INTERNATIONAL JOURNAL OF RESEARCH AND INNOVATION IN APPLIED SCIENCE (IJRIAS)
ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

# Eigenvalue Degeneracy and the Limits of Scalar Spectral Parametrizations on Compact Homogeneous Spaces

**Ajay Minj***

**Department of Mathematics, St. Xavier's College, Ranchi, Jharkhand 834001, India**

***Corresponding Author**

**DOI:** https://doi.org/10.51584/IJRIAS.2026.11080018

**Received: 16 August 2026; Accepted: 21 August 2026; Published: 31 August 2026**

## ABSTRACT

Spectral parametrizations in operator learning fall into two classes. One learns an arbitrary multiplier on the frequency lattice, as the Fourier neural operator does on the torus; the other learns a scalar function of the Laplace–Beltrami eigenvalue, as recent manifold constructions do. For spectral graph networks the second class is limited by the number of distinct eigenvalues, a limitation reported as rare on irregular graphs. On a compact homogeneous space, the situation reverses. Freudenthal's formula confines the Casimir eigenvalues to a fixed lattice, so the number of distinct eigenvalues below $\Lambda^2$ is at most $N_G\Lambda^2 + 1$ whatever the dimension or rank, while the retained modes grow like $\Lambda^n$. On the $d$-torus the deficit is exact: order $(\log\Lambda)^{1/2}$ when $d = 2$ by Landau's theorem, a power of $\Lambda$ when $d \geq 3$ by the three- and four-square theorems. The scalar class therefore omits the Fourier neural operator on every torus of dimension at least two. The bound persists for bundle-valued fields, where multiplicities enter quadratically; a numerical study separates capacity from approximation error, and the joint functional calculus of the invariant differential operators is identified as the constructive remedy.

**Keywords:** spectral multiplier, Fourier multiplier, compact homogeneous space, Gelfand pair, eigenvalue multiplicity, neural operator, invariant differential operator.

## INTRODUCTION

Surrogate models for parametric partial differential equations are increasingly built by learning the solution operator between function spaces rather than individual solutions [1]. Among the spectral constructions in use, two parametrizations should be distinguished, and the distinction is easily lost because both are described as learning a spectral multiplier. On the torus $\mathbb{T}^d$, the Fourier neural operator [2] learns a band-limited multiplier

$$(\mathcal{K}v)(x) = \sum_{|k|\leq\Lambda} R_\varphi(k)\,\hat{v}(k)\,e^{2\pi i k\cdot x} \tag{1}$$

in which $R_\varphi$ is an arbitrary function on the retained lattice points. On a Riemannian manifold, where no frequency lattice is available, the natural substitute is the functional calculus of the Laplace–Beltrami operator: one fixes an eigen basis $-\Delta\psi_j = \lambda_j\psi_j$ and learns a scalar function $m$, giving

$$\mathcal{T}v = \sum_{\lambda_j\leq\Lambda^2} m(\lambda_j)\,\langle v, \psi_j\rangle\,\psi_j \tag{2}$$

This is the form taken by intrinsic and gauge-equivariant operator architectures, for which approximation, stability under metric perturbation, and consistency under mesh refinement have been established [3], and by the Laplace-spectral constructions of [4].

The relation between (1) and (2) has been settled in a neighbouring field. Wang and Zhang [5] prove that a linear spectral graph network is universal precisely when the graph Laplacian has no repeated eigenvalues and no frequency component is missing from the node features, the mechanism being those two components sharing an eigenvalue $\lambda$ are scaled by the same number $m(\lambda)$ and so retain their ratio. Lu et al. [6] make the count explicit: with $k$ distinct eigenvalues a polynomial filter produces at most $k$ distinct filter coefficients, so the dimension

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

of the scalar class equals the number of distinct eigenvalues rather than the number of modes. Wang and Zhang also relate degeneracy to symmetry, showing that a graph whose normalised Laplacian has simple spectrum admits only automorphisms of order less than three. Hordan et al. [7] refine the picture further, stratifying spectral graph networks by the largest eigenvalue multiplicity and showing that incompleteness survives even on graphs with simple spectrum.

Those papers draw an empirical conclusion from the degeneracy criterion. Real graphs are irregular; across ten standard benchmarks fewer than one per cent of eigenvalues are repeated and no frequency component is missing, so the universality conditions are, as [5] puts it, largely satisfied in practice. The purpose of the present note is to record that on the domains of interest in operator learning the conclusion reverses, and to say by how much.

A compact homogeneous space $M = G/K$ carries a transitive action of a positive-dimensional group. The condition that Wang and Zhang show to be benign for irregular graphs is therefore violated maximally and unavoidably, on every such domain. What appears not to have been observed is that the resulting degeneracy admits an a priori bound. The eigenvalues of the Laplace–Beltrami operator on a compact symmetric space are values of a quadratic form on the weight lattice, hence lie in a fixed discrete subset of the line; consequently, the number of distinct eigenvalues below $\Lambda^2$ is $O(\Lambda^2)$ no matter what the dimension or the rank of $M$ may be, while the number of retained modes grows like $\Lambda^n$ with $n = \dim M$. Theorem 3.1 states the bound and Theorem 4.1 evaluates the resulting deficit on the torus, where three classical theorems on sums of squares separate three regimes. Corollary 5.1 draws the consequence for (1) and (2).

