





MR electrical properties mapping using vision transformers and canny edge detectors

Ilias I. Giannakopoulos¹   | Giuseppe Carluccio²  | Mahesh B. Keerthivasan³  |
Gregor Koerzdoerfer³ | Karthik Lakshmanan¹ | Hector L. De Moura¹  |
José E. Cruz Serrallés¹  | Riccardo Lattanzi¹ 

¹The Bernard and Irene Schwartz Center for Biomedical Imaging and Center for Advanced Imaging Innovation and Research (CAI²R), Department of Radiology, New York University Grossman School of Medicine, New York, New York, USA

²Universita di Napoli Federico II, Napoli, Italy

³Siemens Medical Solutions, New York, New York, USA

Correspondence

Ilias I. Giannakopoulos, Center for Biomedical Imaging, Department of Radiology, New York University Grossman School of Medicine, New York, NY 10016, USA.

Email:

ilias.giannakopoulos@nyulangone.org

Funding information

National Institutes of Health, Grant/Award Numbers: P41 EB017183, R01 EB024536

Abstract

Purpose: We developed a 3D vision transformer-based neural network to reconstruct electrical properties (EP) from magnetic resonance measurements.

Theory and Methods: Our network uses the magnitude of the transmit magnetic field of a birdcage coil, the associated transceive phase, and a Canny edge mask that identifies the object boundaries as inputs to compute the EP maps. We trained our network on a dataset of 10 000 synthetic tissue-mimicking phantoms and fine-tuned it on a dataset of 11 000 realistic head models. We assessed performance in-distribution simulated data and out-of-distribution head models, with and without synthetic lesions. We further evaluated our network in experiments for an inhomogeneous phantom and a volunteer.

Results: The conductivity and permittivity maps had an average peak normalized absolute error (PNAE) of 1.3% and 1.7% for the synthetic phantoms, respectively. For the realistic heads, the average PNAE for the conductivity and permittivity was 1.8% and 2.7%, respectively. The location of synthetic lesions was accurately identified, with reconstructed conductivity and permittivity values within 15% and 25% of the ground-truth, respectively. The conductivity and permittivity for the phantom experiment yielded 2.7% and 2.1% average PNAEs with respect to probe-measured values, respectively. The in vivo EP reconstruction truthfully preserved the subject's anatomy with average values over the entire head similar to the expected literature values.

Conclusion: We introduced a new learning-based approach for reconstructing EP from MR measurements obtained with a birdcage coil, marking an important step towards the development of clinically-usable in vivo EP reconstruction protocols.

KEYWORDS

electrical properties mapping, electromagnetic simulations, transfer learning, vision transformers

1 | INTRODUCTION

Electrical properties (EP) govern the interactions between electromagnetic (EM) fields and biological tissues.^{1,2} Relative permittivity ϵ_r reflects how a tissue stores electrical energy in an electric field. Electric conductivity σ_e measures the ability of the tissue to conduct electrical current. The detailed interactions between these properties and external EM field sources can be modeled with Maxwell's or Helmholtz's equations.³ Combined with 3D full-wave EM simulations, access of the EP distribution in tissue^{4,5} could enable precise estimation of the local specific absorption rate, which is especially useful at ultra-high field MRI.⁶ Knowledge of tissue EP could also be used to improve radiofrequency (RF) or thermal-based treatments.⁷⁻⁹ Finally, EP could be utilized as biomarkers for diagnosis and treatment monitoring.¹⁰⁻¹³

Traditional MR-based EP reconstruction methods solve the differential form of Helmholtz equation using a few simplifying assumptions¹⁴ for the transmit magnetic field (\mathbf{B}_1), such as a negligible \mathbf{B}_{1z} component. With this assumption, one can solve the Helmholtz equation using only the right circularly polarized component of the transmit field $\mathbf{B}_1^{(+)} (= \mathbf{B}_{1x} + i\mathbf{B}_{1y})$.¹⁵ The advantage is that the absolute value of $\mathbf{B}_1^{(+)}$ can be measured in MRI,¹⁶ while its phase can be approximated as half of the measured transceive phase¹⁷ when transmission and reception are performed in quadrature and anti-quadrature mode with a birdcage coil.¹⁸ These assumptions hold mainly for the central region of a homogeneous sample and at lower MRI field strengths (≤ 1.5 T),¹⁹ but have been used to map EP at higher field strengths, yielding sub-optimal EP reconstructions.²⁰ Instead, MR-based EP reconstruction methods that rely on the iterative inversion of the integral form of Maxwell's equations²¹⁻²³ do not rely on the above assumptions and can lead to more accurate EP reconstructions. However, integral methods are slow and can require days to converge even for the coarse voxel resolutions. In addition, the performance of these methods depends on how closely the experimental coil can be modeled in simulation.

Learning-based methods are emerging as a new trend in the EP reconstruction community. The method detailed in Reference 24 introduced neural networks into the EP reconstruction field. It leverages a conditional generative adversarial network,²⁵ using the $\mathbf{B}_1^{(+)}$ map, the transceive phase, and a pseudo-Spin Echo MRI image (to provide tissue contrast information) as the input to compute EP maps. The model was trained using slices from 20 numerical realistic head models. The studies in References 26,27 introduced machine learning EP reconstruction methods that utilizes the correlations between EP, water content, and T1 relaxation times. In Reference 28, T2-weighted

spin echo images were employed alongside transceive phase maps as inputs to a neural network (NN) to reconstruct σ_e . Another study performed σ_e reconstruction for the pelvic region using a convolutional neural network (CNN) and roughly 200 models.²⁹ While these methods pioneered the use of NN for MR-based EP reconstruction, they all suffer from limited performance when applied to anatomies that differ from those in the training datasets. The work in Reference 30 highlights several underlying issues responsible for this limitation. Firstly, there is limited availability of diverse and realistic tissue models, resulting in NNs being trained on relatively small datasets of a few hundred volumes, which are insufficient to fully capture the anatomical variability of actual tissue structures. Secondly, these NNs are often two-dimensional to prevent memory overflows, making it difficult to preserve the inherently three-dimensional nature of Maxwell's equations. Thirdly, incorporating MR images alongside $\mathbf{B}_1^{(+)}$ maps as additional NN inputs to provide anatomical context can lead to the networks learning correlations between relaxation times and EP, potentially overshadowing the direct influence of EP on the distribution of the measured magnetic fields.

In this work, we simulated the EM fields of roughly 10 000 synthetic tissue-mimicking phantoms and roughly 11 000 realistic numerical head models to train a novel NN for EP reconstruction. We employed a birdcage coil to simulate the synthetic training set, which means that the network is optimized for data generated using this type of coil and would need to be retrained if the geometry of the coil used in experiments were considerably different. To incorporate anatomical information as input to the network, instead of including MR images, we generated a binary mask of tissue boundaries using Canny edge detection.³¹ A preliminary version of this work was presented at the 2024 meeting of the International Society for Magnetic Resonance in Medicine.³² The PyTorch code for our model and the generation of the Canny edge masks is available at <https://github.com/GiannakopoulosIlias/vision-transformer-network-for-mr-electrical-properties-tomography>.

2 | METHODS

2.1 | Dataset

2.1.1 | Simulated data

We generated 10 160 tissue-mimicking phantoms, discretized with a voxel resolution of 5 mm³. The phantoms had either an ellipsoidal (8160) or a cylindrical geometry (2000). The ellipsoids had random lengths for the

principal semi-axes, which were constrained to be at least 7, 7, and 8.5 cm and at most 9.5, 12, and 11.5 cm in the x , y , and z directions, respectively. Inside every ellipsoid, we randomly placed either 0, 1, or 23 smaller ellipsoids, varying in both size and position. To introduce inhomogeneity, each ellipsoid was randomly assigned tissue-mimicking EP, ranging from 11 to 120 for the relative permittivity and 0.07 to 2.5 S/m for the electric conductivity. The cylindrical phantoms had random length (between 17 and 23 cm) and radius (between 7 and 9.5 cm). They were either homogeneous or inhomogeneous with an additional cylindrical compartment, using random EP values in the same range as for the ellipsoids. The inner cylindrical compartment had the same length as the phantom, whereas its position and radius varied. All models were enclosed in the same cuboid domain of dimensions $19 \times 23.5 \times 23 \text{ cm}^3$ that corresponded to $38 \times 47 \times 46$ voxels. This dataset will be referred to as *synthetic phantoms*.

