NEURAL OPERATORS FOR VARYING GEOMETRY IN THE FORWARD EMG MODEL

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

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

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

ABSTRACT

Forward electromyography(EMG) simulation requires repeated PDE solves in anatomically structured volume conductors, which becomes expensive when exploring different geometries or running inversion. We study neural operators as fast surrogates that map volume-conductor geometry to the full extracellular potential field. Using a simplified multi-geometry forearm setting (layered cylinders with varying tissue thickness), we evaluate operator models across both geometry changes and grid resolutions, and assess their impact on motor unit action potentials(MUAP) generation. The results show that neural operators can provide accurate fields and enable rapid, interactive workflows. This workshop paper is a proof of concept; future work will extend to MRI-derived forearm geometries and broader sources of variability.

1 INTRODUCTION

Biophysics simulations provide a way of conducting experiments and training models in the absence of sufficient experimental data. EMG research is one such case, where experimental recordings are not only expensive and time-consuming to obtain, but also incomplete as all parameters are not able to be measured (Maksymenko et al., 2023). The simulations require solving volume conduction partial differential equations (PDEs) within complex domains representing human anatomy.

Current state-of-the-art models (Maksymenko et al., 2023) (Ma et al., 2024) are based on computationally-expensive finite element methods (FEM) for solving PDEs. Although FEM is highly-accurate, it is restricted to static conditions. This is a significant limitation because, during dynamic, naturalistic movement, the volume conductor deforms as fibres contract, muscles move and physiological parameters change. This requires recalculations as the system changes, which is infeasible given the computationally demanding nature of FEM solvers. Current methods are therefore restricted to single static states.

By training on FEM solutions rather than using them directly, data-driven deep learning alternatives remove the need for FEM recalculations. In this work, we investigate the use of neural operators, building on the preliminary work using physics-informed neural networks (PINNs) for a simple cylindrical geometry (Halatsis et al.). Unlike PINNs, neural operators are resolution-invariant and able to learn a full solution operator (Kovachki et al.) rather than just the single solution, which makes them ideal for EMG simulation.

In this work, we study geometry variation in a simplified setting: a layered cylindrical forearm model where tissue thicknesses vary. This is a proof of concept designed to isolate how geometry

*Equal contribution.

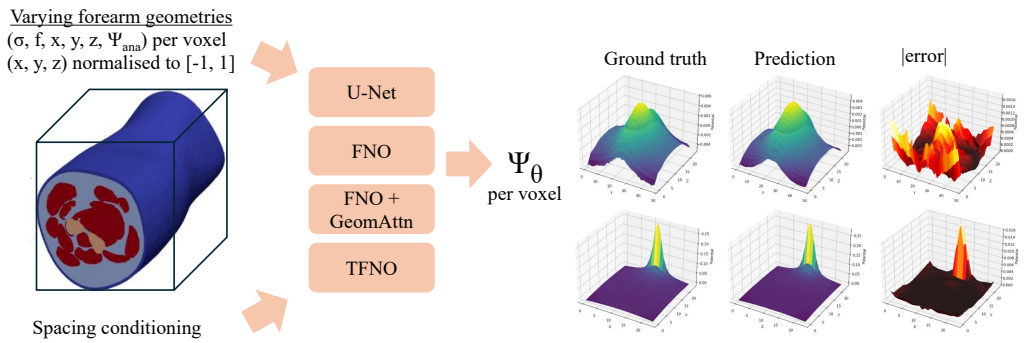

Figure 1: Different volume-conductor geometries are used to train a neural operator to predict the full electric potential field. Once trained, the model enables interactive tools—such as a live volume-conductor editor—that update MUAP properties immediately as the geometry changes.

and resolution affect neural-operator surrogates. In future work, we will extend the same framework to subject-specific forearm geometries derived from MRI.

2 METHOD

2.1 BIOPHYSICAL MODEL

EMG is the result of an electric field being propagated due to current source density flowing through muscle fibres. Simulating dynamic EMG signals therefore requires solving the steady-state Maxwell equations that govern the process. These reduce to the Poisson equation in the quasi-static approximation (Farina et al., 2004b):

$$\nabla \cdot (\sigma \nabla \varphi) = -I \quad \text{in} \quad \boldsymbol{\Omega} \tag{1}$$

where $\boldsymbol{\Omega}$ is the three-dimensional domain of the volume conductor with boundary $\partial\boldsymbol{\Omega}$, $\varphi(\boldsymbol{r})[V]$ is the electric field potential, σ the conductivity tensor, and $I(\boldsymbol{r})[\boldsymbol{A}/m^3]$the current source density. Given that there is assumed to be no current flow between the skin and air, Neumann boundary conditions apply:

$$\nabla\varphi \cdot \mathbf{n} = 0 \quad \text{on} \quad \partial\boldsymbol{\Omega} \tag{2}$$

where $\mathbf{n}$ is the normal to the boundary. Quasi-static conditions means that the conductivity tensor does not depend on the potential, making the equations linear.

