

GREEN’S NEURAL OPERATOR WITH NEUMANN CONDITIONS FOR EMG VOLUME CONDUCTOR MODELLING

Noura Ezaz-Nikpay*

Department of Computing
Imperial College London
London, UK
n.ezaz-nikpay24@imperial.ac.uk

Dimitrios Halatsis*

Department of Bioengineering
Imperial College London
London, UK
d.chalatsis22@imperial.ac.uk

Dario Farina

Department of Bioengineering
Imperial College London
London, UK
d.farina@imperial.ac.uk

ABSTRACT

Simulating electromyography (EMG) signals has a computational bottleneck: repeatedly solving the anisotropic Poisson equation on patient-specific volume conductor geometries. We learn a Green’s function representation using a latent neural operator (LNO), extending prior neural Green’s approaches from Dirichlet to Neumann boundary conditions. In 5-fold cross-validation over held-out geometry-conductivity configurations, the Green’s formulation outperforms a FiLM-conditioned standard LNO formulation with fewer parameters and $2.5\times$ lower variance. However, spatial holdout experiments reveal a vulnerability of the Green’s formulation: when entire regions are devoid of training sources, the Green’s model fails. We trace this to asymmetric gradient flow through the learned eigenfunctions and show that the failure does not occur when training sources have adequate spatial coverage. These findings delineate when the Green’s formulation offers advantages over direct prediction approaches. This workshop paper is a proof of concept; future work will investigate more complex anatomies.

1 INTRODUCTION

Simulating electromyography (EMG) signals is central to building digital twins for neuromuscular diagnostics and prosthetic control, yet each evaluation requires solving the governing partial differential equation (PDE) on a geometry-and-conductivity-specific domain. Classical solvers such as the Finite Element Method (FEM) are accurate but computationally expensive, motivating the use of learned surrogate models. Neural operators have emerged as efficient surrogates for parametric PDEs (Li et al., 2021; Lu et al., 2021), learning mapping between function spaces (Kovachki et al.). However, a limitation shared by standard neural operator formulations is that they learn a direct map from forcing to solution, entangling the geometry-dependent operator with the source-dependent response. Green’s function approaches decouple these roles; the Green’s function $G(x, x')$ encodes only the geometry and material properties of the domain. Recent work has demonstrated the efficacy of learning Green’s functions with neural networks. Boullé et al. (2022) showed that rational neural networks can discover interpretable Green’s functions from input-output data, while Melchers et al. (2026) introduced Neural Green’s Operators that embed symmetry, spectral, and conservation properties directly into the architecture. Yoo et al. (2025) demonstrated that neural Green’s functions achieve superior generalisation across diverse irregular geometries and boundary functions under Dirichlet conditions. Here we extend the Green’s function formulation to Neumann boundary conditions, applying it to the EMG volume conductor problem with sparse monopolar source forcing.

*Equal contribution.

We present proof-of-concept results demonstrating that the Green’s formulation achieves superior geometry and conductivity generalisation compared to a standard neural operator, while also revealing a limitation: when entire spatial regions are devoid of training sources, the spectral solver degrades, owing to an asymmetry in how the learned eigenfunctions are supervised during training. We characterise this failure mode and discuss mitigation strategies for future work toward application on realistic, MRI-derived geometries with FEM-generated training data.

2 METHOD

We compare two formulations built on the same Latent Neural Operator (LNO) backbone (Wang & Wang, 2024): a standard model that directly predicts the potential u , and a Green’s-adapted one, as shown in Figure 1. Data is generated analytically according to Appendix B (Farina et al., 2004a), using a multi-layered infinite cylinder as the volume conductor. The dataset contains 82 geometry-conductivity configurations, each with 100 samples of varying source locations and magnitudes.




Figure 1: The difference between standard vs Green’s formulations in terms of inputs and outputs.

2.0.1 LATENT NEURAL OPERATOR BACKBONE

The LNO follows a perceiver-style architecture (Jaegle et al., 2021). Spatial coordinates are embedded via Fourier positional encoding into \mathbb{R}^{192} , then mapped through trunk and branch networks to produce keys and values of dimension $d = 128$. A set of $M = 256$ learnable latent tokens attend to these N spatial tokens via physics-informed cross-attention (PhCA), compressing the spatial dimension. The latent tokens are then processed by 6 transformer layers with 8-head self-attention, before a symmetric PhCA decoder maps back to N spatial points. Both formulations condition the backbone via concatenation to the branch input, with 4 geometry parameters describing the domain shape and 5 conductivity parameters specifying the heterogeneous material properties.