To generate a dataset of realistic heads, we used six numerical human head models: Duke and Ella (Foundation for Research on Information Technologies in Society, Switzerland,³³), Naomi and Norman (National Radiological Protection Board, UK,³⁴), and Male and Female (Remcom, The Pennsylvania State University,³⁵). For data augmentation, all models were stretched in x , y , and z directions as in Reference 36. As a result, we obtained 216 unique variations for each model, except for Norman which had 120 variations. To further enlarge our dataset and introduce additional diversity in the anatomies, we performed structural deformations to all these variations, except for the 216 derived from Duke. In particular, we shifted all homogeneous regions within the head (e.g., cerebrospinal fluid, cerebellum, etc.) by one, two, or three voxels in the x , y , and z directions (the amount of shift in each direction was chosen randomly). The gaps resulting from the shifting were filled with a randomly selected EP value from the original models. As a result, each homogeneous region was reassembled into a new inhomogeneous region that included the original EP values and a layer with a different EP value. When the randomly shifted areas overlapped with one another, we randomly included only one of them. We performed 10 structural deformations for each head. In addition, we included synthetic uniform lesions in 240 of the deformed heads, using literature values for their EP.³⁷ The total number of variations for Ella, Naomi, Norman, Male, and Female, including stretching and structural deformations, was 11,040. We also generated two variations of the Duke head with synthetic lesions but without structural deformation. All models were enclosed within the same domain as the synthetic phantoms and were discretized using the same voxel resolution (5 mm^3). Head models larger than the computational domain were cropped to fit into it. The final dataset

had 11 256 head models and will be referred to as *realistic heads*.

We modeled a 3 tesla (T) high-pass birdcage coil based on an in-house coil.³⁸ The coil had a radius of 12.3 cm, 8 legs, a length of 22 cm, and a copper width of 1 cm. It was discretized using 2990 triangular elements. The capacitor values were set to 11.634 pF, following.³⁹ We simulated the $B_1^{(+)}$ in quadrature mode and the $B_1^{(-)}$ (receive sensitivity) in anti-quadrature mode for all phantoms and heads using our in-house fast full-wave EM simulator.⁵ In particular, we utilized first-order, piecewise polynomials⁴⁰ to approximate the fields and improve the accuracy of the simulation. The simulation for each case required at most 40 seconds using an NVIDIA A100 PCIe GPU with 40 GB of memory in MATLAB 9.10. All data for this work were generated in less than 10 days. The complex $B_1^{(+)}$ and $B_1^{(-)}$ were corrupted with independent and identically distributed Gaussian noise. The peak SNR was chosen randomly and ranged between 50 and 200. We computed the transceive phase as the sum of the phases of $B_1^{(+)}$ and $B_1^{(-)}$, divided by 2. For each case, we computed edge masks outlining the boundaries between different compartments by applying a Canny edge detection filter³¹ on every axial plane of the ground-truth 3D conductivity map. We assembled all edge masks to a 3D volumetric mask. In our implementation of the Canny algorithm, we set the low threshold to 0.01, the high threshold to 0.05, the kernel size to 1 (due to the coarse resolution), and the σ (standard deviation of the Gaussian kernel) to 1.5. These parameters were found to be the most effective in preserving relevant edges for our datasets. The algorithm proceeds through the standard Canny steps: applying Gaussian smoothing, computing gradient magnitudes and directions using Sobel filters, performing non-maximum suppression to thin edges, and finally, applying double thresholding followed by edge tracking via hysteresis to produce the edge map.

2.1.2 | Phantom experiment

We collected experimental data using a 3T scanner (MAGNETOM Biograph mMR, Siemens Healthineers, Erlangen, Germany) and an in-house head birdcage coil loaded with an inhomogeneous cylindrical phantom with two compartments. The geometry of the coil was similar to the design used in the simulation. The phantom was 21.5 cm long and had a radius of 6.875 cm. The inner compartment was also a cylinder of the same length with a radius 3.1 cm and it was displaced from the center by roughly 0.5 cm. The EP of the outer and inner compartments were similar to those of white matter (WM) and cerebral spinal fluid (CSF), respectively. Both mixtures were constructed using the method in Reference 41 and the EP values

(WM $\epsilon_r = 64$, WM $\sigma_e = 0.67$ S/m, CSF $\epsilon_r = 72$, CSF $\sigma_e = 1.97$ S/m) were confirmed with a dielectric probe (Agilent Technologies, Santa Clara, CA). We used the double-angle method (DAM)⁴² to generate maps of the $|B_1^{(+)}|$. In particular, we performed multi-slice scanning at $2 \times 2 \times 2$ mm³ voxel resolution and resized to 5 mm voxel isotropic resolution to match the computational domain of the simulated data. We applied the Canny filter to one of the two gradient echo MR images used in DAM to generate the edge mask. We computed the transceive phase map using the RF_map pulse sequence,⁴³ which is based on a spin echo acquisition, yielding better performance than gradient echo for the transceive phase reconstructions.⁴⁴ For the RF map acquisition, we scanned at 1 mm² pixel resolution. Each slice was 5 mm thick, with 5 mm gaps between slices to avoid cross-talk. We interpolated between the missing slices. Finally, we resized to 5 mm voxel isotropic resolution.

2.1.3 | In vivo experiment

We imaged a human volunteer without known brain pathologies using the same birdcage coil and scanner. The study was conducted following institutional and national ethical standards, and informed consent was obtained from the participant before imaging. We acquired two sets of $|B_1^{(+)}|$ maps using the RF_map sequence⁴³ scanning at 1 mm² pixel resolution. Each slice was 5 mm thick, with 5 mm gaps between slices to avoid cross-talk. We interpolated between the missing slices. The scan time was 10 min for each sequence. We also generated the corresponding transceive phases to assess repeatability. We also acquired a 3D MPRAGE⁴⁵ dataset ($1 \times 1 \times 5$ mm³ voxel resolution, and 15 min scan time), which provided brain images with sharper internal edges than the RF_map sequence. The MPRAGE images were regridded to match the images of the RF_map acquisitions⁴⁶ to ensure that they aligned with the $|B_1^{(+)}|$ map and the transceive phase. We interpolated the data to 5 mm³ isotropic voxel resolution. We created three distinct input configurations for the network: Configuration C1 used the $|B_1^{(+)}|$ map and the transceive phase from the first RF_map acquisition and the Canny edge mask generated from images of the same RF_map acquisition; configuration C2 used the $|B_1^{(+)}|$ map and the transceive phase from the first RF_map acquisition and the Canny edge mask generated from the MPRAGE images registered to the images of the first RF_map acquisition; and configuration C3 used the $|B_1^{(+)}|$ map and the transceive phase from the second RF_map acquisition and the Canny edge mask generated from the MPRAGE images registered to the images of the second RF_map acquisition.

2.2 | Network architecture

The $|B_1^{(+)}|$, the transceive phase, and the edge mask were used as the input to the neural network. The output of the network was the predicted volumetric maps of the relative permittivity and conductivity. The $|B_1^{(+)}|$ and transceive phase maps were normalized by their respective maximum value. The EP were normalized by dividing them by 135 (for the relative permittivity) and 2.8 S/m (conductivity), respectively, so they would range between 0 and 1. Our network's architecture (Figure 1) consisted of three 3D TransUNets⁴⁷ connected sequentially through 64 channels. TransUNets rely on convolutions for the analysis and synthesis steps, like UNets, and also employ a vision transformer to connect them. We used two pooling layers for the 3D TransUNets. The convolutions in the convolutional blocks of the network were followed by instance normalization (IN), Gaussian error linear unit (GELU) activation functions,⁴⁸ and dropout.⁴⁹ The dropout probability was set to 0.1 through the network. We implemented a residual connection in all convolutional blocks. We incorporated 6 vision transformer layers, each equipped with 16 heads per transformer. We also included a Feature Wise Linear Modulation (FiLM) layer⁵⁰ after the first convolution block of each TransUNet. The parameters of the FiLM's affine transformation (β and γ) were learned using a FiLM generator network. The FiLM generator network consisted of a set of four convolutional blocks, followed by average pooling to retrieve $\beta \in \mathbb{R}^{64 \times 3}$ and $\gamma \in \mathbb{R}^{64 \times 3}$ (3 is used to separate the parameters for each cascade). The architecture of each block was a convolution followed by IN, ReLU with 0.2 slope coefficient,⁵¹ and dropout. The input to the FiLM generator was the Canny edge mask. The network was trained end-to-end.

2.2.1 | Model ablations

We performed an ablation study to evaluate the performance of our proposed model. In particular, we implemented a model with a single 3D UNet, one with a single 3D TransUNet, and one with three 3D TransUNets connected sequentially. For all three architectures, we used IN instead of FiLM. The hyperparameters and the training and testing conditions were the same for all models. We compared the results with those of our proposed architecture.