Sections 6 to 8 take the argument past the counting statement. Section 6 reports a numerical study on the two-torus which separates two quantities that the earlier sections do not distinguish: the dimension deficit, which grows without bound, and the approximation error for a fixed target, which saturates at a level fixed by the anisotropy of the target and not by the truncation. Section 7 extends the bound to fields valued in a homogeneous vector bundle, where the multiplicities enter the count quadratically and the deficit is strictly larger than in the scalar case. Section 8 examines what can be done about it, and identifies the joint functional calculus of the algebra of invariant differential operators as the remedy that restores the full class exactly; the same analysis shows that replacing $\Delta$ by any polynomial in $\Delta$ achieves nothing, and that symmetry-breaking perturbations buy distinct eigenvalues at the cost of an eigen basis that is no longer stable.

We claim no novelty for the criterion itself, which is the continuum transcription of [5] and [6]. What is new here is the observation that on a homogeneous space the degeneracy is structural rather than incidental — it cannot be corrected away, as it can on a graph [6] — together with the exact rates that follow, the bundle-valued extension, and the identification of the constructive remedy.

## SETTING AND THE TWO FAMILIES

Throughout, $G$ is a compact connected Lie group with normalized Haar measure, $K \leq G$ a closed subgroup such that $(G, K)$ is a Gelfand pair, and $M = G/K$ the associated compact homogeneous space, carrying the $G$-invariant metric induced by a bi-invariant metric on $G$ and its Laplace–Beltrami operator $-\Delta_M$. We write $\hat{G}$ for the unitary dual, $d_\xi = \dim\mathcal{H}_\xi$, and $\hat{G}_K \subset \hat{G}$ for the spherical dual, the set of classes admitting a nonzero $K$-fixed vector. The Gelfand hypothesis is exactly that each such class contributes a one-dimensional space of $K$-fixed vectors [8], so that

$$L^2(M) = \oplus_{[\xi]\in\hat{G}_K}\mathcal{E}_\xi, \qquad -\Delta_M|_{\mathcal{E}_\xi} = \lambda_\xi\,\mathrm{id}, \qquad \dim\mathcal{E}_\xi = d_\xi \tag{3}$$

each spherical class occurring exactly once. We write $P_\xi$ for the orthogonal projection onto $\mathcal{E}_\xi$ and, following [9, Def. 10.3.18], set $\langle\xi\rangle = \left(1+\lambda_\xi\right)^{1/2}$. The truncation at level $\Lambda$ is $\Pi_\Lambda = \sum_{\langle\xi\rangle\leq\Lambda} P_\xi$, which by (3) is the spectral projector of $-\Delta_M$ onto $[0, \Lambda^2 - 1]$.

Fix $\Lambda > 0$ and put $\hat{G}_K(\Lambda) = \{[\xi] \in \hat{G}_K : \langle\xi\rangle \leq \Lambda\}$, a finite set. Define

$$P(\Lambda) = \#\hat{G}_K(\Lambda), \qquad Q(\Lambda) = \#\{\lambda_\xi : [\xi] \in \hat{G}_K(\Lambda)\}, \qquad N(\Lambda) = \sum_{\langle\xi\rangle \le \Lambda} d_\xi \tag{4}$$

so that $P(\Lambda)$ counts spherical classes, $Q(\Lambda)$ counts the distinct eigenvalues they carry, and $N(\Lambda) = \dim \Pi_\Lambda L^2(M)$. Evidently $Q(\Lambda) \le P(\Lambda)$. Let

$$\mathcal{M}_\Lambda = \left\{\sum \sigma(\xi) P_\xi : \sigma : \hat{G}_K(\Lambda) \to \mathbb{C}\right\}, \qquad \mathcal{S}_\Lambda = \{m(-\Delta_M)\Pi_\Lambda : m : [0,\infty) \to \mathbb{C}\} \tag{5}$$

denote the truncated multipliers of the two kinds. The following is the continuum form of the criterion established for graphs in [5,6]; we record it because everything below is measured against it.

**Proposition 2.1.** *For every $\Lambda > 0$ one has $\mathcal{S}_\Lambda \subseteq \mathcal{M}_\Lambda$, with $\dim \mathcal{S}_\Lambda = Q(\Lambda)$ and $\dim \mathcal{M}_\Lambda = P(\Lambda)$, so that the codimension is $P(\Lambda) - Q(\Lambda)$. The inclusion is an equality if and only if the map $[\xi] \mapsto \lambda_\xi$ is injective on $\hat{G}_K(\Lambda)$.*

*Proof.* By (3) the operator $m(-\Delta_M)\Pi_\Lambda$ acts on $\mathcal{E}_\xi$ as multiplication by $m(\lambda_\xi)$ when $\langle\xi\rangle \le \Lambda$ and as 0 otherwise, so it is the element of $\mathcal{M}_\Lambda$ with symbol $\sigma = m \circ \lambda$. Conversely an element of $\mathcal{M}_\Lambda$ lies in $\mathcal{S}_\Lambda$ precisely when its symbol is constant on the fibres of $[\xi] \mapsto \lambda_\xi$, since a function on $\hat{G}_K(\Lambda)$ constant on those fibres extends to a function on $[0,\infty)$, arbitrarily off the finite set of values attained. Hence $\mathcal{S}_\Lambda$ is the subspace of fibre-constant symbols, of dimension equal to the number of fibers. Equality of the two spaces is equality of the two dimensions.

