

Unifying spherical convolutions through lifting

Jason D. McEwen^{a)}

¹⁾*Mullard Space Science Laboratory, University College London, Surrey RH5 6NT, United Kingdom*

²⁾*The Alan Turing Institute, Euston Road, London NW1 2DB, United Kingdom*

(Dated: 9 October 2026)

We develop a unified description of convolution on the sphere by lifting to the rotation group. A number of definitions of spherical convolution have been proposed, differing in whether the kernel is axisymmetric or directional, whether the output is on the sphere or the rotation group, whether the signals carry spin, and in their computational cost. No unified treatment exists. Regarding the sphere as the homogeneous space $S^2 \cong \text{SO}(3)/\text{SO}(2)$, we lift a spin s field to the rotation group, convolve on $\text{SO}(3)$, and project back to the sphere through the adjoint of the lift for a potentially different spin s' . We recover every rotation-equivariant convolution considered previously through constraints on the kernel, spin choices and inclusion or omission of the lift and projection. For maps between fields of fixed input and output spin, we prove that directional kernels, output on the sphere and equivariance cannot be achieved simultaneously. We make explicit that spin fields are sections of associated line bundles, linking the physics and bundle literatures, and express the spin rotation phase as the character of the twist of the section. We present harmonic-space algorithms for the exact and efficient computation of lifted spherical convolutions of band-limited fields by appealing to sampling theorems on the sphere and rotation group. All operators are implemented in the open-source `s2conv` code, and numerical experiments confirm the reductions and the equivariance of the operators to machine precision. Finally, we illustrate the operators on fields measured on the Earth and on the celestial sky.

I. INTRODUCTION

In many fields of science and engineering data are measured on the sphere. Applications where data are defined inherently on the sphere are found in cosmology (*e.g.* the cosmic microwave background and its polarisation [1, 2], and gravitational lensing [3]), geophysics and planetary science, meteorology and climate science (*e.g.* global weather forecasting [4]), computer vision [5, 6] and biomedical imaging (*e.g.* diffusion tensor imaging [7]), to name only a few. In many of these applications a harmonic analysis of the data is insightful and convolution is often the fundamental operation: it underlies filtering [8], wavelet analysis [9–12] and, more recently, convolutional neural networks and neural operators defined natively on the sphere [4–6, 13–16]. Moreover, many of the signals of interest are not only scalar but carry spin: tangent vector fields such as horizontal wind are spin ± 1 fields, the polarisation of the cosmic microwave background and cosmic shear of weak gravitational lensing are spin ± 2 fields, and the second flexion of weak gravitational lensing is a spin ± 3 field [17] (Fig. 1); the values of all of these are defined relative to a local frame and acquire a phase when that frame is rotated.

A number of definitions of convolution on the sphere have been proposed in the literature, differing in the class of kernel supported, the domain on which the convolved signal is defined, whether spin is accommodated, and the cost of computation. Axisymmetric convolution (*e.g.* [4, 6, 9, 18–20]) restricts the kernel to be invariant under azimuthal rotation when centred on the North pole, so that the output remains on the sphere and may

^{a)}Corresponding author: jason.mcewen@ucl.ac.uk




FIG. 1. Spin fields on the sphere. From left to right: a scalar (spin-0) field f , shown as colour, and the spin-1, spin-2 and spin-3 fields ∂f , $\partial^2 f$ and $\partial^3 f$ obtained from it by the spin raising operator, drawn in the tangent plane with glyphs whose symmetry is their spin. Temperature, wind, polarisation or shear, and the second flexion of weak lensing are physical examples of spin 0, 1, 2 and 3 respectively.

be computed simply by a product in harmonic space; directional kernels are, however, not supported. Directional convolution (e.g. [5, 9, 16, 19, 21–24]) relaxes this restriction to arbitrary kernels; however, the output is then defined on the rotation group $\text{SO}(3)$, with greater computational cost when all azimuthal orders are retained. Convolution on the rotation group itself [5, 13, 14, 25] takes both field and kernel to be functions on $\text{SO}(3)$. Spin convolutions extend the axisymmetric and directional constructions to spin signals [12, 15, 26–29], while a further family of anisotropic constructions [20, 30–32] recovers an output on the sphere for directional kernels, but is not in general equivariant under all rotations. Roddy & McEwen [32] compare isotropic, left, directional, commutative anisotropic and sifting constructions by support for directional kernels, spherical output and efficient computation, motivating sifting by the limitations of the preceding constructions in that comparison. Several families are related by scalar or axisymmetric specialisations, but their differing conventions and representations obscure the full set of connections. We make those connections explicit and distinguish the original criteria from full rotation equivariance.

We seek a common description in the language of homogeneous spaces. The sphere is the quotient $\text{SO}(3)/\text{SO}(2)$ of the rotation group by the stabiliser of a point, and fields on the sphere carrying spin are sections of the vector bundles associated to this quotient, which may be encoded as functions on the group, the so-called Mackey functions, on which the group acts by the representation induced from the stabiliser [33, 34]. In this language, equivariant linear maps between fields admit group-convolution representations [35, 36]. Cohen, Geiger & Weiler [36] and Xu *et al.* [37] already construct layers by lifting from G/H to G , convolving, and returning to G/H , with possibly different stabiliser representations at input and output. Related local-frame formulations are developed in [38, 39].

In this article we show that all of the rotation-equivariant spherical convolutions listed above are special cases of a single construction, which we call the *generalised lifted spherical convolution*: a spin- s field on the sphere is lifted to the rotation group through the irreducible representation $\rho^{(s)}(\gamma) = e^{is\gamma}$ of $\text{SO}(2)$, convolved on $\text{SO}(3)$ with an arbitrary kernel Ψ , and projected back to the sphere through the adjoint of the lift for a possibly different spin s' . By decomposing the lifted field and the kernel into their Wigner expansions and noting the orthogonality of the Wigner D -functions, the construction is found to act in spin spherical harmonic space by multiplication by the conjugate of a single Wigner coefficient $\Psi_{-s',-s}^\ell$ of the kernel, up to a sign. Each existing equivariant convolution is then recovered by constraining the support of the Wigner coefficients of Ψ , by the choice of spins s and s' , and by whether the lift or the projection is omitted, as we elaborate in Sec. VI. The constants in the literature arise from the different ways in which a kernel on the sphere is identified with a kernel on the rotation group. The adjoint is a bounded equivariant projection on all group functions and selects the prescribed output spin for arbitrary kernels. Consequently, the axisymmetric, directional, group and spin convolutions and others, and the analysis and synthesis operators of wavelet transforms on the sphere, become instances

of one operator whose properties, equivariance in particular, are established once.

The construction also makes explicit what equivariance permits. We collect the four cases of bounded equivariant maps between spin fields and group functions: scalar multipliers between fixed spins, vectors for spin-to-group and group-to-spin maps, and matrices acting on group multiplicities, with exact boundedness conditions. Their kernels are generally distributions, with square integrable representatives precisely for Hilbert–Schmidt maps. Azimuthally band-limited group functions form tuples of spin fields, so a directional convolution with a steerable kernel is a finite tuple of spin-changing convolutions. A key result of the article is that, for maps between fields of fixed input and output spin, directional kernels, output on the sphere and equivariance cannot be achieved simultaneously.

The geometric description yields two further key results. We make explicit that the spin-weighted functions of physics [40, 41], defined by the phase that they acquire under a rotation of the local frame, are the sections of the line bundles associated to the quotient $\text{SO}(3)/\text{SO}(2)$; making this identification explicit links the physics and the bundle literatures. We also express the phase that accompanies the rotation of a spin field through the geometry of the bundle: the residual-frame phase developed by the author and collaborators [12] is the character $\rho^{(s)}$ evaluated on the inverse twist of the bundle section, so that the rotation of a spin field is the representation of the rotation group induced from that character. A forthcoming article [29] applies the construction to spin spherical Fourier neural operators.

The remainder of this article is structured as follows. Harmonic analysis on the sphere and rotation group is reviewed concisely in Sec. II, to present the mathematical preliminaries and conventions required subsequently. In Sec. III we review the spherical convolutions of the literature, restating each as a numbered definition in a single convention together with its harmonic form. Fields on homogeneous spaces, the lift and its adjoint, and induced representations are developed in Sec. IV. The generalised lifted spherical convolution is defined in Sec. V, where its harmonic form, equivariance, kernel constraint and composition rule are derived and the classification of equivariant maps is proved. In Sec. VI each rotation-equivariant construction treated here is recovered as a special case, and in Sec. VII the constructions that fall outside the framework are analysed. Harmonic-space algorithms, their computational cost, the `s2conv` code that implements them and the numerical tests of the reductions and of equivariance are presented in Sec. VIII, and the operators are illustrated on measured fields of the Earth and of the sky in Sec. IX. Concluding remarks are made in Sec. X. The identities for Wigner functions and spin harmonics used in the proofs are collected in Appendix A, and the proofs in Appendix B.

II. HARMONIC ANALYSIS ON THE SPHERE AND ROTATION GROUP

We first describe the sphere as a quotient of the rotation group, before reviewing scalar and spin spherical harmonics, the Wigner D -functions and the relation between them, following the conventions of [12, 42, 43]. Spin fields and their rotation are then discussed.

A. The sphere and the rotation group

We consider the space of square integrable functions $L^2(S^2)$ on the two-sphere $S^2 = \{\omega \in \mathbb{R}^3 : |\omega| = 1\}$, parameterised by spherical coordinates $\omega = (\theta, \varphi)$ with colatitude $\theta \in [0, \pi]$ and longitude $\varphi \in [0, 2\pi)$, with the inner product of $f, g \in L^2(S^2)$ defined by $\langle f, g \rangle = \int_{S^2} f(\omega) g^*(\omega) d\mu(\omega)$, where $d\mu(\omega) = \sin \theta d\theta d\varphi$ is the usual invariant measure on the sphere and g^* denotes complex conjugation. We write $\eta = (0, 0)$ for the North pole.

Rotations $\rho \in \text{SO}(3)$ are parameterised by zyz Euler angles $\rho = (\alpha, \beta, \gamma)$, with $\alpha \in [0, 2\pi)$, $\beta \in [0, \pi]$ and $\gamma \in [0, 2\pi)$, corresponding to the rotation matrix

$$R(\alpha, \beta, \gamma) = R_z(\alpha) R_y(\beta) R_z(\gamma). \quad (1)$$

We adopt the *zyz* Euler convention corresponding to the rotation of a physical body in a *fixed* coordinate system about the z , y and z axes by γ , β and α , respectively. The inner product of $F, G \in L^2(\text{SO}(3))$ is defined by $\langle F, G \rangle = \int_{\text{SO}(3)} F(\rho) G^*(\rho) d\mu(\rho)$, where $d\mu(\rho) = \sin \beta d\alpha d\beta d\gamma$ is the usual invariant measure on the rotation group, of total mass $8\pi^2$, which we do not normalise. Note that $\langle \cdot, \cdot \rangle$ is used to denote inner products over both the sphere and the rotation group (the case adopted can be inferred from the context). The rotation of a function on the sphere is defined by $(\mathcal{R}_\rho f)(\omega) = f(R_\rho^{-1}\omega)$, where R_ρ is the rotation matrix corresponding to ρ , and for a point $\omega = (\theta, \varphi)$ we adopt the shorthand notation $\mathcal{R}_\omega = \mathcal{R}_{(\varphi, \theta, 0)}$.

The map $p : \text{SO}(3) \rightarrow S^2$, $p(\rho) = R_\rho \eta$, sends the rotation $R(\alpha, \beta, \gamma)$ to the point with colatitude β and longitude α , and is invariant under right multiplication by the subgroup $\text{SO}(2) = \{R_z(\gamma)\}$ of rotations about the z axis, which is the stabiliser of the North pole. Hence $S^2 \cong \text{SO}(3)/\text{SO}(2)$, the coset $\rho \text{SO}(2)$ corresponding to the point $\rho\eta$. Throughout, a point of the sphere is written (θ, φ) when it is considered in its own right and (β, α) when it arises as the image of the North pole under a rotation with Euler angles (α, β, γ) ; the two descriptions denote the same point.

B. Spherical harmonics, spin spherical harmonics and Wigner functions

The spherical harmonics $Y_{\ell m} : S^2 \rightarrow \mathbb{C}$ form the canonical orthonormal basis of $L^2(S^2)$ and are defined by

$$Y_{\ell m}(\theta, \varphi) = \sqrt{\frac{2\ell+1}{4\pi} \frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos \theta) e^{im\varphi}, \quad \ell \in \mathbb{N}, \quad |m| \leq \ell, \quad (2)$$

where P_ℓ^m are the associated Legendre functions, for natural degree ℓ and integer order m . We adopt the Condon–Shortley phase convention, with the $(-1)^m$ phase factor included in the definition of the associated Legendre functions, to ensure that the conjugate symmetry relation $Y_{\ell m}^* = (-1)^m Y_{\ell, -m}$ holds. Due to the orthogonality and completeness of the spherical harmonics, any $f \in L^2(S^2)$ may be represented by its spherical harmonic expansion $f = \sum_{\ell m} f_{\ell m} Y_{\ell m}$, where the spherical harmonic coefficients are given by the usual projection onto each basis function, $f_{\ell m} = \langle f, Y_{\ell m} \rangle$.

Spin spherical harmonics ${}_s Y_{\ell m}$, for integer spin s with $|s| \leq \ell$, are defined through the action of the spin raising and lowering operators $\tilde{\partial}$ and $\bar{\partial}$ on the scalar spherical harmonics [40, 41],

$$\tilde{\partial} = -\sin^s \theta \left(\partial_\theta + \frac{i \partial_\varphi}{\sin \theta} \right) \sin^{-s} \theta, \quad \bar{\partial} = -\sin^{-s} \theta \left(\partial_\theta - \frac{i \partial_\varphi}{\sin \theta} \right) \sin^s \theta, \quad (3)$$

which act on a spin- s quantity, by

$${}_s Y_{\ell m} = \begin{cases} \sqrt{\frac{(\ell-s)!}{(\ell+s)!}} \tilde{\partial}^s Y_{\ell m}, & 0 \leq s \leq \ell, \\ (-1)^s \sqrt{\frac{(\ell+s)!}{(\ell-s)!}} \bar{\partial}^{-s} Y_{\ell m}, & -\ell \leq s \leq 0, \end{cases} \quad (4)$$

so that

$$\tilde{\partial} {}_s Y_{\ell m} = \sqrt{(\ell-s)(\ell+s+1)} {}_{s+1} Y_{\ell m}, \quad \bar{\partial} {}_s Y_{\ell m} = -\sqrt{(\ell+s)(\ell-s+1)} {}_{s-1} Y_{\ell m}. \quad (5)$$

For each fixed s the spin spherical harmonics form an orthonormal basis of the space $L_s^2(S^2)$ of square integrable spin- s fields on the sphere (Sec. II C), reducing to the scalar spherical harmonics for $s = 0$. Any ${}_s f \in L_s^2(S^2)$ may thus be represented by its spin spherical harmonic expansion ${}_s f = \sum_{\ell m} {}_s f_{\ell m} {}_s Y_{\ell m}$, with coefficients ${}_s f_{\ell m} = \langle {}_s f, {}_s Y_{\ell m} \rangle$. Note that ${}_s Y_{\ell m}$ exists only for $\ell \geq |s|$, so that a spin- s field has no harmonic content below degree $|s|$.

The Wigner D -functions $D_{mn}^\ell : \text{SO}(3) \rightarrow \mathbb{C}$, for $|m|, |n| \leq \ell$, are the matrix elements of the irreducible unitary representations of the rotation group [44], and may be decomposed as

$$D_{mn}^\ell(\alpha, \beta, \gamma) = e^{-im\alpha} d_{mn}^\ell(\beta) e^{-in\gamma}, \quad (6)$$

where the Wigner d -functions d_{mn}^ℓ are real. Since the D^ℓ are representation matrices they satisfy $D^\ell(\rho\rho') = D^\ell(\rho)D^\ell(\rho')$ and $D^\ell(\rho^{-1}) = D^\ell(\rho)^\dagger$, and the orthogonality relation of their matrix elements reads $\langle D_{mn}^\ell, D_{m'n'}^{\ell'} \rangle = \frac{8\pi^2}{2\ell+1} \delta_{\ell\ell'} \delta_{mm'} \delta_{nn'}$, where δ_{ij} is the Kronecker delta symbol. Together with completeness from the Peter–Weyl theorem [34], this makes the conjugate D -functions an orthogonal basis of $L^2(\text{SO}(3))$, and any $F \in L^2(\text{SO}(3))$ admits the Wigner expansion

$$F(\rho) = \sum_{\ell=0}^{\infty} \frac{2\ell+1}{8\pi^2} \sum_{m,n=-\ell}^{\ell} F_{mn}^\ell D_{mn}^{\ell*}(\rho), \quad F_{mn}^\ell = \langle F, D_{mn}^{\ell*} \rangle = \int_{\text{SO}(3)} F(\rho) D_{mn}^\ell(\rho) d\mu(\rho), \quad (7)$$

where we adopt the conjugate D -functions as basis elements since this convention simplifies connections to convolutions on the sphere. Parseval's identity reads $\|F\|^2 = \sum_{\ell} \frac{2\ell+1}{8\pi^2} \sum_{mn} |F_{mn}^\ell|^2$. The Wigner D -functions are related to the spin spherical harmonics by [41, 42]

$$D_{mn}^\ell(\alpha, \beta, \gamma) = (-1)^n \sqrt{\frac{4\pi}{2\ell+1}} {}_{-n}Y_{\ell m}^*(\beta, \alpha) e^{-in\gamma}, \quad (8)$$

or equivalently

$$e^{-is\gamma} {}_sY_{\ell m}(\beta, \alpha) = (-1)^s \sqrt{\frac{2\ell+1}{4\pi}} D_{m,-s}^{\ell*}(\alpha, \beta, \gamma). \quad (9)$$

Consequently, the spin spherical harmonics of spin s are, up to normalisation and conjugation, the column $n = -s$ of the Wigner D -functions restricted to $\gamma = 0$. This relation underlies the lift developed in Sec. IV: a spin- s field lifted to the rotation group occupies a single column of Wigner space.

Under rotation the spin spherical harmonic coefficients of a field transform with the same Wigner matrices for every spin [12],

$$(\mathcal{R}_\rho {}_s f)_{\ell m} = \sum_{n=-\ell}^{\ell} D_{mn}^\ell(\rho) {}_s f_{\ell n}, \quad (10)$$

where the rotation $\mathcal{R}_\rho$ of a spin field is defined in (12) below. Similarly, the Wigner coefficients of a left-translated function $F(\rho^{-1} \cdot)$ read $\sum_{m'} D_{mm'}^\ell(\rho) F_{m'n}^\ell$, *i.e.* left translation acts on the first index of the Wigner coefficients and leaves the second untouched.

Throughout, we say that a field on the sphere is band-limited at L if its harmonic coefficients vanish for all $\ell \geq L$, and that a function on $\text{SO}(3)$ is in addition azimuthally band-limited at N if $F_{mn}^\ell = 0$ for all $|n| \geq N$. The harmonic band-limit plays no role in the theory developed in Secs. IV–VII, which is formulated in L^2 ; it enters only in Sec. VIII, where it permits all transforms to be computed exactly and efficiently by appealing to sampling theorems on the sphere [42] and rotation group [45], as implemented in [43]. The azimuthal band-limit, in contrast, is structural: it is what makes a function on the rotation group a finite tuple of spin fields (Corollary IV.12 and Proposition VI.6).

C. Spin fields

A spin- s field ${}_s f$ on the sphere is a complex function whose value at ω is defined relative to a choice of orthonormal frame (e_θ, e_φ) in the tangent plane and transforms as

$${}_s f'(\omega) = e^{-is\chi} {}_s f(\omega) \quad (11)$$

under a right-handed rotation of the frame by χ [1, 40, 41].[46] It is important to note that the rotation considered here is *not* a global rotation on the sphere but rather a rotation by χ in the tangent plane centred on ω . Scalar functions such as temperature are spin 0; the complex representation $v_\theta \pm i v_\phi$ of a tangent vector field such as wind has spin ± 1 ; the Stokes combination $Q \pm iU$ of linear polarisation, and likewise cosmic shear, has spin ± 2 [1]; and the second flexion of weak lensing, the spin-raised gradient $\bar{\partial}\gamma$ of the shear, which bends images into three-fold arcs, has spin ± 3 [17]. In Fig. 1 we plot a scalar field together with the spin-1, spin-2 and spin-3 fields obtained from it by repeated application of the spin raising operator $\bar{\partial}$.

Figure 2 shows seven such fields measured on the Earth and on the sky, which span spins 0, 1 and 2 and to which we return in Sec. IX. The relief of the Earth is a scalar field. The deflection of the vertical, the angle by which a plumb line departs from the normal to the reference ellipsoid, the horizontal part of the geomagnetic field, which aligns a compass needle, the wind near the surface and the currents at the surface of the ocean are tangent vector fields, of spin ± 1 , drawn as arrows. The polarisation of the cosmic microwave background and the shear of weak gravitational lensing, the stretching of the images of distant galaxies, are fields of spin ± 2 , drawn as headless sticks, since a polarisation angle or a direction of stretching is unchanged by a half turn. In each case the frame law (11) is a physical statement: a bearing, a polarisation angle or a direction of stretching is a number only once a reference direction has been chosen at each point, and it changes by a known phase when that choice changes. The constructions of this article act on such fields without a global choice of frame, which no choice of coordinates on the sphere can provide.

Because the frame is fixed by the coordinate lines, a global rotation of a spin field must be accompanied by a frame correction. The rotation of a spin field is defined by [12]

$$(\mathcal{R}_\rho \, {}_s f)(\omega) = e^{-is\vartheta} \, {}_s f(R_\rho^{-1}\omega), \quad (12)$$

where the angle $\vartheta \in [0, 2\pi)$ is the third Euler angle of the rotation $R_\rho^{-1}R_\omega$, with $R_\omega = R(\varphi, \theta, 0)$ the rotation matrix corresponding to the Euler angles $(\varphi, \theta, 0)$. For the scalar setting where $s = 0$, (12) reduces to the usual rotation operator defined solely through a rotation of the coordinate system. The exponential factor compensates the residual rotation of the local frame so that the field retains its spin, as described in [12, 47], and yields the harmonic law (10). We write $\mathcal{R}_\rho$ for the rotation of scalar and spin fields alike, the spin being that of the field on which it acts; in Sec. IV D we show that $\mathcal{R}_\rho$ is the representation $\pi_s(R_\rho)$ of the rotation group induced from the character $\rho^{(s)}$, with the phase evaluated on the inverse twist. The Euler-angle description holds away from the coordinate singularities, hence almost everywhere.

III. SPHERICAL CONVOLUTIONS IN THE LITERATURE

In this section we review the definitions of convolution on the sphere proposed in the literature, restating each as a numbered definition in the single convention of Sec. II, together with its harmonic form. We first recall the Euclidean convolution and the difficulty of translating it to the sphere, before reviewing axisymmetric convolutions, directional convolutions, convolution on the rotation group, spin convolutions, anisotropic convolutions whose output remains on the sphere, and finally the projections from the rotation group back to the sphere that arise in wavelet synthesis. The definitions are summarised in Table I.

A. Convolution on the plane and translation on the sphere

On the plane $\mathbb{R}^2$ the convolution $(f \star g)(x) = \int_{\mathbb{R}^2} f(y) g(x - y) dy$ is a bilinear pairing of f with a translated and reflected g , or an inner product when g is real. It commutes with translations, it is commutative, and every translation-equivariant bounded linear operator is of this form, with a distributional kernel in general. On the sphere translations




FIG. 2. Fields of spin 0, 1 and 2 on the Earth and on the celestial sky (the examples of Sec. IX) in the Mollweide projection. The relief (a) is a scalar field. The deflection of the vertical (b), the small angle by which the direction of gravity differs from the normal to the reference ellipsoid because of the uneven distribution of mass within the Earth, the horizontal geomagnetic field (c), which aligns a compass needle, the wind (d) and the currents (e) are spin-1 fields, with the modulus shown as colour and the direction drawn as arrows. The polarisation of the cosmic microwave background (f) and the shear of weak gravitational lensing (g) are spin-2 fields, drawn as headless sticks over the temperature anisotropy and the convergence, and are smoothed so that the sticks follow their large-scale patterns.

are replaced by rotations. However, as Kennedy *et al.* [20] emphasise, the rotation group is not the analogue of the plane: a rotation carries three parameters while a point on the sphere carries two, and a kernel translated to a point by a rotation acquires an orientation that has no counterpart in the Euclidean setting. Every definition of spherical convolution must therefore decide how to treat this third degree of freedom, namely whether to integrate it out, to retain it in the output, or to fix it by a rule. The definitions below treat this orientation in different ways. Roddy & McEwen [32] compare them by three criteria, namely support for directional kernels, an output that remains on the sphere, and computational cost no greater than that of a fast spherical harmonic transform, to which we add equivariance under $\text{SO}(3)$, *i.e.* that convolution commute with the rotation of the signal.

B. Axisymmetric convolutions with output on the sphere

The simplest approach is to restrict the kernel to be axisymmetric, so that its orientation is immaterial.

Definition III.1 (Axisymmetric convolution). For $f, \psi \in L^2(S^2)$ with axisymmetric kernel ψ , *i.e.* $\psi_{\ell m} = \psi_{\ell 0} \delta_{m0}$, the axisymmetric convolution is

$$(f \odot \psi)(\omega) = \langle f, \mathcal{R}_\omega \psi \rangle = \int_{S^2} f(\omega') (\mathcal{R}_\omega \psi)^*(\omega') d\mu(\omega'), \quad \mathcal{R}_\omega = \mathcal{R}_{(\varphi, \theta, 0)}. \quad (13)$$

For an axisymmetric kernel, *i.e.* a kernel invariant under azimuthal rotation when centred on the North pole so that $\psi_{\ell m} = \psi_{\ell 0} \delta_{m0}$, the harmonic form of (13) reads [9, 32]

$$(f \odot \psi)_{\ell m} = \sqrt{\frac{4\pi}{2\ell+1}} f_{\ell m} \psi_{\ell 0}^*, \quad (14)$$

which follows directly from (10) and (8). The output remains on the sphere and the convolution is simply a product in harmonic space, hence it may be computed exactly and efficiently at the cost of two spherical harmonic transforms. This operator underlies both the axisymmetric continuous wavelets of [9], constructed by stereographic dilation, and the needlet and scale-discretised constructions of [11, 48]. Driscoll & Healy [18] instead integrate over the full rotation group.

Definition III.2 (Left convolution of Driscoll & Healy). For $f, h \in L^2(S^2)$, the left convolution is

$$(f \star_L h)(\omega) = \int_{\text{SO}(3)} f(\rho\eta) h(R_\rho^{-1}\omega) d\mu(\rho). \quad (15)$$

Its harmonic form reads [6, 18, 20]

$$(f \star_L h)_{\ell m} = 2\pi \sqrt{\frac{4\pi}{2\ell+1}} f_{\ell m} h_{\ell 0}. \quad (16)$$

Note that the kernel h may be arbitrary but only its zonal part $h_{\ell 0}$ enters, since the integral over the third Euler angle averages h about its own pole (Theorem 3 of [20]); note also that the pairing is bilinear, with no complex conjugate, and that Kennedy *et al.* [20] divide (15) by 2π . The left convolution is the convolution of the spherical convolutional networks of Esteves *et al.* [6], who observe that it “only makes sense to learn zonal filters”, and of the spherical Fourier neural operator of Bonev *et al.* [4].

C. Directional convolutions with output on the rotation group

The second approach retains the orientation of the kernel in the output.

Definition III.3 (Directional convolution). For $f, \psi \in L^2(S^2)$, the directional convolution is

$$(f \circledast \psi)(\rho) = \langle f, \mathcal{R}_\rho \psi \rangle = \int_{S^2} f(\omega) (\mathcal{R}_\rho \psi)^*(\omega) d\mu(\omega), \quad \rho \in \text{SO}(3). \quad (17)$$

By decomposing f and ψ into their spherical harmonic expansions and noting the orthogonality of the spherical harmonics and their rotation by (10), the directional convolution may be written

$$(f \circledast \psi)(\rho) = \sum_{\ell mn} f_{\ell m} \psi_{\ell n}^* D_{mn}^{\ell*}(\rho), \quad (f \circledast \psi)_{mn}^\ell = \frac{8\pi^2}{2\ell+1} f_{\ell m} \psi_{\ell n}^*, \quad (18)$$

where the second expression follows by comparison with (7). Consequently, the output is a function on $\text{SO}(3)$ whose Wigner coefficients are the outer product of the harmonic coefficients of signal and kernel [5, 9, 19, 23, 24]. Wandelt & Górski [19] parameterise the rotation by the scan geometry of a cosmic microwave background experiment and evaluate (18) on a torus, and Challinor *et al.* [21] give the polarised, *i.e.* spin-2, version. Fast algorithms by factorisation of rotations [9, 19, 24] cost $\mathcal{O}(L^4)$ in general, or $\mathcal{O}(NL^3)$ for kernels azimuthally band-limited at N , which is the case of steerable kernels [22]. Cohen *et al.* [5] write the same operator for real signals with the Haar measure normalised to unity, and the discrete–continuous convolution of Ocampo *et al.* [16] is a quadrature of (17) on a sampling of the sphere. Directional convolution supports directional kernels but yields an output that does not remain on the sphere, at greater computational cost than a spherical harmonic transform.