2.3 | Training and validation

We split the synthetic phantoms' dataset into 8065 (training), 1463 (validation), and 632 (testing). Figure 2 (top)

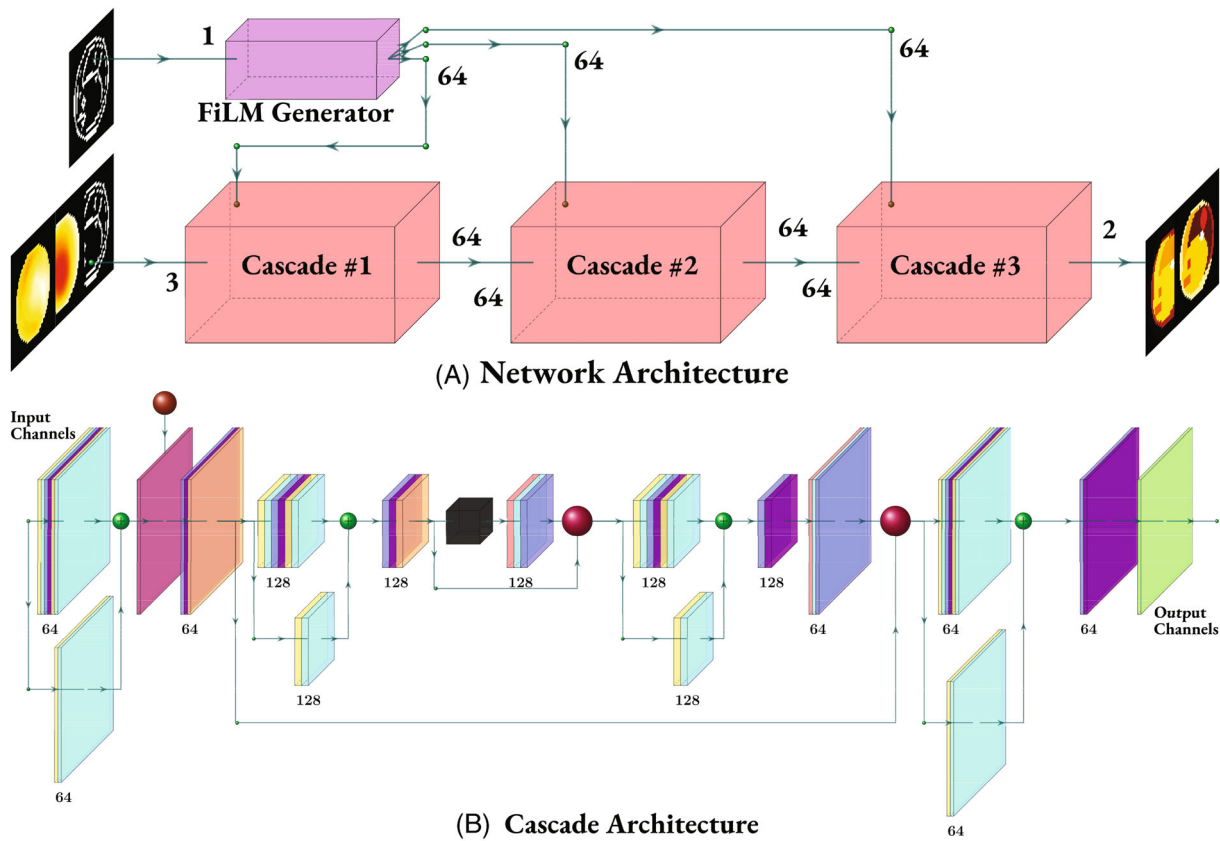


FIGURE 1 (A) Overall network architecture used for the synthetic phantoms. The inputs are the magnitude of the 3D $B_1^{(+)}$, the 3D transceive phase map, and the 3D edge mask (only middle axial slice is shown). The 2 outputs are the corresponding maps of relative permittivity and conductivity. The three inputs are passed to Cascade #1. The edge mask is also passed to the FiLM Generator network. Each Cascade is a TransUNet. The numbers show the input and output channels of each sub-network. (B) Detailed architecture of each cascade. Yellow and green layers represent convolutions with kernel size $3 \times 3 \times 3$ and $1 \times 1 \times 1$, respectively. Red and orange layers represent transpose convolutions and average pooling with kernel size $2 \times 2 \times 2$. Cyan, magenta, and blue layers represent instance normalization, dropout, and GELU activation functions, respectively. The pink layer represents the affine transformation based on the FiLM β and γ parameters. The green and red circles represent residual connections and concatenations, respectively. The transformer, shown in black, is equipped with 16 attention heads.

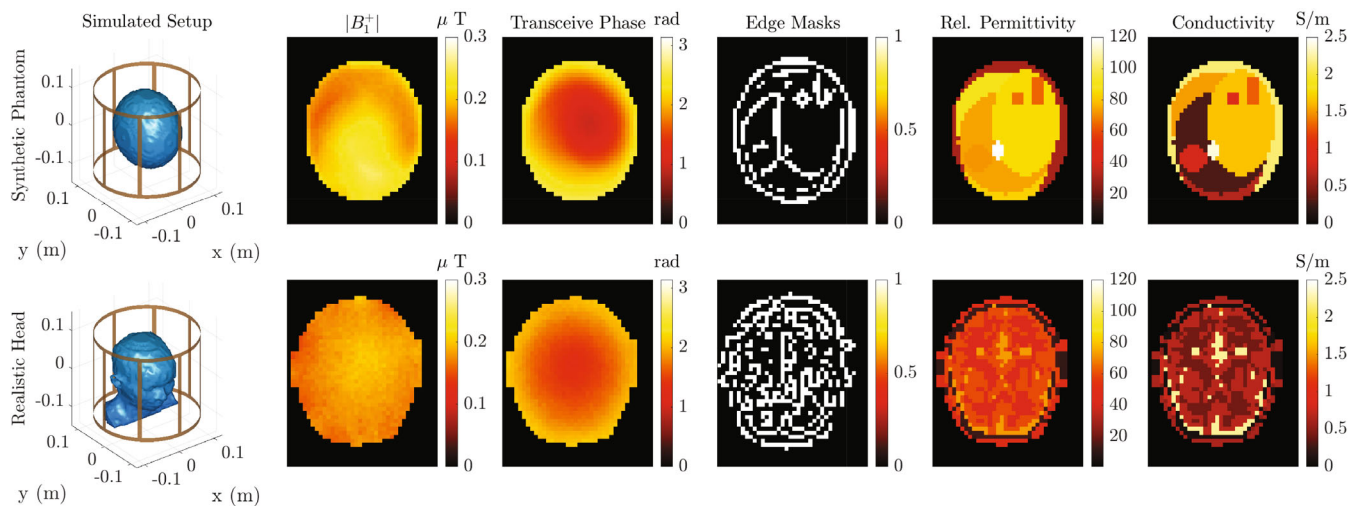


FIGURE 2 Simulations and ground truth EP for two representative cases from the synthetic phantoms (top) and realistic heads (bottom) training datasets. The magnitude of the coil's transmit field ($|B_1^{(+)}$), the transceive phase, and the edge mask were used as inputs to the neural network. Ground truth relative permittivity and electric conductivity were used as outputs. Data are shown for a representative axial slice.

shows data from a representative synthetic phantom. The network was trained for 100 epochs using the adamW optimizer⁵² and a cyclical learning rate (lr) policy⁵³ with an initial lr of 0.01. We then fine-tuned the network on the realistic heads dataset. In particular, we used the final weights of the first training as initial weights to retrain the network and included an additional TransUNet (two input and two output channels) after Cascade #3 (Figure 1) to further boost performance. We did not include a FiLM layer for the additional TransUNet. We did not freeze any weights when fine-tuning. Ella, Naomi, Norman, Male, and Female and their variations were used for fine-tuning (11 040 data), while the Duke models were split into 108 cases for validation, and 108 cases (plus the 2 cases with the lesions) for testing, avoiding data leakage (Figure 2). For fine-tuning, we trained for 100 epochs and reduced the initial learning rate to 0.003.⁵⁴ We used a batch size of 4 for training and 1 for validation. The network was written in PyTorch Lightning and was trained using a single NVIDIA V100 PCIe GPU with 32 GB of memory. The initial training required roughly 5 days, and the fine-tuning roughly 6 days. The loss function was the equally weighted sum of the mean squared error (MSE) value for the σ_e and the MSE value for the ϵ_r . We also recorded the structural similarity index measure (SSIM)⁵⁵ scores for both EP during validation. The three ablated models were trained using the synthetic phantoms dataset.

2.4 | Evaluation

For evaluation, we selected the network state that exhibited the lowest validation MSE to mitigate the risk of assessing performance with a network potentially overfitting the training dataset. We computed the SSIM and the peak normalized absolute error (PNAE) for the test dataset. The PNAE was defined as $\overline{|y - x| / \max\{x\}}$, where x is the ground-truth, and y is the prediction, and $\overline{\cdot}$ indicates averaging. For the phantom experiment, we

treated the probe-measured EP values as ground-truth. For the in vivo experiment, since there was no available ground-truth, we qualitatively compared the reconstructed conductivity against the one obtained with H-EPT (Helmholtz-EPT)⁵⁶ available from the open-source EPTlib.⁵⁷ We used Savitzky–Golay filters⁵⁸ with a kernel size of 1, 2, and 3 and ellipsoid shape to compute numerical derivatives.