2.1 STANDARD FORMULATION (LNO-FILM)

The standard model predicts the solution u directly from a source-conditioned network. Since the forcing is extremely sparse ($<0.01\%$ nonzero nodes), passing the raw forcing field f as input is ineffective. Instead, source parameters $(r_j, z_j, \theta_j, I_j)$ for each active source j are encoded independently by an MLP, masked for inactive sources, and summed to produce a global conditioning vector $\mathbf{c} \in \mathbb{R}^d$. This vector modulates each transformer layer via adaptive layer normalisation (Perez et al., 2017) where learned affine transformations of \mathbf{c} produce per-layer scale and shift parameters for both the attention and MLP sub-layers. The decoder output is mapped to $\hat{u} \in \mathbb{R}^{N \times 1}$ by an MLP head, followed by explicit mean subtraction to enforce the Neumann compatibility condition.

2.2 GREEN’S FORMULATION (LNO-GREEN’S)

The Green’s function model learns a geometry-dependent spectral representation of the solution operator rather than predicting u directly (Boullé et al., 2022). Crucially, the network receives no information about the forcing, neither f nor source parameters. Two MLP heads map the decoded features to learned eigenfunctions $\Phi \in \mathbb{R}^{N \times R}$ and eigenvalues $\Lambda \in \mathbb{R}^{N \times R}$, with spectral rank $R = 32$. The solution is then obtained via a differentiable spectral solver, see Appendix A.

3 EXPERIMENTS

3.1 GEOMETRY GENERALISATION

To evaluate generalisation to unseen geometries, we perform 5-fold cross-validation at the configuration level. The 82 geometry-conductivity configurations are deterministically shuffled and partitioned into 5 folds of 16-17 configurations each. For each fold k , fold k serves as the test set, fold $(k+1) \bmod 5$ as the validation set, and the remaining 3 folds as the training set. This yields approximately 4800-5000 training, 1600-1700 validation, and 1600-1700 test samples per fold, with the test set always containing entirely unseen geometries. Both models are trained for 50 epochs per fold with identical hyper-parameters (learning rate 2×10^{-4} , batch size 4, gradient clipping at 1.0). Performance is measured by relative ℓ_2 error on the mean-centered solution.

Table 1: 5-fold cross-validation results (relative ℓ_2 error) on held-out geometry-conductivity configurations. Each fold tests on 16-17 entirely unseen configurations.

| Method | Fold 0 | Fold 1 | Fold 2 | Fold 3 | Fold 4 | Mean \pm Std |
|-------------------|--------|--------|--------|--------|--------|---------------------------------------|
| LNO-FiLM Standard | 0.3204 | 0.3870 | 0.3289 | 0.4066 | 0.3749 | 0.3636 ± 0.0335 |
| LNO-Green’s | 0.2540 | 0.2719 | 0.2731 | 0.2412 | 0.2749 | 0.2630 ± 0.0133 |

Table 1 shows that LNO-Green’s consistently outperforms LNO-FiLM Standard across all five folds, with a relative improvement of 27.7%. These results suggest that the spectral decomposition acts as a strong inductive prior which not only improves accuracy but also reduces sensitivity to the particular choice of held-out geometries, as evidenced by the $2.5\times$ lower cross-validation variance. Notably, LNO-Green’s achieved this with fewer parameters (1.5M vs. 1.9M) and without access to source information during the forward pass through the network.

3.2 SOURCE LOCATION GENERALISATION

In the geometry generalisation experiment above, source locations are drawn randomly across the full spatial domain in both training and test sets. Each of the 8200 samples has unique source positions, so the model routinely generalises to exact source coordinates never encountered during training. To probe the limits of this generalisation, we design a more challenging spatial holdout in which an entire contiguous region of the domain is devoid of sources during training. We split along the longitudinal axis $z \in [-0.5, 0.5]$ and consider two complementary scenarios:

Extrapolation. The model is trained on samples whose active sources all satisfy $z \leq 0.20$ and tested on samples where every source lies at $z > 0.35$. A validation set drawn from the intermediate band $(0.20, 0.35]$ is used for early stopping. This yields a 70%/15%/15% train/val/test split.

