Towards Real-Time Simulations Of Induced Electric Fields During Brain Stimulation Using Conditioned Transformers

Fabian Greifeneder
Research Department Medical Informatics
RISC Software GmbH, Hagenberg, Austria

Dominik Freinberger
Research Department Medical Informatics
RISC Software GmbH, Hagenberg, Austria

Philipp Moser
Research Department Medical Informatics
RISC Software GmbH, Hagenberg, Austria
philipp.moser@risc-software.at

Abstract

Real-time simulations of the induced electric fields during transcranial magnetic stimulation play an important role in guiding and optimizing the coil positioning. In this paper, we present our ongoing work on a deep learning-based surrogate model designed to rapidly predict the induced electric field distribution across the entire cortex, offering a much faster alternative to traditional numerical solvers. Leveraging (conditioned) transformer architectures, our approach operates directly on mesh-based head geometries, achieving highly accurate simulations in just 0.08 seconds on consumer hardware. While we continue to improve the neural surrogate, its current accuracy and efficiency have already enabled integration into an augmented reality platform, demonstrating a promising foundation for live electric field-guided brain stimulation applications.

1 Introduction

Transcranial magnetic stimulation (TMS) is a non-invasive medical procedure where a coil placed on a subject's scalp induces intracortical electric currents thereby enhancing or inhibiting neuronal activity. TMS is primarily used to treat neuropsychiatric disorders but is also being explored in preoperative functional mapping in tumor patients to minimize the risk of post-surgical deficits. Simulating the induced electric fields can help in efficiently and effectively positioning the coil during treatment to optimize targeted brain stimulation [1]. For live visualizations of the electric field in neuronavigation systems (e.g., via augmented reality) while moving the coil, these simulations necessitate (near) real-time computations.

In modern deep learning (DL), neural surrogates have emerged as an attractive alternative to traditional numerical solvers for approximating the solutions of complex partial differential equations (PDEs) that govern the behavior of physical systems. While numerical solvers offer high accuracy, they are often computationally prohibitive for real-time applications, especially with nonlinear physical models and high-resolution geometries. In contrast, neural PDE surrogates learn input-output relationships from precomputed numerical data, enabling both fast and accurate predictions. Although such surrogate approaches have been applied to e-field modeling in TMS [2], they primarily used convolution-based neural networks. This requires converting the 3D brain mesh to a regular grid, applying the surrogate model, and then mapping the predicted e-fields back onto the mesh, which is cumbersome, increases latency, and is inefficient in terms of memory.

The Third Austrian Symposium on AI, Robotics, and Vision (AIROV26).

To overcome this detour, this paper presents our ongoing work on rapid e-field simulations using neural PDE surrogates that operate directly on mesh geometries. We leverage recent (diffusion) transformer architectures to jointly process mesh-based subject geometries with positional information of the TMS coil, achieving both highly accurate and fast e-field predictions, even on consumer hardware.

2 Methods

2.1 Neural PDE Surrogate: Model Architecture and Training Details

Neural operators are a popular concept to approximate solutions to PDEs by learning a mapping $\mathbf{G}$ between an input space $\mathbf{U}$ and an output space $\mathbf{V}$ as $\mathbf{G} : \mathbf{U} \to \mathbf{V}$. Our work builds on *Universal Physics Transformers* (UPT) [3], a learning approach that approximates $\mathbf{G}$ via three maps: $\mathbf{G} \approx \hat{\mathbf{G}} := \mathbf{D} \circ \mathbf{A} \circ \mathbf{E}$, with encoder $\mathbf{E}$, approximator $\mathbf{A}$, and decoder $\mathbf{D}$. In a supervised learning fashion, the network's parameters are optimized during training by repeatedly evaluating the N input-output pairs $(u_i, v_i) = (u_i, \mathbf{G}(u_i)), i = 1..., N$.