3 | RESULTS

3.1 | Network validation and testing

Figure 3 presents the MSE and SSIM throughout the validation for both the network trained on the synthetic phantoms and the network fine-tuned on the realistic heads. The observed MSE oscillations with the synthetic phantom dataset during the early training phase were expected due to high lr . During fine-tuning, the lr was smaller, therefore the oscillations were reduced. Since only the MSE was used for training, the oscillations in the SSIM reflect the fundamental differences between SSIM and MSE.⁵⁹ The larger SSIM oscillations during fine-tuning as opposed to the original training were expected since the Duke's variations used for validation were out-of-distribution (OOD) examples.

The mean and standard deviation (SD) of the PNAE (SSIM) for the ϵ_r and σ_e of the synthetic phantoms dataset were $1.7\% \pm 1.7\%$ (0.94 ± 0.05) and $1.3\% \pm 1.6\%$ (0.93 ± 0.09), respectively. For the realistic heads, the PNAE and SSIM for σ_e were $1.8\% \pm 0.4\%$ and 0.93 ± 0.01 , respectively. The corresponding values for ϵ_r were $2.7\% \pm 0.5\%$ and 0.86 ± 0.03 , respectively. Figure 4 shows histograms for the PNAE and SSIM value distributions for all volumes in the test datasets.

Figure 5 presents the reconstructed EP and associated PNAE for the middle axial slice of two representative cases from the synthetic phantoms test dataset.

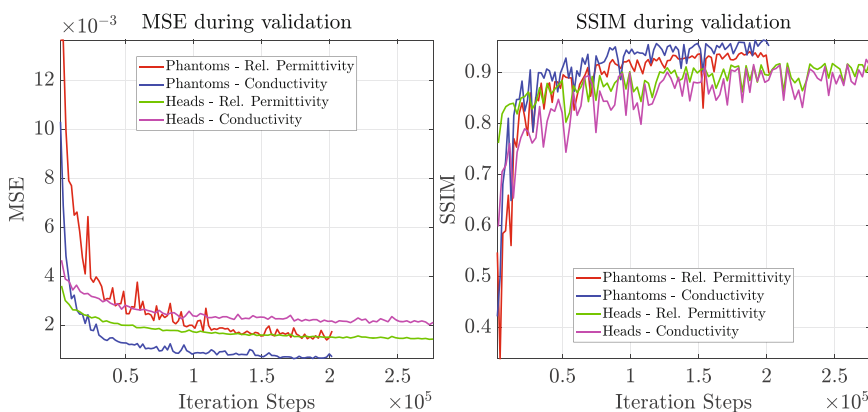


FIGURE 3 Average MSE (left) and SSIM (right) validation error for both the network trained on synthetic phantoms and the network fine-tuned on realistic heads. The loss function used for training was the equally weighted sum of the MSE value for the conductivity and the MSE value for the relative permittivity.

FIGURE 4 Histograms of the PNAE (left) and SSIM (right) for (top) the network trained and tested on synthetic phantoms and (bottom) the network fine-tuned on realistic heads and tested only on Duke's variations. The results account for all volumes in the test datasets.

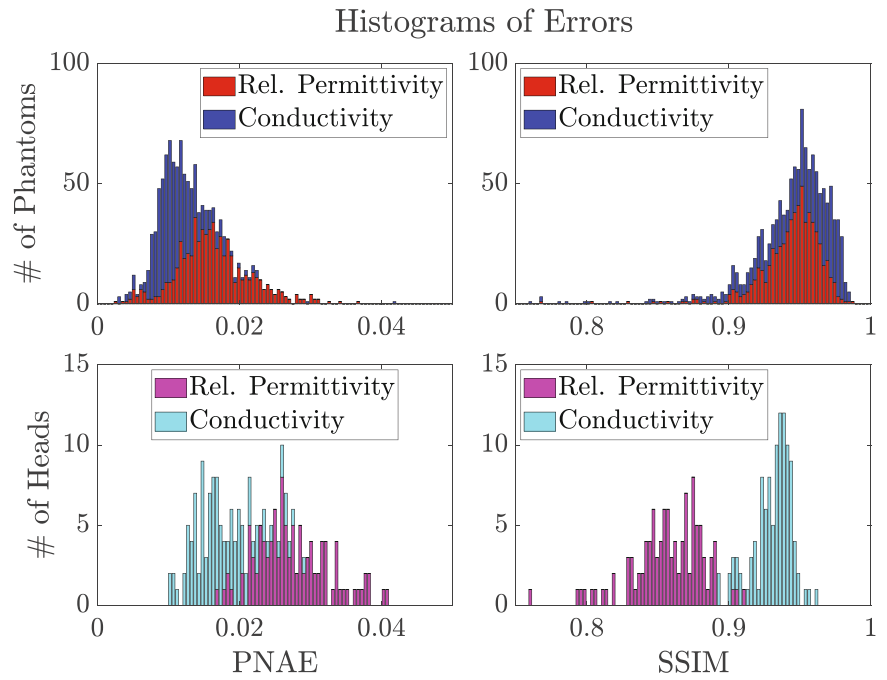
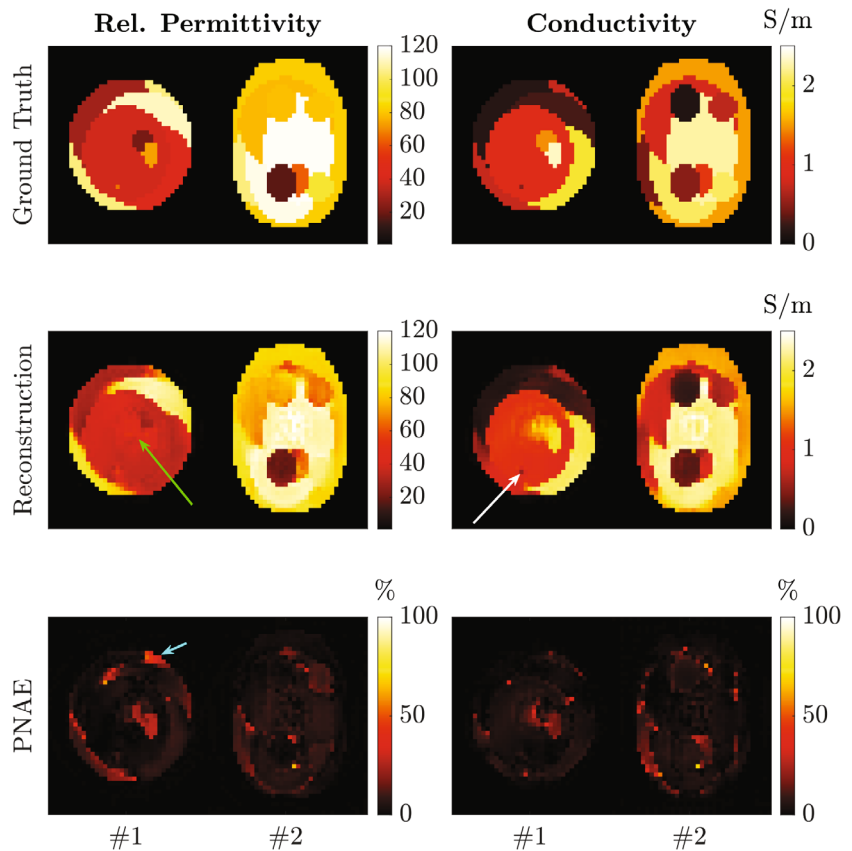


FIGURE 5 (Left) Relative permittivity and (right) conductivity maps for the middle axial slice of two representative examples from the synthetic phantoms test dataset. PNAE maps of the reconstructed EP (middle) with respect to the ground-truth (top) are plotted at the bottom. Single voxel regions are preserved in the reconstructions (white arrow). The PNAE increases near boundaries for ϵ_r (blue arrow). A checkerboard-like artifact is visible in the middle of the phantoms (green arrow).



The reconstructed EP have excellent qualitative similarity with the ground truth values. Blurring can be observed across certain boundaries, and a checkerboard-like artifact appeared at the center of the field of view for some reconstructions (evident in the first example).

Figure 6 shows the results of the fine-tuned network for two representative examples from the realistic heads test dataset. While the internal structure of the brain is retained in the EP maps, a degree of blurring can be observed between WM and GM. Higher errors were found




FIGURE 6 (Left) Relative permittivity and (right) conductivity maps for an axial slice through the brain of two representative examples from the realistic heads test dataset. The PNAE maps of the reconstructed EP (middle) with respect to the ground-truth (top) are presented at the bottom. The PNAE is higher near boundaries, but the internal structures of the head are preserved in the reconstructions. The white arrow indicates a region where blurring between WM and GM occurs in the reconstructions. A region of high PNAE between the brain and the skull is indicated with the blue arrow.

in regions between the brain and skull, particularly where the EP values are small.

