.42]. \square

At this point, we know that global symmetry manifests itself in feature maps and data points through the induced representation, and we used this knowledge to define G -equivariant layers. We also defined convolutional layers and showed that these are special cases of G -equivariant layers, but the converse problem is much more subtle: When can a G -equivariant layer be expressed as a convolutional layer? The answer, as we have just seen, is directly related to the concept of RKHS.

The next result is our main theorem. It extends [12, Theorem 6.1] to layers which are not a priori given by integral operators, and it generalizes [25, Theorem 1] beyond compact groups as discussed below.

Theorem 3.22 *Let G be a unimodular Lie group, let $K \leq G$ be a compact subgroup, and consider homogeneous vector bundles E_ρ, E_σ over $\mathcal{M} = G/K$. Suppose that*

$$\phi : L^2(G; \rho) \rightarrow L^2(G; \sigma), \quad (3.44)$$

is a G -equivariant layer. If ϕ maps into a left-invariant RKH subspace $\mathcal{H} \subseteq L^2(G; \sigma)$, then ϕ is a convolutional layer.

Proof Fix orthonormal bases in V_ρ, V_σ . For $i = 1, \dots, \dim \sigma$, consider the functionals

$$\mathcal{E}^i : L^2(G; \rho) \rightarrow \mathbb{C}, \quad \mathcal{E}^i(f) = (\phi f)^i(e), \quad (3.45)$$

composing ϕ with evaluation at the identity element $e \in G$ and projection onto the i 'th component. As ϕ maps into a left-invariant RKH subspace $\mathcal{H} \subseteq L^2(G; \sigma)$, (3.45) is a bounded linear functional: $|\mathcal{E}^i(f)| \leq \|(\phi f)(e)\|_\sigma \leq \|\phi f\|_{L^2(G; \sigma)} \leq \|\phi\| \|f\|_{L^2(G; \rho)}$. By the Riesz representation theorem, there is a unique $\varphi_i \in L^2(G; \rho)$ such that

$$\mathcal{E}^i(f) = \int_G \langle f(g), \varphi_i(g) \rangle_\rho \, dg = \int_G \sum_{j=1}^{\dim \rho} f^j(g) \overline{\varphi_i^j(g)} \, dg, \quad (3.46)$$

and proceeding as in (3.39)–(3.42) with $\kappa := \varphi_e^*$ yields the desired relation

$$(\phi f)(g) = \int_G \kappa(g^{-1}g') f(g') \, dg'. \quad (3.47)$$

\square

The implications of Theorem 3.22 are not immediately obvious. It is not yet clear whether the left-invariant RKH subspace criterion is satisfied for most G -equivariant layers and how it depends on the group G . Such questions are discussed in Appendix A, where we discuss the concept of *bandwidth* and its close connection to the theory of RKHS. This discussion will give a Fourier analytic perspective on Theorem 3.22 and helps us draw stronger conclusions about G -equivariant layers for certain groups. The following corollary of Theorem 3.22 is proven at the end of Appendix A.

Corollary 3.23 *If G is a discrete abelian or finite group, then any G -equivariant layer is a convolutional layer.*

In particular, setting $G = \mathbb{Z}^2$ shows that convolutional layers are the only possible translation equivariant layers in the ordinary CNN setting.

The next corollary could be compared to the generalized convolutions described in [25, Sect. 4.1] and may simplify numerical computations of convolutional layers. Integrals over the homogeneous space G/K are often faster to compute than integrals over the (larger) group G . For example, rotationally equivariant networks on spherical data ($G = SO(3)$, $\mathcal{M} = S^2$) compute convolutions over $SO(3)$, which are computationally more expensive than convolutions over the sphere S^2 . The expensive computations of convolutional layers is one of the main drawbacks of GCNNs.

