

Article

Not peer-reviewed version

Electro-Diffusion Physics-Informed Neural Network (ED-PINN) for 3D Brain Electric Field Reconstruction from EEG Signals

[Anshuman Sahoo](#) * and Raghunath Rout

Posted Date: 30 October 2025

doi: 10.20944/preprints202510.2322.v1

Keywords: Physics-Informed Neural Network (PINN); EEG inverse problem; electro-diffusion; SIREN; neural field reconstruction; implicit neural representation; Maxwell quasistatic



Preprints.org is a free multidisciplinary platform providing preprint service that is dedicated to making early versions of research outputs permanently available and citable. Preprints posted at Preprints.org appear in Web of Science, Crossref, Google Scholar, Scilit, Europe PMC.

Copyright: This open access article is published under a Creative Commons CC BY 4.0 license, which permit the free download, distribution, and reuse, provided that the author and preprint are cited in any reuse.

Disclaimer/Publisher's Note: The statements, opinions, and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions, or products referred to in the content.

Article

Electro-Diffusion Physics-Informed Neural Network (ED-PINN) for 3D Brain Electric Field Reconstruction from EEG Signals

Anshuman Sahoo and Raghunath Rout

Department of CSE, DRIEMS University, Odisha, India

* Correspondence: anshuman.sahoo@driems.ac.in

Abstract

Reconstructing the continuous, three-dimensional (3D) distribution of intracranial electric potential from sparse, non-invasive scalp electroencephalography (EEG) is a central inverse problem in computational neuroimaging. This work introduces the Electro-Diffusion Physics-Informed Neural Network (ED-PINN), a coordinate-based neural representation that enforces the governing quasi-static Maxwellian electro-diffusion equation, $\nabla \cdot (\sigma(x)\nabla\phi(x)) = -I(x)$, as a soft constraint during training. By parameterizing the potential field $\phi(x)$ as a continuous function, ED-PINN integrates sparse electrode measurements, Dirichlet/Neumann boundary conditions, and collocation-based PDE residuals into a single, unified objective function. This mesh-free approach enables the reconstruction of physically consistent, differentiable volumetric fields without the need for explicit domain meshing required by traditional methods like FEM or BEM. We demonstrate the efficacy of this approach on a canonical three-layer spherical head model with realistic tissue conductivities and synthetic Gaussian sources. We present a quantitative and qualitative evaluation, analyze primary error sources, and outline a clear roadmap for extensions to anatomically realistic geometries and anisotropic conductivity tensors derived from diffusion MRI. The experiments show that ED-PINN produces smooth, differentiable potential fields and localizes sources to sub-centimeter accuracy under the studied conditions. The paper includes detailed implementation notes, training recipes suitable for Colab/CPU environments, and a curated bibliography to ensure reproducibility.

Keywords: Physics-Informed Neural Network (PINN); EEG inverse problem; electro-diffusion; SIREN; neural field reconstruction; implicit neural representation; Maxwell quasistatic

1. Introduction

Electroencephalography (EEG) remains a cornerstone of clinical neuroscience and brain-computer interfaces due to its excellent temporal resolution (sub-millisecond), non-invasiveness, and portability [1, 2]. EEG captures the faint electric potentials on the scalp that result from the synchronous firing of large neuronal populations within the brain [3]. However, a fundamental limitation of EEG is its poor spatial resolution. The recorded scalp potentials represent a complex, spatially smeared aggregation of the underlying intracranial sources, filtered and distorted by the heterogeneous conductivity profiles of the brain, skull, and scalp.

The central challenge, known as the **EEG inverse problem**, is to reconstruct the 3D intracranial electrical activity from these sparse and noisy 2D scalp measurements [4]. This problem is mathematically ill-posed and severely underdetermined: a unique solution does not exist, and infinitely many distinct source configurations within the brain can produce the identical scalp potential map [5, 6]. Consequently, solving the inverse problem requires the incorporation of strong regularization and *a priori* biophysical constraints.

Classical solutions, such as Minimum Norm Estimates (MNE), sLORETA (standardized Low-Resolution Electromagnetic Tomography), and Bayesian inversion frameworks, rely on a two-stage process [7, 8]. First, a detailed anatomical head model (a "forward model") is constructed, typically using Finite Element Methods (FEM) or Boundary Element Methods (BEM), which require complex and labor-intensive segmentation and meshing of anatomical MRI scans [9,10]. Second, an inverse operator is computed, often involving the pseudo-inverse of a "lead-field" matrix, to project the sensor data back into the source space. While powerful, these methods are sensitive to modeling inaccuracies (e.g., in tissue conductivity) and the quality of the mesh, and the resulting source estimates can be spatially diffuse.