The formulation has a reading worth stating, because it names what the scalar parametrization silently assumes. By (5), $\mathcal{M}_\Lambda$ is the truncation of the commutant of the $G$-action, while $\mathcal{S}_\Lambda$ is the truncation of the algebra generated by $\Delta_M$. Proposition 2.1 says these coincide exactly when $\Delta_M$ separates the spherical dual — that is, adopting (2) amounts to assuming the Laplacian generates the whole commutant. Sections 3 and 4 measure how far that fails, and Section 8 returns to the generation question, where the remedy lies.

# AN A PRIORI BOUND ON THE SCALAR FAMILY

On a graph, eigenvalue degeneracy is a property of the particular Laplacian, and can be mitigated by perturbing the spectrum [6]. On a compact symmetric space, it is not: the eigenvalues are constrained to a lattice, and no such correction is available. This is the content of the following bound, which involves neither the dimension nor the rank of $M$.

**Theorem 3.1.** *Let $M = G/K$ be a compact symmetric space, and normalise the invariant metric so that the inner product on the dual of the Cartan subalgebra takes integer values on the weight lattice. Then there is $N_G \in \mathbb{N}$, depending on $G$ alone, with $\lambda_\xi \in N_G^{-1}\mathbb{Z}$ for every $[\xi] \in \hat{G}$, and consequently*

$$Q(\Lambda) \le N_G \Lambda^2 + 1 \qquad \text{for all } \Lambda > 0 \tag{6}$$

*Proof.* Let $\lambda$ denote the highest weight of $\xi$ and $\rho$ the half-sum of the positive roots. Freudenthal's formula [8, Ch. V, §1] evaluates the Casimir eigenvalue as $\lambda_\xi = \langle\lambda+\rho, \lambda+\rho\rangle - \langle\rho,\rho\rangle = |\lambda|^2 + 2\langle\lambda,\rho\rangle$. Under the stated normalisation $|\lambda|^2$ is an integer, and $2\langle\lambda,\rho\rangle$ is an integer because $2\rho$ is a sum of roots; allowing for the finitely many denominators introduced by a different choice of normalisation gives $\lambda_\xi \in N_G^{-1}\mathbb{Z}$. The values not exceeding $\Lambda^2$ therefore number at most $N_G\Lambda^2 + 1$, and $Q(\Lambda)$ counts a subset of them.

A scalar spectral multiplier can thus never carry more than $O(\Lambda^2)$ parameters, whatever the dimension or the rank of $M$. Since the number of retained modes satisfies $N(\Lambda) \asymp \Lambda^n$ with $n = \dim M$, by Weyl's asymptotic law [10], the scalar family is short by a factor of at least order $\Lambda^{n-2}$ relative to the space it acts on, on every compact symmetric space of dimension exceeding two. Combining (6) with Proposition 2.1 gives the codimension bound

$$\dim \mathcal{M}_\Lambda - \dim \mathcal{S}_\Lambda = P(\Lambda) - Q(\Lambda) \ \ge\ P(\Lambda) - N_G\Lambda^2 - 1.$$

The bound (6) is not always binding. On a compact rank-one symmetric space the spherical dual is indexed by a single non-negative integer $j$, and Freudenthal gives $\lambda_\xi(j) = j^2|\mu|^2 + 2j\langle\mu,\rho\rangle$ for a fixed dominant $\mu$, a quadratic with positive coefficients, hence strictly increasing and injective. There $P(\Lambda) \asymp \Lambda$, (6) permits no conclusion, and by Proposition 2.1 the scalar family loses nothing: on the sphere $S^n$ the familiar $\lambda_\ell = \ell(\ell+n-1)$ is injective and (2) is a faithful parametrization, which is consistent with the practice of spherical spectral architectures [11]. The interesting case is the opposite one. Section 8 shows that rank one is exactly the case in which a single operator generates the invariant algebra, which explains the coincidence.

**The torus: exact rates**

Take $G = \mathbb{T}^d$ with $K$ trivial, an abelian and therefore Gelfand pair, and normalise the flat metric so that $\widehat{G}_K = \mathbb{Z}^d$ with $\lambda_k = |k|^2$ and every $d_\xi = 1$. Then $P(\Lambda)$ is a lattice-point count in a ball and $Q(\Lambda)$ counts the integers below $\Lambda^2$ representable as a sum of $d$ squares. Writing $\omega_d = \pi^{d/2}/\Gamma(d/2+1)$ for the volume of the unit ball, the evaluation of $Q$ changes character three times.

**Theorem 4.1.** *With the notation above, as $\Lambda \to \infty$,*

$$\frac{P(\Lambda)}{Q(\Lambda)} \sim \begin{cases} \pi\sqrt{2}\,\mathcal{K}^{-1}(\log\Lambda)^{1/2}, & d=2, \\ (8\pi/5)\,\Lambda, & d=3, \\ \omega_d\,\Lambda^{d-2}, & d \geq 4, \end{cases} \tag{7}$$

*where $\mathcal{K} = 0.76422\ldots$ is the Landau–Ramanujan constant. In particular the ratio is unbounded for every $d \geq 2$, and $\mathcal{S}_\Lambda$ is a proper subspace of $\mathcal{M}_\Lambda$ for every $d \geq 1$ and every $\Lambda > 1$.*