Figure 7 shows the reconstructed EP maps for the two variations of Duke with the synthetic lesions. While the network successfully detected the lesions, it underestimated their EP. Specifically, ϵ_r in the vicinity of the small lesion was 62 ± 5 , compared to the ground truth value of 68, while the reconstructed σ_e was 0.94 ± 0.42 S/m, compared to the ground truth value of 1.1 S/m. Note that the large SD for σ_e is attributable to errors at the boundary of the lesion, where the network incorrectly interpreted some voxels as part of the adjacent CSF tissue. For the large lesion, ϵ_r was 60 ± 2 and σ_e was 0.72 ± 0.09 S/m, compared with ground truth values of 80 and 0.8 S/m, respectively.

3.2 | Phantom experiment

Figure 8 shows experimental results for the two-compartment cylindrical phantom. The reconstructed EP maps preserve the expected homogeneity across the two compartments. The conductivity was 1.58 ± 0.11 S/m and 0.55 ± 0.06 S/m for the CSF- and WM-mimicking compartments, respectively. The relative

permittivity was 74 ± 5.7 and 60 ± 1.3 for the two compartments. The relative difference with respect to the probe-measured values was 20% and 18.3% for the average conductivity in the CSF and WM compartments, respectively. For the permittivity, the differences were 3% for the CSF compartment and 5.7% for the WM compartment.

3.3 | In vivo experiment

Figure 9 presents a qualitative comparison between the EP reconstructed using the C1, C2, and C3 input configurations for representative sagittal, coronal, and axial slices through the brain of the volunteer. The average EP values and the corresponding SD in the entire head for the C1, C2, and C3 configurations were 0.62 ± 0.37 S/m (C1), 0.62 ± 0.36 S/m (C2), and 0.61 ± 0.47 S/m (C3) for σ_e , and 48 ± 13 (C1), 48 ± 13 (C2), and 47.39 ± 13 (C3) for ϵ_r . For comparison, the corresponding values averaged across all Duke variations in the test dataset were 0.61 ± 0.47 S/m for σ_e and 45 ± 19 for ϵ_r .

Figure 10 shows conductivity maps reconstructed from the same in vivo data using H-EPT at 5 mm voxel isotropic resolution. We tested three different Savitzky–Golay (SG) kernel sizes for each method.

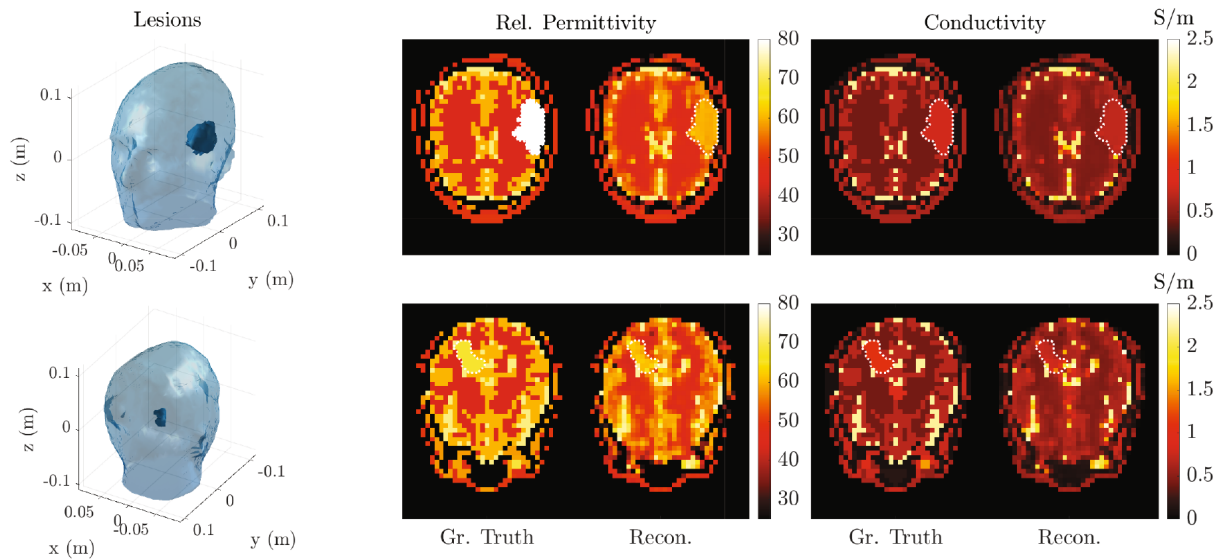


FIGURE 7 (Left) Duke's head with a large (top) and small (bottom) homogeneous synthetic lesion inserted in the brain; associated relative permittivity (middle) and conductivity (right) maps for a representative axial cut through the lesions. The reconstructed relative permittivity of the smaller lesion had a 9% average relative error with respect to ground-truth (within the lesion), whereas the reconstruction error for the larger lesion was 25%. For the conductivity, the error was 15% and 10% for the small and large lesions, respectively. The lesions are contoured with dotted white lines for enhanced visualization.




FIGURE 8 Experimental phantom results. The three inputs of the neural network (normalized $|B_1^+|$ map in dB, MR image obtained with the DAM, and corresponding Canny edge mask) are shown for a representative axial cut (left). Probe measured ground-truth, reconstructed EP, and PNAE are shown for relative permittivity (middle) and conductivity (right). The reconstruction preserves the homogeneity of both compartments. The relative permittivity for both compartments was close to the ground-truth. The average conductivity was underestimated by up to 25% in comparison to the probe-measured value.

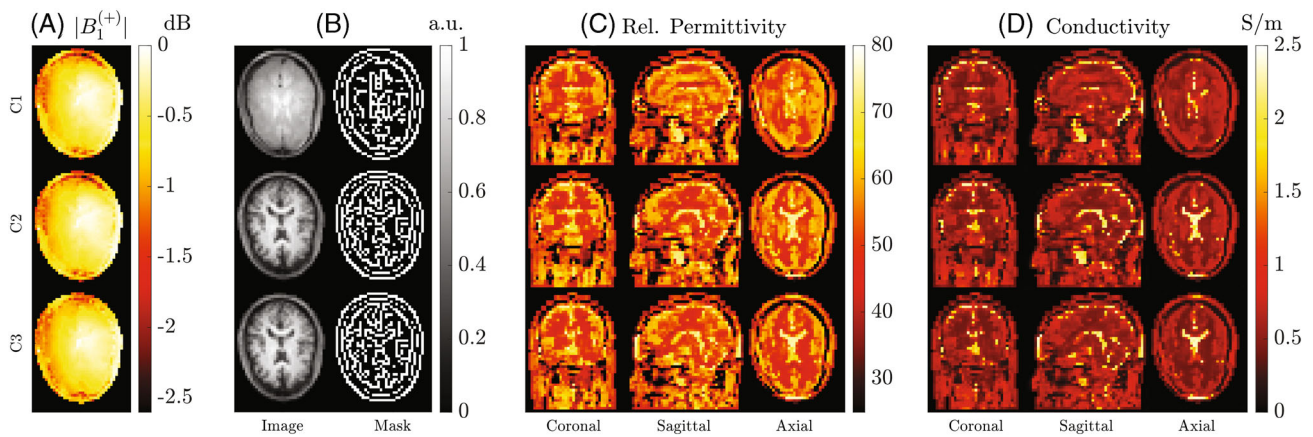


FIGURE 9 In vivo results. (A) Normalized $|B_1^{(+)}|$ maps in dB, (B) MR images, Canny edge masks, (C) reconstructed relative permittivity, and (D) conductivity for each input configuration. EP maps are shown for representative coronal, sagittal, and axial planes.

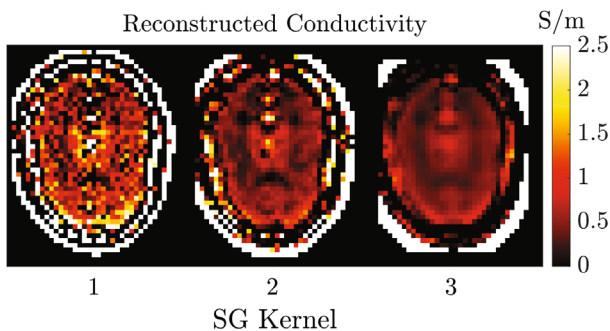


FIGURE 10 Conductivity reconstructions using the H-EPT for 5 mm voxel isotropic resolution for the in vivo experiment. The results are presented for the same axial plane as in Figure 9 and for three different Savitzky-Golay (SG) kernel sizes (1, 2, and 3).