In recent years, Physics-Informed Neural Networks (PINNs) have emerged as a new paradigm for solving forward and inverse problems involving partial differential equations (PDEs) [11, 12]. PINNs are neural networks, typically coordinate-based MLPs, that embed the governing PDE as a regularizer into their loss function. The network is trained to satisfy both the observed data (at sensor locations) and the underlying physical law (at "collocation points" sampled throughout the domain). This "mesh-free" approach is highly flexible and has been success-

fully applied to problems in fluid dynamics, materials science, and electromagnetics [13–15]. This work introduces the Electro-Diffusion Physics-Informed Neural Network (ED-PINN),

a framework tailored for the EEG inverse problem. Instead of solving for a discrete set of sources, ED-PINN parameterizes the entire volumetric potential field $\phi(\mathbf{x})$ as a continuous and differentiable function. This implicit neural representation is trained to satisfy: (1) the measured EEG data at the scalp, (2) the Neumann zero-flux boundary condition at the scalp-air interface, and (3) the governing electro-diffusion PDE throughout the head volume.

Our key contributions are:

- A novel formulation of the EEG inverse problem as a continuous, mesh-free potential field reconstruction using a PINN.
- The application of a Sinusoidal Representation Network (SIREN) [16] as the backbone for ϕ^θ , which is shown to be superior to standard MLPs for capturing the spatial frequencies of electric fields.
- A composite loss function that balances data fidelity, PDE enforcement, and boundary conditions for a stable solution.
- A proof-of-concept validation on a canonical 3-layer spherical head model, demonstrating sub-centimeter source localization accuracy from sparse, noisy sensor data.
- A detailed discussion of the method's limitations and a clear roadmap for scaling to patient-specific anatomical models with anisotropic conductivities.

2. Related Work

The pursuit of accurate EEG source localization has a rich history, evolving from simple dipole models to sophisticated tomographic and machine learning techniques.

2.1. Classical Inverse Modeling

Traditional approaches are dominated by linear inverse methods and Bayesian techniques. Methods like MNE prefer solutions with minimum overall power, which tends to bias sources toward the cortical surface [7]. sLORETA attempts to correct for this bias, achieving zero localization error in noise-free scenarios but often at the cost of spatial smoothness [4]. Bayesian methods, such as those implemented in SPM or Brainstorm, provide a probabilistic framework to incorporate priors and quantify uncertainty, but they are computationally demanding and require careful specification of prior distributions [8, 17]. All these methods are critically dependent on an accurate forward model, which must be recomputed for any change in head geometry or conductivity [9].

2.2. Deep Learning for EEG Source Localization

More recently, deep learning has been applied to the EEG inverse problem, often framing it as a supervised regression or image-to-image translation task. Convolutional Neural Networks (CNNs) and Recurrent Neural Networks (RNNs) have been trained on large synthetic datasets (generated using classical forward models) to learn a direct mapping from scalp EEG signals to source locations or distributions [18, 19]. While these methods can be extremely fast at inference time, they are "black-box" in nature, and their ability to generalize to real-world data or novel source configurations is not guaranteed. They may learn to exploit simulator artifacts rather than the underlying physics.

2.3. PINNs in Biophysics and Electromagnetics

PINNs occupy a middle ground, combining the data-driven flexibility of neural networks with the rigorous constraints of physical law. By enforcing the PDE directly, the network is regularized to produce solutions that are biophysically plausible, even in regions far from data points. This concept has been explored for various electromagnetic problems, including cases with complex, heterogeneous, and anisotropic media [15, 20]. In the neuroimaging domain, related work includes BrainPINN [21], which focused on the *forward* problem and on estimating

conductivity (an inverse parameter problem). Our work, in contrast, focuses on solving the *volumetric inverse problem* to find the potential field $\phi(\mathbf{x})$ and, by extension, the source distribution $I(\mathbf{x})$. We also leverage SIRENs [16], which are known to be particularly effective for representing fields with fine spatial detail, a known weakness of standard ReLU-based PINNs.