*Proof.* Properness for $d \geq 1$ is immediate from Proposition 2.1, since $k$ and $-k$ are distinct elements of $\mathbb{Z}^d$ carrying the same eigenvalue. For the rates, $P(\Lambda) \sim \omega_d \Lambda^d$ is the lattice-point count in a ball of radius $\Lambda$. For $d \geq 4$, Lagrange's four-square theorem makes every non-negative integer representable, so $Q(\Lambda) = \lfloor \Lambda^2 \rfloor + 1$ exactly and the quotient is $\omega_d \Lambda^{d-2}$. For $d = 3$, Legendre's three-square theorem identifies the complement of the representable integers as those of the form $4^a(8b+7)$, of natural density $1/6$; hence $Q(\Lambda) \sim (5/6)\Lambda^2$ and the quotient is $(4\pi/3)(6/5)\Lambda = (8\pi/5)\Lambda$. For $d = 2$, Landau's theorem gives an asymptotic count $\mathcal{K}x(\log x)^{-1/2}$ for the integers up to $x$ that are sums of two squares; putting $x = \Lambda^2$ yields $Q(\Lambda) \sim \mathcal{K}\Lambda^2(2\log\Lambda)^{-1/2}$ and the quotient is $\pi\sqrt{2}\,\mathcal{K}^{-1}(\log\Lambda)^{1/2}$.

The three densities have been confirmed by direct enumeration of the representable integers up to $2.5 \times 10^5$: the observed density is 0.2298 for $d = 2$, consistent with a vanishing limit; 0.8333 for $d = 3$, against the predicted $5/6$; and 1.0000 for $d \geq 4$. The leading constants in (8) are likewise borne out. At $\Lambda = 60$ direct enumeration gives ratios 301.1 for $d = 3$ against the predicted 301.6, and 17762.0 for $d = 4$ against 17765.0.

For $d = 2$ the agreement is looser, 10.61 counted against 11.76 predicted at $\Lambda = 60$, and 14.52 against 15.28 at $\Lambda = 1000$. The discrepancy is the well-known slow convergence in the Landau–Ramanujan asymptotic and not a failure of the estimate. Since $d = 2$ carries much of the reported work on the Fourier neural operator, and since the deficit there is only logarithmic, the exact counts are worth recording; they are collected in Table 1 and show the shortfall already exceeds an order of magnitude at every usable truncation level.

Exact counts on the two-torus, obtained by direct enumeration of $\mathbb{Z}^2 \cap \{|k| \leq \Lambda\}$. The deficit is unbounded but grows only like the square root of $\log\Lambda$.

| $\Lambda$ | 30 | 60 | 120 | 250 | 500 | 1000 |
|---|---|---|---|---|---|---|
| $P(\Lambda) = \#\{k \in \mathbb{Z}^2 : \lvert k\rvert \leq \Lambda\}$ | 2821 | 11289 | 45225 | 196321 | 785349 | 3141549 |
| $Q(\Lambda) = \#\{\lvert k\rvert^2 : \lvert k\rvert \leq \Lambda\}$ | 300 | 1064 | 3860 | 15389 | 57445 | 216342 |
| $P(\Lambda)/Q(\Lambda)$ | 9.40 | 10.61 | 11.72 | 12.76 | 13.67 | 14.52 |

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

### Consequence for the Fourier neural operator

The band-limited layer (1) is the general element of $\mathcal{M}_\Lambda$ for $G = \mathbb{T}^d$, its multiplier $R_\varphi$ ranging over all functions on the retained lattice. Combining this with Proposition 2.1 and Theorem 4.1 gives the statement that motivated the note.

**Corollary 5.1.** *A Fourier neural operator layer on* $\mathbb{T}^d$ *is realisable as a scalar spectral multiplier of the form* (2) *if and only if its multiplier is radial, that is* $R_\varphi(k) = m(|k|^2)$*. The layers with this property form a subspace of codimension* $P(\Lambda) - Q(\Lambda)$*, and the proportion of the parametrisation they occupy tends to* $0$ *as* $\Lambda \to \infty$ *for every* $d \geq 2$.

So, the scalar construction does not contain the Fourier neural operator, on any torus, at any truncation level beyond the trivial one. Since $d = 2$ and $d = 3$ carry most of the reported work on that architecture, the restriction is not confined to a regime of merely theoretical interest. What Theorem 4.1 adds to the bare statement of properness is the rate: on the two-torus the deficit is a factor of order $(\log \Lambda)^{1/2}$, unbounded but slow, while from $d = 3$ onward it is a power of $\Lambda$.

### A numerical illustration on the two-torus

Corollary 5.1 is a statement about dimensions. It says nothing about how badly a given target is approximated, and the two quantities behave differently. This section reports a small computation which makes the difference explicit and gives the expressivity gap a numerical value.

Fix $\Lambda$ and a target multiplier $R^*: \mathbb{Z}^2 \to \mathbb{R}$. In the Hilbert–Schmidt norm on $\mathcal{M}_\Lambda$ the two families are subspaces of a Euclidean space and the closest element of $\mathcal{S}_\Lambda$ to $R^*$ is the orthogonal projection, obtained by averaging $R^*$ over each fibre of $k \mapsto |k|^2$. Writing $\bar{R}^*(\lambda)$ for that average, the quantity

$$E(\Lambda; R^*) = \left(\frac{\sum_{|k|\leq\Lambda} |R^*(k) - \bar{R}^*(|k|^2)|^2}{\sum_{|k|\leq\Lambda} |R^*(k)|^2}\right)^{1/2} \tag{8}$$

is the exact infimum of the relative error over the whole scalar class. It is therefore a lower bound on what any training procedure can reach for a single layer of the form (2), independent of the optimizer, the initialization, and the number of epochs.