D. Convolution on the rotation group

Once the output is a function on $\text{SO}(3)$, subsequent convolutions take place on the rotation group itself.

Definition III.4 (Group convolution on $\text{SO}(3)$). For $\Psi, F \in L^2(\text{SO}(3))$, the group convolution is

$$(\Psi \star F)(\rho) = \langle F, \pi_G(\rho)\Psi \rangle = \int_{\text{SO}(3)} \Psi^*(\rho^{-1}\rho') F(\rho') d\mu(\rho'), \quad (\pi_G(\rho)\Psi)(\rho') = \Psi(\rho^{-1}\rho'), \quad (19)$$

where π_G is the left regular representation of $\text{SO}(3)$.

In Wigner space the group convolution is a matrix product for each degree [5, 14, 25],

$$(\Psi \star F)_{mn}^\ell = \sum_{m'=-\ell}^{\ell} F_{mm'}^\ell \Psi_{nm'}^{\ell*}, \quad (20)$$

as we show in Appendix B. Definition III.4 is written in correlation form, with the kernel conjugated and translated by the left regular representation, so that (19) is manifestly the inner product of F with a rotated copy of Ψ ; this is the $\text{SO}(3)$ layer of Cohen *et al.* [5], the correlation of Kostelec & Rockmore [25], the linear part of the Clebsch–Gordan networks of Kondor *et al.* [13] and, with channel mixing, the generalised convolution of Cobb *et al.* [14]. Our convention differs from that of Xu *et al.* [37] and Kondor & Trivedi [35], who write $\int \kappa(\rho'^{-1}\rho) F(\rho') d\mu(\rho')$, by $\kappa(\rho) = \Psi^*(\rho^{-1})$, and from that of Cohen, Geiger & Weiler [36], who write $\int \kappa(\rho^{-1}\rho') F(\rho') d\mu(\rho')$, by conjugation alone.

E. Spin convolutions

We follow the directional spin wavelet transform of McEwen *et al.* [12], which defines convolutions of spin fields by the same three constructions as in the scalar case: a directional convolution with output on the rotation group, an axisymmetric convolution with output on the sphere, and a synthesis integral back to the sphere. All three rest on the rotation (12) of a spin field.

a. Rotation of spin fields. A coordinate rotation alone, ${}_sf \mapsto {}_sf(R_\rho^{-1} \cdot)$, does not map a spin- s field to a spin- s field, because the coordinate frame at $R_\rho^{-1}\omega$ is not the image under R_ρ^{-1} of the coordinate frame at ω . The operator (12) corrects this with the phase $e^{-is\vartheta}$, where ϑ is the third Euler angle of $R_\rho^{-1}R_\omega$. In [12] the phase is justified in harmonic space: from the composition law $D^\ell(\rho_1\rho_2) = D^\ell(\rho_1)D^\ell(\rho_2)$ and (9) one finds

$$({}_\rho {}_s Y_{\ell m})(\omega) = \sum_{n=-\ell}^{\ell} D_{nm}^\ell(\rho) {}_s Y_{\ell n}(\omega), \quad (21)$$

equivalently (10), so that $\mathcal{R}_\rho$ maps spin- s harmonics to spin- s harmonics with the same Wigner matrices as in the scalar case, which is what makes the constructions below consistent. The residual frame rotation is described in [12, 47]. Here we express it through the section twist, connecting the adopted spin convention to the induced representation of Cohen, Geiger & Weiler [36] (Proposition IV.10 and Remark IV.11). The lift then intertwines $\mathcal{R}_\rho$ with left translation on the rotation group, giving (21).

Definition III.5 (Directional spin convolution). For ${}_sf, {}_s\psi \in L_s^2(S^2)$, the directional spin convolution is

$$({}_sf \otimes {}_s\psi)(\rho) = \langle {}_sf, \mathcal{R}_\rho {}_s\psi \rangle = \int_{S^2} {}_sf(\omega) (\mathcal{R}_\rho {}_s\psi)^*(\omega) d\mu(\omega), \quad \rho \in \text{SO}(3). \quad (22)$$

These are the wavelet coefficients of [12] for a spin wavelet ${}_s\psi$. By decomposing both fields into their spin spherical harmonic expansions and noting (21), the directional spin convolution may be written

$$({}_sf \otimes {}_s\psi)(\rho) = \sum_{\ell mn} {}_sf_{\ell m} {}_s\psi_{\ell n}^* D_{mn}^{\ell*}(\rho), \quad ({}_sf \otimes {}_s\psi)_{mn}^\ell = \frac{8\pi^2}{2\ell+1} {}_sf_{\ell m} {}_s\psi_{\ell n}^*, \quad (23)$$

the spin generalisation of (18). From (23) it is apparent that the output is a scalar function on $\text{SO}(3)$ whatever the spin of the input: the third Euler angle records the orientation of the kernel about each point, and the spin phases of field and kernel cancel in the inner product. The transform may thus be computed via an inverse Wigner transform of the coefficients (23), as we elaborate in Sec. VIII.

Definition III.6 (Axisymmetric spin convolution). For ${}_sf \in L_s^2(S^2)$ and an axisymmetric kernel ${}_s\Phi$, *i.e.* $\mathcal{R}_{(0,0,\gamma)} {}_s\Phi = {}_s\Phi$ or equivalently ${}_s\Phi_{\ell m} = {}_s\Phi_{\ell 0} \delta_{m0}$, the axisymmetric spin convolution is

$$({}_sf \odot {}_s\Phi)(\omega) = \langle {}_sf, \mathcal{R}_\omega {}_s\Phi \rangle, \quad \mathcal{R}_\omega = \mathcal{R}_{(\varphi, \theta, 0)}. \quad (24)$$

These are the scaling coefficients of [12], whose harmonic form reads

$$({}_sf \odot {}_s\Phi)_{\ell m} = \sqrt{\frac{4\pi}{2\ell+1}} {}_sf_{\ell m} {}_s\Phi_{\ell 0}^*, \quad (25)$$

where the coefficients on the left are those of the scalar spherical harmonics: the output is a scalar field on the sphere although the input has spin s . The field is recovered from its wavelet and scaling coefficients $W^{\psi^{(j)}} = {}_sf \otimes {}_s\psi^{(j)}$ and $W^\Phi = {}_sf \odot {}_s\Phi$ by the synthesis [12]

$${}_sf(\omega) = \int_{S^2} W^\Phi(\omega') (\mathcal{R}_{\omega'} {}_s\Phi)(\omega) d\mu(\omega') + \sum_j \int_{\text{SO}(3)} W^{\psi^{(j)}}(\rho) (\mathcal{R}_\rho {}_s\psi^{(j)})(\omega) d\mu(\rho), \quad (26)$$

provided the admissibility condition

$$\frac{4\pi}{2\ell+1} |{}_s\Phi_{\ell 0}|^2 + \frac{8\pi^2}{2\ell+1} \sum_j \sum_{n=-\ell}^{\ell} |{}_s\psi_{\ell n}^{(j)}|^2 = 1 \quad \text{for all } \ell \geq |s| \quad (27)$$

holds. The second integral in (26) is the spin version of the synthesis operator of Definition III.10 below and the first its axisymmetric counterpart; both project back to a spin- s field. As we show in Sec. VI, each is the adjoint of the corresponding analysis operator, which is why admissibility takes the form of a resolution of the identity. The wavelets are steerable, *i.e.* azimuthally band-limited, so that their directional coefficients at all orientations are determined by finitely many [12, 22]; Proposition VI.6 identifies such coefficients with a tuple of spin fields.

Other spin constructions in the literature admit the same description. The spin-weighted spherical convolutional networks of Esteves *et al.* [15] define, for a set of input spins W_F and filter spins W_K , the convolution ${}_s(\widehat{F * K})_m^\ell = \sum_{i \in W_F} i \hat{f}_m^\ell {}_s \hat{k}_i^\ell$ for $s \in W_K$, in their notation, so that the output carries the spins of the filter, and a cross-correlation ${}_s(\widehat{F \star K})_m^\ell = \sum_i i \hat{f}_m^\ell i \hat{k}_s^{\ell*}$, which is a sum of spin-changing multipliers. Their explicit D convention and section harmonics agree with ours; the exact kernel conversion, rather than an unchanged column reading of Definition III.5, is given in Proposition VI.5. The spin needlets of Geller & Marinucci [26, 27] use isotropic two-point kernels $\sum_{\ell m} b(\sqrt{e_{\ell s}}/B^j) {}_s Y_{\ell m}(x) {}_s Y_{\ell m}^*(y)$ with $e_{\ell s} = (\ell - |s|)(\ell + |s| + 1)$, mapping spin s to spin s , and their mixed needlets replace one of the two harmonics by a scalar harmonic, connecting spin s with spin 0 only. The spin-weighted spherical convolution introduced by Tessore *et al.* [28] for making maps of cosmic shear from discrete galaxy catalogues, $F(x) = \int_{S^2} f(y) e^{iS\alpha} K(\theta) e^{-is\alpha'} d\mu(y)$, pairs a radial profile K of the separation θ with relative-frame phases and maps spin s to a freely chosen spin S . They state the theorem $F_{\ell m} = K_\ell f_{\ell m}$ for $K(\theta) = \sum_\ell (2\ell+1) K_\ell d_{Ss}^\ell(\theta)/(4\pi)$. Proposition VI.5 translates the angles explicitly and distinguishes the literal printed convention for odd output spins. Their cap profile $K(\theta) = \cos^{2s}(\theta/2)$ for $\theta \leq \beta$ has closed-form coefficients for $S = s \geq 0$. Finally, a number of operators of mathematical physics change the spin of a field while commuting with rotations: the spin raising and lowering operators (5); the E/B decomposition of a spin- ± 2 polarisation field, $a_{E,\ell m} = -(a_{2,\ell m} + a_{-2,\ell m})/2$ and $a_{B,\ell m} = i(a_{2,\ell m} - a_{-2,\ell m})/2$ [1, 2]; the spherical Kaiser–Squires relation ${}_2\gamma_{\ell m} = \mathcal{D}_\ell \kappa_{\ell m}$, with $\mathcal{D}_\ell = -\sqrt{(\ell+2)(\ell-1)/(\ell(\ell+1))}$, between the convergence (spin 0) and shear (spin 2) of weak gravitational lensing [3]; and the all-sky convolution of a polarised beam [21]. All are recovered in Sec. VI.

F. Anisotropic convolutions with output on the sphere

The third approach to address the orientation degree of freedom is to fix it by a rule, which yields an output on the sphere for a directional kernel.

Definition III.7 (Standard correlation). For $f, \psi \in L^2(S^2)$ with ψ not necessarily axisymmetric, the standard correlation is the directional convolution restricted to the rotations with vanishing third Euler angle,

$$(f \circledast \psi)(R_\omega) = \langle f, \mathcal{R}_\omega \psi \rangle, \quad R_\omega = R(\varphi, \theta, 0). \quad (28)$$

This is the “standard correlation” of Wiaux *et al.* [22, 23], and is formula (13) applied to a directional kernel. Its harmonic coefficients are not a pointwise product [23].

Definition III.8 (Commutative anisotropic convolution). For $f, h \in L^2(S^2)$, the commutative anisotropic convolution of Sadeghi *et al.* [30, 31] is

$$(h \odot_S f)(\theta, \varphi) = \int_{S^2} (\mathcal{R}_{(\varphi, \theta, \pi - \varphi)} h)(\omega') f(\omega') d\mu(\omega'). \quad (29)$$

The pairing is bilinear. Sadeghi *et al.* [30] prove that their bilinear directional correlation evaluated at rotations $(\varphi, \theta, \chi(\varphi))$ is commutative if and only if $\chi = \pi - \varphi$, and give its harmonic form and an $\mathcal{O}(L^3 \log L)$ algorithm; they show commutativity, not equivariance.

Definition III.9 (Harmonic multiplication and sifting convolution). For $f, g \in L^2(S^2)$, harmonic multiplication [20] and the sifting convolution [32] are defined in harmonic space by

$$(g \odot_K f)_{\ell m} = g_{\ell m} f_{\ell m}, \quad (f \odot g)_{\ell m} = f_{\ell m} g_{\ell m}^*, \quad (30)$$

the latter arising in real space as $(f \odot g)(\omega) = \langle \mathcal{T}_\omega f, g \rangle$ for the translation operator $(\mathcal{T}_{\omega'} f)(\omega) = \sum_{\ell m} f_{\ell m} Y_{\ell m}(\omega') Y_{\ell m}(\omega)$.

Both operators are diagonal in the spherical harmonic basis. Kennedy *et al.* [20] characterise isotropic operators by independence of the multiplier from m . Harmonic multiplication is commutative; sifting is commutative up to coefficientwise conjugation, not pointwise conjugation of the field. Sifting has the $\mathcal{O}(L^3)$ cost of spherical harmonic transforms and meets the three criteria of [32]. Both families are always polar-equivariant but not fully rotation-equivariant in general; additional symmetries and equivariant special cases are discussed in Sec. VII.

G. Projections from the rotation group to the sphere

Finally, the adjoint of directional analysis synthesises a signal from group coefficients by an integral over $\text{SO}(3)$. It becomes an inverse only with appropriate admissibility or dual-frame conditions.

Definition III.10 (Synthesis or transpose convolution). For $W \in L^2(\text{SO}(3))$ and $\psi \in L^2(S^2)$, the synthesis is

$$(W \circledast^\dagger \psi)(\omega) = \int_{\text{SO}(3)} W(\rho) (\mathcal{R}_\rho \psi)(\omega) d\mu(\rho), \quad (W \circledast^\dagger \psi)_{\ell m} = \sum_{n=-\ell}^{\ell} W_{mn}^\ell \psi_{\ell n}, \quad (31)$$

where the harmonic form follows by expanding W and ψ and noting (10).

This is Haar-measure synthesis, distinct from the flat-scan “transpose convolution” of Wandelt & Górski [19] as explained in Sec. VIA, and is the synthesis step of directional wavelet transforms on the sphere [9, 10, 12, 24], where the kernel is either the wavelet itself, with the constant $8\pi^2/(2\ell+1)$ carried by the admissibility condition as in (27) [12], or the dual wavelet $\psi_{\ell n} \mapsto \frac{2\ell+1}{8\pi^2} \psi_{\ell n}$, with an admissibility condition free of constants [10, 24]. The left convolution of Kennedy *et al.* [20] and Esteves *et al.* [6] is obtained instead by marginalising the bilinear directional correlation over the third Euler angle; with the sesquilinear convention (17), the spherical kernel is conjugated.

H. Summary

The convolutions reviewed are summarised in Table I. In the rotation-equivariant harmonic formulas, the signal order m is a spectator index: the kernel acts on the remaining indices. Dependence of a diagonal multiplier on m , or an unconstrained restriction to selected rotations, instead gives families that are not fully equivariant in general. Their special cases and symmetry conditions are established in Secs. V–VII.

TABLE I. Spherical convolutions in the notation of Sec. II. Symmetry entries describe the family: “yes” guarantees $\text{SO}(3)$ equivariance, “polar” guarantees polar rotations without excluding other symmetries of particular kernels, and “no” means full equivariance is not guaranteed. Harmonic multiplication is commutative; ^a denotes coefficientwise conjugation for sifting, with longitude reflection in real space. “_” means no numbered definition or, in the commutativity column, not applicable. Costs are algorithmic upper bounds for band-limit L and azimuthal band-limit N , with standard correlation referring to grid values.

Convolution	Def.	Input $\rightarrow$ output	Kernel	Equiv. / comm.	Cost
Axisymmetric $\odot$	III.1	$S^2 \rightarrow S^2$	axisymmetric	yes / –	L^3
Left $\star_L$	III.2	$S^2 \rightarrow S^2$	any (zonal part)	yes / no	L^3
Directional $\otimes$	III.3	$S^2 \rightarrow \text{SO}(3)$	any	yes / no	L^4, NL^3
Group $\star$	III.4	$\text{SO}(3) \rightarrow \text{SO}(3)$	any	yes / no	L^4
Directional spin $\otimes$	III.5	$L^2(S^2) \rightarrow \text{SO}(3)$	spin s	yes / no	L^4, NL^3
Axisymmetric spin $\odot$	III.6	$L^2_s(S^2) \rightarrow L^2_0(S^2)$	axisymmetric spin s	yes / –	L^3
Spin-weighted [15]	–	$L^2_s(S^2) \rightarrow L^2_{s'}(S^2)$	spin s'	yes / no	L^3
Spin-weighted, isotropic [28]	–	$L^2_s(S^2) \rightarrow L^2_s(S^2)$	radial, with frame phases	yes / –	L^3
Synthesis $\otimes^\dagger$	III.10	$\text{SO}(3) \rightarrow S^2$	any	yes / –	L^4, NL^3
Standard correlation	III.7	$S^2 \rightarrow S^2$	directional	polar / no	L^3
Commutative anisotropic $\odot_S$	III.8	$S^2 \rightarrow S^2$	directional	no / yes	$L^3 \log L$
Harmonic mult. / sifting $\odot$	III.9	$S^2 \rightarrow S^2$	any	polar / yes ^a	L^3

IV. FIELDS ON HOMOGENEOUS SPACES: LIFTING AND PROJECTION

In this section we develop the machinery on which the unification rests. We first describe the rotation group as a principal bundle over the sphere and introduce the twist that measures the failure of a rotation to preserve a section, before describing spin fields as sections of associated bundles and as Mackey functions on the group. The lift of a spin field to the rotation group and its adjoint, the projection, are then defined and their structure established, and the induced representation of $\text{SO}(3)$ on spin fields is identified with the rotation (12). Finally, the decomposition of an arbitrary function on the rotation group into a tuple of spin fields is derived. The material of the first two subsections is classical [33, 34] and was brought to equivariant networks in [36, 37, 39]; we present it in the concrete setting of the sphere and in the conventions of Sec. II.

A. The principal bundle $\text{SO}(3) \rightarrow S^2$ and the twist

Let G be a compact Lie group and H a closed subgroup, and consider the homogeneous space $X = G/H$ with base point $x_0 = H$, so that H is its stabiliser and gH corresponds to gx_0 . The quotient map $p : G \rightarrow G/H$ exhibits G as a principal H -bundle over G/H : the fibre $p^{-1}(x) = gH$ over $x = gH$ is a copy of H on which H acts freely by right multiplication. For $G = \text{SO}(3)$, $H = \text{SO}(2)$ and $x_0 = \eta$ the fibre over a point ω of the sphere is the circle of unit tangent vectors at ω : a rotation ρ with $\rho\eta = \omega$ is determined by where it sends the North pole and a reference tangent vector there, and the third Euler angle γ parameterises the fibre. The bundle is illustrated in Fig. 3.

A *section* is a map $s : G/H \rightarrow G$ with $p \circ s = \text{Id}$, *i.e.* a choice of coset representative for every point. On the sphere we adopt

$$s(\theta, \varphi) = R(\varphi, \theta, 0) \quad (0 < \theta < \pi), \quad s(\eta) = e, \quad s(-\eta) = R_y(\pi), \quad (32)$$

which chooses the frame aligned with the coordinate lines away from the poles and is smooth there. The letter s denotes the section in $s(x)$ and in compositions $F \circ s$, and the integer spin in representation and operator labels; we write $\sigma_0 = s(\cdot)$ when an argument-free section symbol is needed for the evaluation operators below.

Remark IV.1 (No global section). The bundle $\text{SO}(3) \rightarrow S^2$ is the unit tangent bundle of the sphere: a section is a choice of unit tangent vector at every point, and by the hairy-ball




FIG. 3. The rotation group as a principal $SO(2)$ -bundle over the sphere S^2 . The fibre over a point x is drawn as the rim of a cone standing on the sphere at x ; its points are the rotations $R(\alpha, \beta, \gamma)$ with $R(\alpha, \beta, \gamma)\eta = x$, coloured by the fibre coordinate γ . The red dot is the section $s(x) = R(\alpha, \beta, 0)$, and a general element of the fibre is $g = s(x)R_z(\gamma)$ (inset).

theorem no such choice is continuous. The section (32), which follows the meridians away from the poles, is measurable but not continuous. The intrinsic spin-field constructions below are independent of this choice, with their coordinate expressions related by (36); the mixed-spin section evaluations of Sec. VII A instead depend on the chosen section and do not transform by one common spin phase. For continuous spin fields, the limiting coordinate values at the poles vanish or carry the windings $e^{\mp i s \varphi}$ [42, 47]; bare L^2 classes have no intrinsic point values there.

The action of G does not preserve a section: in general $g s(x) \neq s(gx)$. The discrepancy lies in H .

Definition IV.2 (Twist). The twist $\hbar : G/H \times G \rightarrow H$ is defined by

$$g s(x) = s(gx) \hbar(x, g), \quad \hbar(x, g) = s(gx)^{-1} g s(x). \quad (33)$$

We write $\hbar(g) = \hbar(x_0, g) = s(gx_0)^{-1} g$ for the twist at the base point.

Lemma IV.3 (Composition of twists). $\hbar(x, g_1 g_2) = \hbar(g_2 x, g_1) \hbar(x, g_2)$ for all $x \in G/H$ and $g_1, g_2 \in G$. On the sphere, with the section (32), $\hbar(R(\alpha, \beta, \gamma)) = R_z(\gamma)$ for $0 < \beta < \pi$, hence almost everywhere.

Proof. The composition rule follows by inserting $s(g_2 x) s(g_2 x)^{-1}$ into $s(g_1 g_2 x)^{-1} g_1 g_2 s(x)$ and regrouping (it is Eq. (20) of [36]). On the sphere, for $0 < \beta < \pi$, $s(R(\alpha, \beta, \gamma)\eta) = R(\alpha, \beta, 0)$, hence $\hbar(R(\alpha, \beta, \gamma)) = R(\alpha, \beta, 0)^{-1} R(\alpha, \beta, \gamma) = R_z(\gamma)$. The details are given in Appendix B. $\square$

Geometrically, a rotation g carries the frame chosen by the section at x to a frame at gx that differs from the frame chosen by the section there by a rotation $\hbar(x, g)$ about the local normal, as illustrated in Fig. 4. This residual rotation is what a spin field must compensate for, and it is the origin of the phase in (12).

B. Associated bundles and fields

The irreducible unitary representations of $H = SO(2)$ are the characters

$$\rho^{(s)}(R_z(\gamma)) = e^{i s \gamma}, \quad s \in \mathbb{Z}, \quad (34)$$




FIG. 4. The twist of the section under a rotation, $g s(x) = s(gx) h(x, g)$. The rotation g carries the section point $s(x)$ in the fibre over x to a point of the fibre over gx that differs from $s(gx)$ by the right action of the twist $h(x, g) \in \text{SO}(2)$.

and we call s the spin. Given a representation ρ of H on a vector space V , the associated vector bundle is $A = G \times_{\rho} V = (G \times V)/H$, whose elements are the orbits $[g, v] = \{(gh, \rho(h^{-1})v) : h \in H\}$ of pairs of a pose g and a feature vector v , with projection $[g, v] \mapsto gH$. A *field of type ρ* is a section of A . Scalar fields have type $\rho^{(0)}$; a real tangent vector field is the pair of fields $v_{\theta} \pm i v_{\varphi}$ of types $\rho^{(\pm 1)}$; the trace-free symmetric part of a rank-two tensor, such as linear polarisation or shear, is the pair $Q \pm iU$ of types $\rho^{(\pm 2)}$. A real field of spin s is thus a pair of complex fields of spins $\pm s$ related by conjugation.

A field admits two equivalent encodings.

Definition IV.4 (Mackey functions). A Mackey function of type ρ is a map $F : G \rightarrow V$ with $F(gh) = \rho(h^{-1})F(g)$ for all $g \in G$, $h \in H$. For $\rho = \rho^{(s)}$ the condition reads $F(gR_z(\gamma)) = e^{-is\gamma}F(g)$, and we write

$$\mathcal{M}_s = \{F \in L^2(\text{SO}(3)) : F(\cdot R_z(\gamma)) = e^{-is\gamma}F \text{ in } L^2 \text{ for every fixed } \gamma\}. \quad (35)$$

For $F \in \mathcal{M}_s$ we use a fibrewise equivariant representative when restricting to a section: its value on the section is the coefficient of $e^{-is\gamma}$ in its fibre Fourier series, not the restriction of an arbitrary L^2 representative to a measure-zero set. The second encoding evaluates F along the section, $f(x) = \sqrt{2\pi} F(s(x))$, the factor $\sqrt{2\pi}$ being chosen to match the normalisation of the lift and projection operators introduced in Sec. IV C below, under which f and F determine each other (Lemma IV.7). Changing the section to $s'(x) = s(x)R_z(\chi(x))$ changes this function to

$$f'(x) = \sqrt{2\pi} F(s(x)R_z(\chi(x))) = e^{-is\chi(x)} f(x), \quad (36)$$

which is precisely the frame dependence (11) of a spin- s field. Thus spin- s fields in the sense of Sec. II C are the coordinate expressions, with respect to the section (32), of type- $\rho^{(s)}$ Mackey functions.

Remark IV.5 (Spin fields are sections). The identification just made is the bridge between the physics and the bundle literatures. The spin-weighted quantities of Newman & Penrose [40] and Goldberg *et al.* [41], defined by their behaviour (11) under a rotation of the local frame and adopted for the polarisation of the cosmic microwave background in [1, 2], are the sections of the line bundle associated to $\text{SO}(3) \rightarrow S^2$ by the character $\rho^{(s)}$, and their coordinate expression depends on the choice of frame exactly as (36) prescribes. Boyle [47] draws the same conclusion from the impossibility of a global frame, noting that spin-weighted functions are not functions on the sphere alone but functions of a point together

with a frame at that point, *i.e.* functions on the rotation group; the general theory of such coordinate-independent fields on manifolds is developed by Weiler *et al.* [39]. On the sphere the oriented orthonormal frame bundle is $\text{SO}(3)$, which is distinct from a chosen local frame field. Note that the encoding as a Mackey function is global and involves no choice of section, whereas the encoding as a spin field does; the former is therefore the natural setting in which to establish properties such as equivariance, and the latter the natural setting in which to compute.

C. Lift and projection

The two encodings are related by the lift and its adjoint.

Definition IV.6 (Lift and projection). For $s \in \mathbb{Z}$ define the lift $\Lambda_s : L^2(S^2) \rightarrow L^2(\text{SO}(3))$ and its adjoint, the projection $\Lambda_s^\dagger : L^2(\text{SO}(3)) \rightarrow L^2(S^2)$, by

$$(\Lambda_s f)(\alpha, \beta, \gamma) = \frac{1}{\sqrt{2\pi}} e^{-is\gamma} f(\beta, \alpha), \quad (\Lambda_s^\dagger F)(\beta, \alpha) = \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} e^{is\gamma} F(\alpha, \beta, \gamma) d\gamma. \quad (37)$$

In the language of Definition IV.2, $(\Lambda_s f)(g) = \rho^{(s)}(\hbar(g)^{-1}) f(gH)/\sqrt{2\pi}$, and the projection is the integral over the fibre against the character. The normalisation by $\sqrt{2\pi}$, the square root of the volume of the fibre, makes the lift an isometry. For a general compact Lie group and closed stabiliser H , the analogous normalisation is $\sqrt{\text{vol}(H)}$.

Lemma IV.7 (Structure of the lift). *Let $s \in \mathbb{Z}$.*

- (a) Λ_s is an isometry of $L^2(S^2)$ onto $\mathcal{M}_s$; $\Lambda_s^\dagger \Lambda_s = \text{Id}$ and $\Lambda_s \Lambda_s^\dagger$ is the orthogonal projector onto $\mathcal{M}_s$. The subspaces $\mathcal{M}_s$ are mutually orthogonal and $L^2(\text{SO}(3)) = \widehat{\bigoplus_{s \in \mathbb{Z}} \mathcal{M}_s}$.
- (b) On $\mathcal{M}_s$ the projection is evaluation on the section: $\Lambda_s^\dagger F = \sqrt{2\pi} F \circ s$.
- (c) In harmonic space,

$$(\Lambda_s f)_{mn}^\ell = (-1)^s \sqrt{\frac{8\pi^2}{2\ell+1}} {}_s f_{\ell m} \delta_{n,-s}, \quad (\Lambda_s^\dagger F)_{\ell m} = (-1)^s \sqrt{\frac{2\ell+1}{8\pi^2}} F_{m,-s}^\ell, \quad (38)$$

where the coefficients on the sphere are taken with respect to the spin- s harmonics.