Traditional approaches involved solving for each fiber at each instant, making the process of simulating realistic forearms with thousands of fibers impractically costly. More recent methods (Maksymenko et al., 2023) adopt a more efficient strategy of solving the equations at a set of static points once, and later integrate fibers as potentials along a subset of these pre-solved points. However, this solution has to be expensively recalculated for each new geometry. Our objective therefore is to solve the parametric volume conduction equation for a single point source, overcoming the bottleneck of current simulation efforts.

We solve a "source-at-electrode" forward problem: we place a monopole source at the electrode location and compute the resulting potential field in the volume conductor. We then use this field as a transfer map to obtain the potential at any point along a muscle fiber. This formulation is consistent with reciprocity: the field generated by a source at the electrode can be used to compute electrode measurements from sources along the fiber.

2.2 DATASET

Dataset generation. We generate ground-truth potential fields using FEM simulations of current flow in a parametric, layered cylindrical volume conductor with five concentric regions: cancellous bone, cortical bone, muscle, fat, and skin. Each geometry is defined by five interface radii $(r_1, \ldots, r_5)$ (equivalently, layer thicknesses) that set the tissue boundaries, with $0 < r_1 < r_2 <$

$r_3 < r_4 < r_5$; we sample these parameters within predefined ranges using Latin Hypercube Sampling (LHS) and enforce validity by ensuring strictly increasing radii (e.g., by sampling thicknesses and converting to cumulative radii). For each sampled geometry, we assign fixed tissue conductivities, place a monopole source at the electrode location on the skin surface, and solve the forward problem to obtain the 3D extracellular potential field on the chosen grid. We train on 901 base-resolution samples comprising 501 original low-resolution simulations ($48\times48\times96$ voxels, spacing tag 2.1) and 400 downsampled high-resolution simulations ($48\times48\times96$ voxels, spacing tag 1.6); for mixed-resolution training we additionally include 50 native high-resolution samples ($96\times96\times192$ voxels, spacing tag 0.8). Here, "spacing" is used only as a resolution tag: all coordinates are normalized to $[-1, 1]$, so the physical domain size is fixed and does not change across resolutions. The base-resolution data are split 70/15/15 into train/validation/test sets using a fixed random seed, and we reserve an independent set of 50 high-resolution samples exclusively for super-resolution evaluation.

2.3 MODELS

We compared four architectures: 3D U-Net with 4 levels and 32 base channels; FNO (Fourier Neural Operator)(Li et al., 2021) with 6 layers, 8 Fourier modes per dimension, and width 32; FNO with geometry cross-attention (Li et al., 2023), using lightweight attention (54 tokens, 32-dim) at layers 4-5; and TFNO (Tensorised FNO) (Kossaifi et al., 2023) with instance normalization. All models except U-Net received an analytical monopole solution as auxiliary input and used spacing-conditioned MLPs for resolution awareness. The analytical monopole solution was calculated using the average of the anisotropic muscle conductivity.

Models were trained for up to 100 epochs with early stopping (patience=10), using Adam optimiser (learning rate=0.001, weight decay=1e-5) with cosine annealing. The loss function combined MSE (with 3-voxel singularity masking the source) and total variation regularisation (weight=0.01). Each configuration was trained with 10 random seeds for uncertainty quantification.

We reported relative L2 error, the L2 norm ratio and gradient energy ratio on the muscle region. L2 norm ratio indicated scale accuracy, and the gradient energy ratio quantified smoothness. Models were evaluated on both the low-resolution test set and the held-out high-resolution test set to assess super-resolution capability.