Interpolation. The model is trained on samples whose sources lie on the edges of the domain ($|z| > 0.07$) and tested on samples where all sources are concentrated near the centre ($|z| \leq 0.07$). The non-test (edge) data is split randomly into 85%/15% train/val subsets, so both contain samples from both edges. This yields a 73%/13%/14% train/val/test split.

In both cases, geometry-conductivity configurations are shared across splits; only the sources differ. Table 2 reveals that when an entire spatial region is held out from training sources, LNO-Green’s

Table 2: Relative ℓ_2 error to quantify performance on unseen source regions.

| Holdout | LNO-FiLM | LNO-Green’s |
|--|--------------|-------------|
| Source extrapolation ($z > 0.35$) | 0.572 | 1.082 |
| Source interpolation ($ z \leq 0.07$) | 0.382 | 2.048 |

degrades, exceeding a relative error of 1.0 (i.e. worse than predicting the trivial zero solution) in both settings. In the interpolation setting, LNO-Green’s achieves a validation error of 0.198 on edge data, comparable to its geometry holdout performance, yet produces a test error of 2.048 on centre

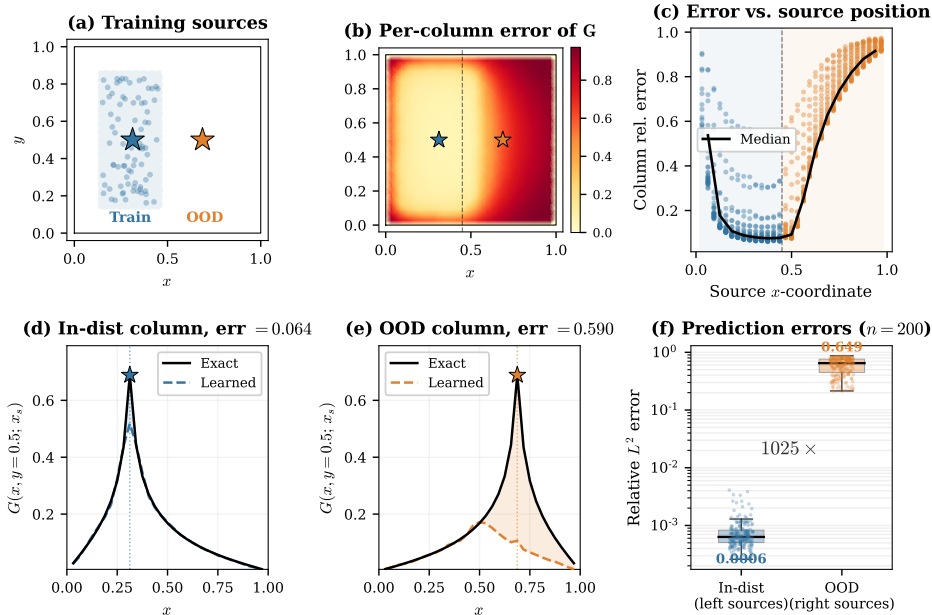


Figure 2: Source extrapolation on $[0, 1]^2$. A direct parametrization (rank 256) is trained on 500 samples with sources in $x \in [0.15, 0.45]$. (a) Training source locations and probe points (stars). (b) Per-column relative error of \mathbf{G}_θ , showing accurate learning only within the training region. (c) Column error vs. source x -coordinate. (d)–(e) Cross-section at $y = 0.5$ of exact (solid) and learned (dashed) Green’s function columns for an in-distribution ($x_s=0.3$) and OOD ($x_s=0.7$) source; shaded area shows discrepancy. (f) Prediction error over 200 test samples per region ($\sim 1000\times$ degradation).

sources. This confirms that the model learns the operator well when training sources provide spatial coverage of the full domain, but fails when an entire region lacks source-side supervision. Although LNO-FiLM also degrades relative to the geometry holdout, it does not do so to the same extent.

It is important to note that this failure mode is distinct from generalisation to individual unseen source locations: in the geometry cross-validation (Table 1), every test sample has unique source coordinates never seen during training, yet LNO-Green’s achieves its best performance. The critical difference is that the training source distribution in that experiment covers the full spatial domain.