- (d) $\Lambda_s {}_s Y_{\ell m} = (-1)^s \sqrt{\frac{2\ell+1}{8\pi^2}} D_{m,-s}^{\ell*}$: lifted spin harmonics are unit-norm Wigner functions.
- (e) $\pi_G(g) \Lambda_s = \Lambda_s \pi_s(g)$ and $\Lambda_s^\dagger \pi_G(g) = \pi_s(g) \Lambda_s^\dagger$ for all $g \in \text{SO}(3)$, where π_G is the left regular representation of $\text{SO}(3)$ on $L^2(\text{SO}(3))$, $(\pi_G(g)F)(g') = F(g^{-1}g')$, and π_s is the representation of $\text{SO}(3)$ on spin- s fields induced from $\rho^{(s)}$; both are developed in Sec. IV D, where (e) is proved as Proposition IV.10.

Proof. We outline the proof; the details are given in Appendix B. Part (a) follows from Fubini's theorem and the Fourier series of F in the fibre coordinate γ , of which $\Lambda_s \Lambda_s^\dagger F$ is the term in $e^{-is\gamma}$. Part (b) is immediate from (35). Part (d) is the relation (9) between spin spherical harmonics and Wigner functions, together with their orthogonality, and (c) follows from (d) by linearity: by (8) and the orthogonality of the complex exponentials, the Wigner coefficients of $\Lambda_s f$ vanish unless $n = -s$, and $\int_0^{2\pi} e^{-i(s+n)\gamma} d\gamma = 2\pi \delta_{n,-s}$ supplies the constant. Part (e) is Proposition IV.10 below. $\square$

Part (c) shows that the lift of a spin- s field occupies the single column $n = -s$ of Wigner space, so that the Wigner transform restricted to that column is simply a spin- s spherical harmonic transform. This is classical [41, 47] (for scalar fields it is Proposition 1 of Kondor & Trivedi [35]), and it underlies Wigner transforms through spin spherical harmonic transforms [45], implemented in [43].

Remark IV.8 (Evaluation versus adjoint). Cohen, Geiger & Weiler [36] and Xu *et al.* [37] define the inverse of the unnormalised lift on Mackey functions by evaluation on the section, $F \mapsto F \circ s$, which by (b) equals $\Lambda_s^\dagger / \sqrt{2\pi}$ on $\mathcal{M}_s$. Xu *et al.* also define a general convolutional projection P^κ on group functions and implement a fixed fibre integral after the nonlinearity; with their unnormalised stabiliser measure this integral is $\sqrt{2\pi} \Lambda_s^\dagger$ and is not restricted to Mackey functions. Kondor & Trivedi [35] use the scalar fibre average. We adopt the adjoint convention throughout: it is bounded on all of $L^2(\text{SO}(3))$, selects a definite output spin equivariantly for arbitrary kernels, and, with the isometric lift, satisfies $\Lambda_s^\dagger \Lambda_s = \text{Id}$. The degree-dependent factors also cancel under section evaluation on a constrained range; the adjoint provides a consistent normalisation, with the constants of the classical reductions tracked explicitly in Sec. VI. The forthcoming article [29] adopts the same convention.

D. Induced representations

The left regular representation $(\pi_G(g)F)(g') = F(g^{-1}g')$ of G on $L^2(\text{SO}(3))$ commutes with right multiplication and therefore preserves every $\mathcal{M}_s$. Transporting it to the sphere gives the representation induced from $\rho^{(s)}$.

Definition IV.9 (Induced representation). $\pi_s = \text{Ind}_H^G \rho^{(s)}$ is the representation of $G = \text{SO}(3)$ on $L_s^2(S^2)$ given by $\pi_s(g) = \Lambda_s^\dagger \pi_G(g) \Lambda_s$.

Proposition IV.10 (Intertwining and explicit form). π_s is unitary, $\pi_G(g) \Lambda_s = \Lambda_s \pi_s(g)$, and

$$(\pi_s(g)f)(x) = \rho^{(s)}(h(x, g^{-1})^{-1}) f(g^{-1}x). \quad (39)$$

On the sphere, $\pi_s(R_\rho)$ is the rotation $\mathcal{R}_\rho$ of a spin field defined in (12): $h(x, R_\rho^{-1}) = R_z(\vartheta)$ with ϑ the third Euler angle of $R_\rho^{-1}R_x$ away from the coordinate singularities, hence almost everywhere.

Proof. Writing $(\Lambda_s f)(g') = \rho^{(s)}(h(g')^{-1}) f(g'H) / \sqrt{2\pi}$, we have

$$(\pi_G(g) \Lambda_s f)(g') = \frac{1}{\sqrt{2\pi}} \rho^{(s)}(h(g^{-1}g')^{-1}) f(g^{-1}g'H), \quad (40)$$

and Lemma IV.3 with $x = x_0$, $g_1 = g^{-1}$ and $g_2 = g'$ splits $h(g^{-1}g')$ into $h(g'H, g^{-1}) h(g')$, from which (39) and the intertwining relation follow. Unitarity follows since π_G is unitary and Λ_s is an isometry onto an invariant subspace. On the sphere, $R_\rho^{-1}R_x = s(R_\rho^{-1}x)R_z(\vartheta)$ by the definition of ϑ , hence $h(x, R_\rho^{-1}) = R_z(\vartheta)$ and (39) reduces to (12). The details are given in Appendix B. $\square$

Remark IV.11 (Geometric meaning of the spin rotation phase). The rotation R_ρ^{-1} carries the reference frame $s(x)$ at x to a frame at $R_\rho^{-1}x$ that differs from the chosen frame there by $h(x, R_\rho^{-1}) = R_z(\vartheta)$ (Fig. 4). The phase $e^{-is\vartheta}$ is the character $\rho^{(s)}$ evaluated on the inverse twist. This residual frame rotation is already described in the spin-wavelet treatment [12] and by Boyle [47]; Proposition IV.10 connects those conventions to the induced-action formula of Cohen, Geiger & Weiler [36]. The common harmonic law (10) then follows because Λ_s intertwines $\mathcal{R}_\rho$ with left translation on $\text{SO}(3)$, which acts on the first Wigner index by the same D^ℓ for every spin.

E. Fibre decomposition and azimuthal band-limits

The decomposition in Lemma IV.7(a) expresses every function on the rotation group as a sum of lifted spin fields.

Corollary IV.12 (Functions on $\text{SO}(3)$ as tuples of spin fields). *Every $F \in L^2(\text{SO}(3))$ decomposes uniquely as $F = \sum_{s \in \mathbb{Z}} \Lambda_s \Lambda_s^\dagger F$, its Fourier series in γ converging in L^2 , with $(\Lambda_s \Lambda_s^\dagger F)_{mn}^\ell = F_{mn}^\ell \delta_{n,-s}$; the component $\Lambda_s^\dagger F$ is a spin- s field whose harmonic coefficients are the Wigner column $F_{m,-s}^\ell$ up to the constant in (38). In particular F is azimuthally band-limited at N , $F_{mn}^\ell = 0$ for $|n| \geq N$, if and only if $F = \sum_{|s| < N} \Lambda_s f_s$ for spin fields $f_s = \Lambda_s^\dagger F$.*

Consequently, a function on the rotation group azimuthally band-limited at N is represented by a tuple with at most $2N - 1$ spin slots, one for each $|s| < N$, some of which may vanish. At degree ℓ there are only $2 \min(\ell, N - 1) + 1$ admissible slots. The Wigner transform of such a function is therefore a stack of at most $2N - 1$ spin spherical harmonic transforms. This is the identity by which Wigner transforms are computed through spin spherical harmonic transforms in the sampling theorem on the rotation group of McEwen *et al.* [45], as implemented in [43]. The identity is used in Sec. VI to identify directional convolutions with azimuthally band-limited, *i.e.* steerable, kernels with tuples of spin-changing convolutions, and in Sec. VII to decompose the anisotropic constructions of Sec. III F into their spin components. Figure 5 (bottom) shows the column structure of a lifted field and Fig. 6 the part of a kernel read by each of the operators defined in the next section.

V. THE GENERALISED LIFTED SPHERICAL CONVOLUTION

We now define the construction that unifies the convolutions of Sec. III. We first define the generalised lifted spherical convolution in its four forms, before deriving its harmonic form and equivariance. The constraint on the kernel under which the projection discards nothing is then characterised, the closure of the construction under composition is established, and finally all rotation-equivariant linear maps between spin fields and functions on the rotation group are classified, with a corollary on the incompatibility of directional kernels, spherical output and equivariance.

A. Definition

Definition V.1 (Generalised lifted convolution). For a kernel $\Psi \in L^2(\text{SO}(3))$ and spins $s, s' \in \mathbb{Z}$ define the four operators

$$\begin{aligned} \mathcal{C}_{s \rightarrow s'}[\Psi] &= \Lambda_{s'}^\dagger (\Psi \star \Lambda_s \cdot) : L_s^2(S^2) \rightarrow L_{s'}^2(S^2), \\ \mathcal{C}_{s \rightarrow G}[\Psi] &= \Psi \star \Lambda_s \cdot : L_s^2(S^2) \rightarrow L^2(\text{SO}(3)), \\ \mathcal{C}_{G \rightarrow s'}[\Psi] &= \Lambda_{s'}^\dagger (\Psi \star \cdot) : L^2(\text{SO}(3)) \rightarrow L_{s'}^2(S^2), \\ \mathcal{C}_{G \rightarrow G}[\Psi] &= \Psi \star \cdot : L^2(\text{SO}(3)) \rightarrow L^2(\text{SO}(3)), \end{aligned} \quad (41)$$

with $\star$ the group convolution of Definition III.4.

The first operator lifts a spin- s field to $\text{SO}(3)$, convolves it there and projects the result to a spin- s' field; there is no requirement that $s' = s$. The other three omit the lift, the projection, or both, and are included so that convolutions with input or output on the rotation group are covered by the same statements. Figure 5 shows the construction and its harmonic-space form, which we now derive.

B. Harmonic form and equivariance

Theorem V.2 (Harmonic form). *Let $\Psi \in L^2(\text{SO}(3))$, $s, s' \in \mathbb{Z}$, and let $\tilde{\Psi}(g) = \Psi^*(g^{-1})$. For $|m|, |n| \leq \ell$, with $\ell \geq |s|$ whenever the input is a spin- s field and $\ell \geq |s'|$ whenever the*




FIG. 5. The generalised lifted spherical convolution $\mathcal{C}_{s \rightarrow s'}[\Psi] = \Lambda_{s'}^\dagger (\Psi \star \Lambda_s \cdot)$, drawn for $s = 2$ and $s' = 1$. Top: a spin- s field on the sphere is lifted to a type- s Mackey function on $\text{SO}(3)$, convolved on $\text{SO}(3)$ and projected to a spin- s' field. Cones are the $\text{SO}(2)$ fibres, with the red dot at $\gamma = 0$; rims are coloured by the phase of the function on the fibre, which winds as $e^{-is\gamma}$ after the lift and mixes all windings after the convolution. Bottom: the same at one degree ℓ . The lift fills the Wigner column $n = -s$, the convolution multiplies by the conjugate transpose of the kernel block, and the projection reads the column $n = -s'$, so that only the entry $\Psi_{-s', -s}^\ell$ reaches the output, (43).

output is a spin- s' field,

$$(\mathcal{C}_{s \rightarrow G}[\Psi] {}_s f)_{mn}^\ell = (-1)^s \sqrt{\frac{8\pi^2}{2\ell+1}} {}_s f_{\ell m} \Psi_{n, -s}^{\ell*}, \quad (42)$$

$$(\mathcal{C}_{s \rightarrow s'}[\Psi] {}_s f)_{\ell m} = (-1)^{s+s'} {}_s f_{\ell m} \Psi_{-s', -s}^{\ell*}, \quad (43)$$

$$(\mathcal{C}_{G \rightarrow s'}[\Psi] F)_{\ell m} = (-1)^{s'} \sqrt{\frac{2\ell+1}{8\pi^2}} \sum_{m'=-\ell}^{\ell} F_{mm'}^\ell \Psi_{-s', m'}^{\ell*}, \quad (44)$$

$$(\mathcal{C}_{G \rightarrow G}[\Psi] F)_{mn}^\ell = \sum_{m'=-\ell}^{\ell} F_{mm'}^\ell \Psi_{nm'}^{\ell*}, \quad (45)$$

and the output of $\mathcal{C}_{s \rightarrow s'}[\Psi]$ has no component of degree $\ell < \max(|s|, |s'|)$. All four operators are bounded and $\text{SO}(3)$ -equivariant for every Ψ : $\mathcal{C}_{s \rightarrow s'}[\Psi] \pi_s(g) = \pi_{s'}(g) \mathcal{C}_{s \rightarrow s'}[\Psi]$ and correspondingly with π_G in place of π_s or $\pi_{s'}$. Their adjoints are $(\mathcal{C}_{s \rightarrow G}[\Psi])^\dagger = \mathcal{C}_{G \rightarrow s}[\tilde{\Psi}]$, $(\mathcal{C}_{s \rightarrow s'}[\Psi])^\dagger = \mathcal{C}_{s' \rightarrow s}[\tilde{\Psi}]$ and $(\mathcal{C}_{G \rightarrow G}[\Psi])^\dagger = \mathcal{C}_{G \rightarrow G}[\tilde{\Psi}]$.

Proof. By Lemma IV.7(c) the Wigner coefficients of the lifted field are supported on the single column $m' = -s$, so that the sum over m' in (20) collapses to a single term and (42) follows; applying $\Lambda_{s'}^\dagger$, by (38) reads the column $n = -s'$ and multiplies by $(-1)^{s'} \sqrt{(2\ell+1)/8\pi^2}$, so that the degree-dependent constants cancel and (43) follows, while (44) is $\Lambda_{s'}^\dagger$ applied to (20). Equivariance follows since Λ_s , $\Lambda_{s'}^\dagger$ and $\Psi \star$ each intertwine the relevant representations, by Lemma IV.7(e) and the left invariance of the Haar measure, and the adjoint identities follow from Fubini's theorem, the adjoint of $\Psi \star$ being $\tilde{\Psi} \star$. The details are given in Appendix B. $\square$

Equation (43) is the central formula of the article: a spin- s to spin- s' lifted convolution multiplies each spin- s harmonic coefficient by a single Wigner coefficient of the kernel, conjugated and up to a sign, $(-1)^{s+s'} \Psi_{-s',-s}^{\ell*}$, with no degree-dependent constant. The equivariant families listed in Table I arise from the different ways in which the literature identifies spherical and group kernels, as we show in Sec. VI. Figure 6 in Sec. VI shows the part of the kernel block read by each of the four operators.

Proposition V.3 (Composition). *The four families of Definition V.1 are closed under composition, and composing multiplies kernel blocks: for kernels $\Psi_1, \Psi_2 \in L^2(\text{SO}(3))$,*

$$\mathcal{C}_{G \rightarrow G}[\Psi_2] \mathcal{C}_{G \rightarrow G}[\Psi_1] = \mathcal{C}_{G \rightarrow G}[\Psi_3], \quad \Psi_3^\ell = \Psi_2^\ell \Psi_1^\ell, \quad (46)$$

and likewise $\mathcal{C}_{G \rightarrow s''}[\Psi_2] \mathcal{C}_{s \rightarrow G}[\Psi_1] = \mathcal{C}_{s \rightarrow s''}[\Psi_3]$, $\mathcal{C}_{G \rightarrow G}[\Psi_2] \mathcal{C}_{s \rightarrow G}[\Psi_1] = \mathcal{C}_{s \rightarrow G}[\Psi_3]$ and $\mathcal{C}_{G \rightarrow s''}[\Psi_2] \mathcal{C}_{G \rightarrow G}[\Psi_1] = \mathcal{C}_{G \rightarrow s''}[\Psi_3]$ with the same Ψ_3 . When the intermediate stage is a spin- s' field only the intermediate index $-s'$ is retained: $\mathcal{C}_{s' \rightarrow s''}[\Psi_2] \mathcal{C}_{s \rightarrow s'}[\Psi_1] = \mathcal{C}_{s \rightarrow s''}[\Psi_3]$ with $(\Psi_3)_{-s'',-s}^\ell = (\Psi_2)_{-s'',-s'}^\ell (\Psi_1)_{-s',-s}^\ell$, so that the multipliers of (43) multiply, $c_\ell^{(3)} = c_\ell^{(2)} c_\ell^{(1)}$; correspondingly $\mathcal{C}_{s' \rightarrow G}[\Psi_2] \mathcal{C}_{s \rightarrow s'}[\Psi_1] = \mathcal{C}_{s \rightarrow G}[\Psi_3]$ with $(\Psi_3)_{n,-s}^\ell = (\Psi_2)_{n,-s'}^\ell (\Psi_1)_{-s',-s}^\ell$.

Proof. By (20), $(\Psi_2 \star (\Psi_1 \star F))^\ell = F^\ell \Psi_1^{\ell\dagger} \Psi_2^\ell = F^\ell (\Psi_2^\ell \Psi_1^\ell)^\dagger$, which is (46); inserting $\Lambda_{s'} \Lambda_{s'}^\dagger$ between the two convolutions, whose action in Wigner space is the projector onto the column $n = -s'$ by Corollary IV.12, retains only the intermediate index $-s'$, and the remaining cases follow by omitting the outer lift or projection. See Appendix B. $\square$

Consequently, consecutive sphere-to-sphere convolutions require no intermediate transforms, since their multipliers simply multiply, whereas an intermediate stage on the rotation group between two spin fields sums over all intermediate spins, which is the content of Proposition VI.6 below. The forthcoming article [29] relies on this rule to stack layers.

C. Lossless projection and the kernel constraint

Only one column of Ψ is read by $\mathcal{C}_{s \rightarrow G}[\Psi]$ and only one entry by $\mathcal{C}_{s \rightarrow s'}[\Psi]$. The projection discards nothing precisely when the intermediate function on $\text{SO}(3)$ is already a Mackey function of the target type.

Corollary V.4 (Lossless projection and kernel constraint). *The following are equivalent:*

- (i) $\Psi \star \Lambda_s f \in \mathcal{M}_{s'}$ for all $f \in L_s^2(S^2)$;
- (ii) $\Psi_{n,-s}^\ell = 0$ for all $n \neq -s'$ and all $\ell \geq |s|$;
- (iii) the type- s right-isotypic part $\Psi^{(s)}(g) = \frac{1}{2\pi} \int_0^{2\pi} \Psi(gR_z(\gamma)) e^{is\gamma} d\gamma$ of the kernel satisfies the two-sided constraint

$$\Psi^{(s)}(R_z(\chi_1) g R_z(\chi_2)) = e^{-is'\chi_1} \Psi^{(s)}(g) e^{-is\chi_2} \quad \text{for all } \chi_1, \chi_2. \quad (47)$$

In that case $\mathcal{C}_{s \rightarrow s'}[\Psi]f = \sqrt{2\pi} (\Psi \star \Lambda_s f) \circ s$, the constrained part of the kernel is

$$\Psi^{(s)}(\alpha, \beta, \gamma) = \sum_{\ell \geq \max(|s|, |s'|)} \frac{2\ell+1}{8\pi^2} \Psi_{-s',-s}^\ell e^{-is'\alpha} d_{-s',-s}^\ell(\beta) e^{-is\gamma}, \quad (48)$$

where the series converges in L^2 . For every kernel lifted from a spin- s field, $\Psi = \Lambda_s \psi$, whether or not projection is lossless,

$$(\mathcal{C}_{s \rightarrow s'}[\Lambda_s \psi] s f)_{\ell m} = (-1)^{s'} \sqrt{\frac{8\pi^2}{2\ell+1}} s f_{\ell m} s \psi_{\ell, -s'}^*. \quad (49)$$

For such a lifted kernel, losslessness is equivalent to $s \psi_{\ell m} = s \psi_{\ell, -s'} \delta_{m, -s'}$.

Proof. The equivalence of (i) and (ii) follows from (42) and the characterisation of $\mathcal{M}_{s'}$ as the column $n = -s'$ of Wigner space (Lemma IV.7); that of (ii) and (iii) by expanding $\Psi^{(s)}$ in Wigner functions and noting that $D_{mn}^{\ell*}(R_z(\chi_1)gR_z(\chi_2)) = e^{im\chi_1}D_{mn}^{\ell*}(g)e^{in\chi_2}$. The evaluation formula is Lemma IV.7(b), (48) is the Wigner expansion restricted to the entry $(-s', -s)$, and (49) is (43) with Lemma IV.7(c). See Appendix B. $\square$

D. Classification of equivariant maps

Theorem V.2 shows that every generalised lifted spherical convolution between spin fields is a multiplier in spin spherical harmonic space. We now prove the converse.

Theorem V.5 (Classification). *Let $T : L_s^2(S^2) \rightarrow L_{s'}^2(S^2)$ be a bounded linear map with $T\pi_s(g) = \pi_{s'}(g)T$ for all $g \in \text{SO}(3)$. Then there is a unique bounded sequence $(c_\ell)_{\ell \geq \max(|s|, |s'|)}$ such that*

$$(T {}_s f)_{\ell m} = c_\ell {}_s f_{\ell m} \quad \text{for all } \ell \geq \max(|s|, |s'|), \quad |m| \leq \ell, \quad (50)$$

and T annihilates the input components with degrees absent from the target; conversely every bounded sequence defines such a T , with $\|T\| = \sup_\ell |c_\ell|$. Moreover:

- (a) $T = \mathcal{C}_{s \rightarrow s'}[\Psi]$ for some $\Psi \in L^2(\text{SO}(3))$ if and only if $\sum_\ell (2\ell + 1)|c_\ell|^2 < \infty$, i.e. if and only if T is Hilbert–Schmidt. The kernel is not unique, but the representative whose entire kernel satisfies the two-sided law (47), with Ψ in place of $\Psi^{(s)}$, is unique: $\Psi_{T-s', -s}^\ell = (-1)^{s+s'} c_\ell^*$ and all other Wigner coefficients are zero.
- (b) Every bounded T is $\mathcal{C}_{s \rightarrow s'}[\Psi_T]$ in the sense of Fourier multipliers, with Ψ_T the distribution on $\text{SO}(3)$ with the Wigner coefficients of (a).
- (c) Set $d_\ell = 2\ell + 1$ and $a_\ell = d_\ell/(8\pi^2)$. The equivariant maps of the other three types act on coefficients as

$$(T_{s \rightarrow G} f)_{mn}^\ell = v_n^\ell {}_s f_{\ell m}, \quad (T_{G \rightarrow s'} F)_{\ell m} = \sum_k v_k^\ell F_{mk}^\ell, \quad (T_{G \rightarrow G} F)^\ell = F^\ell B^\ell. \quad (51)$$

They are bounded if and only if the respective bounds

$$\begin{aligned} \|T_{s \rightarrow G}\| &= \sup_{\ell \geq |s|} \sqrt{a_\ell} \|u^\ell\|_2 < \infty, \\ \|T_{G \rightarrow s'}\| &= \sup_{\ell \geq |s'|} \frac{\|v^\ell\|_2}{\sqrt{a_\ell}} < \infty, \quad \|T_{G \rightarrow G}\| = \sup_{\ell \geq 0} \|B^\ell\|_{\text{op}} < \infty \end{aligned} \quad (52)$$

hold. Degrees absent from the input or output spin space contribute zero. These maps are realised by the column, row and full block of a kernel, respectively, with

$$u_n^\ell = (-1)^s a_\ell^{-1/2} \Psi_{n, -s}^{\ell*}, \quad v_k^\ell = (-1)^{s'} a_\ell^{1/2} \Psi_{-s', k}^{\ell*}, \quad B^\ell = \Psi^{\ell\dagger}. \quad (53)$$

Setting all unread entries to zero gives a canonical distributional kernel, which is square integrable if and only if the map is Hilbert–Schmidt. The corresponding summability conditions are given in Appendix B 5.

Proof. By Frobenius reciprocity the multiplicity of the degree- ℓ irreducible representation of $\text{SO}(3)$ in $\text{Ind}_H^G \rho^{(s)}$ equals the multiplicity of $\rho^{(s)}$ in its restriction to $\text{SO}(2)$, which is one if $\ell \geq |s|$ and zero otherwise, so that $L_s^2(S^2)$ decomposes without multiplicity into the subspaces spanned by the ${}_s Y_{\ell m}$, $|m| \leq \ell$, on which $\text{SO}(3)$ acts by the same matrices D^ℓ for every spin by (10). Schur’s lemma then forces T to map each such subspace to its counterpart of the same degree by a scalar c_ℓ , and to annihilate degrees absent from the target; boundedness of T is equivalent to $\sup_\ell |c_\ell| < \infty$. Part (a) follows by comparing (43)

with (50) and computing $\|\Psi_T\|^2 = \sum_\ell \frac{2\ell+1}{8\pi^2} |c_\ell|^2$, part (b) by extending $\Psi_\star$ to distributional kernels with bounded Wigner coefficients, and part (c) by the same argument applied to the left regular representation, in which the degree- ℓ representation occurs with multiplicity $2\ell+1$, one copy per column. Parseval's identity weights each Wigner block by a_ℓ , whereas the spin harmonic coefficients have unit weight, giving the norms (52). The details are given in Appendix B. $\square$

For scalar fields Theorem V.5 is Theorem 1 and Lemma 10 of Kondor & Trivedi [35] and, for $s = s' = 0$, the characterisation of isotropic operators of Kennedy *et al.* [20]; for general types it is contained in Theorems 3.2–3.4 of Cohen, Geiger & Weiler [36], which for the sphere describe the kernels as functions on $[0, \pi]$ unconstrained except at the poles, and in Proposition 1 and Corollary 1 of Xu *et al.* [37]. Xu *et al.* give the character lift, the single-column harmonic representation, the single-entry input/output kernel and a completeness argument explicitly, while Esteves *et al.* [15] parameterise spin-changing multipliers directly. Here we collect the four cases in the conventions of Sec. II, including the degree ranges, the exact norm bounds and the distinction between square integrable and distributional kernels. The identity already requires the distributional interpretation of part (b).

The differential operators of Sec. VI E require a separate extension beyond bounded maps. For a possibly unbounded sequence (c_ℓ) , set $c_\ell = 0$ on input degrees absent from the target and define

$$\mathcal{D}(T) = \left\{ f \in L_s^2(S^2) : \sum_{\ell m} |c_\ell|^2 |f_{\ell m}|^2 < \infty \right\}. \quad (54)$$

The coefficient formula (50) defines a closed, densely defined operator on this rotation-invariant domain and commutes there with rotations. Polynomial growth of c_ℓ gives a distributional kernel; an arbitrary sequence still acts on finite harmonic sums but need not define a distribution. For $\tilde{\partial}$ and $\tilde{\partial}$ the domain is the first-order spin Sobolev space, and for the second-order operators below it is the second-order space, with the finitely many zero low-degree multipliers imposing no additional restriction.

Corollary V.6 (Equivariant projections and lifts). *Every bounded $\text{SO}(3)$ -equivariant linear map $L^2(\text{SO}(3)) \rightarrow L_{s'}^2(S^2)$ is $\mathcal{C}_{G \rightarrow s'}[\Psi]$ for a kernel (in general a distribution) whose only nonzero row is $n = -s'$ in each degree $\ell \geq |s'|$. If the map acts by $(TF)_{\ell m} = \sum_{m'} c_{m'}^\ell F_{mm'}^\ell$, then $c_{m'}^\ell = (-1)^{s'} \sqrt{(2\ell+1)/(8\pi^2)} \Psi_{-s', m'}^{\ell*}$, subject to (52). The adjoint projection $\Lambda_s^\dagger$ is the case $\Psi_{-s', m'}^{\ell*} = \delta_{m', -s'}$, which reads one Wigner column. Dually, every bounded equivariant map $L_s^2(S^2) \rightarrow L^2(\text{SO}(3))$ is $\mathcal{C}_{s \rightarrow G}[\Psi]$ for a kernel with a single column $m' = -s$, and Λ_s is the case $\Psi_{n, -s}^{\ell*} = \delta_{n, -s}$.*

Proof. This is Theorem V.5(c) with the normalisations of Lemma IV.7(c). $\square$

The learned projecting convolution of the forthcoming article [29] uses the family in Corollary V.6.