Two targets are considered. The first is the time-$t$ solution operator of the anisotropic heat equation $\partial_t u = (\partial_1^2 + \rho\, \partial_2^2)u$ on $\mathbb{T}^2$, with multiplier $R^*(k) = \exp\big(-t(k_1^2 + \rho k_2^2)\big)$; we take $t = 0.02$ and let the anisotropy $\rho$ vary. The second is $R^*(k) = k_1^2/|k|^2$ for $k \neq 0$ and $R^*(0) = 0$, the symbol of $-\partial_1^2(-\Delta)^{-1}$, a bounded operator of order zero which does not smooth. Table 2 records $E(\Lambda; R^*)$ in both cases.

Relative error (9) of the best scalar spectral approximation on $\mathbb{T}^2$. Rows one to four: anisotropic heat semigroup at $t = 0.02$ with anisotropy ρ. Last row: the order-zero symbols $k_1^2/|k|^2$. Entries are rounded to four decimals; a dash indicates that the lattice was not enumerated at that level.

| **target** | $\Lambda = 15$ | $\Lambda = 30$ | $\Lambda = 60$ | $\Lambda = 120$ | $\Lambda = 250$ | $\Lambda = 500$ |
|---|---|---|---|---|---|---|
| heat, $\rho = 1$ | 0.0000 | 0.0000 | 0.0000 | 0.0000 | 0.0000 | — |
| heat, $\rho = 2$ | 0.1696 | 0.1697 | 0.1697 | 0.1697 | 0.1697 | — |
| heat, $\rho = 4$ | 0.3200 | 0.3200 | 0.3200 | 0.3200 | 0.3200 | — |
| heat, $\rho = 10$ | 0.4727 | 0.4727 | 0.4727 | 0.4727 | 0.4727 | — |
| $k_1^2/\lvert k\rvert^2$ | 0.5750 | 0.5765 | 0.5771 | 0.5772 | 0.5773 | 0.5773 |

Three features of the table deserve comment. The isotropic case $\rho = 1$ gives error zero to machine precision, as Corollary 5.1 requires, since the multiplier is then radial. For $\rho \neq 1$ the error is strictly positive and essentially independent of $\Lambda$: the heat multiplier decays so fast that the truncation level is irrelevant beyond the first few shells, and the error is fixed by the anisotropy alone. The dimension deficit of Theorem 4.1 continues to grow across the same range of $\Lambda$, from 9.4 to 12.8 in Table 1, while the approximation error does not move. Capacity and accuracy are distinct quantities, and the counting results of Sections 3 to 5 bound the first without controlling the second.

The order-zero symbol behaves differently in one respect and identically in another. Its error rises slightly with $\Lambda$ and settles at 0.5773, which is $1/\sqrt{3}$ to four decimals. The continuum computation explains the value: writing $k = |k|(\cos\theta, \sin\theta)$, the target is $\cos^2\theta$, its average over each circle is $1/2$ in the aggregate, and the ratio of the variance of $\cos^2\theta$ to its second moment is $(3/8 - 1/4)/(3/8) = 1/3$. That the enumerated value matches to four digits indicates that the lattice points on circles of radius at most $\Lambda$ equidistribute in aggregate, which is all the argument needs; individual circles need not, and the identification of the limit is offered as a numerical observation rather than a theorem. The error again does not decay with $\Lambda$: refining the truncation does not help a scalar multiplier reproduce a target that is not radial.

At the level of individual functions, the same gap appears in the output. At $\Lambda = 120$, with $R^*(k) = k_1^2/|k|^2$, applying the exact multiplier and its best scalar approximation to a random field with $|\hat{v}(k)| = |k|^{-3/2}$ and uniform phases gives a relative $L^2$ error of 0.6304 in the output; for the deterministic input $v(x) = \exp(\sin 2\pi x_1)\cos 2\pi x_2 + \sin(4\pi x_1 + 2\pi x_2)$ the error is 0.7744. Figure 1 collects the two effects: the parameter deficit of Theorem 4.1 on the left, the saturating approximation error on the right.