3.4 | Model ablations

For the synthetic phantoms test dataset, the single UNet with IN achieved 0.849 ± 0.093 SSIM for ϵ_r and 0.901 ± 0.102 for σ_e . The single TransUNet using IN improved the SSIM to 0.906 ± 0.057 for ϵ_r and 0.915 ± 0.106 for σ_e . The model using a cascade of three TransUNets with IN achieved 0.918 ± 0.090 SSIM for ϵ_r and 0.945 ± 0.090 for σ_e . Our proposed model that incorporates FiLM in the architecture with three TransUNets yielded 0.939 ± 0.046 SSIM for ϵ_r and 0.930 ± 0.091 for σ_e .

For the phantom experiment, the architecture with a cascade of three TransUNets with IN achieved SSIM equal to 0.508 for ϵ_r and 0.913 for σ_e , whereas the same architecture using FiLM instead of IN achieved 0.77 for ϵ_r and 0.965 for σ_e .

4 | DISCUSSION

We introduced a novel NN architecture to compute EP maps from MR measurements and evaluated its

performance in simulations and experiments. The integration of edge masks generated using Canny filtering is a robust way to provide anatomical information without including MR images directly as input to the network. This training strategy prevents the network from learning nonphysical correlations between particular MR image contrasts and EP maps,³⁰ therefore, ensuring that the network learns the EP distribution only from magnetic field measurements, to which they are related by Maxwell's equations. Further, the inclusion of Canny edge masks as additional inputs enhances boundary preservation between tissues as shown in Reference 32. For simulated data, we derived edge masks from the available ground truth conductivity maps. For phantom and in vivo experiments, the masks were produced by applying the Canny filter to actual MR images, without any manual processing.

The main component of our architecture is the TransUNet.⁴⁷ TransUNet combines the strengths of vision transformers⁶⁰ and the expressive power of CNNs.⁶¹ This hybrid model enables capturing both global (vision transformer) and local (convolutions) information from the input tensors. In particular, vision transformers are capable of learning the global interactions between EP and the input magnetic field maps, resulting in improved reconstruction performance over purely CNN-based models that rely solely on local interactions.⁶² Our findings are consistent with findings in MRI reconstruction, where transformers and attention-based techniques have demonstrated notable performance improvements.^{63–66} The CNN components of TransUNets can learn the local boundaries between tissues through the utilization of the Canny edge masks, further boosting the model's capability to accurately reconstruct EP maps. In particular, the use of a TransUNet instead of a UNet improved the SSIM scores for the case of the synthetic phantoms dataset by 0.057 for ϵ_r and by 0.014 for σ_e . The use of a cascaded architecture with

three TransUNet instead of a single one was implemented to increase the overall learning capacity of the model following the work in Reference 67. In fact, for the model with three TransUNets, the SSIM for the synthetic phantoms test improved by 0.012 for ϵ_r and by 0.03 for σ_e . FiLM, as an adaptive normalization technique, modulates feature maps across network layers, allowing the network to more effectively control and guide information.⁵⁰ This adaptive capability is particularly beneficial in handling variations and unexpected inputs, which are common in experimental measurements (OOD data). When we replaced the IN with FiLM in our ablation study, the SSIM for the phantom experiment improved by 0.262 for ϵ_r and 0.059 for σ_e . We used the Canny edge masks as the input to the FiLM generator to better preserve the anatomical information during training; although different inputs could also be used, such as a segmentation mask. Other normalization techniques such as Adaptive Instance Normalization,⁶⁸ might help improve our network's performance and they will be explored in future work.

NN architectures that do not integrate prior knowledge of the physical laws are prone to overfitting, especially when trained on small datasets. To address this, we trained our network using two large datasets (8000 synthetic phantoms and 11 000 head models) generated using random augmentations to ensure robust generalization during inference with previously unseen data. Our network was first trained on the synthetic phantoms and then fine-tuned on the realistic heads. We followed this strategy because the synthetic phantoms had a simpler interior structure, which allowed the network to initially learn in a more controlled and less complex environment. Fine-tuning enabled the network to adapt its learned patterns to the more complex internal structures of the realistic heads.⁶⁹ The network provided excellent EP reconstructions for both simulated phantoms (Figure 5) and OOD numerical head models (Figure 6) that were not included in the training dataset. Moreover, the network's performance remained consistent for all simulated test data, as shown by the histograms of Figure 4. Specifically, the SD of the PNAE and SSIM was less than 2% and 0.09, respectively, for all cases. A small degree of blurring occurred between the boundaries of WM and GM (Figure 6). This effect was anticipated, since these fine anatomical details minimally affect the $B_1^{(+)}$ and $B_1^{(-)}$ field distributions and can be challenging to detect with the Canny edge masks at a voxel resolution of 5 mm³. The reconstruction quality was worse for the permittivity compared to the conductivity for the realistic heads dataset, which was reported also in previous work.¹⁵ The observed checkerboard-like artifact in the central region of the sample for some cases (visible, e.g., in Figure 5) can be attributed to the use of transpose convolutions in the TransUNets.⁷⁰ Future work

will focus on different techniques⁷¹ to avoid the emergence of this artifact.

The network was successful in locating synthetic lesions in OOD cases. Although the shape of the pathology was accurately reconstructed (Figure 7), the EP of the lesions were underestimated. Since the network was predominantly fine-tuned on healthy tissue data (only ~2% of the cases in the realistic heads dataset had pathologies), it tends to reconstruct values typical of healthy brain tissues, even in areas where pathologies are present. This can explain the observed underestimation of the EP values for the lesions. Future work will focus on overcoming this issue by including a larger number of head models with realistic lesions as in Reference 72.

Our approach yielded greater EP uniformity within phantom compartments (Figure 8) compared to previous methods,⁷³ where EP reconstructions were often corrupted by noise amplifications associated with numerical derivatives. The observed error in the permittivity was small and within the accuracy and repeatability of dielectric probe measurements. The observed error was larger for conductivity. This will be addressed in future work by better modeling the experimental variability in the simulated training set.

The in vivo EP maps did not show artifacts that are typically observed with other approaches.¹⁵ In Figure 9C,D, both input configurations C2 and C3 yielded EP reconstructions that preserved the anatomy of the brain, while the reconstructions based on C1 did not preserve the structure of the choroid plexus, highlighting the importance of good edge masks for network performance. On the other hand, the different $B_1^{(+)}$ maps (Figure 9A) used for C2 and C3 had a negligible effect on the quality of the EP reconstruction. We assessed the repeatability of our network by comparing the average in vivo EP for the three input configurations. Both the mean and the SD values over the entire head were in the same range for all three configurations and were similar to the values for Duke's head, which, although cannot be assumed as the ground truth, are based on established literature values. The SD of the estimated EP values were lower compared to the SD of the EP for the Duke model. These differences could be attributed to the interpolation into a coarse resolution of the in vivo data, which could have resulted in certain regions being removed due to partial volume effects.⁷⁴ This could be the case of the CSF, which can be as small as 1.2 mm in specific regions in the head while having an impact on the overall SD due to its high EP values. We anticipate better matching between in vivo and simulated data by training our network with data of finer spatial resolutions, for example at 2 mm instead of 5 mm voxel isotropic resolution. Although generating such a dataset would require approximately two months of simulation time, it would

allow (a) higher accuracy of the simulated magnetic fields, (b) capacity for deeper TransUNets with 4 pooling layers per cascade instead of 2, and (c) sharper Canny edge masks for better preservation of the anatomy.

The quality of the reconstruction with our network was superior compared to the H-EPT (Figure 10) implementation provided in Reference 57. In particular, the H-EPT reconstructions with a SG kernel equal to 1 or 2 appeared noisy with a few values in the brain anatomy approaching 0, while larger SG kernel led to blurred reconstructions to the point that different tissues were not easily distinguishable in the EP maps. Other deep learning-based methods, such as the ones presented in References 24,28, yielded accurate EP reconstructions. However, both techniques strongly rely on the contrast in MR images, which is not directly associated with EP, whereas our network learns the EP distribution from magnetic field measurements, which are related to the underlying conductivity and permittivity via Maxwell's equations.

The network proposed here could be alternatively used as an initial guess generator for other EP reconstruction methods^{21,23} that rely on iterative optimization strategies as in Reference 75, to considerably reduce their iteration count, thus making them feasible for real-time reconstructions. In addition, approaches like the one presented in Reference 76 could allow weight initialization for physics-informed models^{77,78} using the reconstructions produced from our proposed supervised network.

One particular constraint of our approach is that it used a local birdcage coil for MR excitation in the training datasets. If samples were scanned with a different coil, for example, the scanner body coil or a multi-channel transmit array, the accuracy of the reconstructed EP maps would worsen unless the network is re-trained with data generated using the new coil setup. Alternatively, the Barlow twins method for redundancy reduction⁷⁹ offers a solution to this problem and could potentially enforce the network to learn the same EP reconstructions from magnetic field maps acquired through different coils. Such an approach could broaden the network's adaptability to different coil configurations and will be explored in future work.