The spectral solver computes u according to Equation A.7. Although $\tilde{\Phi}$ appears symmetrically on both sides, enforcing $G(x, x') = G(x', x)$ by construction, the training gradients are asymmetric. On the output side ($u = \tilde{\Phi} \tilde{c}$), the reconstruction loss constrains $\tilde{\Phi}$ at all N grid points whereas on the source side ($c = \tilde{\Phi}^\top f$), gradients flow only through the handful of nonzero entries in f . $\tilde{\Phi}$ is therefore directly supervised as a source-projection basis only at locations where sources appear during training. When training sources are distributed across the full domain, each region receives at least some source-side gradient signal, and the eigenfunctions are adequately calibrated everywhere, explaining the strong performance in the geometry generalisation experiment. However, when an entire region lacks training sources, the eigenfunction values there are shaped solely by the output-side objective and may not produce correct coefficients when queried by the source-side projection.

To further investigate this behavior, we design a controlled experiment on a 2D unit square with Dirichlet boundary conditions. We train the model using only source functions supported in the left half of the domain ($x \in [0.15, 0.45]$) and evaluate on sources placed in the unseen right half ($x \in [0.55, 0.85]$). As shown in Fig. 3.2, the learned Green’s function \mathbf{G}_θ accurately captures columns corresponding to source locations within the training region, but fails entirely for out-of-distribution positions — with a median relative L^2 error degradation of over $1000\times$. This is expected: each training sample (f_i, u_i) constrains \mathbf{G} only along directions excited by f_i , so columns of \mathbf{G} corresponding to unseen source locations remain effectively unconstrained.

3.3 DISCUSSION AND FUTURE WORK

The Green’s formulation provides a strong inductive bias for generalising across geometries and readily handles novel source positions, provided that training sources are spatially distributed across the domain. It is only when entire regions lack source-side coverage that the spectral solver fails, due to the asymmetric gradient flow through the learned eigenfunctions. In practice, this limitation is benign: randomly placed training sources naturally provide full spatial coverage, as demonstrated by the geometry cross-validation results. Nevertheless, for applications where source coverage may be systematically non-uniform, architectural modifications may be needed to address this vulnerability. Future work will use this proof-of-concept to apply neural green’s operators to increasingly realistic geometries. Numerical solvers such as FEM will be used to generate ground truth data, with the intent to condition the models on MRI geometries.

ACKNOWLEDGMENTS

NEN was supported by UK Research and Innovation [UKRI AI Centre for Doctoral Training in Digital Healthcare number EP/Y030974/1]. DH was supported by the Imperial-META Wearable Neural Interfaces Research Centre and the Onassis Foundation under Scholarship ID: F ZT 012-1/2023-2024.

REFERENCES

- Nicolas Boullé, Christopher J. Earls, and Alex Townsend. Data-driven discovery of Green’s functions with human-understandable deep learning. *Scientific Reports*, 12(1):4824, March 2022. ISSN 2045-2322. doi: 10.1038/s41598-022-08745-5. URL <https://www.nature.com/articles/s41598-022-08745-5>.
- D. Farina, L. Mesin, S. Martina, and R. Merletti. A surface EMG generation model with multilayer cylindrical description of the volume conductor. *IEEE Transactions on Biomedical Engineering*, 51(3):415–426, March 2004a. ISSN 1558-2531. doi: 10.1109/TBME.2003.820998. URL <https://ieeexplore.ieee.org/document/1268212>.
- Dario Farina, Luca Mesin, and Sidayne Martina. Advances in surface electromyographic signal simulation with analytical and numerical descriptions of the volume conductor. *Medical & biological engineering & computing*, 42:467–76, August 2004b. doi: 10.1007/BF02350987.
- Andrew Jaegle, Felix Gimeno, Andrew Brock, Andrew Zisserman, Oriol Vinyals, and Joao Carreira. Perceiver: General Perception with Iterative Attention, June 2021. URL <http://arxiv.org/abs/2103.03206>. arXiv:2103.03206 [cs].
- Nikola Kovachki, Zongyi Li, Burigede Liu, Kamyar Azizzadenesheli, Kaushik Bhattacharya, and Andrew Stuart. Neural Operator: Learning Maps Between Function Spaces With Applications to PDEs.
- Zongyi Li, Nikola Kovachki, Kamyar Azizzadenesheli, Burigede Liu, Kaushik Bhattacharya, Andrew Stuart, and Anima Anandkumar. Fourier Neural Operator for Parametric Partial Differential Equations, May 2021. URL <http://arxiv.org/abs/2010.08895>. arXiv:2010.08895 [cs].
- Lu Lu, Pengzhan Jin, and George Em Karniadakis. DeepONet: Learning nonlinear operators for identifying differential equations based on the universal approximation theorem of operators. *Nature Machine Intelligence*, 3(3):218–229, March 2021. ISSN 2522-5839. doi: 10.1038/s42256-021-00302-5. URL <http://arxiv.org/abs/1910.03193>. arXiv:1910.03193 [cs].
- Kostiantyn Maksymenko, Alexander Kenneth Clarke, Irene Mendez Guerra, Samuel Deslauriers-Gauthier, and Dario Farina. A myoelectric digital twin for fast and realistic modelling in deep learning. *Nature Communications*, 14(1):1600, March 2023. ISSN 2041-1723. doi: 10.1038/s41467-023-37238-w. URL <https://www.nature.com/articles/s41467-023-37238-w>.