2.4 FROM POTENTIAL TO SFAP

The ultimate goal of biophysical simulation is to generate the motor unit action potential (MUAP), which is the impulse response of a motor unit (i.e., a group of muscle fibers). To compute a MUAP, we first generate single-fiber action potentials (SFAPs): we sample the extracellular potential field along each fiber trajectory and then compute the second spatial derivative along the fiber direction (Farina et al., 2004a). This derivative-based transformation amplifies high-frequency components, making SFAP generation highly sensitive to prediction errors and noise in the potential field. The challenge is particularly acute for fibres located at greater depths or oblique angles relative to the electrode, where the potential magnitude decays substantially—following approximately an inverse to the distance. In these low-amplitude regions, even small absolute errors in the predicted field translate to large relative errors in the resulting SFAP waveform. Figure 2 illustrates this effect: prediction accuracy remains high ($r > 0.95$) for shallow fibres but degrades at the deepest positions ($r = 20$) where signal amplitude is minimal. This spatial dependence of prediction quality has direct implications for simulating realistic motor units, which contain fibres distributed across a range of depths and angles within the muscle volume.

3 EXPERIMENTS

Runtime improvements: For a single geometry and electrode configuration, computing the transfer field with FEM takes 13 s. Once the field is available, evaluating potentials at 1000 query points takes 100 ms with FEM versus 1.5 ms with the neural operator (a $\sim67\times$ speedup for point queries). The main advantage appears when geometry varies: FEM must be re-solved (13 s per new geometry), whereas the trained neural operator requires only a forward pass (no additional PDE solve), enabling $\sim8700\times$ faster end-to-end evaluation for 1000-point queries. Training on 1000 FEM ge-

Table 1: Performance of neural operator architectures on base-resolution (48×48×96) and high-resolution (96×96×192) test sets. Results are mean ± std over 10 random seeds.

		Base-res test data			High-res test data		
Model	**Training data**	**Rel L2**	**L2 norm ratio**	**Grad energy**	**Rel L2**	**L2 norm ratio**	**Grad energy**
U-Net	Base-res	0.530 ± 0.106	1.016 ± 0.137	1.138 ± 0.0952	1.25 ± 0.388	0.746 ± 0.504	2.964 ± 3.34
FNO	Base-res	$\mathbf{0.152 \pm 0.020}$	0.980 ± 0.0237	0.913 ± 0.0410	0.495 ± 0.159	0.752 ± 0.171	0.885 ± 0.0753
FNO + GeomAttn	Base-res	0.171 ± 0.023	0.994 ± 0.0276	0.939 ± 0.0608	0.504 ± 0.147	0.869 ± 0.230	0.938 ± 0.205
TFNO	Base-res	0.166 ± 0.017	0.978 ± 0.0408	1.01 ± 0.0682	0.447 ± 0.165	1.14 ± 0.196	1.23 ± 0.142
FNO	Mixed-res	$\mathbf{0.144 \pm 0.0155}$	0.985 ± 0.0377	0.895 ± 0.0392	$\mathbf{0.128 \pm 0.0226}$	0.996 ± 0.0610	0.970 ± 0.0353
FNO + GeomAttn	Mixed-res	0.165 ± 0.0188	0.972 ± 0.0327	0.938 ± 0.0682	0.155 ± 0.0306	0.965 ± 0.0499	1.05 ± 0.0713
TFNO	Mixed-res	0.170 ± 0.0213	0.971 ± 0.0321	1.05 ± 0.121	0.1623 ± 0.0402	0.923 ± 0.0554	1.15 ± 0.184

Table 2: MUAP results (mean ± std) on low- and high-resolution test sets.

Test Set	**Correlation**	**Rel L2**	**Rel MSE**	**P2P Ratio**
Low-res (48×48×96)	0.9927 ± 0.0186	0.2034 ± 0.1623	0.0677 ± 0.1373	0.9743 ± 0.2514
High-res (96×96×192)	0.9372 ± 0.0328	0.3655 ± 0.0869	0.1412 ± 0.0630	0.8265 ± 0.2125

ometries takes ∼5 h; since generating those FEM solutions already costs ∼3.6 h, the additional overhead to obtain the surrogate is ∼1.4 h, after which repeated evaluations amortize quickly while maintaining strong MUAP agreement (mean $r \approx 0.95$).

Interpolation Performance: Table 1 shows that among models trained exclusively on base resolution data, the standard FNO achieved the best performance with a relative L2 error of 0.152 ± 0.020, outperforming the more complex variants and the U-Net baseline. However, all base-resolution-trained models exhibited significant performance degradation when evaluated on high-resolution data, with relative L2 errors increasing by 2–3×.

Mixed-resolution training, where 50 high-resolution samples were included alongside 901 base-resolution samples, dramatically improved generalisation to unseen resolutions. The FNO trained with mixed-resolution data achieved a high-resolution relative L2 error of 0.128 ± 0.023, which was a 3.9× improvement over the same architecture trained on base resolution alone (0.495 ± 0.159). Notably, the L2 norm ratio improved from 0.752 to 0.996, indicating that mixed training helped correct the systematic scale errors observed in base-resolution-only models.