5 | CONCLUSION

We introduced a novel NN, based on vision transformers and FiLM, to reconstruct EP from MR measurements obtained with a head birdcage coil. The use of Canny edge masks as input improved the NN ability to preserve boundaries between compartments with different EP, reducing blurring in the reconstructed EP maps. We demonstrated excellent results for numerical simulations with OOD realistic brain models along with promising in

vivo reconstructions. Our approach is a promising direction towards clinically-usable in vivo EP reconstructions. Future work will focus on refining the resolution of our datasets and improving the network's architecture to further boost performance.

ACKNOWLEDGMENTS

The authors are thankful to Dr Batool Abbas, Mary Bruno, Jerzy Walczyk, and Xinling Yu for valuable discussions.

FUNDING INFORMATION

This work was supported in part by National Institutes of Health R01 EB024536 and was performed under the rubric of the Center for Advanced Imaging Innovation and Research (CAI²R, www.cai2r.net), an NIBIB National Center for Biomedical Imaging and Bioengineering (NIH P41 EB017183).

CONFLICT OF INTEREST STATEMENT

GK and MBK are employees of Siemens Medical Solutions.

ORCID

Ilias I. Giannakopoulos  <https://orcid.org/0000-0003-2180-5898>

Giuseppe Carluccio  <https://orcid.org/0000-0001-5376-3843>

Mahesh B. Keerthivasan  <https://orcid.org/0000-0002-7841-9333>

Hector L. De Moura  <https://orcid.org/0000-0002-6620-9814>

José E. Cruz Serrallés  <https://orcid.org/0000-0002-3323-5688>

Riccardo Lattanzi  <https://orcid.org/0000-0002-8240-5903>

TWITTER

Ilias I. Giannakopoulos  [IliasGiannakop3](https://twitter.com/IliasGiannakop3)

REFERENCES

- Hand JW. Modelling the interaction of electromagnetic fields (10 MHz–10 GHz) with the human body: methods and applications. *Phys Med Biol*. 2008;53:R243.
- Collins Christopher M, Zhangwei W. Calculation of radiofrequency electromagnetic fields and their effects in MRI of human subjects. *Magn Reson Med*. 2011;65:1470-1482.
- Jackson John David. *Classical electrodynamics*. John Wiley & Sons; 2021.
- Eugene M, Georgy G, Angel T-C, et al. Individualized SAR calculations using computer vision-based MR segmentation and a fast electromagnetic solver. *Magn Reson Med*. 2021;85:429-443.
- Giannakopoulos Ilias I, Guryev Georgy D, Serrallés José EC, et al. A hybrid volume-surface integral equation method for rapid electromagnetic simulations in MRI. *IEEE Trans Biomed Eng*. 2022;70:105-114.

6. Riccardo L, Sodickson Daniel K, Grant Aaron K, et al. Electrodynamics constraints on homogeneity and radiofrequency power deposition in multiple coil excitations. *Magn Reson Med.* 2009;61:315-334.
7. Christian R, Dieter H. Review of temperature dependence of thermal properties, dielectric properties, and perfusion of biological tissues at hyperthermic and ablation temperatures. *Crit Rev Biomed Eng.* 2014;42:467-492.
8. Hall Sheldon K, Ooi Ean H, Payne SJ. Cell death, perfusion and electrical parameters are critical in models of hepatic radiofrequency ablation. *Int J Hyperth.* 2015;31:538-550.
9. Edmond B, Petra KH, Gerben S, et al. Hyperthermia treatment planning for cervical cancer patients based on electrical conductivity tissue properties acquired in vivo with EPT at 3 T MRI. *Int J Hyperth.* 2016;32:558-568.
10. Holder DS. Detection of cerebral ischaemia in the anaesthetised rat by impedance measurement with scalp electrodes: implications for non-invasive imaging of stroke by electrical impedance tomography. *Clin Phys Physiol Meas.* 1992;13:63.
11. Antoni I, Bassim A-S, Boris R, et al. In vivo electrical conductivity measurements during and after tumor electroporation: conductivity changes reflect the treatment outcome. *Phys Med Biol.* 2009;54:5949.
12. Edmond B, Lier Astrid LHMW, Hans C, et al. Feasibility of electric property tomography of pelvic tumors at 3T. *Magn Reson Med.* 2015;73:1505-1513.
13. Jaewook S, Jung KM, Joonsung L, et al. Initial study on in vivo conductivity mapping of breast cancer using MRI. *J Magn Reson Imaging.* 2015;42:371-378.
14. Xiaotong Z, Jiaen L, Bin H. Magnetic-resonancebased electrical properties tomography: a review. *IEEE Rev Biomed Eng.* 2014;7:87-96.
15. Stefano M, Petrov Petar I, Vink Jord JT, et al. Brain tissue conductivity measurements with MR-electrical properties tomography: an in vivo study. *Brain Topogr.* 2021;34:56-63.
16. Narae C, Joonsung L, Min-Oh K, et al. A modified multi-echo AFI for simultaneous B1+ magnitude and phase mapping. *Magn Reson Imaging.* 2014;32:314-320.
17. Van Lier ALHMW, Brunner David O, Pruessmann Klaas P, et al. B phase mapping at 7 T and its application for in vivo electrical conductivity mapping. *Magn Reson Med.* 2012;67:552-561.
18. Collins Christopher M, Shizhe L, Smith MB. SAR and B1 field distributions in a heterogeneous human head model within a birdcage coil. *Magn Reson Med.* 1998;40:847-856.
19. Lier Astrid LHMW, Alexander R, Tobias V, et al. Electrical properties tomography in the human brain at 1.5, 3, and 7T: a comparison study. *Magn Reson Med.* 2014;71:354-363.
20. Stefano M, Alessandro S, Ulrich K, et al. Error analysis of helmholtz-based MR-electrical properties tomography. *Magn Reson Med.* 2018;80:90-100.
21. Serrallés José EC, Giannakopoulos Ilias I, Bei Z, et al. Noninvasive estimation of electrical properties from magnetic resonance measurements via Global Maxwell Tomography and match regularization. *IEEE Trans Biomed Eng.* 2019;67:3-15.
22. Ilias G, Serrallés Jos EC, Luca D, et al. Magnetic-resonance-based electrical property mapping using Global Maxwell Tomography with an 8-channel head coil at 7 Tesla: a simulation study. *IEEE Trans Biomed Eng.* 2021;68:236-246.
23. Edmond B, Berg Cornelis AT, Johan T, et al. CSI-EPT: a contrast source inversion approach for improved MRI-based electric properties tomography. *IEEE Trans Med Imaging.* 2015;34:1788-1796.
24. Stefano M, Meliàdò Ettore F, Huttinga Niek RF, et al. Opening a new window on MR-based electrical properties tomography with deep learning. *Sci Rep.* 2019;9:8895.
25. Phillip I, Zhu J-Y, Tinghui Z, et al. Image-to-image translation with conditional adversarial networks. *Proceedings of the IEEE conference on computer vision and pattern recognition;* 2017:1125-1134.
26. Daniel H, Kyoung-Nam K. Use of machine learning to improve the estimation of conductivity and permittivity based on longitudinal relaxation time T1 in magnetic resonance at 7 T. *Sci Rep.* 2023;13:7837.
27. Zheng M, Feiyang L, Yiman H, et al. MR-based electrical property tomography using a physics-informed network at 3 and 7 T. *NMR Biomed* 2024;37:e5137.
28. Kyu-Jin J, Stefano M, Chuanjiang C, et al. Data-driven electrical conductivity brain imaging using 3 T MRI. *Hum Brain Mapp.* 2023;44:4986-5001.
29. Soraya G, Berg Cornelis AT, Savenije Mark HF, et al. Deep learning-based reconstruction of in vivo pelvis conductivity with a 3D patch-based convolutional neural network trained on simulated MR data. *Magn Reson Med.* 2020;84:2772-2787.
30. Nils H, Ulrich K, Berg Cornelis AT, et al. Investigating the challenges and generalizability of deep learning brain conductivity mapping. *Phys Med Biol.* 2020;65:135001.
31. Paul B, Zhang Lei W, Xiaolin. Canny edge detection enhancement by scale multiplication. *IEEE Trans Pattern Anal Mach Intell.* 2005;27:1485-1490.
32. Giannakopoulos Ilias I, Xinling Y, Giuseppe C, et al. Electrical property mapping using vision transformers and canny edge detection. *Proc. ISMRM;* 2024:186.
33. Andreas C, Wolfgang K, Hahn Eckhart G, et al. The Virtual Family-development of surface-based anatomical models of two adults and two children for dosimetric simulations. *Phys Med Biol.* 2009;55:N23.
34. Peter D. Development of the female voxel phantom, NAOMI, and its application to calculations of induced current densities and electric fields from applied low frequency magnetic and electric fields. *Phys Med Biol.* 2005;50:1047.
35. Makarov Sergey N, Noetscher Gregory M, Janakinadh Y, et al. Virtual human models for electromagnetic studies and their applications. *IEEE Rev Biomed Eng.* 2017;10:95-121.
36. Giuseppe C, Eros M, Riccardo L, et al. A comparative study of 2D and 3D deep learning networks for human body models temperature prediction. *2023 IEEE EMBS special topic conference on data science and engineering in healthcare, medicine and biology.* IEEE; 2023:133-134.
37. Damijan M, Nataša P, Hart FX. Electric properties of tissues. *Wiley Encyclopedia of Biomedical Engineering.* 2006. doi:10.1002/9780471740360.ebs0403.
38. Karthik L, Martijn C, Wiggins Graham C, et al. Improving B1 excitation in head apex by combining birdcage coil with crossed dipole elements. *Proceedings of the 21st Annual Meeting of ISMRM;* 2013:2754.
39. Chin Chih-Liang, Collins Christopher M, Li Shizhe, et al. Birdcage-Builder: design of specified-geometry birdcage coils