Corollary 3.24 *Let G , K , and $\phi : L^2(G; \rho) \rightarrow L^2(G; \sigma)$ be as in Theorem 3.22 and let κ be the kernel of the resulting convolutional layer (3.47). Then*

$$(\phi f)(g) = \int_{G/K} \kappa(g^{-1}x) f(x) dx. \quad (3.48)$$

Proof In the proof of Theorem 3.22, we constructed the kernel from the components of $\varphi_i \in L^2(G; \rho)$, and unitarity of ρ clearly implies that the expression $\langle f(x), \varphi_i(x) \rangle_\rho$ is well-defined. We may therefore use the unitary structure (3.13) to get the following relation for all component functions $(\phi f)^i$ and all $g \in G$:

$$(\phi f)^i(g) = \langle f, \text{Ind}_K^G \rho(g) \varphi_i \rangle_{L^2(G; \rho)} = \langle s_f, s_{\text{Ind}_K^G \rho(g) \varphi_i} \rangle_{L^2(E_\rho)} \quad (3.49)$$

$$= \int_{G/K} \langle s_f(x), s_{\text{Ind}_K^G \rho(g) \varphi_i}(x) \rangle_x dx \quad (3.50)$$

$$= \int_{G/K} \langle f(x), \text{Ind}_K^G \rho(g) \varphi_i(x) \rangle_\rho dx \quad (3.51)$$

$$= \int_{G/K} \sum_{j=1}^{\dim \rho} \overline{\varphi_i^j(g^{-1}x)} f^j(x) dx. \quad (3.52)$$

We now obtain (3.48) by reconstructing κ from its components $\kappa_{ij} = \overline{\varphi_i^j}$. □

3.4.1 Relation to previous work

Theorem 3.22 is similar to [25, Theorem 1], even though there are also clear differences. For instance, we study individual G -equivariant layers ϕ whereas [25] is concerned

with G -equivariant multi-layer neural networks

$$\sigma_L \circ \phi_L \circ \sigma_{L-1} \circ \phi_{L-1} \circ \cdots \circ \sigma_1 \circ \phi_1, \quad (3.53)$$

where ϕ_ℓ are linear transformations and where σ_ℓ are non-linear activation functions for $\ell = 1, \dots, L$. That being said, most results on G -equivariance in [25] are actually proven for the linear transformations ϕ_ℓ . The pointwise non-linearities σ_ℓ are applied in a G -equivariant way [10, Sect. 6.2], so any result concerning G -equivariance of the linear transformations can be extended to multi-layer networks (3.53) by simple induction on the number of layers.

Another difference is that we consider unimodular Lie groups G while [25] restricts attention to compact groups G . In fact, in the case of a single-layer neural network with index set $\mathcal{X}_0 = \mathcal{X}_1 = G/K$, [25, Theorem 1] states that all G -equivariant layers are convolutional layers. This unconditional result is stronger than Theorem 3.22 for compact groups. However, the proof of [25, Theorem 1] contains a minor flaw. Their proof uses the inverse Fourier transform to construct a convolution kernel κ for each G -equivariant linear layer ϕ , so that ϕ becomes a convolutional layer $\phi f = \kappa \star f$ (in our notation). To be more precise, their proof makes use of a linear relation

$$\widehat{\phi f}(\gamma_i) = \Phi_i \hat{f}(\gamma_i), \quad i = 1, 2, 3, \dots \quad (3.54)$$

between the Fourier coefficients of ϕf and f , for each feature map f and each unitary irreducible representation γ_i of G . Matrices B_i corresponding to the linear maps Φ_i are interpreted as the Fourier coefficients $\hat{\kappa}(\gamma_i)$ of a function κ that is then constructed using the inverse Fourier transform,

$$\kappa(g) = \sum_i \dim(\gamma_i) \operatorname{tr}[B_i \gamma_i(g)]. \quad (3.55)$$

Finally, the convolution theorem $\widehat{\kappa \star f} = \hat{\kappa} \hat{f} = \widehat{\phi f}$ ensures that $\kappa \star f = \phi f$.

When ϕ is the identity map on $L^2(G)$, however, the matrices $B_i \in \mathbb{C}^{\dim(\gamma_i) \times \dim(\gamma_i)}$ become identity matrices and the Parseval equation [18, p. 145] implies that

$$\|\kappa\|_2^2 = \sum_i \dim(\gamma_i) \operatorname{tr}[B_i^* B_i] = \sum_i \dim(\gamma_i)^2. \quad (3.56)$$

Thus, the matrices B_i are only the Fourier coefficients of an element $\kappa \in L^2(G)$ when the compact group G is finite. Otherwise, the right-hand side in (3.56) diverges and the inverse Fourier transform (3.55) cannot be used to construct κ . This is consistent with Example 3.21, which shows that the identity map on $L^2(G)$ is a G -equivariant layer but not a convolutional layer when the compact group G is not finite.