Despite most metrics focus on the potential field reconstruction, the final metric is MUAP reconstruction, as this is used in most biophysical modeling applications such as benchmarking (Mamidanna et al., 2025), biophysical-model inversion (Halatsis et al., 2025) or general prototyping. We report the MUAP metrics in table 2 and some sample MUAPs can be viewed in figure 2. Overall, due to the high accuracy of the field, most high power MUAPs waveforms are very accurate but performance degrades as we move deeper within the volume conductor and the potential field weakens and becomes noisier.

Overall, these results show that neural operators are a useful tool for biophysical simulation. They enable fast, interactive workflows where users can modify a volume conductor and immediately see the resulting signals, and they also support continuous biophysical model inversion by making repeated forward evaluations practical.

The trade-off between speed and accuracy as the estimation error is negligible for most applications such as EMG decomposition and simulation

4 Discussion and future work

In this current iteration, we focus on parametric geometries as a proof of concept. In future work, we will move to MRI-derived forearm geometries and train models that generalize across a wide range of anatomies. The goal is a tool that can predict how signals vary between individuals in a seamless way, while also capturing rare or unusual cases (outliers). The use of neural network approximations is well motivated, the tool must be fast for intuitive live use, the error introduced for approximation is negligible for most applications like EMG simulation or decomposition.

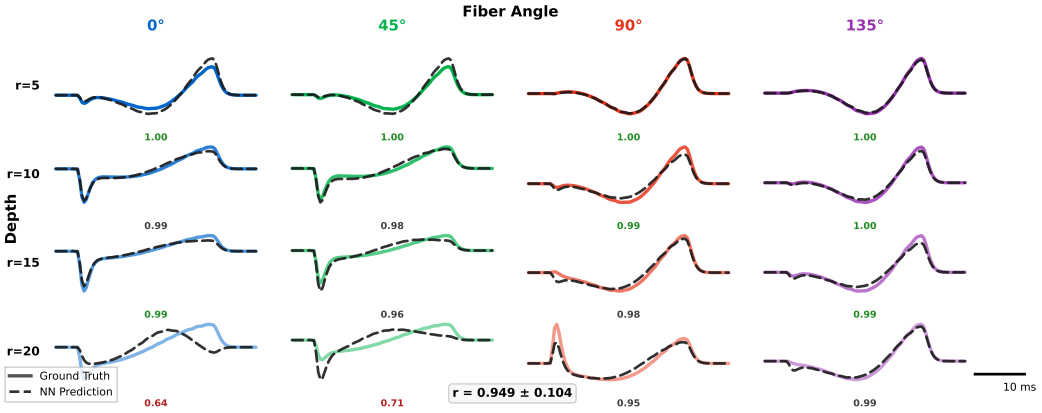

Figure 2: Ground truth (solid) vs. neural network predicted (dashed) MUAPs across fiber depths (rows, r = 5–20 voxels) and angles (columns, 0°–135°). Waveforms are individually normalized; colors encode angle and depth. Correlation coefficients shown below each pair (mean r = 0.949 ± 0.104). Prediction accuracy remains high ($r > 0.95$) except at the deepest positions with lowest signal amplitude.

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

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.

D Halatsis, P Mamidanna, J Pereira, and D Farina. A biophysical-model-informed source separation framework for emg decomposition. In *2025 47th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 1–7. IEEE, 2025.

Dimitrios Halatsis, Alexander Clarke, and Dario Farina. Modelling variation in the forward EMG model.

Jean Kossaifi, Nikola Kovachki, Kamyar Azizzadenesheli, and Anima Anandkumar. Multi-grid tensorized fourier neural operator for high-resolution pdes. *arXiv preprint arXiv:2310.00120*, 2023.

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].

Zongyi Li, Nikola Kovachki, Chris Choy, Boyi Li, Jean Kossaifi, Shourya Otta, Mohammad Amin Nabian, Maximilian Stadler, Christian Hundt, Kamyar Azizzadenesheli, et al. Geometry-informed neural operator for large-scale 3d pdes. *Advances in Neural Information Processing Systems*, 36:35836–35854, 2023.