Corollary V.7 (Incompatibility of directional kernels, spherical output and equivariance). *The kernel, in general a distribution, of a bounded equivariant map from one fixed input spin to one fixed output spin, $T : L_s^2(S^2) \rightarrow L_{s'}^2(S^2)$, is*

$$K(x, y) = (-1)^{s+s'} \sum_\ell c_\ell \frac{2\ell+1}{4\pi} D_{-s', -s}^\ell(s(x)^{-1}s(y)), \quad (55)$$

a radial kernel with the frame phases forced by (36); for $s = s' = 0$ it is $\sum_\ell c_\ell (2\ell+1) P_\ell(x \cdot y)/(4\pi)$. Apart from the angular characters prescribed by the input and output spins, such a map cannot retain independent directional kernel content. Thus, beyond the angular characters prescribed by the spins, directional kernels, a fixed-spin output on the sphere and equivariance cannot be achieved simultaneously for a fixed input spin.




FIG. 6. The part of the kernel block $\Psi_{nm'}^\ell$ read by each of the four operators of Definition V.1 at one degree ℓ : a single entry $(-s', -s)$ for sphere to sphere, the column $m' = -s$ for sphere to $SO(3)$, the row $n = -s'$ for $SO(3)$ to sphere, and the full block for $SO(3)$ to $SO(3)$, together with the convolutions of this section that each case covers and the costs of Sec. VIII.

TABLE II. Taxonomy of spherical convolutions as generalised lifted spherical convolutions. Kernel support is the set of Wigner coefficients $\Psi_{nm'}^\ell$ of the kernel that are read; $\Lambda_s \psi$ denotes a kernel lifted from a spin- s field ψ on the sphere.

Convolution (Prop.)	Operator	Input $\rightarrow$ output	Kernel support
Directional, scalar (VI.1)	$\mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]$	$L_0^2(S^2) \rightarrow SO(3)$	col. $m' = 0$
Directional, spin (VI.1)	$\mathcal{C}_{s \rightarrow G}[\Lambda_s \psi]$	$L_s^2(S^2) \rightarrow SO(3)$	col. $m' = -s$
Synthesis, transpose (VI.2)	$\mathcal{C}_{G \rightarrow s}[\Lambda_s \psi]$	$L^2(SO(3)) \rightarrow L_s^2(S^2)$	row $n = -s$
Axisymmetric (VI.3)	$(2\pi)^{-1/2} \mathcal{C}_{0 \rightarrow 0}[\Lambda_0 \psi]$	$L_0^2(S^2) \rightarrow L_0^2(S^2)$	entry $(0, 0)$
Left (VI.4)	$\mathcal{C}_{0 \rightarrow 0}[\Psi_h]$	$L_0^2(S^2) \rightarrow L_0^2(S^2)$	entry $(0, 0)$
Spin $s \rightarrow s'$ (VI.5)	$\mathcal{C}_{s \rightarrow s'}[\Psi]$	$L_s^2(S^2) \rightarrow L_{s'}^2(S^2)$	entry $(-s', -s)$
Band-limited directional (VI.6)	$\bigoplus_{ s' < N} \mathcal{C}_{s \rightarrow s'}[\Psi]$	$L_s^2(S^2) \rightarrow \bigoplus_{ s' < N} L_{s'}^2(S^2)$	col. $-s$, $ n < N$
Group convolution (VI.7)	$\mathcal{C}_{G \rightarrow G}[\Psi]$	$L^2(SO(3)) \rightarrow L^2(SO(3))$	all
$\bar{\partial}$, E/B , Kaiser–Squires (VI.8)	$\mathcal{C}_{s \rightarrow s'}[\Psi_T]$	$L_s^2(S^2) \rightarrow L_{s'}^2(S^2)$	entry $(-s', -s)$, distrib.

Proof. By (50), $K(x, y) = \sum_\ell c_\ell \sum_{m, s'} Y_{\ell m}(x) {}_s Y_{\ell m}^*(y)$, and the sum over m is evaluated by the addition theorem for spin spherical harmonics, which follows from (9) and $D^\ell(s(x)^{-1} s(y)) = D^\ell(s(x))^\dagger D^\ell(s(y))$; the second Euler angle of $s(x)^{-1} s(y)$ is the geodesic distance between x and y , and the first and third are the phases forced by the frames at x and y . See Appendix B. $\square$

Corollary V.7 adds an equivariance qualification to the criteria of Roddy & McEwen [32]: between two fixed spins the kernel is radial with prescribed frame phases, not necessarily an axisymmetric spin filter in the sense of Definition III.6. For example, a scalar filter $Y_{\ell, -1}$ can produce an equivariant spin-1 output; its prescribed angular character is not independent directionality. Retaining additional directional content from a fixed input spin requires a tuple of output spins, equivalently a function on the rotation group. The restriction does not apply to a group-to-spin synthesis, whose input already contains several spin components.

VI. REDUCTION TO EXISTING SPHERICAL CONVOLUTIONS

In this section we recover the rotation-equivariant convolutions treated in Sec. III as generalised lifted spherical convolutions. For each family we state the input, the part of the Wigner block $\Psi_{nm'}^\ell$ of the kernel that is read, and the output, and we recover the harmonic form together with its constant. The results are collected in Table II, and Fig. 6 shows the part of the kernel block that each operator reads. Throughout, $\Lambda_s \psi$ denotes a kernel lifted from a spin- s field ψ on the sphere; by Lemma IV.7(c) such a kernel has Wigner support in the single column $m' = -s$. The reductions follow by direct calculation from Theorem V.2 and Lemma IV.7; proofs not immediate from the harmonic formulas are supplied in Appendix B.

A. Directional convolutions and synthesis

We begin with the directional convolution, which is the lifted convolution with no projection and a kernel lifted from the sphere.

Proposition VI.1 (Directional convolutions). *For ${}_sf, {}_s\psi \in L_s^2(S^2)$ (including $s = 0$),*

$${}_sf \otimes {}_s\psi = \mathcal{C}_{s \rightarrow G}[\Lambda_s \psi] {}_sf, \quad ({}_sf \otimes {}_s\psi)_{mn}^\ell = \frac{8\pi^2}{2\ell+1} {}_sf_{\ell m} {}_s\psi_{\ell n}^*, \quad (56)$$

exactly, with no residual constant. The kernel $\Lambda_s \psi$ has Wigner support in the column $m' = -s$.

Proof. By the intertwining relation of Lemma IV.7(e) and the isometry of the lift, $(\Lambda_s \psi \star \Lambda_s f)(g) = \langle \Lambda_s f, \pi_G(g) \Lambda_s \psi \rangle = \langle \Lambda_s f, \Lambda_s \pi_s(g) \psi \rangle = \langle f, \pi_s(g) \psi \rangle$, which is Definition III.5 since $\pi_s(g) = \mathcal{R}_g$ by Proposition IV.10; the harmonic form is (42) with Lemma IV.7(c). $\square$

Proposition VI.1 covers the directional convolutions of Wandelt & Górski [19], Challinor *et al.* [21], McEwen *et al.* [9, 12, 24] and Wiaux *et al.* [22, 23], and the S^2 correlation of Cohen *et al.* [5] for real signals with the measure normalised to unity. Note that the absence of a residual constant is a consequence of the isometric normalisation of the lift: the directional convolution on the sphere *is* the group convolution of the lifts.

Proposition VI.2 (Synthesis and transpose convolution). *For $W \in L^2(\text{SO}(3))$ and ${}_s\psi \in L_s^2(S^2)$,*

$$W \otimes {}_s\psi^\dagger = \int_{\text{SO}(3)} W(\rho) (\pi_s(\rho) {}_s\psi) d\mu(\rho) = \mathcal{C}_{G \rightarrow s}[\widetilde{\Lambda_s \psi}] W = (\mathcal{C}_{s \rightarrow G}[\Lambda_s \psi])^\dagger W, \quad (57)$$

$$(W \otimes {}_s\psi^\dagger)_{\ell m} = \sum_n W_{mn}^\ell {}_s\psi_{\ell n}. \quad (58)$$

Synthesis is the adjoint of analysis, and the flipped kernel $\widetilde{\Lambda_s \psi}$ has Wigner support in the row $n = -s$.

Proof. Since $\pi_s(\rho) \psi = \Lambda_s^\dagger \pi_G(\rho) \Lambda_s \psi$ and $\Lambda_s^\dagger$ is bounded, it may be taken outside the integral, and $\int W(\rho) (\pi_G(\rho) \Phi)(g) d\mu(\rho) = (\Phi \star W)(g)$ since $\Phi^*(g^{-1}\rho) = \Phi(\rho^{-1}g)$; the harmonic form follows from (44) and the adjoint identity from Theorem V.2. $\square$

Proposition VI.2 recovers the Haar-measure synthesis step of the directional wavelet transforms of [9, 10, 12, 24], whose synthesis kernels compensate the constant $8\pi^2/(2\ell+1)$ of (23) either through the dual kernel $\frac{2\ell+1}{8\pi^2} \psi_{\ell n}$ [10, 24] or through the admissibility condition [12], so that the sum over scales of analysis followed by synthesis is the identity. Consequently, the admissibility condition (27) is the statement that the sum of the operators $A_j^\dagger A_j$ over the scales, with A_j the analysis operators of Proposition VI.1 and Proposition VI.3 below, is the identity, *i.e.* that the wavelets form a Parseval frame.

B. Axisymmetric convolutions

Proposition VI.3 (Axisymmetric convolution). *For axisymmetric ψ , $\psi_{\ell m} = \psi_{\ell 0} \delta_{m0}$,*

$$f \odot \psi = \frac{1}{\sqrt{2\pi}} \mathcal{C}_{0 \rightarrow 0}[\Lambda_0 \psi] f, \quad (f \odot \psi)_{\ell m} = \sqrt{\frac{4\pi}{2\ell+1}} f_{\ell m} \psi_{\ell 0}^*, \quad (59)$$

and $\Lambda_0 \psi$ satisfies the constraint (47) with $s = s' = 0$. The same holds for the axisymmetric spin convolution of Definition III.6, which is $(2\pi)^{-1/2} \mathcal{C}_{s \rightarrow 0}[\Lambda_s \Phi]$ with harmonic form (25). For a non-axisymmetric ψ the left-hand side is the standard correlation of Definition III.7, which is not of the form $\mathcal{C}_{0 \rightarrow 0}$ (Sec. VII).

Proof. By Proposition VI.1 evaluated at the section, $\langle f, \mathcal{R}_\omega \psi \rangle = (\mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]f)(s(\omega))$, and since the output lies in $\mathcal{M}_0$ when ψ is axisymmetric (Corollary V.4), Lemma IV.7(b) gives $(\mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]f) \circ s = (2\pi)^{-1/2} \Lambda_0^\dagger \mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]f$; the constant follows from (49). The spin case is identical with $s' = 0$. $\square$

Proposition VI.3 recovers the axisymmetric convolution of [9, 20, 32], the axisymmetric wavelets and needlets of [11, 48], and the scaling coefficients of the spin wavelet transform of [12], which are a spin $s \rightarrow 0$ convolution. The factor $(2\pi)^{-1/2}$ is the difference between evaluating a Mackey function on the section and projecting it with the adjoint.

Proposition VI.4 (Left convolution). *For any $h \in L^2(S^2)$ let $\Psi_h(g) = h^*(g^{-1}\eta)$. Then*

$$f \star_L h = \mathcal{C}_{0 \rightarrow 0}[\Psi_h]f, \quad (f \star_L h)_{\ell m} = 2\pi \sqrt{\frac{4\pi}{2\ell+1}} f_{\ell m} h_{\ell 0}. \quad (60)$$

Only the entry $(\Psi_h)_{00}^\ell$ is read, so that only the zonal part of h enters; the factor 2π is the fibre volume that appears because (15) lifts f and h without the normalisation of Definition IV.6. The marginalisation of the bilinear directional correlation over the third Euler angle used in [6, 20] is $\sqrt{2\pi} \Lambda_0^\dagger$. Applied instead to the sesquilinear correlation of Definition III.3, it conjugates the spherical kernel.

Proof. By direct substitution, $(\Psi_h \star \Lambda_0 f)(g) = (2\pi)^{-1/2} \int h(g'^{-1}g\eta) f(g'\eta) d\mu(g')$; applying $\Lambda_0^\dagger$ and noting that $R_z(\gamma)\eta = \eta$ gives (15), and the constant follows by evaluating $(\Psi_h)_{00}^\ell$ with (8). See Appendix B. $\square$

Proposition VI.4 recovers Theorem 3 of Kennedy *et al.* [20], with their factor of $1/2\pi$, the convolution of Esteves *et al.* [6] and the convolution theorem of Bonev *et al.* [4]. The two convolutions in Definitions III.1 and III.2 read the same kernel entry but differ in both normalisation and kernel conjugation: the axisymmetric correlation is sesquilinear, whereas the left convolution is bilinear, with the conjugation absorbed in $\Psi_h(g) = h^*(g^{-1}\eta)$.

C. Spin-changing convolutions and steerable kernels

Proposition VI.5 (Spin-changing convolutions). *Every bounded multiplier ${}_s f_{\ell m} = c_\ell {}_s f_{\ell m}$, for $\ell \geq \max(|s|, |s'|)$ and annihilating input degrees absent from the target, is $\mathcal{C}_{s \rightarrow s'}[\Psi_T]$ in the Fourier-multiplier sense of Theorem V.5(b), with $\Psi_T^\ell_{-s', -s} = (-1)^{s+s'} c_\ell^*$. The canonical kernel is generally a distribution and belongs to L^2 exactly when $\sum_\ell (2\ell+1) |c_\ell|^2 < \infty$. In particular:*

- (a) *the spin-weighted convolution of Esteves et al. [15], Eq. (13), with a single input spin s and output spin s' has $c_\ell = {}_s k_s^\ell$ in their notation, and their multi-spin convolution sums over input spins and stacks output spins. Their cross-correlation, Eq. (15), $g_{\ell m}^t = \sum_i f_{\ell m}^i k_{\ell t}^{i*}$, instead corresponds to $\Psi_{-t, -i}^\ell = (-1)^{i+t} k_{\ell t}^i$;*
- (b) *the spin needlets of Geller & Marinucci [26, 27] are $\mathcal{C}_{s \rightarrow s}$ with $c_\ell = b(\sqrt{e_{\ell s}}/B^j)$, $e_{\ell s} = (\ell - |s|)(\ell + |s| + 1)$, and their mixed needlets are the corresponding $\mathcal{C}_{s \rightarrow 0}$ and $\mathcal{C}_{0 \rightarrow s}$; mixed needlets connect spin s with spin 0 only;*
- (c) *write $s(x)^{-1}s(y) = R_z(a)R_y(b)R_z(c)$. In the translated canonical angles $(a, b, -c)$, the radial-with-phases construction of Tessore et al. [28] has kernel $e^{iS^a} K(b) e^{isc}$, where $K(b) = \sum_\ell (\ell+1) K_\ell d_{S^s}^\ell(b)/(4\pi)$, and is $\mathcal{C}_{s \rightarrow S}[\Psi_T]$ with $c_\ell = K_\ell$ by (A10). The first angle differs by π from the source's printed quadrant convention: its literal angles are $\alpha = a + \pi$, $\alpha' = -c$, giving $c_\ell = (-1)^S K_\ell$ in our conventions with the same radial coefficients;*

- (d) the directional spin convolution of [12], Definition III.5, is not itself of this form, its output being a function on $\text{SO}(3)$, but for a steerable kernel, as the directional spin wavelets are, it is by Proposition VI.6 below the tuple of the spin-changing convolutions $\mathcal{C}_{s \rightarrow s'}[\Lambda_s \psi]$ over $|s'| < N$, with $c_\ell = (-1)^{s'} \sqrt{8\pi^2/(2\ell+1)} {}_s\psi_{\ell, -s'}^*$ as in (49): the azimuthal order $-s'$ of the kernel is read to produce the output of spin s' , so that the directional spin convolution computes every spin-changing convolution of its kernel at once, and its axisymmetric case, the scaling coefficients of Definition III.6, is $(2\pi)^{-1/2} \mathcal{C}_{s \rightarrow 0}[\Lambda_s \Phi]$, with multiplier $c_\ell = \sqrt{4\pi/(2\ell+1)} {}_s\Phi_{\ell 0}^*$ as in Proposition VI.3.
- (e) the spin spherical Fourier neural operator layer of the forthcoming article [29], which uses the same convention as here, learns the multiplier c_ℓ directly, as one harmonic weight per degree for each pair of spins (s, s') .

Proposition VI.5 makes precise the sense in which a single layer may map a spin- s field to a spin- s' field without forming a feature map on $\text{SO}(3)$: the operator is a multiplier, evaluated with one forward spin- s and one inverse spin- s' spherical harmonic transform, and by Theorem V.5 the family of such multipliers is the complete family of bounded equivariant linear maps between the fixed input and output spins. Forming the feature map on $\text{SO}(3)$, as the directional spin convolution does, is conversely the way to obtain all of the spin-changing convolutions of a kernel in one operation.

Proposition VI.6 (Azimuthally band-limited and steerable kernels). *If the column $m' = -s$ of Ψ is supported on $|n| < N$ then*

$$\mathcal{C}_{s \rightarrow G}[\Psi] = \sum_{|s'| < N} \Lambda_{s'} \mathcal{C}_{s \rightarrow s'}[\Psi], \quad (61)$$

i.e. the output is represented by at most $2N-1$ spin-changing components $\mathcal{C}_{s \rightarrow s'}[\Psi]$, $|s'| < N$, with $2\min(\ell, N-1) + 1$ admissible output slots at degree ℓ . Conversely, any finite tuple of bounded equivariant spin-changing maps from the same input spin assembles into a bounded map $\mathcal{C}_{s \rightarrow G}$ with a generally distributional kernel. A square integrable representative exists exactly when the sum of their Hilbert–Schmidt norms squared is finite. At finite band-limit the assembled kernel is always square integrable.

Proof. By (42) the Wigner coefficients of $\mathcal{C}_{s \rightarrow G}[\Psi]f$ vanish for $|n| \geq N$; apply Corollary IV.12 to the output and note that $\Lambda_s^\dagger \mathcal{C}_{s \rightarrow G}[\Psi] = \mathcal{C}_{s \rightarrow s'}[\Psi]$ by definition. $\square$

Proposition VI.6 expresses the finite-spin content behind the steerability result of Wiaux *et al.* [22]: a kernel with azimuthal band-limit N produces at most $2N-1$ spin components, each obtained by a spin-changing convolution on the sphere. The directional spin wavelets of [12] are steerable by construction and their coefficients at each scale form such a tuple. This also relates the $\mathcal{O}(NL^3)$ group layers of Cobb *et al.* [14] to multi-spin layers in the sense of Esteves *et al.* [15], after translating the kernel conventions. In particular, their cross-correlation is obtained at finite band-limit from projected directional convolutions by replacing their spherical kernels with $k_{\ell n}^i = (-1)^n \sqrt{(2\ell+1)/(8\pi^2)} k_{\ell, -n}^i$, rather than using those kernels unchanged.

D. Convolution on the rotation group

Proposition VI.7 (Convolution on the rotation group). *$\mathcal{C}_{G \rightarrow G}[\Psi]$ is the group convolution of Definition III.4, and by Theorem V.5(c) every bounded equivariant map $L^2(\text{SO}(3)) \rightarrow L^2(\text{SO}(3))$ is of this form with a distributional kernel.*

Proposition VI.7 covers the $\text{SO}(3)$ layer of Cohen *et al.* [5], written there for real signals with the measure normalised to unity and the convolution theorem in the form $\hat{f}\hat{\psi}^\dagger$, the correlation of Kostelec & Rockmore [25], the Clebsch–Gordan networks of Kondor *et al.*

[13], whose linear maps are the per-degree blocks of (45) acting across fragments, and the generalised convolution of Cobb *et al.* [14], whose channel mixing is the multi-channel version of (45). A lifted spin- s input reduces $\mathcal{C}_{G \rightarrow G}$ to $\mathcal{C}_{s \rightarrow G}$; it gives the traditional directional convolution of Proposition VI.1 after replacing the kernel by the lifted spherical kernel with the same relevant input column. At infinite band-limit this replacement may require a distributional spherical kernel. For an input supported on $|m'| < N_{\text{in}}$ and an effective kernel supported on $|n| < N_{\text{out}}$ in those input columns, Corollary IV.12 gives an array with at most $(2N_{\text{out}} - 1) \times (2N_{\text{in}} - 1)$ spin-changing entries. Setting $N_{\text{in}} = N_{\text{out}} = N$ gives the square array used in Sec. VIII.

E. Spin-changing operators of mathematical physics

The rotation-equivariant operators considered here are harmonic multipliers, with generally distributional kernels. Bounded examples fall under Theorem V.5; differential examples use the domain (54) instead.

Proposition VI.8 (Spin-changing operators of mathematical physics). *The following classical operators have the multiplier form (50), with (54) understood for unbounded examples:*

- (a) *the spin raising operator $\mathfrak{D} : \mathcal{D}(\mathfrak{D}) \rightarrow L_{s+1}^2(S^2)$, with $c_\ell = \sqrt{(\ell-s)(\ell+s+1)}$, and the spin lowering operator $\bar{\mathfrak{D}} : \mathcal{D}(\bar{\mathfrak{D}}) \rightarrow L_{s-1}^2(S^2)$, with $c_\ell = -\sqrt{(\ell+s)(\ell-s+1)}$, have the first-order spin Sobolev domains $\mathcal{D}(\mathfrak{D}), \mathcal{D}(\bar{\mathfrak{D}}) \subset L_s^2(S^2)$ specified by (54) [40, 41];*
- (b) *the E and B components of a spin- ± 2 polarisation field, $a_{E,\ell m} = -\frac{1}{2}(a_{2,\ell m} + a_{-2,\ell m})$ and $a_{B,\ell m} = \frac{i}{2}(a_{2,\ell m} - a_{-2,\ell m})$ [1, 2], each the sum of a $\mathcal{C}_{2 \rightarrow 0}$ and a $\mathcal{C}_{-2 \rightarrow 0}$ operator with constant c_ℓ ; the scalars $\bar{\mathfrak{D}}^2(Q+iU)$ and $\bar{\mathfrak{D}}^2(Q-iU)$ of [1] are the constituent $\mathcal{C}_{2 \rightarrow 0}$ and $\mathcal{C}_{-2 \rightarrow 0}$ maps with $c_\ell = \sqrt{(\ell+2)!/(\ell-2)!}$;*
- (c) *the spherical Kaiser–Squires relation of weak gravitational lensing between convergence and shear [3]: $\mathcal{C}_{0 \rightarrow 2}$ with $c_\ell = \mathcal{D}_\ell = -\sqrt{(\ell+2)(\ell-1)/(\ell(\ell+1))}$ and its inverse $\mathcal{C}_{2 \rightarrow 0}$ with $c_\ell = \mathcal{D}_\ell^{-1}$, for $\ell \geq 2$; the real-space forward map of [3] is synthesis with an axisymmetric spin-2 kernel, the adjoint of spin-2 to scalar analysis, rather than that forward analysis itself. Its two-point kernel is (55) with $(s, s') = (0, 2)$;*
- (d) *the all-Stokes beam response of Challinor *et al.* [21], i.e. their Eq. (39), corresponds to a sum of $\mathcal{C}_{0 \rightarrow G}$ maps on intensity I and circular polarisation V, and $\mathcal{C}_{\pm 2 \rightarrow G}$ maps on linear polarisation. For an axisymmetric co-polar beam, the intensity kernel has azimuthal order 0 and the linear-polarisation kernels have the prescribed orders ∓ 2 for output spins ± 2 . Their Eqs. (47)–(50) give scalar and spin-2 smoothing multipliers W_ℓ and ${}_2W_\ell$; the response in Eq. (46) depends on the third Euler angle through $\{1, \cos 2\psi, \sin 2\psi\}$, corresponding to the retained spins 0, ± 2 and band-limit $N = 3$ in Proposition VI.6.*

The identity, with $c_\ell \equiv 1$, illustrates the distributional-kernel clause of Theorem V.5. Spin raising instead has an unbounded polynomially growing multiplier and uses the separate domain (54); the second-order differential maps in part (b) have second-order spin Sobolev domains. None of these operators has a square integrable canonical kernel. Finally, the discrete–continuous convolution of Ocampo *et al.* [16] replaces the integral in Definition III.3 by a quadrature on a sampling of the sphere; it is Proposition VI.1 up to quadrature error and inherits its equivariance to that accuracy, while its restricted form, with output on the sphere, is a section evaluation, to which we now turn.

VII. BEYOND THE FRAMEWORK

The families of Sec. III F are not fully $\text{SO}(3)$ -equivariant in general, although they contain equivariant special cases. Their non-equivariant members are not generalised lifted spherical

convolutions, by Theorem V.2, but are described by the same tools: they either evaluate a lifted convolution on a section without the constraint of Corollary V.4, or use an m -dependent diagonal harmonic multiplier. We characterise full equivariance and establish guaranteed residual symmetries, without classifying the exact stabiliser of every kernel.

A. Evaluation on a section

Proposition VII.1 (Evaluation on a section). *Let $\sigma(x) = s(x)R_z(\chi_\sigma(x))$ be a measurable section and let $\Psi \in L^2(\text{SO}(3))$ satisfy $\Psi_{n,-s}^\ell = 0$ for $|n| \geq N$. Define $E_\sigma[\Psi]_s f = (\Psi \star \Lambda_s \star f) \circ \sigma$, using the continuous representative of the $L^2 \star L^2$ correlation. Then*

$$E_\sigma[\Psi]_s f = \sum_{|s'| < N} \frac{e^{-is'\chi_\sigma}}{\sqrt{2\pi}} \mathcal{C}_{s \rightarrow s'}[\Psi]_s f, \quad (62)$$

where N bounds the first index within the effective input column. With the fixed coordinate-frame output action, the zero map is equivariant into every spin space for every section. Otherwise $E_\sigma[\Psi]$ is equivariant into spin- t fields if and only if $\Psi_{n,-s}^\ell = 0$ for $n \neq -t$ and $e^{-it\chi_\sigma}$ is constant almost everywhere. The latter condition is automatic for $t = 0$.

For $\sigma_0 = s(\cdot)$, all polar rotations are symmetries. For $\chi_\sigma = \pi - \varphi$, a polar rotation $R_z(a)$ is a symmetry exactly when $e^{ita} = 1$ for every active output spin t , where active means $\Psi_{-t,-s}^\ell \neq 0$ for some ℓ . This criterion does not classify non-polar symmetries.

Proof. Corollary IV.12 and Definition IV.6 give the finite sum, which is bounded since each component is bounded and the phase has unit modulus. If it is equivariant into spin t , Theorem V.5 applies. Testing each input harmonic and comparing the independent Wigner columns representing the different output spins forces all components except t to vanish and, for a nonzero map, forces its phase to be constant almost everywhere. The converse follows from equivariance of $\mathcal{C}_{s \rightarrow t}$. Polar rotations act by coordinate rotation for every spin; substituting $\chi_\sigma = \pi - \varphi$ gives the stated phase criterion. The degreewise argument and the zero-map case are detailed in Appendix B 7. $\square$

When only spin t is present, $E_\sigma = e^{-it\chi_\sigma} E_{\sigma_0}$ already expresses the output in the σ frame. Returning it to the coordinate frame multiplies by $e^{+it\chi_\sigma}$; reading it in the σ frame does not remove the phase.

The standard correlation of Definition III.7 is $E_{\sigma_0}[\Lambda_0 \psi]$. The restricted discrete-continuous convolution of Ocampo *et al.* [16] has this form up to quadrature. Wiaux *et al.* [22] expand the standard correlation as $\sum_n \sum_{\ell m} (-1)^n \sqrt{4\pi/(2\ell+1)} \psi_{\ell,-n}^* f_{\ell m} Y_{\ell m}$, which is (62) with (49), constants included: it sums the spin components in the coordinate frame. Polar equivariance is guaranteed, but full rotation equivariance is not present in general. Ocampo *et al.* derive the limiting rotation behaviour as $\beta \rightarrow 0$, corresponding to the vanishing twist for polar rotations, away from the coordinate singularities.

The commutative anisotropic convolution of Definition III.8 uses $\chi_\sigma = \pi - \varphi$ and a bilinear pairing. It is commutative [30] but not fully rotation-equivariant in general. Its polar symmetries depend on the active orders: even orders retain a half-turn, and orders divisible by q retain the cyclic rotations through multiples of $2\pi/q$. A zonal kernel retains full scalar equivariance, while non-polar symmetries of other kernels require separate analysis. Proposition VII.1 gives the precise obstruction to combining nonzero-spin components into one equivariant scalar output.