In the encoder $\mathbf{E}$, the k mesh coordinates were first normalized using global min-max scaling and multiplied by 300, following [3], and then embedded into the latent space using multi-scale sine-cosine positional encoding. Further, each coordinate had four input features—tissue conductivity and three Euclidean distances to the coil center and both loop centers. These features were linearly projected via a multi-layer perceptron to the latent dimension and added to the latent positional input. Next, a message-passing layer aggregated information at $n_s \ll k$ randomly sampled supernodes from within a radius r_{sn} around each supernode. The TMS coil position (x, y, z) and orientation (3×3 rotation matrix) were linearly projected to the latent dimension and injected as modulation/conditioning [3, 4] in the encoder's diffusion transformer and perceiver blocks.

Due to our stationary simulation scenario in this work, the approximator $\mathbf{A}$ consists solely of regular transformer blocks and employs no temporal evolution as in the original UPT [3]. The decoder $\mathbf{D}$ also starts with conditioned diffusion transformer blocks (using a separate multi-layer perceptron projecting the 12 coil positioning values to the latent dimension). Perceiver-like cross-attention layers with queries based on positional-encoded coordinates finally generate the induced electric field across the entire cortex, see Figure 1.

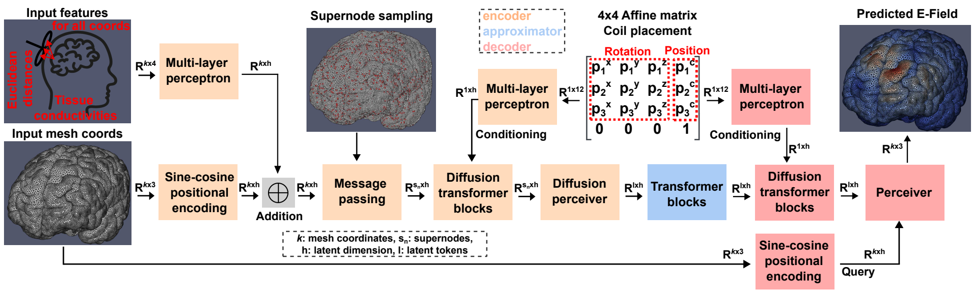

Figure 1: Schematics of the proposed transformer-based neural PDE surrogate.

The following summarizes the selected model hyperparameters: cortical coordinates k: 150000, supernodes n_s: 8000, supernode aggregation radius r_{sn}: 8, latent tokens: 160, latent dimension: 160, attention heads in all transformer blocks: 4, drop path: 0.25, total trainable parameters: 5.36 million. A learning rate of 5×10^{-5}, a batch size of 16, Lion optimizers with weight decay of 0.25, and a weighted mean squared error (MSE) loss (low/high weights for small/large target e-field values) were chosen. All trainings and evaluations were performed on a local workstation (Intel i7-13700 and NVIDIA GeForce RTX 4090 24GB) in PyTorch 2.2.0.

2.2 Generation of training/testing data: Numerical simulations of induced electric fields

A comprehensive dataset for training and evaluation of our proposed neural surrogate was precomputed using SimNIBS v4.01 [5], a popular toolkit for finite element-based simulations of non-invasive brain stimulation. First, realistic head models for 16 randomly selected subjects of the publicly available WU-Minn Human Connectome Project were constructed: Based on 3T T_1w- and T_2w-MR images, SimNIBS's preprocessing pipeline *charm* automatically performed subject-specific whole-head multi-tissue segmentations, assigned conductivities to various tissues types and created computational brain meshes. Second, we randomly selected 24 rotations of a MagVenture-MCF-B65 butterfly coil on the 21 central positions of the EEG-10-10 system, totaling 8064 samples. For each coil placement, SimNIBS was used to simulate the induced electric field on the gray matter cortex, see Figure 2. We used a random train-test split of $75\%/25\%$.