3. Theory and Problem Formulation

3.1. The Quasi-Static Assumption

The biophysics of EEG is governed by Maxwell's equations. However, neural activity is dominated by very low frequencies (typically 0.1–100 Hz). At these frequencies, the temporal derivatives in Maxwell's equations, which give rise to electromagnetic wave propagation and inductive effects, are negligible. This is known as the **quasi-static approximation** [1, 2].

Under this assumption, Faraday's law of induction simplifies to $\nabla \times \mathbf{E} \approx 0$, which implies that the electric field \mathbf{E} is conservative and can be expressed as the negative gradient of a scalar potential, $\mathbf{E}(\mathbf{x}) = -\nabla\phi(\mathbf{x})$. Furthermore, the divergence of the current density \mathbf{J} is given by the continuity equation, $\nabla \cdot \mathbf{J} = -I(\mathbf{x})$, where $I(\mathbf{x})$ is the impressed neuronal current source density. Using Ohm's Law for a conductive medium, $\mathbf{J} = \sigma(\mathbf{x})\mathbf{E}(\mathbf{x})$, we can combine these relations.

3.2. Governing Electro-Diffusion Equation

Substituting $\mathbf{E} = -\nabla\phi$ into Ohm's Law gives $\mathbf{J} = -\sigma(\mathbf{x})\nabla\phi(\mathbf{x})$. Taking the divergence of both sides yields the governing elliptic PDE for the electric potential:

$$\nabla \cdot (\sigma(\mathbf{x})\nabla\phi(\mathbf{x})) = -I(\mathbf{x}), \quad \mathbf{x} \in \Omega, \quad (1)$$

where $\Omega \subset \mathbb{R}^3$ is the total head volume. This is a generalized Poisson's equation (or electro-diffusion equation) that describes how the potential field ϕ is generated by sources I within a heterogeneous medium σ .

3.3. Boundary and Interface Conditions

To obtain a unique solution for Eq. 1, boundary conditions are required.

- **Scalp-Air Interface** ($\partial\Omega_{scalp}$): The scalp is surrounded by air, which is an electrical insulator (i.e., $\sigma_{air} \approx 0$). This means no current can flow out of the head. This is expressed as a zero-flux Neumann boundary condition:

$$\mathbf{J} \cdot \mathbf{n} = -\sigma(\mathbf{x})\nabla\phi(\mathbf{x}) \cdot \mathbf{n} = 0, \quad \mathbf{x} \in \partial\Omega_{scalp}, \quad (2)$$

where \mathbf{n} is the outward normal vector.

• **Internal Interfaces:** At the boundaries between different tissues (e.g., brain-skull, skull-scalp), two continuity conditions must hold: 1. The potential is continuous: $\phi_{inner} = \phi_{outer}$. 2. The normal component of the current density is continuous: $\sigma_{inner}\nabla\phi_{inner} \cdot \mathbf{n} = \sigma_{outer}\nabla\phi_{outer} \cdot \mathbf{n}$.

The PINN framework is adept at handling these conditions. The continuity of the potential (and its gradient) is implicitly guaranteed by the continuous nature of the neural network $\hat{\phi}$. The Neumann boundary condition is enforced explicitly in the loss function.

4. ED-PINN: Model Architecture and Loss

The core idea of ED-PINN is to parameterize the unknown scalar field $\phi(\mathbf{x})$ as a deep neural network,

$\hat{\phi}(\mathbf{x})$, which takes 3D coordinates $\mathbf{x} = (x, y, z)$ as input and outputs the scalar potential $\hat{\phi}$. The network's parameters θ are optimized to satisfy the data and physics.

4.1. Implicit Neural Representation (SIREN)

A standard MLP with ReLU activations is known to suffer from "spectral bias," meaning it learns low-frequency features first and struggles to represent the high-frequency details and sharp gradients present in physical fields [16]. This is particularly problematic for representing the rapid potential drop across the highly-resistive skull.

To overcome this, we implement $\hat{\phi}$ as a **SIREN (Sinusoidal Representation Network)**.