B. Operators diagonal in the spherical harmonic basis

Proposition VII.2 (Harmonic multiplication and sifting). *Let T act on $L^2(S^2)$ by $(Tf)_{\ell m} = g_{\ell m} f_{\ell m}$ with a bounded coefficient sequence g . This is harmonic multiplication,*

or sifting with the coefficientwise-conjugated filter; the filter may be a distribution. Then T commutes with every polar rotation, and is $\text{SO}(3)$ -equivariant if and only if $g_{\ell m}$ is independent of m for each ℓ . In that case $T = \mathcal{C}_{0 \rightarrow 0}[\Psi_T]$ with $c_\ell = g_{\ell 0}$ in the Fourier-multiplier sense of Theorem V.5, with a generally distributional zonal representative. If $g_{\ell m} = g_{\ell, -m}$, T also commutes with $R_y(\pi)$, and thus with the embedded $O(2)$ normaliser of polar rotations.

Proof. The full-equivariance criterion is Theorem V.5 with $s = s' = 0$. Polar rotations act diagonally by e^{-ima} and commute with any diagonal multiplier, while $R_y(\pi)$ exchanges m and $-m$ up to a sign, so commutes when the corresponding multipliers agree. $\square$

Proposition VII.2 is the characterisation of isotropic operators of Kennedy *et al.* [20]. An L^2 sifting filter gives a Hilbert–Schmidt diagonal operator, with norm squared $\sum_{\ell m} |g_{\ell m}|^2$; general bounded diagonal operators allow distributional filters. The sifting filter and the equivalent zonal convolution kernel are different representations of the operator.

C. The desiderata revisited

Roddy & McEwen [32] ask for directional kernels, spherical output and the $\mathcal{O}(L^3)$ cost of spherical harmonic transforms. Their comparison includes isotropic, left, directional, commutative anisotropic and sifting constructions, and motivates sifting by the limitations of the preceding constructions considered there. Sifting meets the three criteria, as does coordinate-section standard correlation, whose $\mathcal{O}(L^3)$ grid evaluation is derived in Appendix B8. The cited commutative anisotropic algorithm costs $\mathcal{O}(L^3 \log L)$ [30] and does not meet the strict third criterion.

Full rotation equivariance adds a different restriction: from one fixed input spin to one fixed output spin, Corollary V.7 permits only a radial kernel with the prescribed frame phases, not independent directionality. Standard correlation and sifting are therefore not fully equivariant in general, although each has equivariant special cases and guaranteed polar symmetry. A tuple of output spins retains directional information equivariantly, at an $\mathcal{O}(NL^3)$ upper bound for at most $2N - 1$ components.

VIII. COMPUTATION

In this section we describe how the operators of the article are computed exactly and efficiently for band-limited signals. We first present harmonic-space algorithms for the four generalised lifted spherical convolutions and their computational complexity, which follow from the part of the kernel that each operator reads and from the representation retained at its output. We then describe `s2conv`, the open-source code in which the operators are implemented. We finally present numerical experiments that verify the reductions of Sec. VI and the equivariance of every operator to machine precision, and that confirm the predicted scaling of the computation time with the band-limit.

A. Harmonic-space algorithms and their complexity

The lifted input need not be formed explicitly: for band-limited signals every operator of Table II is computed in harmonic space by spin spherical harmonic and Wigner transforms, by appealing to the sampling theorems on the sphere [42] and on the rotation group [45], which are exact for band-limited signals. Let L be the band-limit and $1 \leq N \leq L$ an azimuthal band-limit. We assume that the Wigner coefficients of the kernel are available. We also consider the cost of the transforms of the sampled input and output. Working memory is given by the arrays of harmonic coefficients held during the computation, with the Wigner d -functions of the transforms generated on the fly; the kernel parameters are

counted separately, and fully precomputed Wigner d -functions for a transform of fixed spin require $\mathcal{O}(L^3)$ further storage [43]. The four operators of Definition V.1 are computed as follows.

- $\mathcal{C}_{s \rightarrow s'}$: a forward spin- s spherical harmonic transform, multiplication of the coefficients by $c_\ell = (-1)^{s+s'} \Psi_{-s', -s}^{\ell*}$, and an inverse spin- s' transform, at a cost of $\mathcal{O}(L^3)$ [42], with $\mathcal{O}(L^2)$ working memory and at most L kernel parameters.
- $\mathcal{C}_{s \rightarrow G}$: a forward spin- s transform, multiplication by $\Psi_{n, -s}^{\ell*}$ for $|n| < N$, and an inverse Wigner transform, which comprises at most $2N - 1$ inverse spin transforms and a Fourier series in γ [43, 45], at a cost of $\mathcal{O}(NL^3)$, with $\mathcal{O}(NL^2)$ storage for the output and $\mathcal{O}(NL)$ kernel parameters; $N = L$ allows an unrestricted column.
- $\mathcal{C}_{G \rightarrow s'}$: for an input supported on $|m'| < N$, a forward Wigner transform, contraction with the entries $\Psi_{-s', m'}^{\ell*}$ of the row $-s'$, and an inverse spin- s' transform, at a cost of $\mathcal{O}(NL^3)$, with $\mathcal{O}(NL^2)$ storage for the input and $\mathcal{O}(NL)$ kernel parameters.
- $\mathcal{C}_{G \rightarrow G}$: a forward Wigner transform, a matrix product for each degree, and an inverse Wigner transform. With input azimuthal band-limit N_{in} and kernel row bound N_{out} as in Sec. VID, the transforms cost $\mathcal{O}((N_{\text{in}} + N_{\text{out}})L^3)$ and the products $\mathcal{O}(N_{\text{in}}N_{\text{out}}L^2)$, with $\mathcal{O}(N_{\text{in}}N_{\text{out}}L)$ kernel parameters and $\mathcal{O}((N_{\text{in}} + N_{\text{out}})L^2)$ storage; taking both bounds equal to N gives a cost of $\mathcal{O}(NL^3 + N^2L^2)$, and $N = L$ the unrestricted $\mathcal{O}(L^4)$.

These costs are collected in Table III, together with those of the constructions of Sec. VII, which are formed from the same transforms: the evaluation on a general section is computed component by component in $\mathcal{O}(NL^3)$, each spin component being accumulated into the single $\mathcal{O}(L^2)$ output array as it is computed, so that the tuple of $\mathcal{O}(NL^2)$ coefficients need not be held at once, while the standard correlation on grid values, harmonic multiplication and the sifting convolution cost $\mathcal{O}(L^3)$ and the commutative anisotropic convolution $\mathcal{O}(L^3 \log L)$ [30]. Note that a single spin-changing convolution costs $\mathcal{O}(L^3)$, as does any finite tuple of spins, and that further structure of the kernel may permit other algorithms. Furthermore, by Proposition V.3 consecutive operators require no intermediate transforms: a chain of sphere-to-sphere convolutions is a single multiplier, and a tuple of spins is stored as $2N - 1$ arrays of harmonic coefficients between which no transform is required. All of the harmonic transforms required are available with exact quadrature on the sphere [42] and on the rotation group [45], and are implemented differentiably on hardware accelerators, with precomputed or on-the-fly Wigner d -functions, in the `s2fft` code [43] (<https://github.com/astro-informatics/s2fft>).

B. Implementation: the `s2conv` code

We have implemented the algorithms described above in `s2conv`, an open-source code that accompanies this article and is publicly available at <https://github.com/astro-informatics/s2conv> under the MIT licence. The code is written in JAX and relies on the `s2fft` code [43] to compute spin spherical harmonic and Wigner transforms, so that every operator is differentiable, may be compiled, runs on hardware accelerators (*e.g.* GPUs, TPUs) and is exact for band-limited signals. Since the operators are differentiable, their kernel coefficients may also be learned, as in the forthcoming article [29]. The code follows the article closely. A single function implements the four operators of Definition V.1 for a kernel given by its Wigner coefficients, in the harmonic forms of Theorem V.2, reading only the entry, column, row or block of the kernel that the operator uses, so that the complexity of Table III is realised. The lift, the adjoint projection and the decomposition of a function on the rotation group into its spin components (Corollary IV.12) are also provided, together with the kernels of Sec. VI, namely the kernel lifted from a spin field, the kernel Ψ_h of the left convolution, the canonical kernel of a multiplier (Theorem V.5)

TABLE III. Algorithmic complexity for band-limit L and azimuthal band-limit $N \leq L$: the scaling of the computation time, the working memory and the number of kernel parameters of each operator, with the part of the kernel that it reads. Kernel coefficients are supplied and the transforms are those of [42, 45], implemented in [43]. Memory is the working memory of Sec. VIII A, the arrays of input and output coefficients held during the computation with Wigner d -functions generated on the fly; parameter storage is listed separately, and fully precomputed Wigner d -functions increase the memory. Counts per degree are further limited by $2\ell + 1$, and “—” means not applicable.

Operator	Effective kernel support	Compute	Memory	Parameters
$\mathcal{C}_{s \rightarrow s'}$	one entry per ℓ	L^3	L^2	L
$\mathcal{C}_{s \rightarrow G}, n < N$	at most $2N - 1$ entries per ℓ	NL^3	NL^2	NL
$\mathcal{C}_{s \rightarrow G}, N = L$	one column per ℓ	L^4	L^3	L^2
$\mathcal{C}_{G \rightarrow s'}$, input $ m' < N$	row $-s', m' < N$	NL^3	NL^2	NL
$\mathcal{C}_{G \rightarrow G}$, input $ m' < N$	$ n , m' < N$	$NL^3 + N^2L^2$	NL^2	N^2L
$\mathcal{C}_{G \rightarrow G}, N = L$	full block per ℓ	L^4	L^3	L^3
General-section evaluation, accumulated	col. $-s, n < N$	NL^3	L^2	NL
Standard correlation, grid values	one column per ℓ	L^3	L^2	L^2
Commutative anisotropic [30]	one column per ℓ	$L^3 \log L$	—	L^2
Harmonic multiplication / sifting	diagonal in (ℓ, m)	L^3	L^2	L^2

and the flipped kernel $\tilde{\Psi}$ of an adjoint. Furthermore, the classical convolutions of Sec. III are provided in the closed forms of Sec. VI, the operators of mathematical physics of Proposition VI.8 as multipliers, and the constructions of Sec. VII, namely the standard correlation and the commutative anisotropic convolution as evaluations on a section, and harmonic multiplication and the sifting convolution as diagonal operators. The harmonic-space operators act on arrays of coefficients in the layout of `s2fft` and broadcast over leading axes, so that batches, channels or spins may be carried as leading dimensions or mapped over. The transforms may be computed with Wigner d -functions generated on the fly, with the $\mathcal{O}(L^2)$ working memory of Table III, or precomputed and stored, which trades $\mathcal{O}(L^3)$ memory for speed, as `s2fft` allows.

C. Numerical validation

In order to verify the reductions of Sec. VI and the equivariance of the operators, and to assess their computation time, we perform the following numerical experiments with the benchmark scripts distributed with the code. We measure the rotational equivariance of an operator T by the protocol of Cobb *et al.* [14] and Ocampo *et al.* [16]. We simulate band-limited test signals f_i whose harmonic or Wigner coefficients are drawn from the standard normal distribution, draw rotations g_j from the Haar measure of $\text{SO}(3)$, and compute the relative errors

$$\epsilon_{ij} = \frac{\|T \pi(g_j) f_i - \pi'(g_j) T f_i\|}{\|\pi'(g_j) T f_i\|}, \quad (63)$$

where π and π' are the representations acting on the input and on the output, π_s on a spin- s field and π_G on a function on the rotation group, and where the norms are those of $L_s^2(S^2)$ or $L^2(\text{SO}(3))$, which are computed by Parseval’s identity and are exact for band-limited outputs. The numerator is the equivariance error, the difference between rotating the input and rotating the output, and the denominator is the norm of the output, which a rotation does not change. For an evaluation on a section, whose output is not band-limited, the comparison is made through the spin components of Proposition VII.1, which is exact and considerably faster than a comparison point by point. The errors are averaged over 100 random signals and 100 random rotations for each operator.

For each operator we also compute a reduction error, the largest relative difference, over 100 random signals, between its closed form and its form as a generalised lifted spherical convolution, *i.e.* Definition V.1 composed literally from the lift, the group convolution in Wigner space and the adjoint projection. This verifies the reductions of Sec. VI together

TABLE IV. Reduction errors and equivariance errors of the operators of **s2conv** at band-limit $L = 128$, in double precision, from the benchmark script of the code. The equivariance error is the mean of (63) over 100 random signals and 100 random rotations, and the reduction error the largest relative difference between the closed form of the operator and Definition V.1 composed from the lift, the group convolution and the projection, over 100 random signals (“-” where the operator is the definition itself or has no lifted form). The random kernels of the lifted operators, the directional kernels and the functions on the rotation group are azimuthally band-limited at $N = 5$. The rule separates the equivariant operators from the constructions of Sec. VII.

Operator	Reduction error	Equivariance error
$\mathcal{C}_{0 \rightarrow 0}[\Psi]$	1.4×10^{-16}	9.2×10^{-16}
$\mathcal{C}_{2 \rightarrow 1}[\Psi]$	1.4×10^{-16}	9.0×10^{-16}
$\mathcal{C}_{-1 \rightarrow 2}[\Psi]$	1.4×10^{-16}	8.7×10^{-16}
$\mathcal{C}_{0 \rightarrow G}[\Psi]$	1.3×10^{-16}	1.2×10^{-15}
$\mathcal{C}_{2 \rightarrow G}[\Psi]$	1.3×10^{-16}	1.2×10^{-15}
$\mathcal{C}_{G \rightarrow 0}[\Psi]$	2.0×10^{-16}	1.3×10^{-15}
$\mathcal{C}_{G \rightarrow 2}[\Psi]$	2.1×10^{-16}	1.3×10^{-15}
$\mathcal{C}_{G \rightarrow G}[\Psi]$	0	1.6×10^{-15}
axisymmetric convolution	1.7×10^{-16}	8.3×10^{-16}
left convolution	0	7.6×10^{-16}
directional convolution	1.5×10^{-16}	1.1×10^{-15}
directional synthesis	1.6×10^{-16}	1.2×10^{-15}
$\bar{\partial}$	0	9.6×10^{-16}
$\bar{\bar{\partial}}$	0	9.6×10^{-16}
E mode from $Q + iU$	0	9.7×10^{-16}
Kaiser–Squires	0	8.8×10^{-16}
standard correlation	3.8×10^{-16}	1.3×10^0
commutative anisotropic convolution	–	1.3×10^0
sifting convolution	–	1.4×10^0
harmonic multiplication	–	1.4×10^0

with their constants, for example the $\sqrt{4\pi/(2\ell+1)}$ of the axisymmetric convolution and the 2π of the left convolution, and, for the multipliers of Proposition VI.8, the canonical kernel of Theorem V.5.

The operators tested are the four operators of Definition V.1 with random kernels, for the spins $0 \rightarrow 0$, $2 \rightarrow 1$ and $-1 \rightarrow 2$ between spin fields, 0 and 2 to the rotation group, the rotation group to spins 0 and 2 , and the rotation group to itself, with functions on the rotation group azimuthally band-limited at $N = 5$; the classical convolutions of Sec. III, namely the axisymmetric convolution, the left convolution, the directional convolution of a spin-2 field with a steerable kernel and the synthesis; the operators of mathematical physics, namely $\bar{\partial}$ from spin 1 to spin 2, $\bar{\bar{\partial}}$ from spin 1 to spin 0, the E mode of a spin-2 field and the Kaiser–Squires relation; and, for contrast, the standard correlation, the commutative anisotropic convolution, the sifting convolution and harmonic multiplication, which Sec. VII shows are not equivariant under general rotations.

Results at $L = 128$ are presented in Table IV; the results at $L = 32$ and 64 are distributed with the code. All numerical experiments are performed in double precision on an NVIDIA GB10 graphics processing unit. The equivariance error of every equivariant operator is of order 10^{-15} , *i.e.* close to machine precision. The reduction errors are of order 10^{-16} , or vanish identically where the closed form and the lifted form read the same coefficients, which verifies the reductions of Sec. VI with their constants. In contrast, the standard correlation, the commutative anisotropic convolution, the sifting convolution and harmonic multiplication exhibit equivariance errors of order unity, as expected, since Sec. VII shows that these constructions are equivariant only under rotations about the polar axis.

In order to assess the scaling of computation time with band-limit, we time one operator of each of the four kinds, $\mathcal{C}_{2 \rightarrow 1}$ and, with $N = 5$, $\mathcal{C}_{2 \rightarrow G}$, $\mathcal{C}_{G \rightarrow 2}$ and $\mathcal{C}_{G \rightarrow G}$, with random




FIG. 7. Computation time of the four generalised lifted spherical convolutions against band-limit, for $C_{2 \rightarrow 1}$ and, with $N = 5$, $C_{2 \rightarrow G}$, $C_{G \rightarrow 2}$ and $C_{G \rightarrow G}$, with random kernels, each run end to end with its transforms and measured by the median of five compiled runs in double precision on a single Cortex-X925 core of an NVIDIA GB10 with multithreading disabled. As expected, computation time scales as $\mathcal{O}(L^3)$, the complexity of the transforms (Table III), which is shown by the dotted line.

kernels, for band-limits from $L = 16$ to $L = 256$. Each operator is run end to end, including its transforms, after compilation, and the computation time is measured by the median of five runs in double precision. The timings are performed on a single processor core, one Cortex-X925 core of the GB10 with the multithreading of the compiler disabled, so that they count operations rather than the occupancy of a graphics processing unit. Results are plotted in Fig. 7. With N fixed, the complexity of each of the four operators is the $\mathcal{O}(L^3)$ of its transforms (Table III), and, as expected, the measured computation time is indeed found to scale as $\mathcal{O}(L^3)$. The operators on the rotation group are more costly than the operator between spin fields since each of their Wigner transforms runs over $2N - 1 = 9$ azimuthal orders.

IX. PHYSICAL ILLUSTRATIONS

In this section we apply the generalised lifted spherical convolution to fields measured on the Earth and on the sky, in order to illustrate in practice the operators that it provides, which map between fields of different spin exactly and efficiently and which respect the geometry of each field. Many of the fields of geophysics and cosmology are not scalars. Wind, an ocean current, the deflection of the vertical and the horizontal magnetic field are tangent vector fields, of spin ± 1 , while the polarisation of the cosmic microwave background, the shear of weak gravitational lensing and the horizontal part of a gravity-gradient tensor are fields of spin ± 2 , each describing a headless direction at every point. The components of such a field are defined only with respect to the local frame (e_θ, e_φ) of Sec. II C, which turns from point to point and is undefined at the poles of the coordinates. Treating the components as independent scalar fields, for example by smoothing v_θ and v_φ , or Q and U , separately, combines values expressed in different frames, and the result then depends on the choice of coordinates and is incorrect near their poles. In each of the examples we adopt the lifted convolutions whose input and output spins are those of the physical quantities

involved. A tangent vector field with northward and eastward components (v, u) is the spin-1 field $v_\theta + i v_\varphi = -v + iu$, since e_θ points south, and a polarisation with Stokes parameters (Q, U) is the spin-2 field $Q + iU$. The fields are truncated at the band-limit $L = 64$ for these illustrations, and the fields of the Earth are multiplied by the taper $e^{-\ell(\ell+1)/L^2}$, which limits the ringing of the truncation at coastlines and jets. The examples are computed with `s2conv`, in a notebook distributed with the code, from which the figures of this section are taken.

We restrict our attention to the lifted convolution itself, which, since its kernels may be distributions, also includes differential operators such as $\bar{\partial}$ and $\bar{\bar{\partial}}$ and their compositions. Note that convolution is often one component of other analyses, such as wavelet transforms on the sphere, convolutional neural networks and Fourier neural operators, which build on the operators illustrated here but which we do not consider further.

The fields considered are shown in Fig. 2 of Sec. II C. Throughout this section, maps of the whole sphere use the Mollweide projection, centred on the meridian of longitude zero with north at the top, the fields of the Earth in geographic coordinates and the Planck maps in Galactic coordinates, and views of one hemisphere use the orthographic projection, centred on the point given in each caption. A signed quantity is shown as colour on a diverging scale, red for positive and blue for negative values and white for zero, and a non-negative quantity, such as the modulus of a spin field, on a sequential scale from light to dark. Spin fields are drawn in the tangent plane with the glyphs of Fig. 1: a spin-1 field ${}_1f$ as an arrow at the angle $\arg({}_1f)$ from e_θ towards e_φ , and a spin-2 field ${}_2f$ as a headless stick at the angle $\arg({}_2f)/2$, at quasi-uniform points, with lengths that grow with the modulus of the field. Coastlines are drawn as black lines.

A. Topography: slopes and ridges

The elevation of the land and the depth of the ocean floor, the topography and bathymetry f of ETOPO 2022 [49], form a scalar field, but the two quantities that describe the shape of the relief are not scalars: the slope is a vector and the orientation of a mountain range is a headless direction. We obtain the slope by the lifted convolution $\mathcal{C}_{0 \rightarrow 1}$ whose kernel is the multiplier $\sqrt{\ell(\ell+1)}$ of the spin raising operator, Proposition VI.8(a). Since $\bar{\partial}f = -(\partial_\theta + i \sin^{-1} \theta \partial_\varphi)f$ for a scalar field, by (3), the spin-1 field $\bar{\partial}f/R_\oplus$, where $R_\oplus$ is the radius of the Earth, is the vector $-\nabla f$, which points downhill and whose modulus is the slope.

To locate mountain ranges we consider a steerable ridge detector ψ , a kernel elongated along one direction with a zero-mean profile across it, like a second derivative, truncated to the azimuthal orders $|m| < 5$. Its directional convolution with the field, $f \circledast \psi$, is a function on the rotation group that gives the response of the detector at every position and every orientation γ . Since the kernel is steerable, this function is the tuple of the nine spin fields ${}_{s'}g = \mathcal{C}_{0 \rightarrow s'}[\Lambda_0 \psi]f$, $|s'| < 5$, by Proposition VI.6, from which the response at any orientation follows by Proposition VII.1 as $\sum_{s'} (2\pi)^{-1/2} e^{-is'\gamma} {}_{s'}g$. The largest response over γ measures the strength of a ridge, and the spin-2 component ${}_2g$, the part of the response that varies with orientation as $\cos 2\gamma$, like that of a line, is itself the orientation of the ridges as a spin-2 field: the response is largest when the detector lies at the angle $\arg({}_2g)/2$, along the ridge. No sampling of orientations is required, and since both outputs are rotation-equivariant, ranges are detected alike whatever their orientation and position, near the equator or near the poles. We apply the detector to the land elevation $f_+ = \max(f, 0)$, since the continental margins, a step of several kilometres, would otherwise dominate its response. Figure 8 shows the results: the arrows run down the flanks of the Tibetan Plateau, the ridge strength picks out the Himalaya, the Andes, the Rocky Mountains, the highlands of East Africa and the ice sheets of Greenland and Antarctica, and where the ridge strength is large the sticks lie along the ranges, from east to west along the Himalaya and from north to south along the Andes.




FIG. 8. The topography of the Earth, its slope and its ridges, from lifted convolutions of spin 0 to spins 1 and 2. The ridge strength (c) is the largest response of the ridge detector over all orientations, and the ridge orientation (e), (f) is the spin-2 component $_2g$ of the response, drawn as headless sticks at the angle $\arg(_2g)/2$ with its modulus shown as colour. The orthographic views are centred on 32°N , 85°E (d), (e) and on 20°S , 68°W (f).

B. Gravity: one potential and its derivatives of spin 0, 1 and 2

The geoid is the surface of constant gravity potential that coincides with mean sea level over the oceans, and its height N above the reference ellipsoid is a scalar field, which we take from the Earth Gravitational Model EGM2008 [50]. By Bruns' formula [51], N is the disturbing potential T , the difference between the gravity potential of the Earth and that of the ellipsoid, divided by the normal gravity γ , and the quantities that geodesy measures are derivatives of T of increasing spin. The gravity anomaly, measured by gravimeters, is the scalar field with coefficients $\Delta g_{\ell m} = \frac{\gamma}{R_\oplus}(\ell - 1)N_{\ell m}$ in the spherical approximation, *i.e.* the lifted convolution $\mathcal{C}_{0 \rightarrow 0}$ with the multiplier $\frac{\gamma}{R_\oplus}(\ell - 1)$. The deflection of the vertical, the angle between the direction of gravity and the normal to the ellipsoid, measured by astronomical positioning, is the vector $-\nabla N/R_\oplus$, *i.e.* the spin-1 field $\varepsilon = \partial N/R_\oplus$ given by $\mathcal{C}_{0 \rightarrow 1}$ with the multiplier $\sqrt{\ell(\ell + 1)}/R_\oplus$. The horizontal trace-free part of the tensor of second derivatives of T , measured by the gradiometer of the GOCE satellite, is the spin-2 field $\Gamma = \frac{1}{2}\partial^2 T/R_\oplus^2$ given by $\mathcal{C}_{0 \rightarrow 2}$.

The classical integral formulae of physical geodesy are lifted convolutions between these spins. The formula of Vening Meinesz [51], which gives the deflection of the vertical from gravity anomalies, is the spin-changing convolution $\mathcal{C}_{0 \rightarrow 1}$ with the multiplier $\sqrt{\ell(\ell + 1)}/(\gamma(\ell - 1))$. The recovery of the geoid from deflections of the vertical, astro-




FIG. 9. The gravity field of the Earth from EGM2008, as fields of spin 0, 1 and 2 related by lifted convolutions. The gravity anomaly, in mGal (10^{-5} m s^{-2}), is shown on a scale saturated at the 99.5th percentile of its modulus; the arrows of the deflection point down the slope of the geoid; and the sticks of the gradients, in Eötvös (10^{-9} s^{-2}), lie along the direction in which the second derivative of the disturbing potential is largest. The orthographic views are centred on 5°S , 80°E (d), 20°S , 75°W (e) and 5°S , 140°E (f). The deflections point into the geoid low of the Indian Ocean and away from the geoid high over New Guinea, and the gradients are largest along the Andes and the adjacent trench.

geodetic levelling, is the lifted convolution $\mathcal{C}_{1 \rightarrow 0}$ that inverts $\mathcal{C}_{0 \rightarrow 1}$, and the recovery of the geoid from its gradients is the lifted convolution $\mathcal{C}_{2 \rightarrow 0}$ that inverts $\frac{1}{2}\partial^2$, the operator that also underlies the Kaiser–Squires relation of weak lensing. Each of these relations is a single multiplier in harmonic space and holds to machine precision for band-limited fields: in Fig. 9 the formula of Vening Meinesz reproduces ε from Δg to a relative error of 10^{-16} , and the geoid is recovered from ε and from Γ to relative errors of 2×10^{-16} and 10^{-16} , apart from its monopole and dipole, which the derivatives do not determine.

C. The geomagnetic field: a potential field of spin 0 and 1, continued to the core

The main magnetic field of the Earth, generated by the dynamo in the fluid outer core, is the gradient of a potential, $\mathbf{B} = -\nabla V$, which the International Geomagnetic Reference Field provides as spherical harmonic coefficients up to degree 13 [52]. With $v = V/R_\oplus$ on the surface of the Earth, the radial component of the field is the scalar field $B_r = \mathcal{C}_{0 \rightarrow 0}[\ell + 1] v$, and the horizontal component is the spin-1 field $B_\theta + iB_\varphi = \partial v = \mathcal{C}_{0 \rightarrow 1}[\sqrt{\ell(\ell + 1)}] v$, the field that aligns a compass needle. The declination, the angle by which the needle deviates




FIG. 10. The geomagnetic field of IGRF-14 in 2025, as fields of spin 0 and 1, and its downward continuation to the core. The radial field is blue where it points into the Earth and red where it points out, and the arrows of the horizontal field point in the direction of a compass needle. The declination is saturated at $\pm 30^\circ$, with contours every 10° , among which those of 0° are the agonic lines, and the magnetic equator is the contour of zero inclination. In (f) the coastlines are projected radially onto the boundary for orientation.

east of geographic north, is the direction of the horizontal field measured from north, and the inclination, the angle by which the field dips below the horizontal, follows from the ratio of $-B_r$ to the modulus of the horizontal field. Each degree ℓ of a potential field of internal origin decays with radius as $r^{-(\ell+2)}$, so that the field at the boundary between the core and the mantle, at radius $r_c = 3485$ km, follows from that at the surface by the multiplier $(R_\oplus/r_c)^{\ell+2}$, the downward continuation. The multiplier preserves spin, and we apply it to the radial field as $\mathcal{C}_{0 \rightarrow 0}$ and to the horizontal vector field as $\mathcal{C}_{1 \rightarrow 1}$, directly and without reference to its components. Downward continuation amplifies the coefficients of degree 13 by $(R_\oplus/r_c)^{15} \approx 8500$ and reveals the structure of the field at the top of the core (Fig. 10): intense patches of flux at high latitudes, and patches of reversed flux beneath the South Atlantic and southern Africa, which weaken the surface field in the South Atlantic Anomaly.