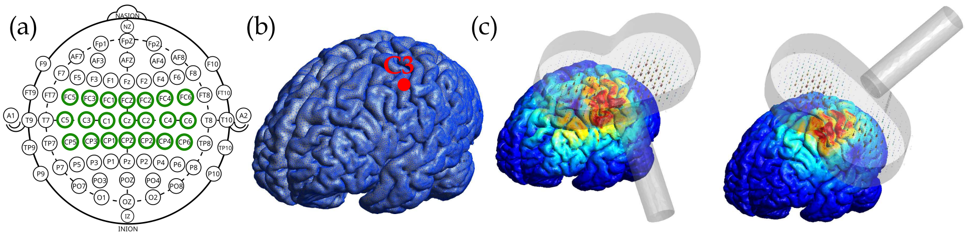

Figure 2: (a) EEG-10-10 system with central positions in green. (b) Head model of one subject with indicated cortical position C3. (c) Two exemplary coil rotations on the C3 position and the corresponding induced electric field magnitude on the cortical gray matter.

3 Results

Evaluating the trained surrogate model demonstrated highly accurate and smooth predictions, characterized by low MSE values and strong similarities of the predicted and ground truth electric field distributions. We achieved mean MSEs of $(1.36 \pm 0.34) \times 10^{-3}$ V/m and $(1.39 \pm 0.37) \times 10^{-3}$ V/m averaged over all train and test samples, respectively (see Figure 3)a. For four test set samples, the exemplary field distributions induced by different coil placements are illustrated in Figure 3b. The training of the neural PDE surrogate with 180 epochs took 82 hours on our consumer hardware.

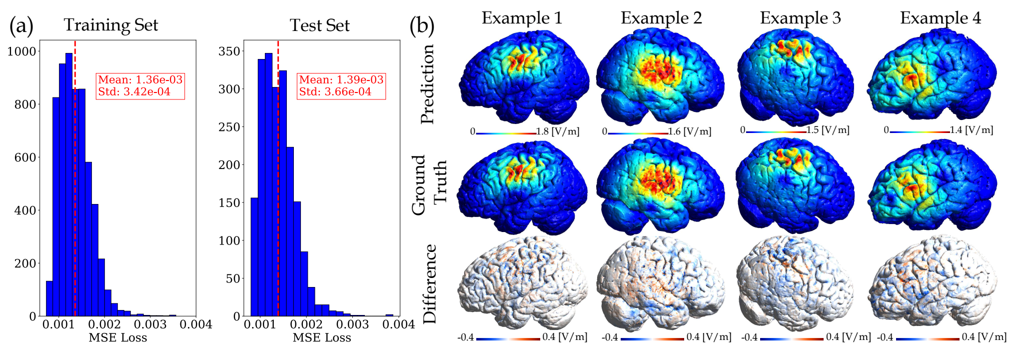

Figure 3: (a) Histogram of the train/test MSE losses. (b) Exemplary e-field magnitude predictions.

3.1 Integration into an Augmented Reality application

For a selected subject from our dataset, we implemented an augmented reality (AR) application in Unity for real-time, electric field-guided TMS, running on a Meta Quest 3. As illustrated in Figure 4, a virtual TMS coil is moved over a virtual, semi-transparent scalp, with the induced electric fields visualized on the gray matter below. Coil orientation and position are tracked by the AR headset and streamed to a Python backend, where a trained surrogate model predicts the corresponding e-fields in real time (0.075 s per simulation including latency, corresponding to the overall coil repositioning time; i.e., 130 times faster compared to the 10 seconds of SimNIBS's default conjugate gradient solver).