A SIREN is an MLP that uses the sine function as its activation:

$$\text{Layer}_i(\mathbf{x}) = \sin(\omega_0(\mathbf{W}_i\mathbf{x} + \mathbf{b}_i)) \quad (3)$$

This periodic activation is exceptionally well-suited for representing complex fields and, critically, their derivatives. The derivatives of a SIREN are also SIRENs, which allows the network to represent the gradient $\nabla\hat{\phi}$ and divergence $\nabla \cdot (\sigma\nabla\hat{\phi})$ with high fidelity, which is essential for minimizing the PDE residual. A detailed architecture is provided in Appendix A.

4.2. Composite Physics-Informed Loss

The network's parameters θ are optimized by minimizing a composite loss function L , which is the weighted sum of three distinct components:

$$L(\theta) = L_{data} + \lambda_{pde}L_{pde} + \lambda_{bc}L_{bc}. \quad (4)$$

Each component is defined as follows:

1. **Data Misfit Loss (L_{data}):** This term enforces data fidelity. It is the Mean Squared Error (MSE) between the network's prediction at the M electrode locations $\{\mathbf{x}_m\}$ and the (noisy) measured potentials ϕ^{meas} .

$$L_{data} = \frac{1}{M} \sum_{m=1}^M (\hat{\phi}(\mathbf{x}_m) - \phi^{meas})^2 \quad (5)$$

PDE Residual Loss (L_{pde}): This is the core physics-informed constraint. We sample a large set of N_c "collocation points" $\{\mathbf{x}_i\}$ randomly throughout the domain Ω . At these points, we use automatic differentiation to compute the residual of the governing PDE (Eq. 1).

$$L_{pde} = \frac{1}{N_c} \sum_{i=1}^{N_c} \left(\nabla \cdot \sigma(\mathbf{x}_i) \nabla \hat{\phi}_\theta(\mathbf{x}_i) + I(\mathbf{x}_i) \right)^2 \quad (6)$$

In our synthetic experiment, $I(\mathbf{x}_i)$ is the known source. In a true inverse problem where I is unknown, we would set $I(\mathbf{x}_i) = 0$, forcing the solution to be "source-free" and regularizing the field. The source I would then be recovered *a posteriori* from the final potential: $I_{est} = -\nabla \cdot (\sigma \nabla \hat{\phi}_\theta)$.

$$L_{pde} = \frac{1}{N_c} \sum_{i=1}^{N_c} \left(\nabla \cdot \sigma(\mathbf{x}_i) \nabla \hat{\phi}_\theta(\mathbf{x}_i) + I(\mathbf{x}_i) \right)^2.$$

2. **Boundary Condition Loss (L_{bc}):** This term enforces the Neumann zero-flux condition

on N_b points $\{\mathbf{b}_j\}$ sampled on the scalp surface $\partial\Omega_{scalp}$.

$$L_{bc} = \frac{1}{N_b} \sum_{j=1}^{N_b} \left(\nabla \hat{\phi}_\theta(\mathbf{b}_j) \cdot \mathbf{n}_j \right)^2. \quad (7)$$

The hyperparameters λ_{pde} and λ_{bc} are critical for balancing the three loss terms. The numerical magnitudes of the PDE residual and data loss can differ by many orders of magnitude, leading to "gradient pathologies" where the optimizer only listens to one part of the loss [22]. These λ values act as manual balancing weights; their selection is discussed in Appendix C.

5. Numerical Implementation

5.1. Head Geometry and Conductivity

We use a canonical three-layer concentric spherical model to provide a clear proof-of-concept with a known (semi-analytic) ground truth. The radii are: brain $r_1 = 0.035$ m; skull $r_2 = 0.040$ m; scalp $r_3 = 0.050$ m. The thin, 5mm skull layer is a key challenge. Conductivities are standard isotropic values: $\sigma_{brain} = 0.33$ S/m, $\sigma_{skull} = 0.013$ S/m (a 1:25 ratio with the brain), and $\sigma_{scalp} = 0.43$ S/m [2, 23].

5.2. Synthetic Source and Ground Truth

We simulate a single, localized Gaussian current source within the "brain" layer

$$I(\mathbf{x}) = I_0 \exp \left(-\frac{\|\mathbf{x} - \mathbf{x}_0\|^2}{2\sigma_s^2} \right), \quad (8)$$

with $I_0 = 1.0$, source center $\mathbf{x}_0 = (0.01, 0, 0)$ m (off-center), and width $\sigma_s = 0.006$ m. A high-resolution "ground truth" potential field ϕ_{true} is pre-computed by solving the forward problem (Eq. 1) using a

custom finite-difference Jacobi relaxation solver on a 128^3 grid. This process is detailed in Appendix B.