Hugo Melchers, Joost Prins, and Michael Abdelmalik. Neural Green's Operators for Parametric Partial Differential Equations, January 2026. URL <http://arxiv.org/abs/2406.01857>. arXiv:2406.01857 [cs].

Ethan Perez, Florian Strub, Harm de Vries, Vincent Dumoulin, and Aaron Courville. FiLM: Visual Reasoning with a General Conditioning Layer, December 2017. URL <http://arxiv.org/abs/1709.07871>. arXiv:1709.07871 [cs].

Tian Wang and Chuang Wang. Latent Neural Operator for Solving Forward and Inverse PDE Problems, December 2024. URL <http://arxiv.org/abs/2406.03923>. arXiv:2406.03923 [cs] version: 5.

Seungwoo Yoo, Kyeongmin Yeo, Jisung Hwang, and Minhyuk Sung. Neural Green's Functions. October 2025. URL <https://openreview.net/forum?id=t2wMLGmcXd>.

A NEURAL GREEN’S FUNCTION FORMULATION

EMG is the result of an electric field being propagated due to current source density flowing through muscle fibres. In the quasi-static approximation (Farina et al., 2004b), the electric field potential $u(x)$ satisfies the anisotropic Poisson equation:

$$\nabla \cdot (\sigma(x)\nabla u(x)) = f(x), \quad x \in D \quad (\text{A.1})$$

where D is the three-dimensional domain of the volume conductor with boundary ∂D , $\sigma(x)$ the conductivity tensor, and $f(x)$, the forcing function, is $-I(x)$, the current source density. Since no current flow between the skin and air, Neumann boundary conditions apply:

$$\sigma(x)\nabla u(x) \cdot n(x) = 0, \quad x \in \partial D \quad (\text{A.2})$$

where $n(x)$ is the normal to the boundary. More generally, with \mathcal{L} as the linear differentiable operator:

$$\mathcal{L}u(x) = f(x), \quad x \in D \quad (\text{A.3})$$

In discrete form this becomes:

$$\mathbf{L}\mathbf{u} = \mathbf{M}\mathbf{f} \quad (\text{A.4})$$

where \mathbf{L} is the discretisation of \mathcal{L} , \mathbf{M} is the mass matrix of mesh \mathbf{D} , and \mathbf{f} is the vector of the discretised scalar-valued source function f . When \mathbf{L} is full-rank and symmetric (e.g. in Dirichlet conditions), $\mathbf{u} = \mathbf{G}\mathbf{M}\mathbf{f}$, with $\mathbf{G} = \mathbf{L}^{-1}$ and the Green’s function admits an eigen-decomposition:

$$\mathbf{G} = \Phi\Lambda^{-1}\Phi^T \quad (\text{A.5})$$

where Φ are the eigenvectors of \mathcal{L} and Λ represents the corresponding eigenvalues.

However, under Neumann conditions, \mathbf{L} is singular, possessing a one-dimensional null space corresponding to arbitrary constant offsets, and so its zero eigenvalue mode $\lambda_0 = 0$ means its inverse is not defined. Instead, we define the Green’s function as the the Moore-Penrose inverse:

$$\mathbf{G} = \mathbf{L}^\dagger = \sum_{i \neq 0} \frac{1}{\lambda_i} \phi_i \phi_i^T \quad (\text{A.6})$$

By removing the zero-eigenvalue mode, the minimum norm solution is found, which corresponds to the mean-centered solution: $\sum_i \phi_i = 0$. Physically this constraint Maksymenko et al. (2023) effectively acts as a distributed current sink, balancing the injected source current uniformly across the volume and rendering the system solvable.