D. Wind: vorticity, divergence and smoothing

We consider the near-surface wind in January, averaged over 1991–2020, from the NCEP/NCAR reanalysis [53]. With northward and eastward components (v, u) , and since e_θ points south, the wind is the spin-1 field ${}_1f = -v + iu$. For a spin-1 field the spin lowering operator gives the divergence and the vorticity, the radial component of the curl,

together,

$$\bar{\partial}_{-1} f = -R_{\oplus} (\nabla \cdot \mathbf{v} + i \hat{\mathbf{r}} \cdot \nabla \times \mathbf{v}), \quad (64)$$

so that the single lifted convolution $\mathcal{C}_{1 \rightarrow 0}$ yields both as the real and imaginary parts of one scalar field, exactly for the band-limited field, without finite differences and without special treatment of the poles. Since $\bar{\partial}\bar{\partial}$ is the Laplacian on spin-1 fields, applying the lifted convolution $\mathcal{C}_{0 \rightarrow 1}$ of $\bar{\partial}\nabla^{-2}$ to the divergence and to the vorticity separately gives the Helmholtz decomposition of the wind into a curl-free part, the flow into and out of regions of convergence, and a divergence-free part, the rotational flow of the trade winds and the westerlies, which sum to the wind to a relative error of 10^{-16} (Fig. 11).

Smoothing the wind by a Gaussian is the spin-preserving lifted convolution $\mathcal{C}_{1 \rightarrow 1}$ whose kernel is the Gaussian multiplier $e^{-\ell(\ell+1)\sigma^2/2}$. A common alternative smooths the components v_{θ} and v_{φ} as though they were two scalar fields. Since the frame $(e_{\theta}, e_{\varphi})$ turns around the poles of the coordinates, this averages vectors expressed in different frames, and near a pole the components vary rapidly even where the wind itself is uniform. The poles of the coordinates are a matter of convention, and to expose the dependence of each method on that convention we place one pole at the location of the strongest smoothed wind and smooth in both systems of coordinates (Fig. 12). The spin-1 smoothing, which is rotation-equivariant, gives the same result in either system to a relative error of 6×10^{-16} . The component-wise smoothing gives a plausible result in geographic coordinates, where the strongest winds lie far from the poles, but with the pole placed in the strong flow it combines the components across the turning frame and destroys the flow near the pole, with a relative error of ten per cent over the sphere: its result depends on the choice of coordinates, which is the error that the spin formalism removes.

E. Ocean currents: gyres and upwelling

The surface currents of the ocean in 2025, averaged over the year at a depth of 5 m, from the Global Ocean Data Assimilation System [54], form a spin-1 field like the wind, which we set to zero on land. The lifted convolution $\mathcal{C}_{1 \rightarrow 0}$ again gives their vorticity and divergence. The stream function $\psi = R_{\oplus}^2 \nabla^{-2} (\hat{\mathbf{r}} \cdot \nabla \times \mathbf{v})$ of the divergence-free part, obtained from the vorticity by the inverse Laplacian, has as its contours the streamlines of that flow, along which the rotational part of the current runs with higher values of ψ to its right. The streamlines trace the subtropical gyres, which circulate clockwise in the northern hemisphere and anticlockwise in the southern hemisphere, and the Antarctic Circumpolar Current (Fig. 13). The divergence of the surface flow marks upwelling, where water rises from below to replace surface water that is carried away. The trade winds drive a transport of the surface water away from the equator on both sides, so that the surface flow diverges along the equator, in the cold tongues of the eastern Pacific and Atlantic, and along the coasts of Peru and north-west Africa, where the winds blow along the coast.

F. Polarisation of the cosmic microwave background

The linear polarisation of the cosmic microwave background, described by the Stokes parameters Q and U , is the spin-2 field $Q + iU$. Its physical content is carried by two scalar fields, the E and B modes, which do not depend on the frame and which separate the contributions of different physical processes: density perturbations in the early Universe generate only E modes, while B modes are generated by the gravitational lensing of E modes and would be generated by primordial gravitational waves. By Proposition VI.8(b), $a_E = -\frac{1}{2}(a_2 + a_{-2})$ and $a_B = \frac{1}{2}(a_2 - a_{-2})$, where $a_{\pm 2}$ are the coefficients of $Q \pm iU$, so that each mode is a sum of the lifted convolutions $\mathcal{C}_{2 \rightarrow 0}$ and $\mathcal{C}_{-2 \rightarrow 0}$. Conversely, the polarisation follows from the two modes by the lifted convolutions $\mathcal{C}_{0 \rightarrow \pm 2}$, as $a_{\pm 2} = -(a_E \pm ia_B)$, so that the decomposition may be inverted exactly.




FIG. 11. The January near-surface wind and its vorticity, divergence and Helmholtz decomposition, from the lifted convolutions $\mathcal{C}_{1 \rightarrow 0}$ and $\mathcal{C}_{0 \rightarrow 1}$. The vorticity and the divergence are shown on scales saturated at the 99.5th percentile of their moduli, red for anticlockwise rotation and for outflow and blue for clockwise rotation and for inflow, and the orthographic views are centred on 20°N , 150°W . The vorticity is positive in the cyclonic flow around the Aleutian and Icelandic lows and negative in the anticyclones of the northern subtropics, and the divergence is negative along the intertropical convergence zone, into which the trade winds of both hemispheres flow.

We consider the Planck 2018 SMICA maps of the temperature and polarisation of the cosmic microwave background [55], in Galactic coordinates, so that the plane of the Galaxy lies along the equator of the maps. We decompose the polarisation into its E and B modes and synthesise the polarisation again from the two modes (Fig. 14). On the large scales considered here, the polarisation measured by Planck combines the signal with noise and residual Galactic emission, most evident along the Galactic plane, which an analysis would mask; the B mode, dominated on these scales by noise and by these residuals, has a root mean square of 79 per cent of that of the E mode. The decomposition and its inverse are single multipliers in harmonic space, $-\frac{1}{2}$ and $\pm\frac{i}{2}$ one way and -1 and $\mp i$ the other, so that the polarisation synthesised from the two modes reproduces the measured polarisation to the precision of the arithmetic. Starting from the map of the polarisation, computing the maps of the E and B modes and synthesising the map of the polarisation again, including all the harmonic transforms, recovers the map to a relative error of 1.3×10^{-13} , set by the machine precision of the transforms.

G. Weak gravitational lensing: shear and convergence

The images of distant galaxies are distorted by the gravitational deflection of their light by the matter along the line of sight. The convergence κ , the projected density of




FIG. 12. Smoothing the wind as a spin-1 field and as two scalar components. Orthographic views centred on the strongest wind after smoothing, at 12°N , 40°W , of the wind smoothed by a Gaussian of width $8/L$ radians, with the speed shown as colour on a common scale, in m s^{-1} , and the wind drawn as arrows. With the pole at the centre, the component-wise smoothing destroys the flow within about the width of the smoothing of the pole, where the frame turns fastest.

matter weighted by the geometry of the lensing, is a scalar field, and the shear γ , the anisotropic stretching of the images, is a spin-2 field. The two are related by the spherical Kaiser–Squires relation [3, 56], which is the lifted convolution $\gamma = \mathcal{C}_{0 \rightarrow 2}[\mathcal{D}]\kappa$ with $\mathcal{D}_\ell = -\sqrt{(\ell+2)(\ell-1)/(\ell(\ell+1))}$, Proposition VI.8(c), and whose inverse $\kappa = \mathcal{C}_{2 \rightarrow 0}[\mathcal{D}^{-1}]\gamma$ recovers the convergence, apart from its monopole and dipole, from the shear, on the full sphere and without the planar approximation of the original method. Since the galaxies are stretched tangentially around concentrations of mass, the sticks of the shear circle the overdense regions and point radially away from the underdense regions.

We simulate the convergence as a shifted lognormal field [57], which captures the skewness of the distribution of matter, with a realistic power spectrum: the Limber integral of the non-linear matter power spectrum computed by CAMB [58] with HMcode [59] for the Planck 2018 cosmology, for a Euclid-like distribution of source galaxies with median redshift 0.9. We compute the shear from the convergence and recover the convergence from the shear




FIG. 13. The surface currents of the ocean in 2025, their stream function and their divergence, from the lifted convolution $\mathcal{C}_{1 \rightarrow 0}$. The divergence is shown on a scale saturated at the 99.5th percentile of its modulus, red where the surface flow diverges and water upwells and blue where it converges, and the orthographic view is centred on 5°S , 110°W . The fastest currents are the westward equatorial currents; the subtropical gyres appear as maxima of ψ in the northern hemisphere and minima in the southern hemisphere; and the divergence is concentrated along the equator in the Pacific and Atlantic and along the coast of Peru.

(Fig. 15); since the simulated convergence has no monopole or dipole, the recovery is exact, to a relative error of 10^{-16} .

In summary, in each example we have applied lifted convolutions $\mathcal{C}_{s \rightarrow s'}$ whose input and output spins are those of the physical quantities concerned. Derivatives raise or lower the spin, as do their compositions, such as the formula of Vening Meinesz and the Kaiser–Squires relation; smoothing and the continuation of potential fields preserve it; the decomposition of polarisation into E and B modes lowers it; and the steerable ridge detector maps a scalar field to the tuple of spin fields that makes up a function on the rotation group. Each operator is a single multiplier of the harmonic coefficients, computed exactly and efficiently for band-limited fields, invertible where the physics allows, as the round trips of the geoid, the polarisation and the convergence confirm to the precision of the arithmetic, and rotation-equivariant, so that its result does not depend on the choice of coordinates. The comparison with the smoothing of the components of the wind as scalar fields shows the error that the spin formalism removes.




FIG. 14. The temperature and polarisation of the cosmic microwave background measured by Planck (top), and the E and B modes of the polarisation, from the lifted convolutions $\mathcal{C}_{2 \rightarrow 0} + \mathcal{C}_{-2 \rightarrow 0}$, with the polarisation synthesised again from the two modes by $\mathcal{C}_{0 \rightarrow 2}$ (bottom), in Galactic coordinates. The sticks are drawn along the direction of polarisation, at the angle $\arg(Q + iU)/2$, after smoothing by the spin-2 convolution $\mathcal{C}_{2 \rightarrow 2}$ with a Gaussian of width $4/L$ radians, so that they follow its large-scale pattern. Both modes show residual Galactic emission along the equator of the maps.

X. CONCLUSIONS

We have developed a unified description of convolution on the sphere by lifting to the rotation group. A spin- s field is lifted to $\text{SO}(3)$ through the irreducible representation of spin s of the stabiliser $\text{SO}(2)$, convolved there with an arbitrary kernel, and projected back to the sphere through the adjoint of the lift for a possibly different spin s' . In spin spherical harmonic space the resulting operator multiplies each coefficient by a single conjugated Wigner coefficient of the kernel, up to a sign. Every rotation-equivariant spherical convolution considered is recovered by constraining the kernel, choosing the input and output spins, and including or omitting the lift and the projection. The constants of the literature follow from the different ways in which a kernel on the sphere is identified with a kernel on the rotation group.

The construction also settles what an equivariant spherical convolution can be. Every bounded rotation-equivariant map between spin fields is a multiplier of this form, and the equivariant maps to and from the rotation group are read from a single column or row of the kernel. A function on the rotation group azimuthally band-limited at N is a tuple of $2N - 1$ spin fields, so that a directional convolution with a steerable kernel is a finite tuple of spin-changing convolutions, computed in $\mathcal{O}(NL^3)$ operations. A key result of the article is that, for maps between fields of fixed input and output spin, directional kernels, output on the sphere and equivariance cannot be achieved simultaneously. The anisotropic constructions of the literature that keep their output on the sphere are evaluations of a




FIG. 15. The convergence and shear of weak gravitational lensing, related by the lifted convolutions $\mathcal{C}_{0 \rightarrow 2}$ and $\mathcal{C}_{2 \rightarrow 0}$. The convergence is red for overdense and blue for underdense regions. The orthographic view (d), centred on 20°N , 34°E in the coordinates of the simulation, shows the convergence smoothed by a Gaussian of width $8/L$ radians, with the shear smoothed in the same way drawn as headless sticks along the direction in which the images are stretched, at the angle $\arg(\gamma)/2$; the smoothing is a multiplier and commutes with $\mathcal{C}_{0 \rightarrow 2}[\mathcal{D}]$, so that the smoothed shear is the shear of the smoothed convergence. The sticks circle the overdense regions and point away from the underdense regions.

lifted convolution on a section of the bundle or multipliers that depend on the azimuthal order, and they are equivariant only in special cases.

The geometric description yields two further key insights. The spin-weighted functions of physics, defined by the phase that they acquire under a rotation of the local frame, are the sections of the line bundles associated to the quotient $\text{SO}(3)/\text{SO}(2)$, or equivalently Mackey functions on the rotation group. Making this identification explicit links the physics and the bundle literatures. The phase that accompanies the rotation of a spin field is thereby expressed through the geometry of the bundle: it is the character of the stabiliser evaluated on the inverse twist of the section under the rotation, so that the rotation of a spin field is precisely the representation induced from that character.

All of the convolutions may be computed exactly and efficiently by spin spherical harmonic and Wigner transforms, at a cost that follows from the part of the kernel that each operator reads. We have implemented every operator of the article in the open-source `s2conv` code, whose numerical tests confirm the reductions and the equivariance of the operators to machine precision, and whose computation time scales as expected. We have applied the

operators to the relief, gravity, magnetic field, winds and currents of the Earth and to the polarisation of the cosmic microwave background and weak gravitational lensing.

In addition to unifying existing spherical convolutions, the framework also suggests new ones. In a forthcoming article [29] we build on the operators presented to construct artificial intelligence (AI) architectures based on spin spherical Fourier neural operators, in which the multipliers are learned and in which a lift, a pointwise non-linearity on the rotation group and a projecting convolution yield equivariant layers for spin fields.

ACKNOWLEDGMENTS

This work was funded under the embedded CSE programme of the ARCHER2 UK National Supercomputing Service (<https://www.archer2.ac.uk>), hosted by the University of Edinburgh and funded by EPSRC, through the GPU eCSE project *Accelerated and differentiable spherical transforms* (GPU-eCSE01-41).

AUTHOR DECLARATIONS

Conflict of interest

The author has no conflicts to disclose.

Use of large language models

Large language models (LLMs) were used in the writing of this manuscript. The author takes full responsibility for the accuracy of the material presented.

DATA AVAILABILITY

No new observational data were generated for this article. The measured fields of Sec. IX are publicly available from the sources cited there, and the only data generated are the simulated convergence and shear of Sec. IX G. The data of all of the examples, the simulation included, are distributed with the open-source **s2conv** code (<https://github.com/astro-informatics/s2conv>), together with the notebook that produces the figures and the benchmark scripts that produce the numerical results of Sec. VIII C.

Appendix A: Wigner functions and spin harmonics

Every constant in the article depends on the choices fixed in Sec. II. This appendix collects the identities for Wigner functions and spin harmonics that the proofs use.

a. Wigner d-functions. The real functions d_{mn}^ℓ in (6) are those of Varshalovich *et al.* [44], fixed by

$$d_{\ell n}^\ell(\beta) = \sqrt{\frac{(2\ell)!}{(\ell+n)!(\ell-n)!}} \left(-\sin \frac{\beta}{2}\right)^{\ell-n} \left(\cos \frac{\beta}{2}\right)^{\ell+n}, \quad d_{m0}^\ell(\beta) = \sqrt{\frac{(\ell-m)!}{(\ell+m)!}} P_\ell^m(\cos \beta), \quad (\text{A1})$$

so that $d_{10}^1(\beta) = -\sin \beta/\sqrt{2}$ and $d_{11}^1(\beta) = \cos^2(\beta/2)$ [43]. They satisfy the symmetries

$$d_{mn}^\ell(\beta) = (-1)^{m-n} d_{nm}^\ell(\beta) = d_{-n,-m}^\ell(\beta) = (-1)^{m-n} d_{mn}^\ell(-\beta) = (-1)^{\ell+m} d_{m,-n}^\ell(\pi - \beta), \quad (\text{A2})$$

for which the $\ell = 1$ entries above provide simple sign checks, and take the values

$$d_{mn}^\ell(0) = \delta_{mn}, \quad d_{mn}^\ell(\pi) = (-1)^{\ell-n} \delta_{m,-n} = (-1)^{\ell+m} \delta_{m,-n} \quad (\text{A3})$$

at the ends of the interval (for $\ell = 1$: $d_{00}^1(\pi) = \cos \pi = -1$ and $d_{1,-1}^1(\pi) = \sin^2(\pi/2) = 1$). Inserting (A2) in (6) gives the conjugation symmetry of the D -functions and, from the homomorphism property with $D_{mn}^\ell(R_z(\chi)) = e^{-im\chi} \delta_{mn}$, their behaviour under the stabiliser,

$$D_{mn}^{\ell*}(\rho) = (-1)^{m-n} D_{-m,-n}^\ell(\rho), \quad D_{mn}^\ell(R_z(\chi_1) \rho R_z(\chi_2)) = e^{-im\chi_1} D_{mn}^\ell(\rho) e^{-in\chi_2}, \quad (\text{A4})$$

the second of which is the identity behind Corollary V.4. At the poles, (A3) gives $D_{mn}^\ell(\alpha, 0, \gamma) = \delta_{mn} e^{-im(\alpha+\gamma)}$ and $D_{mn}^\ell(\alpha, \pi, \gamma) = (-1)^{\ell-n} \delta_{m,-n} e^{-im(\alpha-\gamma)}$.

b. Spin raising and lowering. Acting on a spin- s field, the operators in (5) are [40, 41, 43]

$$\tilde{\partial} = -\sin^s \theta \left(\partial_\theta + \frac{i}{\sin \theta} \partial_\varphi \right) \sin^{-s} \theta, \quad \bar{\partial} = -\sin^{-s} \theta \left(\partial_\theta - \frac{i}{\sin \theta} \partial_\varphi \right) \sin^s \theta, \quad (\text{A5})$$

and the spin harmonics are generated from (2) by

$${}_s Y_{\ell m} = \begin{cases} \sqrt{\frac{(\ell-s)!}{(\ell+s)!}} \tilde{\partial}^s Y_{\ell m}, & 0 \leq s \leq \ell, \\ (-1)^s \sqrt{\frac{(\ell+s)!}{(\ell-s)!}} \bar{\partial}^{-s} Y_{\ell m}, & -\ell \leq s \leq 0, \end{cases} \quad (\text{A6})$$

the sign $(-1)^s$ in the second line being what fixes the phase of the negative-spin harmonics. Setting $\gamma = 0$ in (9) gives the explicit form

$${}_s Y_{\ell m}(\theta, \varphi) = (-1)^s \sqrt{\frac{2\ell+1}{4\pi}} d_{m,-s}^\ell(\theta) e^{im\varphi}, \quad (\text{A7})$$

which reduces to (2) for $s = 0$ by (A1) and shows directly that ${}_s Y_{\ell m}$ exists only for $\ell \geq |s|$.

c. Values at the poles. By (A7) and (A3),

$${}_s Y_{\ell m}(0, \varphi) = (-1)^s \sqrt{\frac{2\ell+1}{4\pi}} \delta_{m,-s} e^{-is\varphi}, \quad {}_s Y_{\ell m}(\pi, \varphi) = (-1)^\ell \sqrt{\frac{2\ell+1}{4\pi}} \delta_{m,s} e^{is\varphi} : \quad (\text{A8})$$

at the north pole only the order $m = -s$ survives and at the south pole only $m = s$, with the windings $e^{\mp is\varphi}$ of Remark IV.1. For the constrained kernel of Corollary V.4 the same values give $d_{-s',-s}^\ell(0) = \delta_{ss'}$ and $d_{-s',-s}^\ell(\pi) = (-1)^{\ell+s} \delta_{s',-s}$, the pole conditions of the constrained kernel.

d. Conjugation. From (A7) and the first two symmetries in (A2),

$${}_s Y_{\ell m}^* = (-1)^{s+m} {}_{-s} Y_{\ell, -m}, \quad (\text{A9})$$

which extends $Y_{\ell m}^* = (-1)^m Y_{\ell, -m}$. Consequently the complex conjugate of a spin- s field ${}_s f$ is the spin- $(-s)$ field with coefficients $({}_s f^*)_{\ell m} = (-1)^{s+m} {}_s f_{\ell, -m}^*$, the relation between the two members of a real field of spin s (Sec. IV B).

e. Addition theorem. For $x, y \in S^2$ and any spins s, s' ,

$$\sum_{m=-\ell}^{\ell} {}_{s'} Y_{\ell m}(x) {}_s Y_{\ell m}^*(y) = (-1)^{s+s'} \frac{2\ell+1}{4\pi} D_{-s',-s}^\ell(s(x)^{-1}s(y)), \quad (\text{A10})$$

with $s(\cdot)$ the section (32). The proof in Appendix B inserts (9) for both harmonics and uses $D^\ell(s(x))^\dagger D^\ell(s(y)) = D^\ell(s(x)^{-1}s(y))$. For noncoincident, nonantipodal points, write the

relative rotation as $R_z(a)R_y(b)R_z(c)$. Then b is the separation, and the bearings measured from e_θ towards e_φ are a at x towards y and $\pi - c$ at y towards x . Using $(-1)^{s+s'}d_{-s',-s}^\ell = d_{s',s}^\ell$, the right-hand side of (A10) becomes $(2\ell+1)e^{is'a}d_{s',s}^\ell(b)e^{isc}/(4\pi)$. The literal quadrant-resolved angles of Tessore *et al.* [28] are instead $\alpha = a + \pi$ and $\alpha' = -c$, the distinction described in Proposition VI.5. For $s = s' = 0$ this is the Legendre addition theorem, and for equal spins and $y = x$ it gives $\sum_m |s Y_{\ell m}(x)|^2 = (2\ell+1)/(4\pi)$.

Appendix B: Proofs

We collect here the proofs deferred from the main text, in the order in which the results appear. Throughout, fields on the sphere are square integrable, kernels are square integrable functions on $\text{SO}(3)$ unless they are explicitly said to be distributions, and identities between functions on $\text{SO}(3)$ or the sphere hold almost everywhere; since $\Psi \star F$ is continuous for $\Psi, F \in L^2(\text{SO}(3))$, by the strong continuity of the regular representation, its evaluation on a section is nevertheless well defined. We use freely that left translation acts on the first Wigner index, $(F(\rho^{-1} \cdot))_{mn}^\ell = \sum_{m'} D_{mm'}^\ell(\rho) F_{m'n}^\ell$, and right translation on the second, $(F(\cdot \rho))_{mn}^\ell = \sum_{n'} F_{mn'}^\ell D_{n'n}^\ell(\rho^{-1})$, both of which follow from (7), the change of variables $g \mapsto \rho g$ or $g \mapsto g\rho^{-1}$ and the representation property $D^\ell(\rho g) = D^\ell(\rho)D^\ell(g)$.

1. The twist

Proof of Lemma IV.3. Inserting $s(g_2x)s(g_2x)^{-1}$ into the definition (33) and regrouping,

$$\begin{aligned} \hbar(x, g_1 g_2) &= s(g_1 g_2 x)^{-1} g_1 g_2 s(x) \\ &= [s(g_1 g_2 x)^{-1} g_1 s(g_2 x)] [s(g_2 x)^{-1} g_2 s(x)] = \hbar(g_2 x, g_1) \hbar(x, g_2), \end{aligned} \quad (\text{B1})$$

where the first bracket is the twist of g_1 at the point $g_2 x$ and the second the twist of g_2 at x . On the sphere, for $0 < \beta < \pi$, the rotation $R(\alpha, \beta, \gamma)$ sends the North pole to the point (β, α) , so that $s(R(\alpha, \beta, \gamma)\eta) = R(\alpha, \beta, 0)$ and

$$\hbar(R(\alpha, \beta, \gamma)) = R(\alpha, \beta, 0)^{-1} R(\alpha, \beta, \gamma) = R_y(\beta)^{-1} R_z(\alpha)^{-1} R_z(\alpha) R_y(\beta) R_z(\gamma) = R_z(\gamma), \quad (\text{B2})$$

as claimed. At $\beta = 0$, the choice $s(\eta) = e$ instead gives $\hbar(R(\alpha, 0, \gamma)) = R_z(\alpha + \gamma)$; at $\beta = \pi$, the choice $s(-\eta) = R_y(\pi)$ gives $R_z(\gamma - \alpha)$. These endpoint choices do not affect any L^2 identity. $\square$

2. The lift, its adjoint and the induced representation

Proof of Lemma IV.7. (a) By Fubini's theorem, $\|\Lambda_s f\|^2 = \frac{1}{2\pi} \int_0^{2\pi} d\gamma \int_{S^2} |f|^2 d\mu = \|f\|^2$, so that Λ_s is an isometry, and

$$\begin{aligned} \langle \Lambda_s f, F \rangle &= \int_{S^2} d\mu(\beta, \alpha) f(\beta, \alpha) \left[\frac{1}{\sqrt{2\pi}} \int_0^{2\pi} e^{is\gamma} F(\alpha, \beta, \gamma) d\gamma \right]^* \\ &= \langle f, \Lambda_s^\dagger F \rangle \end{aligned} \quad (\text{B3})$$

identifies (37) as the adjoint. Next, $\Lambda_s^\dagger \Lambda_s f = f$ since $\frac{1}{2\pi} \int_0^{2\pi} e^{is\gamma} e^{-is\gamma} d\gamma = 1$. Since $R(\alpha, \beta, \gamma)R_z(\gamma') = R(\alpha, \beta, \gamma + \gamma')$, the lift $\Lambda_s f$ satisfies (35); conversely, $F \in \mathcal{M}_s$ has a fibrewise equivariant representative with $F(\alpha, \beta, \gamma) = e^{-is\gamma} F(\alpha, \beta, 0)$ for almost every (α, β, γ) , so that $F = \Lambda_s(\sqrt{2\pi} F \circ s)$ with $\sqrt{2\pi} F \circ s \in L^2(S^2)$ by Fubini's theorem, and the range of Λ_s is all of $\mathcal{M}_s$. Finally,

$$(\Lambda_s \Lambda_s^\dagger F)(\alpha, \beta, \gamma) = e^{-is\gamma} \frac{1}{2\pi} \int_0^{2\pi} e^{is\gamma'} F(\alpha, \beta, \gamma') d\gamma' \quad (\text{B4})$$

is the term in $e^{-is\gamma}$ of the Fourier series of $F(\alpha, \beta, \cdot)$ in γ , which for almost every (α, β) is a square integrable function on the circle. Hence $\Lambda_s \Lambda_s^\dagger$ is the orthogonal projector onto $\mathcal{M}_s$, the subspaces $\mathcal{M}_s$ are mutually orthogonal, and $\sum_s \Lambda_s \Lambda_s^\dagger F = F$ in $L^2(\text{SO}(3))$ by Plancherel's theorem on the circle and Fubini's theorem, which is the decomposition $L^2(\text{SO}(3)) = \widehat{\bigoplus_s \mathcal{M}_s}$.

(b) Using the fibrewise equivariant representative of $F \in \mathcal{M}_s$,

$$\begin{aligned} (\Lambda_s^\dagger F)(\beta, \alpha) &= \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} e^{is\gamma} e^{-is\gamma} F(\alpha, \beta, 0) d\gamma \\ &= \sqrt{2\pi} F(\alpha, \beta, 0). \end{aligned}$$

(d) By (9), $\Lambda_s {}_s Y_{\ell m} = (2\pi)^{-1/2} e^{-is\gamma} {}_s Y_{\ell m}(\beta, \alpha) = (-1)^s \sqrt{(2\ell+1)/8\pi^2} D_{m, -s}^{\ell*}$, which has unit norm by the orthogonality relation of the Wigner D -functions.

(c) By (d) and the orthogonality relation, the Wigner coefficients of $\Lambda_s {}_s Y_{\ell m}$ are $\int \Lambda_s {}_s Y_{\ell m} D_{m' n'}^\ell d\mu = (-1)^s \sqrt{8\pi^2/(2\ell+1)} \delta_{\ell\ell'} \delta_{mm'} \delta_{n', -s}$, and the first relation of (38) follows by linearity and the continuity of Λ_s . For the projection, $(\Lambda_s^\dagger F)_{\ell m} = \langle \Lambda_s^\dagger F, {}_s Y_{\ell m} \rangle = \langle F, \Lambda_s {}_s Y_{\ell m} \rangle = (-1)^s \sqrt{(2\ell+1)/8\pi^2} \int F D_{m, -s}^\ell d\mu$, which is the second relation.