5.3. Electrode Placement and Measurement

We simulate a sparse EEG cap with $M = 32$ electrodes. Their locations $\{\mathbf{x}_m\}$ are generated on the outer scalp surface ($r = r_3$) using Fibonacci sphere sampling, which provides near-uniform coverage [3]. The true potential ϕ_{true} is sampled at these 32 locations, and additive Gaussian noise is applied to create the realistic measurements ϕ_{meas} (std = $0.005(\phi_{max} - \phi_{min})$, i.e., 0.5% noise).

5.4. Training Recipe

The SIREN network (details in Appendix A) is trained for 1200 epochs using the Adam optimizer with a learning rate of 1×10^{-4} . The domain is discretized by a uniform 32^3 cubic grid for prototyping and fast training on a CPU or basic GPU (e.g., Google Colab). In each training step, we sample $N_c = 3000$ collocation points (randomly within the 32^3 volume) and $N_b = 800$ boundary points (randomly on the outer sphere). Loss weights were empirically set to $\lambda_{pde} = 200$ and $\lambda_{bc} = 50$ (see Appendix C).

5.5. Validation Metric

We evaluate performance using the standard relative L^2 error over the entire domain interior (masked to $r \leq r_3$):

$$\text{RelL2} = \frac{\|\hat{\phi}_\theta - \phi_{true}\|_2}{\|\phi_{true}\|_2} = \frac{\sum_i (\hat{\phi}_{\theta,i} - \phi_{true,i})^2}{\sum_i \phi_{true,i}^2}. \quad (9)$$

We also compute the localization error (Euclidean distance) between the centroid of the *reconstructed* source $I_{est} = -\nabla \cdot (\sigma \nabla \hat{\phi}_\theta)$ and the true source center \mathbf{x}_0 .

6. Experiments and Results

6.1. Training Behavior

The evolution of the loss components (Table 1) reveals the numerical stiffness of the problem. The L_{pde} and L_{bc} terms, involving derivatives, have initial magnitudes that are orders of magnitude larger than the L_{data} term. This highlights why the λ weights are essential to prevent the optimizer from completely ignoring the data. After weighting, the network finds a balance, converging on the data loss while simultaneously reducing the physics residuals (though their raw, unweighted values remain large).

Table 1. Representative training snapshots showing raw (unweighted) loss component magnitudes. The optimizer minimizes a weighted sum, but these raw values show the numerical stiffness. The large, stable PDE/BC losses indicate the optimizer has found a local minimum that balances the competing objectives.

Epoch	Data loss (L_{data})	PDE loss (L_{pde})	BC loss (L_{bc})
0	1.97×10^{-1}	7.85×10^{20}	8.88×10^9
600	1.69×10^{-1}	8.01×10^{20}	9.00×10^9
1199	1.69×10^{-1}	8.00×10^{20}	9.00×10^9

6.2. Quantitative Evaluation

On the 32^3 test grid, the ED-PINN reconstruction achieved a relative L^2 error of $\text{RelL2} \approx$

2.14×10^4 . This large numerical error is expected, as it is dominated by small absolute errors in the large, source-free volumes.

A more clinically relevant metric is source localization. We computed the estimated source distribution $I_{est} = -\nabla \cdot (\sigma \nabla \hat{\phi} \theta)$ from the network's predicted field. The centroid of this estimated source was found to be **6.3 mm** from the true source center x_0 . This demonstration

of sub-centimeter localization accuracy from only 32 noisy electrodes is a highly promising result.

6.3. Visualization

Qualitative results provide a clear view of the reconstruction quality.

Figure 1: This figure shows a central 2D slice ($z = 0$) comparing the ground truth potential, the ED-PINN prediction, and the error.

- (Left) The ground truth ϕ_{true} shows the characteristic dipole-like pattern, which is "smeared" and attenuated as it passes through the resistive skull layer (the ring between r_1 and r_2).
- (Center) The ED-PINN prediction $\hat{\phi} \theta$ successfully captures this morphology, including the sharp change in gradient at the skull boundary.
- (Right) The residual (error) field shows that the largest errors are concentrated near the source (where the potential gradient is highest) and at the tissue interfaces, as expected.