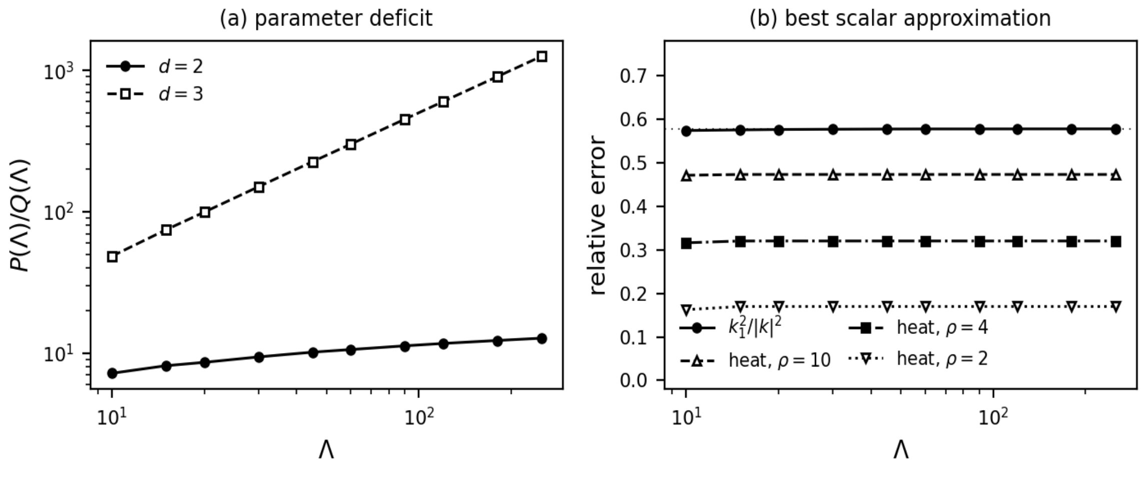


(a) The ratio $P(\Lambda)/Q(\Lambda)$ on $\mathbb{T}^2$ and $\mathbb{T}^3$, by direct enumeration; both axes logarithmic. (b) The relative error (9) of the best scalar spectral approximation for the targets of Table 2; the dotted line marks $1/\sqrt{3}$. The deficit on the left grows without bound; the error on the right does not move.

**Vector- and tensor-valued fields**

The counting so far concerns scalar fields, for which the Gelfand hypothesis makes the symbol of an invariant operator a scalar on the spherical dual. Features in operator learning are frequently vector- or tensor-valued — velocity fields, stresses, magnetic fields — and it is reasonable to ask whether the degeneracy survives. It does, and in a stronger form.

Let $\tau$ be an irreducible unitary representation of $K$ on $V_\tau$ and let $E_\tau = G \times_K V_\tau$ be the associated homogeneous vector bundle, whose sections are the $\tau$-equivariant $V_\tau$-valued functions on $G$. Frobenius reciprocity gives the isotypic decomposition

$$L^2(M, E_\tau) \cong \bigoplus_{[\xi]\in\hat{G}} \mathcal{H}_\xi \otimes \mathrm{Hom}_K\left(\mathcal{H}_\xi|_K, V_\tau\right), \qquad m(\xi,\tau) = \dim\mathrm{Hom}_K\left(\mathcal{H}_\xi\big|_K, V_\tau\right) \tag{9}$$

and the Bochner Laplacian $\Delta_\tau$ built from the canonical connection acts on the $\xi$-isotypic component as the scalar $\lambda_\xi - c_\tau$, where $c_\tau$ is the Casimir eigenvalue of $\tau$ [12]. The Hodge Laplacian on $p$-forms is obtained by taking $\tau = \Lambda^p \mathrm{Ad}^*$ and adding a curvature term which is again constant on each isotypic component, so the same conclusion applies to it.

**Proposition 7.1.** *Let $\mathcal{M}_\Lambda^\tau$ be the space of $G$-equivariant endomorphisms of the truncation of $L^2(M, E_\tau)$ at level $\Lambda$, and $\mathcal{S}_\Lambda^\tau$ the subspace of those of the form $m(\Delta_\tau)$. Then*

$$\dim\mathcal{M}_\Lambda^\tau = \sum_{\langle\xi\rangle\le\Lambda} m\,(\xi,\tau)^2, \qquad \dim\mathcal{S}_\Lambda^\tau = Q_\tau(\Lambda) \le N_G\Lambda^2 + 1 \tag{10}$$

*where $Q_\tau(\Lambda)$ counts the distinct values of $\lambda_\xi - c_\tau$ over the classes retained. In particular the bound of Theorem 3.1 holds verbatim for the bundle-valued scalar family, and the codimension of $\mathcal{S}_\Lambda^\tau$ in $\mathcal{M}_\Lambda^\tau$ is at least as large as in the scalar case whenever the same classes are retained.*

*Proof.* By (10) and Schur's lemma the $G$-equivariant endomorphisms of $\mathcal{H}_\xi \otimes \mathrm{Hom}_K\left(\mathcal{H}_\xi|_K, V_\tau\right)$ are $\mathrm{id} \otimes A$ with $A$ an arbitrary endomorphism of the multiplicity space, giving $m(\xi,\tau)^2$ parameters per class and the first equality. Since $\Delta_\tau$ acts on that component by $\lambda_\xi - c_\tau$, an operator $m(\Delta_\tau)$ is determined by the values of $m$ at the distinct numbers $\lambda_\xi - c_\tau$, which gives the second equality; as $c_\tau$ is a fixed constant, $Q_\tau(\Lambda)$ counts the same fibres as $Q(\Lambda)$ and Theorem 3.1 applies unchanged. Finally, $m(\xi,\tau)^2 \ge m(\xi,\tau) \ge 1$ on the classes that occur, so the first sum dominates the corresponding scalar count.

The torus makes the effect concrete. There $K$ is trivial, every bundle is trivial, and a $\mathbb{C}^r$-valued field carries a matrix-valued multiplier: $\dim\mathcal{M}_\Lambda^\tau = r^2 P(\Lambda)$ against $\dim\mathcal{S}_\Lambda^\tau = Q(\Lambda)$, so the ratio in Theorem 4.1 is multiplied by $r^2$. For 1-forms on $\mathbb{T}^2$ the Hodge Laplacian is the component wise Laplacian, each eigenvalue $|k|^2$ carrying multiplicity 2 in place of 1, and the deficit at $\Lambda = 1000$ rises from a factor of 14.5 to a factor of 58.1. Passing from scalars to vectors therefore aggravates the restriction rather than relieving it, and the architectures that treat vector fields by Hodge decomposition [13] should be read as enlarging $\mathcal{M}_\Lambda^\tau$ by branch structure, not as removing the degeneracy inside a branch.