(e) This is Proposition IV.10, proved next. $\square$

Proof of Proposition IV.10. Since $(\pi_G(g)F)(g'h) = F(g^{-1}g'h) = \rho^{(s)}(h^{-1})(\pi_G(g)F)(g')$ for $F \in \mathcal{M}_s$, the left regular representation preserves $\mathcal{M}_s$, and $\pi_s(g) = \Lambda_s^\dagger \pi_G(g) \Lambda_s$ is therefore a representation: by Lemma IV.7(a), $\pi_s(g)\pi_s(g') = \Lambda_s^\dagger \pi_G(g) \Lambda_s \Lambda_s^\dagger \pi_G(g') \Lambda_s = \Lambda_s^\dagger \pi_G(g)\pi_G(g') \Lambda_s = \pi_s(gg')$, since $\Lambda_s \Lambda_s^\dagger$ is the identity on the range of $\pi_G(g') \Lambda_s$, which lies in $\mathcal{M}_s$. It is unitary since $\pi_s(g)^\dagger = \Lambda_s^\dagger \pi_G(g)^\dagger \Lambda_s = \Lambda_s^\dagger \pi_G(g^{-1}) \Lambda_s = \pi_s(g^{-1})$, and it intertwines, $\Lambda_s \pi_s(g) = \Lambda_s \Lambda_s^\dagger \pi_G(g) \Lambda_s = \pi_G(g) \Lambda_s$, again because $\pi_G(g) \Lambda_s f \in \mathcal{M}_s$; taking adjoints gives $\pi_s(g) \Lambda_s^\dagger = \Lambda_s^\dagger \pi_G(g)$.

For the explicit form, write $(\Lambda_s f)(g) = (2\pi)^{-1/2} \rho^{(s)}(\hbar(g)^{-1}) f(gH)$ with $\hbar(g) = s(gH)^{-1}g$, which is Definition IV.6 by Lemma IV.3. Since $\pi_G(g) \Lambda_s f \in \mathcal{M}_s$, Lemma IV.7(b) gives

$$(\pi_s(g)f)(x) = \sqrt{2\pi} (\Lambda_s f)(g^{-1}s(x)) = \rho^{(s)}(\hbar(g^{-1}s(x))^{-1}) f(g^{-1}x), \quad (\text{B5})$$

and $\hbar(g^{-1}s(x)) = s(g^{-1}x)^{-1}g^{-1}s(x) = \hbar(x, g^{-1})$ by (33), which is (39). The same conclusion follows from Lemma IV.3 with $x = x_0$, $g_1 = g^{-1}$ and $g_2 = g'$, which splits $\hbar(g^{-1}g')$ into $\hbar(g'H, g^{-1})\hbar(g')$.

On the sphere, away from the coordinate singularities, with $g = R_\rho$ and $x = \omega$, the rotation $R_\rho^{-1}R_\omega$ sends the North pole to $R_\rho^{-1}\omega$, so that its Euler angles are $(\alpha', \beta', \vartheta)$ with $(\beta', \alpha') = R_\rho^{-1}\omega$, i.e. $R_\rho^{-1}R_\omega = R(\alpha', \beta', 0)R_z(\vartheta) = s(R_\rho^{-1}\omega)R_z(\vartheta)$. Hence $\hbar(\omega, R_\rho^{-1}) = s(R_\rho^{-1}\omega)^{-1}R_\rho^{-1}s(\omega) = R_z(\vartheta)$, and (39) reads $(\pi_s(R_\rho)f)(\omega) = e^{-is\vartheta} f(R_\rho^{-1}\omega)$, which is (12).

Finally, (10) follows by lifting: by Lemma IV.7(c) and the action of left translation on the first Wigner index,

$$\begin{aligned} (\pi_s(\rho)f)_{\ell m} &= (-1)^s \sqrt{\frac{2\ell+1}{8\pi^2}} (\pi_G(\rho) \Lambda_s f)_{m, -s}^\ell = (-1)^s \sqrt{\frac{2\ell+1}{8\pi^2}} \sum_{m'} D_{mm'}^\ell(\rho) (\Lambda_s f)_{m', -s}^\ell \\ &= \sum_{m'} D_{mm'}^\ell(\rho) {}_s f_{\ell m'}, \end{aligned} \quad (\text{B6})$$

the constants cancelling, which holds with the same matrices $D^\ell(\rho)$ for every spin. $\square$

Proof of Corollary IV.12. The decomposition and its uniqueness are Lemma IV.7(a). By (38), $(\Lambda_s \Lambda_s^\dagger F)_{mn}^\ell = (-1)^s \sqrt{8\pi^2/(2\ell+1)} (\Lambda_s^\dagger F)_{\ell m} \delta_{n, -s} = F_{m, -s}^\ell \delta_{n, -s}$, so that the component of spin s carries exactly the Wigner column $n = -s$; hence $F_{mn}^\ell = 0$ for all $|n| \geq N$ if and only if the components with $|s| \geq N$ vanish. $\square$

3. Convolution on the rotation group

Proof of (20), and of the equivariance, boundedness and adjoint of $\Psi\star$. Substituting $\rho = \rho'h^{-1}$, i.e. $h = \rho^{-1}\rho'$, in the Wigner transform of (19), which is permitted since the Haar measure is bi-invariant, and expanding $D_{mn}^\ell(\rho'h^{-1}) = \sum_k D_{mk}^\ell(\rho') D_{nk}^{\ell*}(h)$ by the representation property and unitarity,

$$\begin{aligned} (\Psi \star F)_{mn}^\ell &= \int_{\text{SO}(3)} \int_{\text{SO}(3)} \Psi^*(h) F(\rho') D_{mn}^\ell(\rho'h^{-1}) d\mu(h) d\mu(\rho') \\ &= \sum_k \left[\int_{\text{SO}(3)} F(\rho') D_{mk}^\ell(\rho') d\mu(\rho') \right] \left[\int_{\text{SO}(3)} \Psi^*(h) D_{nk}^{\ell*}(h) d\mu(h) \right] = \sum_k F_{mk}^\ell \Psi_{nk}^{\ell*}, \end{aligned} \quad (\text{B7})$$

where the interchange of integrals is justified by Fubini's theorem since $\Psi, F \in L^2(\text{SO}(3)) \subset L^1(\text{SO}(3))$. For equivariance,

$$\begin{aligned} (\Psi \star \pi_G(g)F)(\rho) &= \int \Psi^*(\rho^{-1}\rho') F(g^{-1}\rho') d\mu(\rho') \\ &= \int \Psi^*(\rho^{-1}g\rho'') F(\rho'') d\mu(\rho'') = (\Psi \star F)(g^{-1}\rho). \end{aligned} \quad (\text{B8})$$

For boundedness, in the notation F^ℓ for the $(2\ell+1) \times (2\ell+1)$ block of coefficients, (20) reads $(\Psi \star F)^\ell = F^\ell \Psi^{\ell\dagger}$, hence by Parseval's identity $\|\Psi \star F\|^2 = \sum_\ell \frac{2\ell+1}{8\pi^2} \|F^\ell \Psi^{\ell\dagger}\|_F^2 \leq \sup_\ell \|\Psi^\ell\|_{\text{op}}^2 \|F\|^2$, and $\|\Psi^\ell\|_{\text{op}}^2 \leq \|\Psi^\ell\|_F^2 \leq \frac{8\pi^2}{2\ell+1} \|\Psi\|^2$, where $\|\cdot\|_F$ denotes the Frobenius norm. For the adjoint, by Fubini's theorem

$$\langle \Psi \star F, G \rangle = \int_{\text{SO}(3)} F(g') \left[\int_{\text{SO}(3)} \Psi(g^{-1}g') G(g) d\mu(g) \right]^* d\mu(g'), \quad (\text{B9})$$

and $\int \Psi(g^{-1}g') G(g) d\mu(g) = \int \tilde{\Psi}^*(g'^{-1}g) G(g) d\mu(g) = (\tilde{\Psi} \star G)(g')$ with $\tilde{\Psi}(k) = \Psi^*(k^{-1})$, so that $(\Psi\star)^\dagger = \tilde{\Psi}\star$. As a check, the Wigner coefficients of the flipped kernel are $\tilde{\Psi}_{mn}^\ell = \int \Psi^*(g^{-1}) D_{mn}^\ell(g) d\mu(g) = \int \Psi^*(g) D_{nm}^{\ell*}(g) d\mu(g) = \Psi_{nm}^{\ell*}$, i.e. $\tilde{\Psi}^\ell = \Psi^{\ell\dagger}$, and $(\tilde{\Psi} \star G)^\ell = G^\ell \Psi^\ell$ is indeed the adjoint of $F^\ell \mapsto F^\ell \Psi^{\ell\dagger}$ with respect to the Frobenius inner product. $\square$

4. Harmonic form, kernel constraint and composition

Proof of Theorem V.2. By (38), $(\Lambda_s s f)_{mm'}^\ell = (-1)^s \sqrt{8\pi^2/(2\ell+1)} {}_s f_{\ell m} \delta_{m',-s}$, so that the sum over m' in (20) collapses to the single term $m' = -s$, which gives (42). Applying $\Lambda_{s'}^\dagger$ by (38) reads the column $n = -s'$ and multiplies by $(-1)^{s'} \sqrt{(2\ell+1)/8\pi^2}$, so that the degree-dependent constants cancel and (43) follows; (44) is $\Lambda_{s'}^\dagger$, applied to (20), and (45) is (20) itself. Degrees $\ell < |s|$ are absent from $L_s^2(S^2)$ and degrees $\ell < |s'|$ from $L_{s'}^2(S^2)$, whence the restriction on ℓ . Boundedness follows from that of Λ_s , $\Lambda_{s'}^\dagger$ and $\Psi\star$. For equivariance,

$$\begin{aligned} \mathcal{C}_{s \rightarrow s'}[\Psi] \pi_s(g) &= \Lambda_{s'}^\dagger (\Psi \star \Lambda_s \pi_s(g) \cdot) = \Lambda_{s'}^\dagger (\Psi \star \pi_G(g) \Lambda_s \cdot) \\ &= \Lambda_{s'}^\dagger \pi_G(g) (\Psi \star \Lambda_s \cdot) = \pi_{s'}(g) \mathcal{C}_{s \rightarrow s'}[\Psi] \end{aligned} \quad (\text{B10})$$

by Lemma IV.7(e) and the equivariance of $\Psi\star$; the other three cases drop the lift, the projection, or both. The adjoint identities follow from $(\Psi\star)^\dagger = \tilde{\Psi}\star$ and $(\Lambda_s)^\dagger = \Lambda_s^\dagger$: for instance $(\mathcal{C}_{s \rightarrow G}[\Psi])^\dagger = (\Psi \star \Lambda_s)^\dagger = \Lambda_s^\dagger \tilde{\Psi}\star = \mathcal{C}_{G \rightarrow s}[\tilde{\Psi}]$. $\square$

Proof of Corollary V.4. By Lemma IV.7(a) and (c), a function on $\text{SO}(3)$ lies in $\mathcal{M}_{s'}$ if and only if its Wigner coefficients vanish outside the column $n = -s'$. By (42), the coefficients

of $\Psi \star \Lambda_s f$ are $(-1)^s \sqrt{8\pi^2/(2\ell+1)} {}_s f_{\ell m} \Psi_{n,-s}^{\ell*}$, which vanish for all $n \neq -s'$ and all f if and only if $\Psi_{n,-s}^\ell = 0$ for all $n \neq -s'$ and all $\ell \geq |s|$, as is seen by taking $f = {}_s Y_{\ell m}$; this is the equivalence of (i) and (ii). For (iii), the Wigner coefficients of the right-isotypic part are

$$\begin{aligned} (\Psi^{(s)})_{nm'}^\ell &= \frac{1}{2\pi} \int_0^{2\pi} e^{is\gamma} \int_{\text{SO}(3)} \Psi(gR_z(\gamma)) D_{nm'}^\ell(g) d\mu(g) d\gamma \\ &= \frac{1}{2\pi} \int_0^{2\pi} e^{i(s+m')\gamma} d\gamma \Psi_{nm'}^\ell = \delta_{m',-s} \Psi_{nm'}^\ell, \end{aligned} \quad (\text{B11})$$

by the substitution $g' = gR_z(\gamma)$ and $D_{nm'}^\ell(g'R_z(-\gamma)) = D_{nm'}^\ell(g')e^{im'\gamma}$: the right-isotypic part is the column $m' = -s$ of Ψ , which is exactly the part read by (42). Since $D_{nm'}^{\ell*}(R_z(\chi_1)gR_z(\chi_2)) = e^{in\chi_1} D_{nm'}^{\ell*}(g) e^{im'\chi_2}$, the Wigner expansion of $\Psi^{(s)}$ satisfies (47) if and only if its coefficients vanish unless $n = -s'$ and $m' = -s$, which is (ii). In that case $\Psi \star \Lambda_s f \in \mathcal{M}_{s'}$ and Lemma IV.7(b) gives $\mathcal{C}_{s \rightarrow s'}[\Psi]f = \sqrt{2\pi} (\Psi \star \Lambda_s f) \circ s$; the Wigner expansion (7) restricted to the entry $(-s', -s)$, with $D_{-s',-s}^{\ell*}(\alpha, \beta, \gamma) = e^{-is'\alpha} d_{-s',-s}^\ell(\beta) e^{-is\gamma}$, is (48). Finally, for any $\Psi = \Lambda_s \psi$, (38) gives $\Psi_{-s',-s}^\ell = (-1)^s \sqrt{8\pi^2/(2\ell+1)} {}_s \psi_{\ell,-s'}$, and substituting into (43) yields (49). $\square$

Proof of Proposition V.3. In the block notation of the proof of (20), $(\Psi_2 \star (\Psi_1 \star F))^\ell = F^\ell \Psi_1^{\ell\dagger} \Psi_2^{\ell\dagger} = F^\ell (\Psi_2^\ell \Psi_1^\ell)^\dagger$, which is (46) provided a kernel $\Psi_3 \in L^2(\text{SO}(3))$ with blocks $\Psi_2^\ell \Psi_1^\ell$ exists; it does, namely $\Psi_3 = \tilde{\Psi}_1 \star \Psi_2$, since $(\tilde{\Psi}_1 \star \Psi_2)^\ell = \Psi_2^\ell \tilde{\Psi}_1^{\ell\dagger} = \Psi_2^\ell \Psi_1^\ell$ by $\tilde{\Psi}_1 = \Psi_1^{\ell\dagger}$, and $\Psi_3 \in L^2(\text{SO}(3))$ by the boundedness of $\tilde{\Psi}_1 \star$. The cases with an outer lift or projection, or both, follow by composing with Λ_s on the right or $\Lambda_{s'}^\dagger$ on the left. When the intermediate stage is a spin- s' field, the operator $\Lambda_{s'} \Lambda_{s'}^\dagger$ is inserted between the two convolutions; by Corollary IV.12 its action on Wigner coefficients is $F^\ell \mapsto F^\ell E_{-s'}$, with $E_{-s'}$ the diagonal matrix with a single unit entry at $-s'$, so that the composed coefficients are $F^\ell \Psi_1^{\ell\dagger} E_{-s'} \Psi_2^{\ell\dagger} = F^\ell (\Psi_2^\ell E_{-s'} \Psi_1^\ell)^\dagger$, and the block $\Psi_2^\ell E_{-s'} \Psi_1^\ell$ has entries $(\Psi_2)_{n,-s'}^\ell (\Psi_1)_{-s',m'}^\ell$. Reading the entry $(-s'', -s)$ gives the spin-to-spin case, with $c_\ell^{(3)} = (-1)^{s+s''} [(\Psi_2)_{-s'',-s'}^\ell (\Psi_1)_{-s',-s}^\ell]^* = c_\ell^{(2)} c_\ell^{(1)}$ by (43), and reading the column $m' = -s$ gives the remaining case. $\square$

5. Classification

Proof of Theorem V.5. By Lemma IV.7(a) and (d), Λ_s is a unitary equivalence between $L_s^2(S^2)$ and $\mathcal{M}_s = \bigoplus_{\ell \geq |s|} V_\ell^{(s)}$, where $V_\ell^{(s)} = \text{span}\{D_{m,-s}^{\ell*} : |m| \leq \ell\}$ is the image of the degree- ℓ spin harmonics, and by the action of left translation on the first Wigner index, $\pi_G(\rho) D_{m,-s}^{\ell*} = \sum_k D_{km}^\ell(\rho) D_{k,-s}^{\ell*}$, so that each $V_\ell^{(s)}$ is invariant and carries the irreducible representation of $\text{SO}(3)$ of degree ℓ , in the same matrices for every s ; this is (10). Thus $L_s^2(S^2)$ decomposes into irreducible representations without multiplicity, each degree $\ell \geq |s|$ occurring exactly once. This is Frobenius reciprocity [33, 34] in the present case: the multiplicity of the degree- ℓ representation in $\text{Ind}_{\text{SO}(2)}^{\text{SO}(3)} \rho^{(s)}$ equals the multiplicity of $\rho^{(s)}$ in the restriction of the degree- ℓ representation to $\text{SO}(2)$, which is one for $|s| \leq \ell$ and zero otherwise since $D_{mn}^\ell(R_z(\gamma)) = e^{-in\gamma} \delta_{mn}$ contains each character with $|n| \leq \ell$ exactly once.

Let T be bounded and equivariant. Since irreducible representations of different degree are inequivalent, T maps the degree- ℓ component of $L_s^2(S^2)$ into the degree- ℓ component of $L_{s'}^2(S^2)$, and annihilates it if $\ell < |s'|$. For $\ell \geq \max(|s|, |s'|)$ the two components are equivalent irreducible representations, realised in the bases ${}_s Y_{\ell m}$ and ${}_{s'} Y_{\ell m}$ by the same matrices $D^\ell(\rho)$, so that by Schur's lemma the restriction of T is a scalar c_ℓ times the identity in these bases, which is (50); the sequence is unique, and $\|T\| = \sup_\ell |c_\ell|$ by Parseval's identity, so that T is bounded if and only if the sequence is. Conversely, a bounded sequence defines a bounded T by (50), which commutes with π_s by (10).

(a) If $T = \mathcal{C}_{s \rightarrow s'}[\Psi]$ with $\Psi \in L^2(\text{SO}(3))$, then $c_\ell = (-1)^{s+s'} \Psi_{-s', -s}^{\ell*}$ by (43), and $\sum_\ell \frac{2\ell+1}{8\pi^2} |c_\ell|^2 \leq \|\Psi\|^2 < \infty$ by Parseval's identity. Conversely, if $\sum_\ell (2\ell+1) |c_\ell|^2 < \infty$, the function Ψ_T with Wigner coefficients $\Psi_{-s', -s}^\ell = (-1)^{s+s'} c_\ell^*$ and all others zero lies in $L^2(\text{SO}(3))$, with $\|\Psi_T\|^2 = \sum_\ell \frac{2\ell+1}{8\pi^2} |c_\ell|^2$, and $\mathcal{C}_{s \rightarrow s'}[\Psi_T] = T$ by (43). Since the Hilbert–Schmidt norm of T is $\|T\|_{\text{HS}}^2 = \sum_{\ell m} |c_\ell|^2 = \sum_\ell (2\ell+1) |c_\ell|^2$, the condition is that T be Hilbert–Schmidt. Any Ψ with $\mathcal{C}_{s \rightarrow s'}[\Psi] = T$ has the same entries $\Psi_{-s', -s}^\ell$, all other entries being unread, and Ψ_T is the unique such kernel whose entire Wigner block is supported on that entry, equivalently whose entire kernel satisfies the two-sided law of (47). Constraining only $\Psi^{(s)}$ would still permit arbitrary unread input columns and would not give uniqueness.

(b) For a bounded sequence, Ψ_T is a distribution on $\text{SO}(3)$ with bounded Wigner coefficients, and $\mathcal{C}_{s \rightarrow s'}[\Psi_T]$ is defined by (43) on the dense subspace of band-limited fields, on which it coincides with T ; since $|c_\ell|$ is bounded it extends by continuity to $L_s^2(S^2)$, and the extension is T . This is the sense in which every bounded equivariant map is a lifted convolution with a distributional kernel.

(c) The left regular representation decomposes as $L^2(\text{SO}(3)) = \widehat{\bigoplus}_\ell \bigoplus_{|n| \leq \ell} V_\ell^{(n)}$, with $V_\ell^{(n)} = \text{span}\{D_{mn}^{\ell*} : |m| \leq \ell\}$ the n -th Wigner column, each column carrying the degree- ℓ representation in the same matrices; the degree- ℓ representation thus occurs with multiplicity $2\ell+1$, one copy per column. Schur's lemma gives a vector of coefficients u^ℓ from a spin field into the Wigner columns, a vector v^ℓ from those columns into a spin field, and an arbitrary matrix B^ℓ acting on the multiplicity index for a group-to-group map, as in (51). Conversely these coefficient maps commute degree by degree with the action on the first index. For the first case Parseval's identity gives

$$\|T_{s \rightarrow G} f\|^2 = \sum_{\ell \geq |s|} a_\ell \|u^\ell\|_2^2 \sum_m |{}_s f_{\ell m}|^2. \quad (\text{B12})$$

For the second, Cauchy–Schwarz gives

$$\|T_{G \rightarrow s'} F\|^2 \leq \sup_{\ell \geq |s'|} \frac{\|v^\ell\|_2^2}{a_\ell} \sum_\ell a_\ell \|F^\ell\|_{\text{F}}^2, \quad (\text{B13})$$

with equality on a single row of a single block proportional to $(v^\ell)^*$; taking degrees approaching the supremum gives the exact norm. Similarly, $\|F^\ell B^\ell\|_{\text{F}} \leq \|F^\ell\|_{\text{F}} \|B^\ell\|_{\text{op}}$, with equality attained on a suitable row, gives the third norm in (52). Thus those uniform bounds are necessary and sufficient for extension from finite harmonic sums to the Hilbert direct sums; entrywise bounds alone do not suffice.

Comparing with Theorem V.2 gives (53). The norm bounds make the canonical kernel blocks uniformly bounded in operator norm, and hence define distributions. Summing the squared images of orthonormal bases, with the factor $a_\ell^{-1/2}$ for a unit-norm Wigner coefficient, gives

$$\begin{aligned} \|T_{s \rightarrow s'}\|_{\text{HS}}^2 &= \sum_\ell d_\ell |c_\ell|^2, & \|T_{s \rightarrow G}\|_{\text{HS}}^2 &= \sum_\ell d_\ell a_\ell \|u^\ell\|_2^2, \\ \|T_{G \rightarrow s'}\|_{\text{HS}}^2 &= \sum_\ell \frac{d_\ell}{a_\ell} \|v^\ell\|_2^2, & \|T_{G \rightarrow G}\|_{\text{HS}}^2 &= \sum_\ell d_\ell \|B^\ell\|_{\text{F}}^2. \end{aligned} \quad (\text{B14})$$

Each sum is over the degrees present in both spaces. By Parseval's identity each is $8\pi^2 \|\Psi_T\|^2$ for the canonical kernel with unread entries set to zero, so a square integrable representative exists exactly when the map is Hilbert–Schmidt. Adding unread coefficients cannot remove a divergence in the canonical part. $\square$

For completeness, the unbounded extension (54) is dense because it contains all finite harmonic sums. If $f_j \rightarrow f$ and $Tf_j \rightarrow h$ in the respective L^2 spaces, convergence of every coefficient gives $h_{\ell m} = c_\ell f_{\ell m}$; since $h \in L^2$, $f \in \mathcal{D}(T)$ and $Tf = h$, proving closedness.

Rotations preserve each degree norm and hence the domain, and commute with its scalar multiplier. Polynomial growth gives a distributional kernel by the rapid decay of smooth test-function coefficients; arbitrary growth need not do so. The first- and second-order Sobolev assertions follow from the respective growth rates $|c_\ell| \asymp \ell$ and ℓ^2 , with finitely many low modes immaterial to the domain.

Proof of Corollary V.6. This is Theorem V.5(c), the row and column cases. By (38), $\Lambda_s^\dagger$ reads the column $n = -s'$ with the constant $(-1)^{s'} \sqrt{(2\ell+1)/8\pi^2}$, which is (44) with $\Psi_{-s',m'}^{\ell*} = \delta_{m',-s'}$, and Λ_s is (42) with $\Psi_{n,-s}^{\ell*} = \delta_{n,-s}$. $\square$

Proof of Corollary V.7. By (50),

$$\begin{aligned} (Tf)(x) &= \sum_{\ell} c_{\ell} \sum_m s' Y_{\ell m}(x) \langle f, {}_s Y_{\ell m} \rangle = \int_{S^2} K(x, y) f(y) d\mu(y), \\ K(x, y) &= \sum_{\ell} c_{\ell} \sum_m s' Y_{\ell m}(x) {}_s Y_{\ell m}^*(y), \end{aligned} \quad (\text{B15})$$

the series converging in the sense of distributions. By (9) at $\gamma = 0$, ${}_s Y_{\ell m}(x) = (-1)^{s'} \sqrt{(2\ell+1)/4\pi} D_{m,-s'}^{\ell*}(s(x))$ and ${}_s Y_{\ell m}^*(y) = (-1)^s \sqrt{(2\ell+1)/4\pi} D_{m,-s}^{\ell}(s(y))$, so that, by $D_{mn}^{\ell*}(g) = D_{nm}^{\ell}(g^{-1})$ and the representation property,

$$\begin{aligned} \sum_m s' Y_{\ell m}(x) {}_s Y_{\ell m}^*(y) &= (-1)^{s+s'} \frac{2\ell+1}{4\pi} \sum_m D_{-s',m}^{\ell}(s(x)^{-1}) D_{m,-s}^{\ell}(s(y)) \\ &= (-1)^{s+s'} \frac{2\ell+1}{4\pi} D_{-s',-s}^{\ell}(s(x)^{-1} s(y)), \end{aligned} \quad (\text{B16})$$

which is the addition theorem for spin spherical harmonics and gives (55). The rotation $s(x)^{-1} s(y)$ sends the North pole to $s(x)^{-1} y$, the point y expressed in the frame of x , whose colatitude is the geodesic distance β_{xy} between x and y ; writing its Euler angles as $(\alpha_{xy}, \beta_{xy}, \gamma_{xy})$, the kernel reads $D_{-s',-s}^{\ell}(\alpha_{xy}, \beta_{xy}, \gamma_{xy}) = e^{is'\alpha_{xy}} d_{-s',-s}^{\ell}(\beta_{xy}) e^{is\gamma_{xy}}$, in which only the d -function depends on the distance and the two exponentials are the phases of a kernel with output spin s' in x and dual input spin $-s$ in y : under a frame change, $K'(x, y) = e^{-is'\chi(x)} K(x, y) e^{+is\chi(y)}$ by (36). For $s = s' = 0$, $d_{00}^{\ell}(\beta) = P_{\ell}(\cos \beta)$ and $\cos \beta_{xy} = x \cdot y$. $\square$

6. Reductions

Proof of Proposition VI.1. By Definition III.4, $(\Lambda_s \psi \star \Lambda_s f)(g) = \langle \Lambda_s f, \pi_G(g) \Lambda_s \psi \rangle = \langle \Lambda_s f, \Lambda_s \pi_s(g) \psi \rangle = \langle f, \pi_s(g) \psi \rangle$, by Lemma IV.7(e) and the isometry of the lift, and $\pi_s(g) = \mathcal{R}_g$ by Proposition IV.10, so that this is (22). In harmonic space, (42) with $\Psi = \Lambda_s \psi$ and (38) give $(-1)^s \sqrt{8\pi^2/(2\ell+1)} {}_s f_{\ell m} [(-1)^s \sqrt{8\pi^2/(2\ell+1)} {}_s \psi_{\ell n}]^* = \frac{8\pi^2}{2\ell+1} {}_s f_{\ell m} {}_s \psi_{\ell n}^*$. $\square$