4 Discussion

In this paper, we have investigated the mathematical foundations of G -equivariant convolutional neural networks (GCNNs), which are designed for deep learning tasks exhibiting global symmetry. We presented a basic framework for equivariant neural networks that include both gauge equivariant neural networks and GCNNs as special cases. We also demonstrated how GCNNs can be obtained from homogeneous vector bundles, when G is a unimodular Lie group and $K \leq G$ is a compact subgroup.

In Theorem 3.22, we gave a precise criterion for when a given G -equivariant layer is, in fact, a convolutional layer. This criterion uses reproducing kernel Hilbert spaces (RKHS) and cannot be circumvented, as demonstrated in Example 3.21. When the group G is either discrete abelian or finite, we showed that the criterion in our main theorem is automatically satisfied. In other words, there are no G -equivariant layers other than convolutional layers for these groups. This implies that implementations of GCNNs using convolutional layers are maximally expressive - at least with regards to the linear transformations, since we did not discuss non-linear activation functions. Finally, we proved that convolutional layers can be computed by integrating over the homogeneous space G/K rather than integrating over the group G . Convolutional layers are expensive to compute, especially for large groups G . This result has the potential of speeding up numerical computations of convolutional layers and making GCNNs even more viable in practical applications.

One limitation of the current paper, compared to [12, 25], is that the homogeneous space G/K does not change between layers. This restriction was made in order to limit the scope of our analysis and may be relaxed in future work.

Appendix A. RKHS and bandlimited functions

The significance of our main theorem 3.22 depends on how common left-invariant RKH subspaces of $L^2(G; \sigma)$ are. It turns out that there is a deep connection between RKHS and the concept of bandwidth in Fourier analysis which we will make use of. To our knowledge, this connection has not been noted in previous works on GCNNs, so we hope that this appendix may aid future research.

Left-invariant RKH subspaces $\mathcal{H} \subseteq L^2(G)$ have been fully characterized for groups of type I [8, 19], which essentially means groups that have manageable representation theory. Most unimodular Lie groups are already of type I, so adding this property to our list of assumptions about G barely reduces the generality of our results. Using known results about the scalar-valued case $\mathcal{H} \subseteq L^2(G)$ will also provide information about the vector-valued case. This is due to the unitary equivalence in Lemma 3.13,

$$\begin{aligned} A : L^2(G) \otimes V &\rightarrow L^2(G, V) \\ f \otimes v &\mapsto fv \end{aligned}, \quad (\text{A.1})$$

which shows that $A(\mathcal{H} \otimes V_\sigma)$ and its closed subspace

$$A(\mathcal{H} \otimes V_\sigma) \cap L^2(G; \sigma), \quad (\text{A.2})$$

are left-invariant RKH subspaces of $L^2(G, V_\sigma)$ and $L^2(G; \sigma)$, respectively, whenever \mathcal{H} is a left-invariant RKH subspace of $L^2(G)$.

Recall the construction of the reproducing kernel $\kappa : G \rightarrow \text{Hom}(V)$ for arbitrary left-invariant RKH subspaces $\mathcal{H} \subseteq L^2(G, V)$, below Definition 3.20. After choosing an orthogonal basis in V and fixing $g \in G$, we used the Riesz representation theorem to find functions $\varphi_{g,i} \in \mathcal{H}$ such that

$$f(g) = \sum_{i=1}^{\dim V} \langle f, \varphi_{g,i} \rangle e_i, \quad f \in \mathcal{H}. \quad (\text{A.3})$$

We then applied left-invariance to see that $f(g) = (\Lambda(g^{-1})f)(e)$, which enabled us to construct κ as a matrix-valued function, in the given basis, from the components of $\varphi_{e,i} \in \mathcal{H}$. Left-invariant RKH subspaces $\mathcal{H} \subseteq L^2(G)$ correspond to the case $V = \mathbb{C}$, so the same construction yields a single function $\varphi_e \in \mathcal{H}$ used to construct the kernel:

$$f(g) = \langle \lambda(g^{-1})f, \varphi_e \rangle = \int_G \overline{\varphi_e(g^{-1}g')} f(g') \, dg' = (f * \varphi_e^*)(g), \quad (\text{A.4})$$

with $\varphi_e^*(g) = \overline{\varphi_e(g^{-1})}$ denoting involution. So the reproducing kernel $\kappa = \varphi_e$ is not only a complex-valued function in this case, it is even an element of \mathcal{H} .