## ARCHITECTURAL MITIGATIONS

The preceding sections diagnose a restriction. Four routes present themselves, of which one is exact, one is empty, and two are partial.

*(i) Joint functional calculus of the invariant algebra.* Let $\mathbb{D}(M)$ denote the algebra of $G$-invariant differential operators on $M = G/K$. For a compact symmetric space of rank $r$, the Harish-Chandra isomorphism identifies $\mathbb{D}(M)$ with the algebra of Weyl-invariant polynomials on the dual of a maximal abelian subspace, a polynomial algebra in $r$ generators $D_1, \dots, D_r$ [8]. Each spherical class is an eigen class for all of them simultaneously, and the joint eigenvalue $\left(\lambda_\xi^{(1)}, \dots, \lambda_\xi^{(r)}\right)$ determines the highest weight, hence the class. Replacing (2) by

$$\mathcal{T}v = \sum_{\langle\xi\rangle\le\Lambda} m\left(\lambda_\xi^{(1)}, \dots, \lambda_\xi^{(r)}\right) P_\xi v \tag{11}$$

with $m$ a learnable function of $r$ variables therefore yield exactly $\mathcal{M}_\Lambda$: the parametrisation is complete, equivariance is retained, and the parameter count is $P(\Lambda)$ rather than $Q(\Lambda)$. On $\mathbb{T}^d$ the generators are $-i\,\partial_1, \dots, -i\,\partial_d$, the joint eigenvalue of the mode $e^{2\pi i k\cdot x}$ is $k$ itself, and (12) is the Fourier neural operator layer (1); on a rank-one space the single generator is $\Delta_M$ and (12) reduces to (2), which is why nothing was lost there in Section 3. The obstruction of Theorem 3.1 is thus specific to $r = 1$, and its removal costs the ability to diagonalize $r$ commuting operators rather than one — that is, knowledge of the isotypic decomposition rather than of the Laplace spectrum alone. On a mesh this is the difference between an eigen solver and an eigen solver followed by a symmetry-adapted refinement of each eigenspace.

*(ii) Higher-order and polynomial Laplacians.* A natural first attempt is to replace $\Delta_M$ by $\Delta_M^p$, by a polynomial $p(\Delta_M)$, or by a fractional power. This achieves nothing whatever: all such operators have exactly the eigenspaces

of $\Delta_M$, so $m\big(p(\Delta_M)\big)\Pi_\Lambda$ is again fibre-constant and $\mathcal{S}_\Lambda$ is unchanged as a set. The point is worth stating because the reverse is true for the higher-order Laplacians in the sense of Hodge theory, which act on different bundles and do enlarge the model class, though by the mechanism of Proposition 7.1 — through the multiplicity spaces — and not by lifting the degeneracy of the scalar part. Only operators outside the algebra generated by $\Delta_M$ can help, which returns the discussion to route (i).

*(iii) Symmetry-breaking perturbations.* One may perturb the metric, or the domain, so that the $G$-action is destroyed and the spectrum becomes simple; this is the continuum analogue of the eigenvalue correction of [6]. Generic metrics do have simple spectrum, so the perturbation succeeds in principle. Two costs attach to it. First, the perturbed operator is no longer $G$-invariant, so the model class is no longer contained in the commutant, and whatever equivariance motivated (2) is lost along with the degeneracy. Second, and more seriously, the correction is numerically self-defeating: a perturbation of size $\varepsilon$ splits an eigenvalue of multiplicity $d_\xi$ into at most $d_\xi$ values separated by gaps of order $\varepsilon$, and by the Davis–Kahan bound the eigenvectors are then determined only to accuracy of order $\varepsilon/\delta$ with $\delta$ the resulting gap. Here $\delta \asymp \varepsilon$, so the eigenvector error is of order one and the learned coefficients are attached to a basis that the discretization does not reproduce stably. Distinct eigenvalues are obtained; a stable eigen basis is not. The same tension is what the rotation-equivariant construction of [7] is designed to avoid in the graph setting.

*(iv) Equivariant action on eigenspaces.* Rather than assign one scalar per eigenvalue, one may act on each eigenspace by a map equivariant under its residual symmetry, which for an eigenspace of dimension $m$ is a subgroup of $O(m)$. On $M = G/K$ with $(G, K)$ a Gelfand pair this route terminates immediately and favourably: by Schur's lemma $\mathrm{End}_G\big(\mathcal{E}_\xi\big) = \mathbb{C}$, so the equivariant maps on the eigenspace are exactly the scalars $\sigma(\xi)$, and the class they generate is $\mathcal{M}_\Lambda$ itself. In other words, $\mathcal{M}_\Lambda$ is not a larger and less structured family than $\mathcal{S}_\Lambda$; it is the full equivariant family, and $P(\Lambda) - Q(\Lambda)$ measures exactly what the scalar parametrisation discards relative to equivariance, with no equivariance gained in return. In the bundle-valued case of Section 7 the same argument gives $m(\xi, \tau)^2$ parameters per class. Route (iv) and route (i) are the same remedy described from the two sides.

The practical form of the remedy on a general mesh, where no group is available, is a clustering of the computed spectrum: eigenvalues that agree to within the discretization tolerance are grouped, and a full matrix rather than a single scalar is learned within each group. The parameter cost is $\sum_c m_c^2$ with $m_c$ the cluster sizes, which interpolates between $Q(\Lambda)$ and $P(\Lambda)$ as the tolerance is varied, and the construction reduces to (12) when the mesh resolves a homogeneous geometry.