Therefore, in our Green’s Neural Operator, we learn Φ and Λ and formulate the spectral solver as:

$$\hat{u} = \tilde{\Phi}\bar{\Lambda}^{-1}\tilde{\Phi}^T(M \odot f) \quad (\text{A.7})$$

where $\tilde{\Phi} = \Phi - \langle \Phi \rangle$ enforces zero-mean eigenfunctions (removing the infinite mode), $\bar{\Lambda} = \text{softplus}(\langle \Lambda \rangle)$ are positive global eigenvalues, and M is the finite-element mass matrix. The physics constraint of $\sum_i \phi_i = 0$ is therefore enforced explicitly on the network output, strictly ensuring the resulting solution to lie in the full-rank, solvable subspace orthogonal to the null space of the Neumann Laplacian. By projecting exactly onto the orthogonal complement, the network is prevented from struggling to learn the arbitrary floating constant associated with the Neumann null space. Also, Green’s functions for linear PDEs must be symmetric and this processing ensured $G_{ij} = G_{ji}$. This approach is related to the framework of Melchers et al. (2026), who similarly embed symmetry and spectral properties into the learned Green’s function representations.

B DATA GENERATION

We generate data using the infinite cylinder model, a standard starting point for simulating EMG for which an analytical solution is available (Farina et al., 2004a). The dataset consists of 8200 samples, with 82 unique geometry-conductivity configurations, each with 100 different source configurations. Each sample has $N = 49,152$ nodes, with one near-Dirac point source of varying location and magnitude. For 16 distinct geometries, conductivity values are fixed, with geometry parameters sampled uniformly from physiologically plausible ranges: bone radius $r_b \in [15, 25]$ mm, muscle radius $r_m \in [40, 50]$ mm, with fat and skin layer thicknesses $t_f \in [2, 5]$ mm and $t_s \in [1, 3]$