Shihan Ma, Alexander Kenneth Clarke, Kostiantyn Maksymenko, Samuel Deslauriers-Gauthier, Xinjun Sheng, Xiangyang Zhu, and Dario Farina. Conditional Generative Models for Simulation of EMG During Naturalistic Movements. *IEEE Transactions on Neural Networks and Learning Systems*, pp. 1–14, 2024. ISSN 2162-2388. doi: 10.1109/TNNLS.2024.3438368. URL https://ieeexplore.ieee.org/abstract/document/10636282. Conference Name: IEEE Transactions on Neural Networks and Learning Systems.

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.

Pranav Mamidanna, Thomas Klotz, Dimitrios Halatsis, Agnese Grison, Irene Mendez Guerra, Shihan Ma, Arnault H. Caillet, Simon Avrillon, Robin Rohlén, and Dario Farina. MUniverse: A simulation and benchmarking suite for motor unit decomposition. In *The Thirty-ninth Annual Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2025. URL https://openreview.net/forum?id=Slrp3l7aYo.

A Extended Discussion and Future Work

Discussion. Our results suggest that neural operators are a practical surrogate for forward EMG simulation in settings where many evaluations are needed. In the base-resolution setting, operator-based models consistently predict the full potential field with low error, and the same trained model can be evaluated on different grids when trained with mixed-resolution data. This matters because EMG pipelines often need repeated forward solves (e.g., sweeping fiber locations, electrode placement, or tissue thickness), and classical FEM becomes a bottleneck in interactive or optimization-based workflows.

A key observation is that resolution changes are not handled well by default. Models trained only at one resolution degrade noticeably on higher-resolution grids, even when the geometry family is unchanged. Mixed-resolution training greatly reduces this gap and leads to substantially better high-resolution performance. This supports the view that "resolution robustness" is mostly a data-coverage issue in practice: the model must see multiple discretizations during training to avoid learning resolution-specific shortcuts.

We also find that geometry-aware conditioning (e.g., providing geometry information explicitly) can improve robustness, but its benefit depends on how geometry is represented. In our current proof-of-concept setting, geometries are generated by varying tissue thicknesses in a layered cylindrical model. This allows controlled experiments and makes it easy to generate training data, but it does not capture the full anatomical variability of real forearms. As a result, our conclusions mainly apply to parametric geometry changes and should not yet be interpreted as full generalization across subject-specific anatomy.

Finally, while our evaluation focuses on field-level metrics, the downstream goal is accurate signal generation. MUAPs (and SFAPs) depend on sampling the field along fiber trajectories and computing second spatial derivatives, which can amplify small field errors. This motivates measuring not only global field error but also signal-level metrics. In this workshop version we provide initial MUAP results and use them to highlight where improved field fidelity translates into improved waveform quality.

Limitations. This study has several limitations. (1) The geometry family is simplified (layered cylinders) and does not include realistic, irregular tissue boundaries. (2) Source placement is fixed in our dataset, so we have not tested generalization to varying electrode positions or multiple electrodes. (3) Conductivity parameters are fixed, so we do not yet study uncertainty or variability in tissue properties. (4) "Spacing" is treated as a resolution tag with normalized coordinates, so we do not capture changes in physical scale. (5) We do not report runtime and memory comparisons in detail, although these are important for the intended interactive use cases.

Future work: MRI-based geometries. A main next step is to move from parametric cylinders to forearm geometries derived from MRI segmentations. This will introduce realistic anatomical variation (e.g., non-circular boundaries, varying muscle shapes, bone positioning) and will provide a stronger test of whether the surrogate can generalize across individuals. A practical goal is to train a model that can take an MRI-derived conductivity map (or a compact representation of it) and predict the potential field without re-training per subject.

Future work: richer variability and better generalization. Beyond geometry, we plan to expand the dataset to include variability in electrode placement, fiber trajectories, and tissue conductivities. This is important for building surrogates that remain accurate when used inside larger optimization or inversion loops. We also plan to explore geometry encodings that scale to high-resolution segmentations (e.g., signed distance fields or multi-channel tissue labels) and to test which conditioning strategies transfer best across subjects.

Future work: signal-level evaluation and interactive tools. We ultimately aim to support fast exploration and inversion of biophysical models. On the evaluation side, we will add more signal-level metrics (e.g., MUAP amplitude and timing errors, waveform correlation, and robustness across fiber depths and orientations) and link them more directly to field-level errors. On the application side, our long-term goal is an interactive "volume conductor editor" where users can modify tis-

sue boundaries or properties and immediately observe predicted MUAP/SFAP changes, making it feasible to study inter-subject variability and detect outlier anatomies in real time.