Proposition A.1 [19, Proposition 2.40] *Let $\mathcal{H} \subseteq L^2(G)$ be a left-invariant RKH subspace. The kernel $\kappa \in \mathcal{H}$ is then a self-adjoint convolution idempotent,² and*

$$\mathcal{H} = L^2(G) * \kappa = \left\{ f * \kappa \mid f \in L^2(G) \right\} \subset C(G). \quad (\text{A.5})$$

*Conversely, if $\kappa \in L^2(G)$ is a self-adjoint convolution idempotent, then $\mathcal{H} = L^2(G) * \kappa$ is a left-invariant RKH subspace of $L^2(G)$.*

Example A.2 Consider the real line $G = \mathbb{R}$ and let $\mathcal{H} \subseteq L^2(\mathbb{R})$ be a left-invariant RKH subspace with kernel $\kappa \in \mathcal{H}$. Then (A.4) becomes

$$f(x) = \int_{-\infty}^{\infty} \overline{\kappa(y-x)} f(y) \, dy = \langle f, \lambda(x)\kappa \rangle = (f * \kappa^*)(x), \quad (\text{A.6})$$

which implies that f is continuous, since the regular representation λ is continuous, hence $\mathcal{H} \subset C(\mathbb{R})$. Furthermore, setting $f = \kappa$ shows that the kernel is a self-adjoint convolution idempotent:

$$\kappa = \kappa * \kappa^* = (\kappa * \kappa^*)^* = \kappa^*. \quad (\text{A.7})$$

Combining the Plancherel transform on $L^2(\mathbb{R})$ (see Theorem A.4 and equation (A.23)) with the convolution theorem in Fourier analysis, we observe that, for all $f \in \mathcal{H}$,

$$\widehat{f} = \widehat{f * \kappa} = \widehat{f} \widehat{\kappa}. \quad (\text{A.8})$$

² That is, $\kappa = \kappa * \kappa^* = \kappa^*$.

In particular, $\hat{\kappa} = \hat{\kappa}^2$, and so $\hat{\kappa}$ is the characteristic function 1_E on a subset $E \subset \mathbb{R}$. The set E has finite Lebesgue measure according to the Plancherel theorem,

$$\text{vol}(E) = \|1_E\|_2^2 = \|\kappa\|_2^2 < \infty, \quad (\text{A.9})$$

and inserting $\hat{\kappa} = 1_E$ in (A.8) implies that $\text{supp}(\hat{f}) \subset E$. It follows that \mathcal{H} is a space of bandlimited functions. \square

This example illustrates that any measurable subset $E \subset \mathbb{R}$ with finite measure induces a left-invariant closed RKH subspace

$$\mathcal{H}_E = \left\{ f \in L^2(\mathbb{R}) \mid \text{supp}(\hat{f}) \subset E \right\} = L^2(\mathbb{R}) * \kappa_E, \quad (\text{A.10})$$

κ_E being the inverse Plancherel transform of 1_E [19, 2.63–2.65]. This relation between left-invariant RKH subspaces $\mathcal{H} \subseteq L^2(G)$ and bandlimited functions generalizes to unimodular Lie groups G of type I. Going into detail on this rather technical subject would distract from the topic at hand, however, so we summarize the general theory and refer curious readers to the relevant literature [19]. We will then restrict attention to two important cases where we can be more explicit: Abelian and compact groups.

First recall that two representations (γ, V_γ) and (η, V_η) of G are *equivalent* if there exists a linear isomorphism $\psi : V_\gamma \rightarrow V_\eta$ such that $\psi \circ \gamma = \eta \circ \psi$. Equivalence of representations is an equivalence relation, and the equivalence classes of unitary irreducible representations (γ, V_γ) form the *unitary dual* \hat{G} of G :

$$\hat{G} = \{ \text{equivalence classes } [\gamma] \text{ of irreducible unitary representations } (\gamma, V_\gamma) \}. \quad (\text{A.11})$$

The unitary dual \hat{G} is a topological space whose topology is, in a sense, well-behaved precisely when G is of type I [18, Theorem 7.6]. This is what we referred to earlier when saying that type I groups have manageable representation theory. Our interest in the unitary dual is due to the following extension of the Fourier transform:

Definition A.3 [19, Sect. 3.5] Let G be a unimodular Lie group of type I. Then the *operator-valued Fourier transform* on G maps each $f \in L^1(G)$ to the family

$$\mathcal{F}(f) = (\hat{f}(\gamma))_{\gamma \in \hat{G}}, \quad (\text{A.12})$$

where each $\hat{f}(\gamma) \in \mathcal{B}(V_\gamma)$ is a bounded operator on V_γ given by the Bochner integral

$$\hat{f}(\gamma) = \int_G f(g) \gamma(g) \, \text{d}g. \quad (\text{A.13})$$

This operator-valued Fourier transform can be extended to a Plancherel transform on $L^2(G)$ that decomposes both the left regular representation λ and its commutant

$$\lambda(G)' = \left\{ T \in \mathcal{B}(L^2(G)) \mid T \circ \lambda(g) = \lambda(g) \circ T \text{ for all } g \in G \right\}, \quad (\text{A.14})$$

into *direct integrals* \int_I^\oplus of Hilbert spaces. Direct integrals generalize direct sums \bigoplus_I of Hilbert spaces to allow nondiscrete index sets I . See [19, §3.3] for an introduction, though detailed knowledge of direct integrals is not necessary for reading this paper. For example, the Plancherel transform (A.15) below decomposes a function $f \in L^2(G)$ into Fourier coefficients $\hat{f}(\gamma)$ for $[\gamma] \in \widehat{G}$, and one may simply view the direct integral notation as a reminder that \widehat{G} need not be discrete.

Theorem A.4 [19, Theorem 3.48] *There is a canonical Plancherel measure ν for the unitary dual \widehat{G} with the following properties:*

(i) \mathcal{F} extends to a unitary operator

$$\mathcal{P} : L^2(G) \rightarrow \int_{\widehat{G}}^\oplus V_\gamma \otimes V_\gamma \, d\nu([\gamma]), \quad (\text{A.15})$$

called the Plancherel transform of G .

(ii) \mathcal{P} implements the following unitary equivalences:

$$\lambda \simeq \int_{\widehat{G}}^\oplus \gamma \otimes \text{Id} \, d\nu([\gamma]), \quad (\text{A.16})$$

$$\lambda(G)' \simeq \int_{\widehat{G}}^\oplus \text{Id} \otimes \mathcal{B}(V_\gamma) \, d\nu([\gamma]). \quad (\text{A.17})$$

Observe that if $\mathcal{H} \subseteq L^2(G)$ is a left-invariant closed subspace, then the projection $P : L^2(G) \rightarrow \mathcal{H}$ commutes with the left regular representation and is thus an element of the commutant $\lambda(G)'$. It therefore has a direct integral decomposition

$$P = \int_{\widehat{G}}^\oplus \text{Id} \otimes \hat{P}_\gamma \, d\nu([\gamma]), \quad (\text{A.18})$$

where $\hat{P}_\gamma \in \mathcal{B}(V_\gamma)$ for each $[\gamma] \in \widehat{G}$.

Theorem A.5 [19, Theorem 4.22, Proposition 2.40] *Suppose that $\mathcal{H} \subseteq L^2(G)$ is a left-invariant closed subspace with equation (A.18) denoting the projection onto \mathcal{H} . Then \mathcal{H} is a RKHS iff*

$$\int_{\widehat{G}} \text{rank}(\hat{P}_\gamma) \, d\nu(\gamma) < \infty. \quad (\text{A.19})$$

We interpret this theorem as a bandwidth restriction, similar to Example A.2. The integrand in (A.19) is an integer-valued function on \widehat{G} , so the integral is finite only if the projection (A.18) is supported on a set $E \subseteq \widehat{G}$ of finite Plancherel measure,

$$\nu\left(\underbrace{\left\{\gamma \in \widehat{G} : \hat{P}_\gamma \neq 0\right\}}_E\right) = \int_{\widehat{G}} 1_E(\gamma) \, d\nu(\gamma) \leq \int_{\widehat{G}} \text{rank}(\hat{P}_\gamma) \, d\nu(\gamma). \quad (\text{A.20})$$

That is, the left-invariant RKH subspaces $\mathcal{H} \subseteq L^2(G)$ are precisely those subspaces whose elements are bandlimited on a set $E \subseteq \widehat{G}$ in the sense that, for all $f \in \mathcal{H}$ and each equivalence class $\gamma \notin E$,

$$\widehat{f}(\gamma) = \widehat{Pf}(\gamma) = \widehat{f}(\gamma) \circ \widehat{P}_\gamma = 0, \quad (\text{A.21})$$

the second equality being [19, Corollary 4.17]. Bandlimited functions are thus central to the theory of RKHS and, by extension, to the mathematical theory of GCNNs.