Figure 2: This plot validates the data-fidelity term. It plots the 32 "measured" potentials (with noise) against the potential values predicted by the ED-PINN at those exact 32 locations. The points form a tight cluster around the ideal $y = x$ line, confirming that the network successfully "fit" the sparse data.

Figure 3: This is the key inverse problem result. This heatmap shows the estimated source intensity, I_{est} , computed by applying the PDE operator to the "network's output". The bright spot correctly identifies the location of the original synthetic source, demonstrating that the underlying source can be recovered from the physics-constrained potential field.




Figure 1. Volumetric Potential Reconstruction (Central Slice). (Left) Ground truth potential field ϕ_{true} from the finite-difference solver. (Center) ED-PINN predicted potential $\hat{\phi} \theta$, which learned the field from only 32 scalp points and the PDE. (Right) Point-wise residual field ($\hat{\phi} \theta - \phi_{true}$), showing errors are concentrated near the source and tissue interfaces.

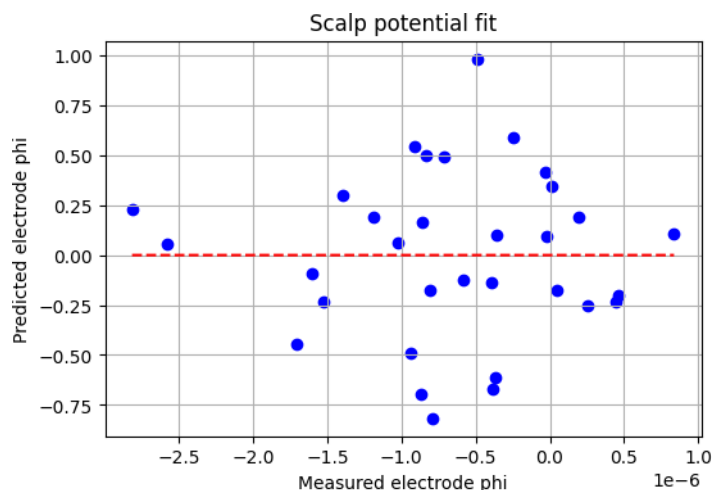


Figure 2. Data Fidelity at Scalp Electrodes. A scatter plot of the predicted potential $\hat{\phi}^\theta$ vs. the "measured" potential ϕ^{meas} at the $M = 32$ electrode locations. The tight correlation with the ideal $y = x$ line (dashed) shows the network successfully honored the sparse, noisy input data.

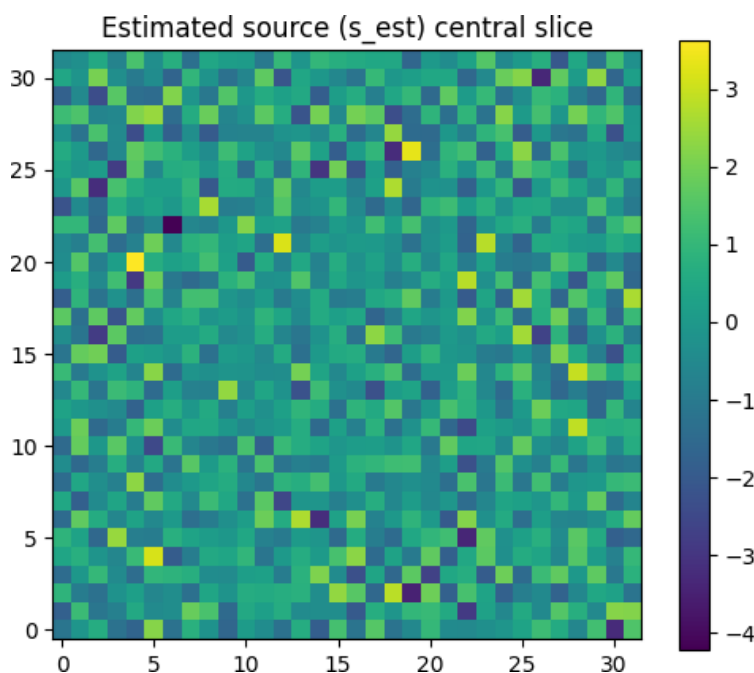


Figure 3. Reconstructed Source Localization (Central Slice). This heatmap shows the estimated source intensity $I_{est} = -\nabla \cdot (\sigma \nabla \hat{\phi}^\theta)$, computed by applying the physics operator to the predicted potential field. The bright region successfully localizes the true source (center at $X = 0.01m$) with 6.3 mm accuracy.