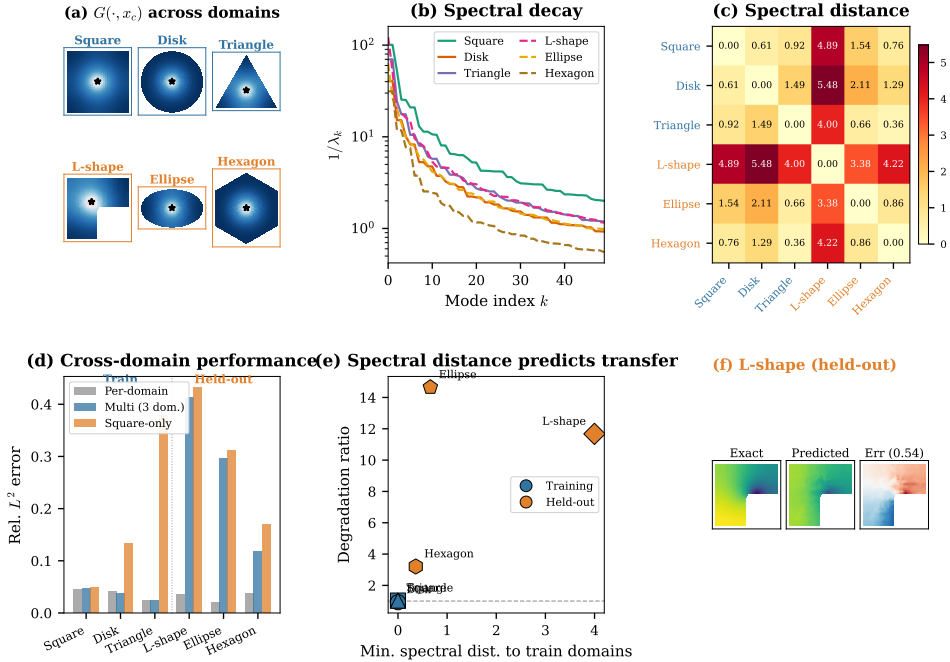


Figure 3: (a) Cross-domain generalization of neural Green’s functions. (a) Point-source response $G(\cdot, x_c)$ for six domains under zero-Neumann boundary conditions, with x_c at each domain’s centroid (black star). Blue labels denote training domains; orange labels denote held-out domains. Each domain produces a geometrically distinct Green’s function. (b) Spectral decay of G , plotted as $1/\lambda_k$ versus mode index. Domains with similar geometry (e.g., Disk and Hexagon) exhibit similar spectral profiles, while the L-shape — with its re-entrant corner — decays markedly slower. Solid lines: training domains; dashed: held-out. (c) Pairwise spectral distance between all domains, computed as the ℓ^2 distance between normalized eigenvalue signatures (first 60 modes). The L-shape is spectrally most distant from all other domains. (d) Test-set relative L^2 error under three training strategies: per-domain (grey, oracle baseline), multi-domain trained on Square+Disk+Triangle (blue), and Square-only (orange). Multi-domain training matches per-domain baselines on training domains but degrades 3–15 \times on held-out domains. (e) Degradation ratio (multi-domain error / per-domain error) versus minimum spectral distance to the training set. Spectral distance is predictive of transfer quality: the Hexagon (3.2 \times) is spectrally closest to the training set, while the Ellipse (14.7 \times) and L-shape (11.7 \times) are most distant. (f) Example prediction on the L-shape (worst held-out domain): exact solution, multi-domain MLP prediction, and pointwise error (relative $L^2 = 0.54$).

mm respectively. For each geometry, 100 samples of varying point source location and magnitude are generated. The same is done for 16 distinct muscle conductivity configurations (with fixed geometry), where muscle conductivities are sampled log-uniformly: transverse $\sigma_t \in [0.05, 0.2]$ S/m and longitudinal $\sigma_\ell \in [0.5, 0.8]$ S/m, while other tissue conductivities remain constant. Finally, we also vary both geometry and muscle conductivity parameters within these ranges to generate a further 50 combinations, again with 100 samples each.

For each sample, the forward problem for a multi-layered anisotropic cylindrical volume conductor was solved analytically using a Fourier Bessel expansion (Farina et al., 2004a). The electric potential $u(r, \theta, z)$ satisfies Equation A.1 with point source forcing $f = \sum_i q_i \delta(\mathbf{r} - \mathbf{r}_i)$, and q_i are source magnitudes. Each sample is evaluated on a cylindrical grid ($n_r \times n_\theta \times n_z = 32 \times 24 \times 64$ points) and stored as a point cloud with coordinates $\mathbf{x} \in \mathbb{R}^{N \times 3}$, forcing function $f \in \mathbb{R}^N$, potential field $u \in \mathbb{R}^N$, and mass matrix diagonal $m \in \mathbb{R}^N$ (accounting for the r -dependent volume element $dV = r dr d\theta dz$). A zero-mean constraint is applied to remove gauge freedom: $u \leftarrow u - \langle u \rangle$.

C GEOMETRY GENERALIZATION

The Green’s function \mathbf{G} is determined entirely by the domain geometry and boundary conditions, independent of the source function f . To understand how \mathbf{G} varies across domains and whether a single model can generalize across geometries, we train an MLP-based model on three domains (Square, Disk, Triangle) and evaluate on three held-out domains (L-shape, Ellipse, Hexagon), all under zero-Neumann conditions.

As shown in Figure C(a), the point-source response of \mathbf{G} is qualitatively different across domains: convex domains produce smooth, roughly radial responses, while the L-shape exhibits anisotropic decay due to its re-entrant corner. These differences are captured quantitatively by the eigenvalue spectrum of the Laplacian (Figure C(b)). We define a spectral distance between domains as the ℓ^2 distance between their normalized eigenvalue signatures and observe that it clusters geometrically similar domains together (Figure C(c)).

The multi-domain MLP matches per-domain baselines on training domains (degradation $\leq 1.1\times$), confirming that a single network can represent multiple Green’s functions simultaneously without loss of capacity. However, transfer to held-out domains is limited: the Hexagon, which is spectrally closest to the training set, degrades only $3.2\times$, while the L-shape and Ellipse degrade $12\text{--}15\times$ (Figure C(d,e)). Crucially, the spectral distance to the nearest training domain is predictive of this degradation, suggesting that transfer success depends on the similarity of the underlying eigenfunction basis rather than superficial geometric resemblance.

This result highlights a fundamental asymmetry: whereas the Green’s function varies smoothly with continuous geometry parameters such as aspect ratio (enabling interpolation between trained configurations), transferring across topologically or spectrally distinct domains requires the model to extrapolate to an entirely different eigenfunction basis — a much harder task.