Abelian groups The irreducible representations $[\gamma] \in \widehat{G}$ of any abelian group G are 1-dimensional and may be identified with their character $\chi_\gamma = \text{tr} \gamma$. There are several useful consequences of this fact.

First, the unitary dual \widehat{G} is now the set of continuous homomorphisms $\chi : G \rightarrow \mathbb{T}$, where \mathbb{T} is the circle group. This is a locally compact group with respect to pointwise multiplication and, as $[\gamma] \in \widehat{G}$ is unitary, we may write $\chi_\gamma = e^{i\xi_\gamma}$ where $\xi_\gamma : G \rightarrow \mathbb{R}$. The Fourier transform then takes the more familiar form

$$\widehat{f}(\gamma) = \int_G f(g) e^{-i\xi_\gamma(g)} \, dg, \quad (\text{A.22})$$

for $f \in L^1(G) \cap L^2(G)$. Moreover, the Haar measure on \widehat{G} can be made to coincide with the Plancherel measure such that (A.15) becomes a unitary equivalence

$$\mathcal{P} : L^2(G) \rightarrow L^2(\widehat{G}). \quad (\text{A.23})$$

Another consequence of the fact that irreducible representations are 1-dimensional, is that the integrand in (A.19) takes values in $\{0, 1\}$ and (A.20) becomes an equality. Using the same arguments as in Example A.2, we see that the left-invariant RKH subspaces $\mathcal{H} \subseteq L^2(G)$ are precisely the spaces of bandlimited functions, $\text{supp}(\widehat{f}) \subset E$, for subsets $E \subset \widehat{G}$ of finite Haar/Plancherel measure. Also, the kernel $\kappa_E \in \mathcal{H}$ is the inverse Plancherel transform of the characteristic function 1_E .

Compact groups The unitary dual \widehat{G} of any compact group G is a discrete space, and its Plancherel measure reduces to the counting measure. Furthermore, all irreducible representations of G are finite-dimensional. The integral (A.19) is thus a discrete sum with finite summands, and converges iff $\widehat{P}_\gamma = 0$ for all but finitely many $[\gamma] \in \widehat{G}$.

We are now ready to prove Corollary 3.23, with which we conclude this appendix.

Proof Let us first consider finite groups G . The unitary dual \widehat{G} is then finite as well [18, Proposition 5.27], which means that the integral (A.19) reduces to a finite sum for all closed, left-invariant subspaces $\mathcal{H} \subseteq L^2(G)$. In particular, $\mathcal{H} = L^2(G)$ is itself a RKHS. The same is thus true of $L^2(G) \otimes V \simeq L^2(G, V)$ and of any closed subspace $L^2(G; \sigma)$. Consequently, setting $\mathcal{H} = L^2(G; \sigma)$ in Theorem 3.22 implies that any G -equivariant layer

$$\phi : L^2(G; \rho) \rightarrow L^2(G; \sigma), \quad (\text{A.24})$$

is a convolutional layer.

When G is discrete abelian, the unitary dual \widehat{G} is compact [14, Proposition 3.1.5]. Since the integrand in (A.19) is at most 1, the integral is bounded by the finite volume of \widehat{G} and therefore converges for all closed, left-invariant subspaces $\mathcal{H} \subseteq L^2(G)$. This proves that $\mathcal{H} = L^2(G)$ is a RKHS when G is discrete abelian, and proceeding as in the finite case shows that any G -equivariant layer is a convolutional layer. \square

Acknowledgements I am grateful for the support from my research group: Oscar Carlsson, Jan Gerken, Hampus Linander, Fredrik Ohlsson, Christoffer Petersson, and last but not least, my advisor Daniel Persson. Thanks also to David Müller for the interesting discussions on gauge equivariance in physics, and to the anonymous reviewers for carefully reading my work and providing detailed, helpful comments. This work was supported by the Wallenberg AI, Autonomous Systems and Software Program (WASP) funded by the Knut and Alice Wallenberg Foundation.