4 Discussion and Work-in-Progress

Simulating the induced electric field during brain stimulation has been described as beneficial for accurate, fast, and reproducible coil handling [1, 5]. To this end, our proposed neural surrogate

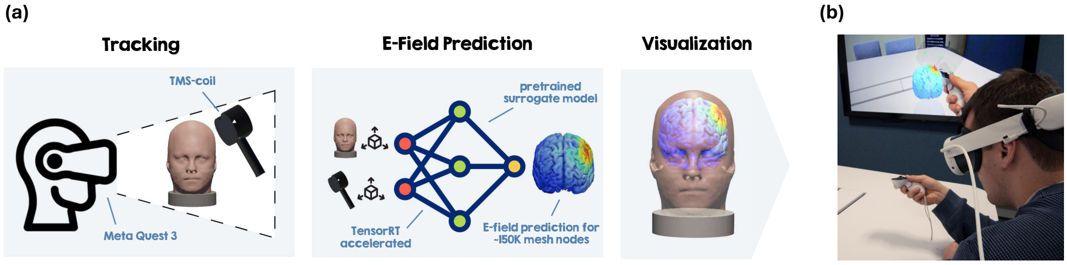

Figure 4: (a) Schematic overview of the AR application and its (b) live use.

leverages transformer-based techniques to rapidly generate highly precise cortical field distributions based on subject-specific head geometries. A key advancement is that we directly operate on mesh geometries, avoiding the need for convolutional architectures that require cumbersome (re-)sampling onto regular grids. Also, diffusion transformers [4] allowed to elegantly merge the TMS coil positional information with the mesh-based geometrical data via model conditioning.

While the model was successfully trained using diverse coil positions across the scalp, careful considerations were necessary to stabilize the training process due to repeated and significant loss spikes (as also mentioned in the original UPT paper [3]). We could control these by employing Lion optimizers with weight decay, a low learning rate, and float32 precision.

In summary, our neural PDE surrogate has already demonstrated promising results and has been successfully integrated into an AR application for electric field-guided TMS, providing real-time insights into brain stimulation. As a neural surrogate, the model is limited by the quality and diversity of the numerical simulations it is trained on. Hence, our ongoing efforts to include a broader range of subjects and cortical positions will ensure robust cross-subject generalizability. Additionally, feedback from medical professionals will help refine the AR application and support its future integration into a clinical TMS study. These advancements pave the way for more personalized, effective, and real-time electric field-guided TMS workflows.

Acknowledgments and Disclosure of Funding

This project is financed by research subsidies granted by the government of Upper Austria (MIMAS.ai, mimas.ng) and by the FFG (nARvibrain, grant no. 894756). www.ffg.at. “ICT of the Future” programme – an initiative of the Federal Ministry for Climate Action, Environment, Energy, Mobility, Innovation and Technology (BMK). RISC Software GmbH is a member of UAR (Upper Austrian Research) Innovation Network.

References

[1] Opitz, A., Zafar, N., Bockermann, V., Rohde, V. & Paulus, W. (2014) Validating computationally predicted TMS stimulation areas using direct electrical stimulation in patients with brain tumors near precentral regions. *NeuroImage: Clinical* **4**:500–507.

[2] Park, T.Y., Franke, L., Pieper, S., Haehn, D. & Ning, L. (2024) A review of algorithms and software for real-time electric field modeling techniques for transcranial magnetic stimulation. *Biomedical Engineering Letters* **14**(3):393–405.

[3] Alkin, B., Fürst, A., Schmid, S., Gruber, L., Holzleitner, M. & Brandstetter, J. (2024) Universal Physics Transformers: A Framework For Efficiently Scaling Neural Operators. In *Proceedings of the 37th Annual International Conference of the Advances in Neural Information Processing Systems (NeurIPS)*, pp. 25152-25194

[4] Peebles, W. & Xie, S. (2023) Scalable Diffusion Models with Transformers. In *Proceedings of 2023 IEEE/CVF International Conference on Computer Vision (ICCV)*, pp. 4172-4182

[5] Thielscher, A., Antunes, A. & Saturnino, G.B. (2015) Field modeling for transcranial magnetic stimulation: A useful tool to understand the physiological effects of TMS? In *Proceedings of the 37th Annual International Conference of the IEEE Engineering in Medicine and Biology Society (EMBC)*, pp. 222–225.