- with desired current pattern and resonant frequency. *Concepts Magn Reson* 2002;15(2):156–163.
40. Georgakis Ioannis P, Giannakopoulos Ilias I, Litsarev Mikhail S, et al. A fast volume integral equation solver with linear basis functions for the accurate computation of EM fields in MRI. *IEEE Trans Antennas Propag*. 2020;69:4020–4032.
 41. Carlotta I, Zwart Jacco A, Qi D, et al. Synthesized tissue-equivalent dielectric phantoms using salt and polyvinylpyrrolidone solutions. *Magn Reson Med*. 2018;80:413–419.
 42. Cunningham Charles H, Pauly John M, Nayak Krishna S. Saturated double-angle method for rapid B1+ mapping. *Magn Reson Med* 2006;55(6):1326–1333.
 43. Akoka S, Florence F, Seguin F, et al. Radiofrequency map of an NMR coil by imaging. *Magn Reson Imaging*. 1993;11:437–441.
 44. Soraya G, Yulia S, Bartels Lambertus W, et al. Transceive phase mapping using the PLANET method and its application for conductivity mapping in the brain. *Magn Reson Imaging*. 2020;83:590–607.
 45. Mugler III, John P, Brookeman JR. Three-dimensional magnetization-prepared rapid gradient-echo imaging (3D MP RAGE). *Magn Reson Med*. 1990;15:152–157.
 46. Eros M, Antonella B, Marco B, et al. A multi-metric registration strategy for the alignment of longitudinal brain images in pediatric oncology. *Med Biol Eng Comput*. 2020;58:843–855.
 47. Jieneng C, Yongyi L, Qihang Y, et al. Transunet: transformers make strong encoders for medical image segmentation. *arXiv preprint:2102.04306*. 2021.
 48. Dan H, Kevin G. Gaussian error linear units (gelus). *arXiv preprint:1606.08415*. 2016.
 49. Nitish S, Geoffrey H, Alex K, et al. Dropout: a simple way to prevent neural networks from overfitting. *J Mach Learn Res*. 2014;15:1929–1958.
 50. Ethan P, Florian S, Harm DV, et al. Film: Visual reasoning with a general conditioning layer. *Proceedings of the AAAI conference on artificial intelligence*. Vol 32; 2018.
 51. Xu B, Naiyan W, Tianqi C, et al. Empirical evaluation of rectified activations in convolutional network. *arXiv preprint:1505.00853*. 2015.
 52. Kingma Diederik P, Jimmy B. Adam: a method for stochastic optimization. *arXiv preprint:1412.6980*. 2014.
 53. Smith LN. Cyclical learning rates for training neural networks. *2017 IEEE winter conference on applications of computer vision (WACV)*. IEEE; 2017:464–472.
 54. Hao L, Pratik C, Hao Y, et al. Rethinking the hyperparameters for fine-tuning. *arXiv preprint:2002.11770*. 2020.
 55. Zhou W, Bovik Alan C, Sheikh Hamid R, et al. Image quality assessment: from error visibility to structural similarity. *IEEE Trans Image Process*. 2004;13:600–612.
 56. Ulrich K, Tobias V, Christian F, et al. Determination of electric conductivity and local SAR via B1 mapping. *IEEE Trans Med Imaging*. 2009;28:1365–1374.
 57. Alessandro A. EPTlib: An open-source extensible collection of electric properties tomography techniques. *Appl Sci*. 2021;11:3237.
 58. Press William H, Teukolsky SA. Savitzky-Golay smoothing filters. *Comput Phys*. 1990;4:669–672.
 59. Umme S, Morium A, Shorif UM. Image quality assessment through FSIM, SSIM, MSE and PSNR—a comparative study. *J Comput Commun*. 2019;7:8–18.
 60. Dosovitskiy Alexey, Beyer Lucas, Kolesnikov Alexander, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint:2010.11929*. 2020.
 61. Olaf R, Philipp F, Thomas B. U-net: Convolutional networks for biomedical image segmentation. *Medical Image Computing and Computer-Assisted Intervention—MICCAI 2015: 18th International Conference, Munich, Germany, October 5–9, 2015, Proceedings, Part III 18*. Springer; 2015:234–241.
 62. Ilias G, José S, Georgy G, et al. On the usage of deep neural networks as a tensor-to-tensor translation between MR measurements and electrical properties. *Proc. Intl. Soc. Mag. Reson. Med*. 28:3193; 2020.
 63. Chun-Mei F, Yunlu Y, Huazhu F, et al. Task transformer network for joint MRI reconstruction and super resolution. *Medical image computing and computer assisted intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24*. 307–317. Springer International Publishing; 2021.
 64. Yilmaz K, Dar Salman UH, Mahmut Y, et al. Unsupervised MRI reconstruction via zero-shot learned adversarial transformers. *IEEE Trans Med Imaging*. 2022;41:1747–1763.
 65. Dianlin H, Yikun Z, Jianfeng Z, et al. Trans-net: Transformer-enhanced residual-error alternative suppression network for MRI reconstruction. *IEEE Trans Instrum Meas*. 2022;71:1–13.
 66. Giannakopoulos Ilias I, Muckley Matthew J, Jesi K, et al. Accelerated MRI reconstructions via variational network and feature domain learning. *Sci Rep*. 2024;14:10991.
 67. Xia Xide, Kulis Brian. W-net: A deep model for fully unsupervised image segmentation. *arXiv preprint: 1711.08506*. 2017.
 68. Jawook G, Chul YJ. AdaIN-based tunable CycleGAN for efficient unsupervised low-dose CT denoising. *IEEE Trans Comput Imaging*. 2021;7:73–85.
 69. Nima T, Shin Jae Y, Gurudu Suryakanth R, et al. Convolutional neural networks for medical image analysis: Full training or fine tuning? *IEEE Trans Med Imaging*. 2016;35:1299–1312.
 70. Augustus O, Vincent D, Chris O. Deconvolution and checkerboard artifacts. *Distill*. 2016;1:e3.
 71. Robin R, Andreas B, Dominik L, et al. High-resolution image synthesis with latent diffusion models. *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*; 2022:10684–10695.
 72. Muhammad UA, Måns L, Ida B, et al. Brain tumor segmentation using synthetic MR images—A comparison of GANs and diffusion models. *Sci Data*. 2024;11:259.
 73. Naohiro E, Motofumi F, Keisuke H, et al. A method for electrical property tomography based on a three-dimensional integral representation of the electric field. *IEEE Trans Med Imaging*. 2021;41:1400–1409.
 74. Gonzalez BMA, Zisserman Andrew P, Michael B. Estimation of the partial volume effect in MRI. *Med Image Anal*. 2002;6:389–405.
 75. Reijer L, Cornelis B, Andrew W, et al. Combining deep learning and 3D contrast source inversion in MR-based electrical properties tomography. *NMR Biomed*. 2022;35:e4211.
 76. Homayoon T, Behzad NS. Weight initialization algorithm for physics-informed neural networks using finite differences. *Eng Comput*. 2024;40:1603–1619.
 77. Garcia IAJ, Ying HS, Nevrez İ, et al. Physics informed neural networks (PINN) for low snr magnetic resonance

- electrical properties tomography (MREPT). *Diagnostics*. 2022; 12:2627.
78. Xinling Y, Serrallés José EC, Giannakopoulos Ilias I, et al. PIFON-EPT: MR-based electrical property tomography using physics-informed Fourier networks. *IEEE J Multiscale Multi-physics Comput Tech*. 2023;9:49-60.
79. Jure Z, Li J, Ishan M, et al. Barlow twins: Self-supervised learning via redundancy reduction. *International Conference on Machine Learning*; 2021:12310-12320.

How to cite this article: Giannakopoulos II, Carluccio G, Keerthivasan MB, et al. MR electrical properties mapping using vision transformers and canny edge detectors. *Magn Reson Med*. 2025;93:1117-1131. doi: 10.1002/mrm.30338