Funding Open access funding provided by Chalmers University of Technology.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Albawi, S., Mohammed, T.A., Al-Zawi, S.: Understanding of a convolutional neural network. In: International Conference on Engineering and Technology (ICET). IEEE, 2017
2. Alpaydin, E.: Introduction to machine learning. MIT Press (2014)
3. Boyda, D.L., Chernodub, M.N., Gerasimeniuk, N.V., Goy, V.A., Liubimov, S.D., Molochkov, A.V.: Finding the deconfinement temperature in lattice yang-mills theories from outside the scaling window with machine learning. *Phys. Rev. D* (2021)
4. Brock, A., De, S., Smith, S.L., Simonyan, K.: High-performance large-scale image recognition without normalization. pages 1059–1071 (2021)
5. Bronstein, M.M., Bruna, J., LeCun, Y., Szlam, A., Vandergheynst, P.: Geometric deep learning: going beyond euclidean data. *IEEE Signal Processing Magazine* (2017)
6. Bronstein, M.M., Bruna, J., Cohen, T., Veličković, P.: Geometric deep learning: Grids, groups, graphs, geodesics, and gauges (2021). arXiv preprint [arXiv:2104.13478](https://arxiv.org/abs/2104.13478)
7. Cao, W., Yan, Z., He, Z., He, Z.: A comprehensive survey on geometric deep learning. *IEEE Access* (2020)
8. Carey, A.L.: Group representations in reproducing kernel hilbert spaces. *Rep. Math. Phys.* (1978)
9. Cheng, M., Anagiannis, V., Weiler, M., de Haan, P., Cohen, T., Welling, M.: Covariance in physics and convolutional neural networks (2019). arXiv preprint [arXiv:1906.02481](https://arxiv.org/abs/1906.02481)
10. Cohen, T., Welling, M.: Group equivariant convolutional networks. In: International conference on machine learning, PMLR (2016)
11. Cohen, T., Geiger, M., Köhler, J., Welling, M.: Spherical CNNs. In: International Conference on Learning Representations (2018)
12. Cohen, T., Geiger, M., Weiler, M.: A general theory of equivariant cnns on homogeneous spaces. *Adv. Neural Inform. Process. Syst.* **32** (2019)
13. Cohen, T., Weiler, M., Kicanaoglu, B., Welling, M.: Gauge equivariant convolutional networks and the icosahedral cnn. In International Conference on Machine Learning. PMLR (2019)
14. Deitmar, A., Echterhoff, S.: Principles of harmonic analysis. Springer International Publishing (2014)