Proof of Proposition VI.2. Since $\pi_s(\rho) \psi = \Lambda_s^\dagger \pi_G(\rho) \Lambda_s \psi$ and $\Lambda_s^\dagger$ is bounded,

$$\int_{\text{SO}(3)} W(\rho) \pi_s(\rho) \psi d\mu(\rho) = \Lambda_s^\dagger \int_{\text{SO}(3)} W(\rho) \pi_G(\rho) \Phi d\mu(\rho), \quad (\text{B17})$$

where $\Phi = \Lambda_s \psi$ and the integrals are Bochner integrals in $L^2(\text{SO}(3))$. Evaluating at g and noting that $\tilde{\Phi}^*(g^{-1} \rho) = \Phi(\rho^{-1} g)$,

$$\int_{\text{SO}(3)} W(\rho) \Phi(\rho^{-1} g) d\mu(\rho) = \int_{\text{SO}(3)} \tilde{\Phi}^*(g^{-1} \rho) W(\rho) d\mu(\rho) = (\tilde{\Phi} \star W)(g), \quad (\text{B18})$$

so that the synthesis is $\mathcal{C}_{G \rightarrow s}[\tilde{\Phi}]W$, which is $(\mathcal{C}_{s \rightarrow G}[\Phi])^\dagger W$ by Theorem V.2. In harmonic space, $\tilde{\Phi}_{-s, m'}^\ell = \Phi_{m', -s}^{\ell*} = (-1)^s \sqrt{8\pi^2/(2\ell+1)} {}_s\psi_{\ell m'}^*$ by (38), so that (44) gives

$$(\mathcal{C}_{G \rightarrow s}[\tilde{\Phi}]W)_{\ell m} = (-1)^s \sqrt{\frac{2\ell+1}{8\pi^2}} \sum_{m'} W_{mm'}^\ell (-1)^s \sqrt{\frac{8\pi^2}{2\ell+1}} {}_s\psi_{\ell m'} = \sum_{m'} W_{mm'}^\ell {}_s\psi_{\ell m'}, \quad (\text{B19})$$

the constants cancelling. $\square$

Proof of Proposition VI.3. By Proposition VI.1 evaluated at $g = s(\omega) = R_\omega$, $\langle f, \mathcal{R}_\omega \psi \rangle = (\mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]f)(s(\omega))$. For an axisymmetric kernel, $\psi_{\ell m} = \psi_{\ell 0} \delta_{m0}$, the kernel $\Lambda_0 \psi$ has, by (38), Wigner support at the single entry $(0, 0)$ of each degree, so that it satisfies condition (ii) of Corollary V.4 with $s = s' = 0$, and hence $\mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]f \in \mathcal{M}_0$; Lemma IV.7(b) then gives $(\mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]f) \circ s = (2\pi)^{-1/2} \Lambda_0^\dagger \mathcal{C}_{0 \rightarrow G}[\Lambda_0 \psi]f = (2\pi)^{-1/2} \mathcal{C}_{0 \rightarrow 0}[\Lambda_0 \psi]f$. The harmonic form follows from (49) with $s = s' = 0$, divided by $\sqrt{2\pi}$: $(2\pi)^{-1/2} \sqrt{8\pi^2/(2\ell+1)} = \sqrt{4\pi/(2\ell+1)}$. The spin case is identical with $s' = 0$: the kernel $\Lambda_s \Phi$ with ${}_s\Phi_{\ell m} = {}_s\Phi_{\ell 0} \delta_{m0}$ satisfies (ii) with target spin 0, and (49) gives (25). If ψ is not axisymmetric the constraint fails and the section evaluation is the standard correlation, treated in Proposition VII.1. $\square$

Proof of Proposition VI.4. With $\Psi_h(g) = h^*(g^{-1}\eta)$ and $k = g^{-1}g'$, $\Psi_h^*(g^{-1}g') = h(k^{-1}\eta) = h(g'^{-1}g\eta)$, so that by Definition III.4

$$(\Psi_h \star \Lambda_0 f)(g) = \frac{1}{\sqrt{2\pi}} \int_{\text{SO}(3)} h(g'^{-1}g\eta) f(g'\eta) d\mu(g'). \quad (\text{B20})$$

Applying $\Lambda_0^\dagger$ at the point $\omega = (\beta, \alpha)$ integrates over γ with $g = R(\alpha, \beta, \gamma)$, and since $g\eta = \omega$ does not depend on γ the integrand is constant, whence $(\mathcal{C}_{0 \rightarrow 0}[\Psi_h]f)(\omega) = \frac{1}{2\pi} \cdot 2\pi \int h(g'^{-1}\omega) f(g'\eta) d\mu(g') = (f \star_L h)(\omega)$, exactly. For the constant, only the entry $(\Psi_h)_{00}^\ell$ is read by (43). Writing $g = R(\alpha, \beta, \gamma)$, the point $g^{-1}\eta = R_z(-\gamma)R_y(-\beta)\eta$ has colatitude β and longitude $\pi - \gamma$, and $D_{00}^\ell(g) = d_{00}^\ell(\beta) = P_\ell(\cos \beta) = \sqrt{4\pi/(2\ell+1)} Y_{\ell 0}(\beta, \cdot)$, so that

$$\begin{aligned} (\Psi_h)_{00}^\ell &= \int_{\text{SO}(3)} h^*(\beta, \pi - \gamma) \sqrt{\frac{4\pi}{2\ell+1}} Y_{\ell 0}(\beta) \sin \beta d\alpha d\beta d\gamma \\ &= 2\pi \sqrt{\frac{4\pi}{2\ell+1}} \int_{S^2} h^*(\omega') Y_{\ell 0}(\omega') d\mu(\omega') = 2\pi \sqrt{\frac{4\pi}{2\ell+1}} h_{\ell 0}^*, \end{aligned} \quad (\text{B21})$$

where the integral over α gives 2π and the substitution $\varphi' = \pi - \gamma$ runs over a full period. Hence (43) with $s = s' = 0$ gives $(f \star_L h)_{\ell m} = f_{\ell m} (\Psi_h)_{00}^{\ell*} = 2\pi \sqrt{4\pi/(2\ell+1)} f_{\ell m} h_{\ell 0}$, which is (16), with no conjugate on h since the pairing (15) is bilinear. Finally, the marginalisation over γ of the bilinear directional convolution $\int f(\omega)(\mathcal{R}_\rho h)(\omega) d\mu(\omega)$, which is a function on $\text{SO}(3)$, is by definition $\sqrt{2\pi}$ times its projection $\Lambda_0^\dagger$. $\square$

Proof of Proposition VI.6. By (42) the Wigner coefficients of $\mathcal{C}_{s \rightarrow G}[\Psi]f$ vanish for $|n| \geq N$ when the column $m' = -s$ of Ψ does. Corollary IV.12 applied to the output gives $\mathcal{C}_{s \rightarrow G}[\Psi]f = \sum_{|s'| < N} \Lambda_{s'} \Lambda_{s'}^\dagger \mathcal{C}_{s \rightarrow G}[\Psi]f$, and $\Lambda_{s'}^\dagger \mathcal{C}_{s \rightarrow G}[\Psi] = \mathcal{C}_{s \rightarrow s'}[\Psi]$ by Definition V.1. Conversely, given at most $2N - 1$ bounded multipliers $c_\ell^{(s')}$, the canonical column $\Psi_{-s', -s}^\ell = (-1)^{s+s'} c_\ell^{(s')*}$ defines a distribution and assembles the tuple by the same identity. The norm bound is $\sup_\ell (\sum_{s'} |c_\ell^{(s')}|^2)^{1/2} < \infty$. By Parseval, the kernel belongs to L^2 exactly when $\sum_{\ell, s'} (2\ell+1) |c_\ell^{(s')}|^2 < \infty$, the sum of the component Hilbert–Schmidt norms squared. This condition holds automatically at finite band-limit. $\square$

Proof of Proposition VI.8(c). Wallis *et al.* [3] write the shear as

$${}_2\gamma(\omega) = \int_{S^2} (\mathcal{R}_{\omega'} {}_2K)(\omega) \kappa(\omega') d\mu(\omega'), \quad (\text{B22})$$

with an axisymmetric spin-2 kernel, ${}_2K_{\ell m} = {}_2K_{\ell 0}\delta_{m0}$. By (10) and (8),

$$(\mathcal{R}_{\omega'} {}_2K)_{\ell m} = D_{m0}^\ell(\varphi', \theta', 0) {}_2K_{\ell 0} = \sqrt{\frac{4\pi}{2\ell+1}} Y_{\ell m}^*(\omega') {}_2K_{\ell 0}, \quad (\text{B23})$$

and integrating against κ gives ${}_2\gamma_{\ell m} = \sqrt{4\pi/(2\ell+1)} {}_2K_{\ell 0} \kappa_{\ell m}$, a multiplier from spin 0 to spin 2; it is $\mathcal{C}_{0 \rightarrow 2}[\Psi_T]$ with $c_\ell = \sqrt{4\pi/(2\ell+1)} {}_2K_{\ell 0}$ by Theorem V.5, and its kernel is (55) with $(s, s') = (0, 2)$: the integral is the adjoint synthesis of the spin-2 to scalar axisymmetric analysis of Proposition VI.3, not that forward analysis. The choice ${}_2K_{\ell 0} = \sqrt{(2\ell+1)/4\pi} \mathcal{D}_\ell$ gives $c_\ell = \mathcal{D}_\ell$, which is the spherical Kaiser–Squires relation. $\square$

7. Beyond the framework

Proof of Proposition VII.1. Let $F = \Psi \star \Lambda_s f$. It is continuous, and $\|F\|_\infty \leq \|\Psi\|_2 \|f\|_2$ by Cauchy–Schwarz and the isometry of the lift, so section evaluation is a bounded map into the coordinate L^2 space. The explicit support hypothesis $\Psi_{n,-s}^\ell = 0$ for $|n| \geq N$ makes its fibre decomposition a finite sum. Evaluating each term at $\sigma(x) = s(x)R_z(\chi_\sigma(x))$ gives (62), since $\Lambda_q^\dagger F = \mathcal{C}_{s \rightarrow q}[\Psi]f$.

Suppose E_σ is equivariant into spin t with the fixed coordinate-frame action. By Theorem V.5 it acts on degree ℓ by a scalar λ_ℓ , or annihilates that degree if $\ell < |t|$. Put $a_{\ell q} = (-1)^{s+q} \Psi_{-q,-s}^{\ell*} / \sqrt{2\pi}$. Testing on every ${}_s Y_{\ell m}$ gives, almost everywhere,

$$\sum_{\substack{|q| < N \\ |q| \leq \ell}} a_{\ell q} e^{-iq\chi_\sigma(x)} {}_q Y_{\ell m}(x) = \lambda_\ell {}_t Y_{\ell m}(x), \quad (\text{B24})$$

with the right-hand side interpreted as zero when $\ell < |t|$. For a fixed x , the vectors $({}_q Y_{\ell m}(x))_{m=-\ell}^\ell$ for distinct admissible q are nonzero multiples of distinct columns of a unitary Wigner matrix, by (A7). They are therefore independent. Since there are countably many tests, a common full-measure set may be used. Comparing columns in (B24) forces $a_{\ell q} = 0$ for $q \neq t$ and all components to vanish when the target degree is absent. If some $a_{\ell t} \neq 0$, its column gives $e^{-it\chi_\sigma(x)} = \lambda_\ell / a_{\ell t}$ almost everywhere. Thus the phase is constant. If every coefficient vanishes, $E_\sigma = 0$, which is equivariant into every target spin independently of the section. Conversely, one active output spin and a constant phase give a scalar multiple of $\mathcal{C}_{s \rightarrow t}$ and hence equivariance.

For polar rotations, $R_z(a)s(x) = s(R_z(a)x)$ away from the poles, so all spins act by coordinate rotation. The phase for σ_0 is constant. For $\chi_\sigma = \pi - \varphi$, one has $\chi_\sigma(R_z(-a)x) = \chi_\sigma(x) + a$ modulo 2π . Testing harmonics and comparing the same independent columns shows that commutation is equivalent to $a_{\ell q}(1 - e^{-iqa}) = 0$ for every ℓ, q . This is $e^{iqa} = 1$ for every active spin q , including the vacuous condition for the zero map; it makes no assertion about non-polar symmetries. $\square$

Proof of Proposition VII.2. For a rotation $R_z(a)$ about the polar axis the twist vanishes, so that $(\pi_0(R_z(a))f)_{\ell m} = e^{-ima} f_{\ell m}$ by (10) and (6); a diagonal multiplier commutes with this action for every $g_{\ell m}$. Equivariance under all of $\text{SO}(3)$ is, by Theorem V.5 with $s = s' = 0$, the condition that T act by a scalar on each degree, *i.e.* that $g_{\ell m}$ be independent of m , in which case T is the axisymmetric convolution $\mathcal{C}_{0 \rightarrow 0}[\Psi_T]$ with $c_\ell = g_{\ell 0}$, understood in the distributional multiplier sense. Moreover, $D^\ell(R_y(\pi))$ exchanges m and $-m$ up to the nonzero sign in (A3), so it commutes with the diagonal multiplier when $g_{\ell m} = g_{\ell, -m}$. Together with the polar rotations this gives the embedded $O(2)$ normaliser. $\square$

8. Coordinate-section standard correlation

The factorisation-of-rotations method [9, 19, 24] gives an $\mathcal{O}(L^3)$ evaluation of standard correlation on an equiangular coordinate grid. Set $\Delta_{km}^\ell = d_{km}^\ell(\pi/2)$. In the conventions of

(6),

$$d_{mn}^\ell(\theta) = i^{m-n} \sum_{k=-\ell}^{\ell} \Delta_{km}^\ell \Delta_{kn}^\ell e^{-ik\theta}. \quad (\text{B25})$$

Substituting this identity in (18) at $\gamma = 0$ and contracting over n before summing over ℓ gives

$$G_k^\ell = \sum_{n=-\ell}^{\ell} i^{-n} \Delta_{kn}^\ell \psi_{\ell n}^*,$$

$$H_{mk} = \sum_{\ell=\max(|m|,|k|)}^{L-1} i^m \Delta_{km}^\ell f_{\ell m} G_k^\ell, \quad (\text{B26})$$

$$\langle f, \mathcal{R}_{(\varphi, \theta, 0)} \psi \rangle = \sum_{m,k=-(L-1)}^{L-1} H_{mk} e^{im\varphi - ik\theta}. \quad (\text{B27})$$

Each contraction costs $\mathcal{O}(L^3)$, and a two-dimensional FFT evaluates the last sum on a grid with $\mathcal{O}(L^2)$ points in $\mathcal{O}(L^2 \log L)$. The Δ blocks may be generated and discarded degree by degree, leaving $\mathcal{O}(L^2)$ working storage; retaining every block instead costs $\mathcal{O}(L^3)$. These are values of a finite bivariate Fourier sum on the coordinate section, not scalar spherical harmonic coefficients of band-limit L . Indeed $D_{01}^1(\varphi, \theta, 0) \propto \sin \theta$ has an infinite scalar harmonic expansion, since an axisymmetric finite harmonic sum is a polynomial in $\cos \theta$.

- ¹M. Zaldarriaga and U. Seljak, “An all-sky analysis of polarization in the microwave background,” *Physical Review D* **55**, 1830–1840 (1997), arXiv:astro-ph/9609170.
- ²M. Kamionkowski, A. Kosowsky, and A. Stebbins, “Statistics of cosmic microwave background polarization,” *Physical Review D* **55**, 7368–7388 (1997), arXiv:astro-ph/9611125.
- ³C. G. R. Wallis, M. A. Price, J. D. McEwen, T. D. Kitching, B. Leistedt, and A. Plouviez, “Mapping dark matter on the celestial sphere with weak gravitational lensing,” *Monthly Notices of the Royal Astronomical Society* **509**, 4480–4497 (2022), arXiv:1703.09233.
- ⁴B. Bonev, T. Kurth, C. Hundt, J. Pathak, M. Baust, K. Kashinath, and A. Anandkumar, “Spherical Fourier neural operators: Learning stable dynamics on the sphere,” in *International Conference on Machine Learning (ICML)* (2023) arXiv:2306.03838.
- ⁵T. S. Cohen, M. Geiger, J. Köhler, and M. Welling, “Spherical CNNs,” in *International Conference on Learning Representations (ICLR)* (2018) arXiv:1801.10130.
- ⁶C. Esteves, C. Allen-Blanchette, A. Makadia, and K. Daniilidis, “Learning SO(3) equivariant representations with spherical CNNs,” in *European Conference on Computer Vision (ECCV)* (2018) pp. 52–68, arXiv:1711.06721.
- ⁷T. Goodwin-Allcock, J. D. McEwen, R. Gray, P. Nachev, and H. Zhang, “How can spherical CNNs benefit ML-based diffusion MRI parameter estimation?” in *Computational Diffusion MRI (CDMRI)* (2022) arXiv:2207.00572.
- ⁸J. D. McEwen, M. P. Hobson, and A. N. Lasenby, “Optimal filters on the sphere,” *IEEE Transactions on Signal Processing* **56**, 3813–3823 (2008), arXiv:astro-ph/0612688.
- ⁹J. D. McEwen, M. P. Hobson, D. J. Mortlock, and A. N. Lasenby, “Fast directional continuous spherical wavelet transform algorithms,” *IEEE Transactions on Signal Processing* **55**, 520–529 (2007), arXiv:astro-ph/0506308.
- ¹⁰Y. Wiaux, J. D. McEwen, P. Vanderghenst, and O. Blanc, “Exact reconstruction with directional wavelets on the sphere,” *Monthly Notices of the Royal Astronomical Society* **388**, 770–788 (2008), arXiv:0712.3519.
- ¹¹B. Leistedt, J. D. McEwen, P. Vanderghenst, and Y. Wiaux, “S2LET: A code to perform fast wavelet analysis on the sphere,” *Astronomy & Astrophysics* **558**, A128 (2013), arXiv:1211.1680.
- ¹²J. D. McEwen, B. Leistedt, M. Büttner, H. V. Peiris, and Y. Wiaux, “Directional spin wavelets on the sphere,” (2015), arXiv:1509.06749.
- ¹³R. Kondor, Z. Lin, and S. Trivedi, “Clebsch–Gordan nets: A fully Fourier space spherical convolutional neural network,” in *Advances in Neural Information Processing Systems (NeurIPS)* (2018) arXiv:1806.09231.
- ¹⁴O. J. Cobb, C. G. R. Wallis, A. N. Mavor-Parker, A. Marignier, M. A. Price, M. d’Avezac, and J. D. McEwen, “Efficient generalized spherical CNNs,” in *International Conference on Learning Representations (ICLR)* (2021) arXiv:2010.11661.
- ¹⁵C. Esteves, A. Makadia, and K. Daniilidis, “Spin-weighted spherical CNNs,” in *Advances in Neural Information Processing Systems (NeurIPS)* (2020) arXiv:2006.10731.

- ¹⁶J. Ocampo, M. A. Price, and J. D. McEwen, “Scalable and equivariant spherical CNNs by discrete-continuous (DISCO) convolutions,” in *International Conference on Learning Representations (ICLR)* (2023) arXiv:2209.13603.
- ¹⁷D. J. Bacon, D. M. Goldberg, B. T. P. Rowe, and A. N. Taylor, “Weak gravitational flexion,” *Monthly Notices of the Royal Astronomical Society* **365**, 414–428 (2006), vERIFY: arXiv:astro-ph/0504478; volume, pages and DOI not checked locally.
- ¹⁸J. R. Driscoll and D. M. Healy, Jr., “Computing Fourier transforms and convolutions on the 2-sphere,” *Advances in Applied Mathematics* **15**, 202–250 (1994).
- ¹⁹B. D. Wandelt and K. M. Górski, “Fast convolution on the sphere,” *Physical Review D* **63**, 123002 (2001), arXiv:astro-ph/0008227.
- ²⁰R. A. Kennedy, T. A. Lamahewa, and L. Wei, “On azimuthally symmetric 2-sphere convolution,” *Digital Signal Processing* **21**, 660–666 (2011).
- ²¹A. Challinor, P. Fosalba, D. Mortlock, M. Ashdown, B. Wandelt, and K. Górski, “All-sky convolution for polarimetry experiments,” *Physical Review D* **62**, 123002 (2000), arXiv:astro-ph/0008228. VERIFY: volume, issue, article number and doi (only title and author list verified locally).
- ²²Y. Wiaux, L. Jacques, P. Vielva, and P. Vanderghenst, “Fast directional correlation on the sphere with steerable filters,” *The Astrophysical Journal* **652**, 820–832 (2006), arXiv:astro-ph/0508516. VERIFY: doi.
- ²³Y. Wiaux, J. D. McEwen, and P. Vielva, “Complex data processing: fast wavelet analysis on the sphere,” *Journal of Fourier Analysis and Applications* **13**, 477–493 (2007), arXiv:0704.3144.
- ²⁴J. D. McEwen, P. Vanderghenst, and Y. Wiaux, “On the computation of directional scale-discretized wavelet transforms on the sphere,” in *Wavelets and Sparsity XV*, Proceedings of SPIE, Vol. 8858 (2013) p. 88580I, arXiv:1308.5706. VERIFY: article number 88580I.
- ²⁵P. J. Kostelec and D. N. Rockmore, “FFTs on the rotation group,” *Journal of Fourier Analysis and Applications* **14**, 145–179 (2008), VERIFY: doi.
- ²⁶D. Geller and D. Marinucci, “Spin wavelets on the sphere,” *Journal of Fourier Analysis and Applications* **16**, 840–884 (2010), arXiv:0811.2935. VERIFY: doi.
- ²⁷D. Geller and D. Marinucci, “Mixed needlets,” *Journal of Mathematical Analysis and Applications* **375**, 610–630 (2011), arXiv:1006.3835. VERIFY: doi.
- ²⁸Euclid Collaboration, N. Tessore, B. Joachimi, A. Loureiro, A. Hall, G. Cañas-Herrera, I. Tutusaus, N. Jeffrey, K. Naidoo, *et al.*, “Euclid preparation. LIX. angular power spectra from discrete observations,” *Astronomy & Astrophysics* **694**, A141 (2025), arXiv:2408.16903. VERIFY: volume, article number and doi; author list truncated to the first eight of the collaboration.
- ²⁹K. Mulder, M. A. Price, E. Leftaki, and J. D. McEwen, “Spin spherical Fourier neural operators: Efficient equivariant learning for spin fields on the sphere,” (2026), in preparation; submitted to the International Conference on Machine Learning (ICML). VERIFY: author list (taken from sfno/main.tex) and status.
- ³⁰P. Sadeghi, R. A. Kennedy, and Z. Khalid, “Commutative anisotropic convolution on the 2-sphere,” *IEEE Transactions on Signal Processing* **60**, 6697–6703 (2012).
- ³¹Z. Khalid, R. A. Kennedy, and P. Sadeghi, “Efficient computation of commutative anisotropic convolution on the 2-sphere,” in *International Conference on Signal Processing and Communication Systems (ICSPCS)* (2012) pp. 1–7, VERIFY: doi (title, venue, pages and year taken from Roddy & McEwen’s bibliography).
- ³²P. J. Roddy and J. D. McEwen, “Sifting convolution on the sphere,” *IEEE Signal Processing Letters* **28**, 304–308 (2021), arXiv:2007.12153. VERIFY: doi (the local bib file carries an MNRAS doi for this entry by mistake).
- ³³G. W. Mackey, “Induced representations of locally compact groups I,” *Annals of Mathematics* **55**, 101–139 (1952).
- ³⁴G. B. Folland, *A Course in Abstract Harmonic Analysis*, 2nd ed., Textbooks in Mathematics (CRC Press, Boca Raton, 2016).
- ³⁵R. Kondor and S. Trivedi, “On the generalization of equivariance and convolution in neural networks to the action of compact groups,” in *International Conference on Machine Learning (ICML)* (2018) pp. 2747–2755, arXiv:1802.03690.
- ³⁶T. S. Cohen, M. Geiger, and M. Weiler, “A general theory of equivariant CNNs on homogeneous spaces,” in *Advances in Neural Information Processing Systems (NeurIPS)* (2019) pp. 9142–9153, arXiv:1811.02017.
- ³⁷Y. Xu, J. Lei, E. Dobriban, and K. Daniilidis, “Unified Fourier-based kernel and nonlinearity design for equivariant networks on homogeneous spaces,” in *International Conference on Machine Learning (ICML)* (2022) arXiv:2206.08362.
- ³⁸T. S. Cohen, M. Weiler, B. Kicanaoglu, and M. Welling, “Gauge equivariant convolutional networks and the icosahedral CNN,” in *International Conference on Machine Learning (ICML)* (2019) pp. 1321–1330, arXiv:1902.04615.
- ³⁹M. Weiler, P. Forré, E. Verlinde, and M. Welling, “Coordinate independent convolutional networks – isometry and gauge equivariant convolutions on Riemannian manifolds,” (2021), arXiv:2106.06020.
- ⁴⁰E. T. Newman and R. Penrose, “Note on the Bondi–Metzner–Sachs group,” *Journal of Mathematical Physics* **7**, 863–870 (1966).
- ⁴¹J. N. Goldberg, A. J. Macfarlane, E. T. Newman, F. Rohrlich, and E. C. G. Sudarshan, “Spin- s spherical harmonics and $\bar{\sigma}$,” *Journal of Mathematical Physics* **8**, 2155–2161 (1967).
- ⁴²J. D. McEwen and Y. Wiaux, “A novel sampling theorem on the sphere,” *IEEE Transactions on Signal Processing* **59**, 5876–5887 (2011).

- ⁴³M. A. Price and J. D. McEwen, “Differentiable and accelerated spherical harmonic and Wigner transforms,” *Journal of Computational Physics* **510**, 113109 (2024), arXiv:2311.14670.
- ⁴⁴D. A. Varshalovich, A. N. Moskalev, and V. K. Khersonskii, *Quantum Theory of Angular Momentum* (World Scientific, Singapore, 1988).
- ⁴⁵J. D. McEwen, M. Büttner, B. Leistedt, H. V. Peiris, and Y. Wiaux, “A novel sampling theorem on the rotation group,” *IEEE Signal Processing Letters* **22**, 2425–2429 (2015), arXiv:1508.03101. VERIFY: volume, number, pages and doi against ref.papers/our_papers/1508.03101.
- ⁴⁶The sign convention adopted for the argument of the complex exponential differs from the original definition [40] but is identical to the convention used in the context of the polarisation of the cosmic microwave background [1, 2].
- ⁴⁷M. Boyle, “How should spin-weighted spherical functions be defined?” *Journal of Mathematical Physics* **57**, 092504 (2016), arXiv:1604.08140.
- ⁴⁸F. J. Narcowich, P. Petrushev, and J. D. Ward, “Localized tight frames on spheres,” *SIAM Journal on Mathematical Analysis* **38**, 574–594 (2006).
- ⁴⁹NOAA National Centers for Environmental Information, “ETOPO 2022 15 arc-second global relief model,” NOAA National Centers for Environmental Information (2022).
- ⁵⁰N. K. Pavlis, S. A. Holmes, S. C. Kenyon, and J. K. Factor, “The development and evaluation of the Earth Gravitational Model 2008 (EGM2008),” *Journal of Geophysical Research: Solid Earth* **117**, B04406 (2012).
- ⁵¹W. A. Heiskanen and H. Moritz, *Physical Geodesy* (W. H. Freeman, San Francisco, 1967).
- ⁵²P. Alken, E. Thébault, C. D. Beggan, *et al.*, “International Geomagnetic Reference Field: the thirteenth generation,” *Earth, Planets and Space* **73**, 49 (2021), the coefficients used are those of the fourteenth generation, IGRF-14, epoch 2025.0, from IAGA. VERIFY: author list truncated.
- ⁵³E. Kalnay, M. Kanamitsu, R. Kistler, W. Collins, D. Deaven, L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, Y. Zhu, M. Chelliah, W. Ebisuzaki, W. Higgins, J. Janowiak, K. C. Mo, C. Ropelewski, J. Wang, A. Leetmaa, R. Reynolds, R. Jenne, and D. Joseph, “The NCEP/NCAR 40-year reanalysis project,” *Bulletin of the American Meteorological Society* **77**, 437–471 (1996).
- ⁵⁴D. W. Behringer and Y. Xue, “Evaluation of the global ocean data assimilation system at NCEP: the Pacific Ocean,” in *Eighth Symposium on Integrated Observing and Assimilation Systems for Atmosphere, Oceans, and Land Surface, AMS 84th Annual Meeting* (Seattle, Washington, 2004) VERIFY: venue and year taken from the data description of the NOAA Physical Sciences Laboratory.
- ⁵⁵Planck Collaboration, “Planck 2018 results. IV. diffuse component separation,” *Astronomy & Astrophysics* **641**, A4 (2020), arXiv:1807.06208.
- ⁵⁶N. Kaiser and G. Squires, “Mapping the dark matter with weak gravitational lensing,” *The Astrophysical Journal* **404**, 441–450 (1993), vERIFY: volume, pages and DOI not checked locally.
- ⁵⁷S. Hilbert, J. Hartlap, and P. Schneider, “Cosmic shear covariance: the log-normal approximation,” *Astronomy & Astrophysics* **536**, A85 (2011), arXiv:1105.3980.
- ⁵⁸A. Lewis, A. Challinor, and A. Lasenby, “Efficient computation of CMB anisotropies in closed FRW models,” *The Astrophysical Journal* **538**, 473–476 (2000), arXiv:astro-ph/9911177.
- ⁵⁹A. J. Mead, S. Brieden, T. Tröster, and C. Heymans, “HMcode-2020: improved modelling of non-linear cosmological power spectra with baryonic feedback,” *Monthly Notices of the Royal Astronomical Society* **502**, 1401–1422 (2021), arXiv:2009.01858.