## REMARKS ON SCOPE

Three limitations should be stated plainly.

First, the criterion of Proposition 2.1 is not new; it is the continuum transcription of results proved for graphs in [5] and [6], and the connection between spectral degeneracy and the order of the automorphism group is likewise established there, with the finer stratification by multiplicity given in [7]. The contribution here is Theorem 3.1, Theorem 4.1 and Proposition 7.1, together with the observation that the empirical conclusion drawn in those papers — that degeneracy is rare, and correctable when it occurs — reverses on a homogeneous space, where transitivity of the group action forces degeneracy and the lattice structure of the spectrum bounds the scalar family a priori.

Second, the counting results bound expressive capacity at a fixed spectral budget and do not by themselves bound approximation error for a given target. Section 6 makes the distinction quantitative on the two-torus: the deficit grows with $\Lambda$ while the relative error of the best scalar approximation saturates, at a level set by the anisotropy of the target and equal to zero exactly in the radial case. A smaller model class may still approximate a particular solution operator adequately, and for a rotationally invariant target on the sphere the two classes coincide, as Section 3 records. Error estimates for the scalar construction are available in [3], and for the Fourier neural operator on the torus in [14] and, more recently, in [15].

ISSN No. 2454-6194 | DOI: 10.51584/IJRIAS |Volume XI Issue VIII August 2026

Third, the treatment is confined to compact homogeneous spaces. On a manifold with no isometries the spectrum is generically simple, $Q(\Lambda) = P(\Lambda)$, and the deficit vanishes in principle; but the relevant quantity in computation is the number of numerically resolved eigenvalue clusters rather than the number of distinct eigenvalues, and a manifold close to a homogeneous one inherits clusters of size $d_\xi$ whose internal gaps are governed by the distance to the symmetric model. The homogeneous case treated here is the extreme point of that continuum rather than an isolated pathology, and quantifying the intermediate regime — clusters, gaps, and the conditioning they induce — is left open. Section 7 removes the earlier restriction to scalar fields for bundles over a homogeneous base; the corresponding count for bundles over a manifold with no transitive symmetry is subject to the same open question.

## REFERENCES


1. Kovachki N, Li Z, Liu B, Azizzadenesheli K, Bhattacharya K, Stuart A, Anandkumar A. Neural operator: learning maps between function spaces with applications to PDEs. J Mach Learn Res. 2023;24(89):1–97.
2. Li Z, Kovachki N, Azizzadenesheli K, Liu B, Bhattacharya K, Stuart A, Anandkumar A. Fourier neural operator for parametric partial differential equations. In: Proceedings of the 9th International Conference on Learning Representations (ICLR); 2021.
3. Cheng P. Gauge-equivariant intrinsic neural operators for geometry-consistent learning of elliptic PDE maps. arXiv:2603.14734 [cs.LG]. 2026.
4. Tang H, Zhu J, Feng Z, Li H, Li C. Geometric Laplace neural operator. arXiv:2512.16409 [cs.LG]. 2025.
5. Wang X, Zhang M. How powerful are spectral graph neural networks. In: Proceedings of the 39th International Conference on Machine Learning (ICML). PMLR. 2022; 162:23341–23362.
6. Lu K, Yu Y, Fei H, Li X, Yang Z, Guo Z, Liang M, Yin M, Chua TS. Improving expressive power of spectral graph neural networks with eigenvalue correction. In: Proceedings of the 38th AAAI Conference on Artificial Intelligence. 2024;38(13):14158–14166.
7. Hordan S, Lifshitz G, Bechler-Speicher M, Dym N. Spectral graph neural networks are incomplete on graphs with a simple spectrum. In: Advances in Neural Information Processing Systems (NeurIPS). 2025. arXiv:2506.05530.
8. Helgason S. Groups and geometric analysis: integral geometry, invariant differential operators, and spherical functions. Orlando: Academic Press; 1984.
9. Ruzhansky M, Turunen V. Pseudo-differential operators and symmetries: background analysis and advanced topics. Pseudo-Differential Operators: Theory and Applications, vol. 2. Basel: Birkhäuser; 2010.
10. Chavel I. Eigenvalues in Riemannian geometry. Pure and Applied Mathematics, vol. 115. Orlando: Academic Press; 1984.
11. Bonev B, Kurth T, Hundt C, Pathak J, Baust M, Kashinath K, Anandkumar A. Spherical Fourier neural operators: learning stable dynamics on the sphere. In: Proceedings of the 40th International Conference on Machine Learning (ICML). PMLR. 2023;202:2806–2823.
12. Wallach NR. Harmonic analysis on homogeneous spaces. Pure and Applied Mathematics, vol. 19. New York: Marcel Dekker; 1973.
13. Zheng D, Zhong T, Allen-Blanchette C. Topology-preserving neural operator learning via Hodge decomposition. arXiv:2605.13834 [cs.LG]. 2026.
14. Kovachki N, Lanthaler S, Mishra S. On universal approximation and error bounds for Fourier neural operators. J Mach Learn Res. 2021;22(290):1–76.
15. Chandramoorthy N, Sanz-Alonso D, Waniorek N. From spectral methods to sample complexity bounds for Fourier neural operators. arXiv:2607.00320 [math.NA]. 2026.