15. Esteves, C., Allen-Blanchette, C., Makadia, A., Daniilidis, K.: Learning $so(3)$ equivariant representations with spherical cnns. In: *Proceedings of the European Conference on Computer Vision (ECCV)* (2018)
16. Esteves, C.: Theoretical aspects of group equivariant neural networks (2020). [arXiv preprint arXiv:2004.05154](https://arxiv.org/abs/2004.05154)
17. Favoni, M., Ipp, A., Müller, D., Schuh, D.: Lattice gauge equivariant convolutional neural networks. *Phys. Rev. Lett.* 128 (2022)
18. Folland, G.B.: *A course in abstract harmonic analysis*. CRC Press (2016)
19. Führ, H.: *Abstract harmonic analysis of continuous wavelet transforms*. Springer Science & Business Media (2005)
20. Fukushima, K., Miyake, S.: *Neocognitron: a self-organizing neural network model for a mechanism of visual pattern recognition*. In: *Competition and cooperation in neural nets*. Springer (1982)
21. Goodfellow, I., Bengio, Y., Courville, A.: *Deep learning*. MIT Press (2016)
22. Jia, C., Yang, Y., Xia, Y., Chen, Y.-T., Parekh, Z., Pham, H., Le, Q.V., Sung, Y., Li, Z., Duerig, T.: Scaling up visual and vision-language representation learning with noisy text supervision. In: *International Conference on Machine Learning*. PMLR (2021)
23. Kolár, I., Michor, P.W., Slovák, J.: *Natural operations in differential geometry*. Springer Science & Business Media (2013)
24. Koller, O., Zargaran, S., Ney, H., Bowden, R.: Deep sign: enabling robust statistical continuous sign language recognition via hybrid cnn-hmms. *Int. J. Comput. Vis.* (2018)
25. Kondor, R., Trivedi, S.: On the generalization of equivariance and convolution in neural networks to the action of compact groups. In: *International Conference on Machine Learning*. PMLR (2018)
26. Krizhevsky, A., Sutskever, I., Hinton, G. E.: ImageNet classification with deep convolutional neural networks. *Adv. Neural Inform. Process. Syst.* (2012)
27. Lang, L., Weiler, M.: A wigner-eckart theorem for group equivariant convolution kernels. In: *International Conference on Learning Representations* (2021)
28. Lawrence, S., Giles, C.L., Tsoi, A.C., Back, A.D.: Face recognition: a convolutional neural-network approach. *IEEE Tran. Neural Netw.* (1997)
29. LeCun, Y., Jackel, L.D., Bottou, L., Cortes, C., Denker, J.S., Drucker, H., Guyon, I., Muller, U.A., Sackinger, E., Simard, P., et al.: Learning algorithms for classification: a comparison on handwritten digit recognition. *Neural Netw. Stat. Mech. Perspect.* (1995)
30. Lee, J.M.: *Introduction to smooth manifolds*. Springer (2013)
31. Li, J., Bi, Y., Lee, G.H.: Discrete rotation equivariance for point cloud recognition. In: *2019 International Conference on Robotics and Automation (ICRA)*. IEEE (2019)
32. Luo, D., Carleo, G., Clark, B., Stokes, J.: Gauge equivariant neural networks for quantum lattice gauge theories. *Bull. Am. Phys. Soc.* (2021)
33. Mehta, D., Sridhar, S., Sotnychenko, O., Rhodin, H., Shafiei, M., Seidel, H.-P., Xu, W., Casas, D., Theobalt, C.: Vnect: Real-time 3d human pose estimation with a single rgb camera. *ACM Trans. Graph. (TOG)* (2017)
34. Monti, F., Boscaini, D., Masci, J., Rodola, E., Svoboda, J., Bronstein, M.M.: Geometric deep learning on graphs and manifolds using mixture model cnns. In: *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017)
35. Müller, D.: *Yang-mills theory, lattice gauge theory and simulations* (2019)
36. Nakahara, M.: *Geometry, topology and physics*. CRC press (2003)
37. Peng, C., Zhang, X., Yu, G., Luo, G., Sun, J.: Large kernel matters—improve semantic segmentation by global convolutional network. In: *Proceedings of the IEEE conference on computer vision and pattern recognition* (2017)
38. Song, S., Huang, H., Ruan, T.: Abstractive text summarization using lstm-cnn based deep learning. *Multimed. Tools Appl.* (2019)
39. Steenrod, N.: *The topology of fibre bundles*. Princeton University Press (1960)
40. Tan, M., Le, Q.V.: Efficientnetv2: Smaller models and faster training (2021)
41. Veeling, B.S., Linmans, J., Winkens, J., Cohen, T., Welling, M.: Rotation equivariant cnns for digital pathology. In: *International Conference on Medical image computing and computer-assisted intervention*. Springer (2018)
42. Wallach, N.R.: *Harmonic analysis on homogeneous spaces*. M. Dekker (2018)
43. Weiler, M., Geiger, M., Welling, M., Boomsma, W., Cohen, T.: 3D Steerable CNNs: learning rotationally equivariant features in volumetric data. *Adv. Neural Inform. Process. Syst.* 31 (2018)

44. Worrall, D., Brostow, G.: Cubenet: Equivariance to 3d rotation and translation. In: Proceedings of the European Conference on Computer Vision (ECCV) (2018)
45. Wu, J.: Introduction to convolutional neural networks. National Key Lab for Novel Software Technology. Nanjing University, China (2017)
46. Wu, H., Xiao, B., Codella, N., Liu, M., Dai, X., Yuan, L., Zhang, L.: Cvt: Introducing convolutions to vision transformers (2021)
47. Xu, Y., Jia, Z., Wang, L-B., Ai, Y., Zhang, F., Lai, M., Eric, I., Chang, C.: Large scale tissue histopathology image classification, segmentation, and visualization via deep convolutional activation features. BMC bioinform. (2017)
48. Zhang, Y., Qiu, Z., Yao, T., Liu, D., Mei, T.: Fully convolutional adaptation networks for semantic segmentation. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (2018)