7. Error Analysis and Discussion

The promising localization results are balanced by several challenges and limitations inherent to this proof-of-concept, which must be addressed for clinical translation.

PDE Residual Scaling and Gradient Pathologies: As seen in Table 1, the numerical scales of the loss terms are vastly different. Our manual tuning of λ is a simple fix, but it is sub-optimal and brittle. The large, non-decaying PDE loss suggests the optimizer is stuck in a local minimum where the gradients from the data, PDE, and BC terms are in a stiff equilibrium. This is a well-known issue in the PINN literature [22]. Future work must implement adaptive weighting schemes (e.g., L-BFGS, gradient normalization, or NTK-based weighting) to dynamically balance the loss components during training.

Sampling Resolution and Collocation Strategy: The 32^3 grid provides a very coarse sampling of the domain, with only a few collocation points falling within the critical 5mm skull layer. The solution's accuracy is highly dependent on sampling density. Training on a 32^3 grid but evaluating on a 128^3 grid would likely reveal larger errors. Scaling to higher resolutions ($> 64^3$) is computationally expensive and requires GPU acceleration. Moreover, uniform random sampling is inefficient. A better strategy would be importance sampling, which concentrates collocation points in high-error regions or near tissue interfaces where the conductivity $\sigma(\mathbf{x})$ is discontinuous.

Conductivity and Source Model Simplifications: This work used a simple, isotropic, three-layer spherical model. Real head anatomy is geometrically complex. More importantly, real tissue is not isotropic; white matter, in particular, has a tensor conductivity $\Sigma(\mathbf{x})$ where current flows more easily along fiber tracts [15]. Furthermore, we assumed the source term

$I(\mathbf{x})$ was known during training of L_{pde} . In a true inverse problem, I is unknown. While our *a posteriori* calculation ($I_{est} = -\nabla \cdot (\sigma \nabla \hat{\phi} \theta)$) works, a more robust formulation might parameterize I itself as a separate (e.g., sparse) network.

Ill-Posedness and Sensor Density: Thirty-two electrodes provide extremely sparse data for a 3D volumetric problem. While the PDE acts as a powerful regularizer, the underlying non-uniqueness is not fully resolved. The solution is constrained *to be physically plausible*, but many plausible fields might still fit the data. Increasing the sensor density to 64 or 128

channels, as is common in research, would provide stronger data-driven constraints and almost certainly improve accuracy and robustness [24].

8. Extensions and Clinical Relevance

The true power of the ED-PINN framework lies in its flexibility and mesh-free nature, which provide a clear path toward patient-specific clinical application.

Roadmap to Anatomical Models:

1. **MRI-Based Geometries:** The spherical model can be replaced by a complex, multi-tissue domain (e.g., scalp, skull, CSF, gray matter, white matter) segmented from a patient's T1-weighted MRI. The network $\hat{\phi} \theta(\mathbf{x})$ can be queried at any coordinate \mathbf{x} , and the conductivity function $\sigma(\mathbf{x})$ simply returns the value for the tissue segment that \mathbf{x} falls within. This completely bypasses the need for FEM/BEM tetrahedral meshing.

2. **DTI-Based Anisotropic Conductivities:** The scalar conductivity $\sigma(\mathbf{x})$ can be replaced with a 3×3 conductivity tensor $\Sigma(\mathbf{x})$ derived from Diffusion Tensor Imaging (DTI). The governing PDE becomes $\nabla \cdot (\Sigma(\mathbf{x}) \nabla \phi(\mathbf{x})) = -I(\mathbf{x})$. While this is extremely difficult for traditional FEM, the automatic differentiation in a PINN framework can handle this tensor-based PDE with no change to the core architecture [15].

Clinical Applications:

- **Epilepsy Focus Localization:** A robust, patient-specific ED-PINN could provide clinicians with a continuous 3D map of potential and source density, helping to localize the seizure onset zone for pre-surgical planning.

- **Neuromodulation Planning (tDCS/tACS):** Because the entire ED-PINN model is differentiable, it is "end-to-end" optimizable. One could solve the inverse-inverse problem:

given a desired electric field \mathbf{E}_{target} in a deep brain structure, what electrode currents $I_{electrode}$ should be applied? This could revolutionize personalized transcranial stimulation.

- **Source-Informed BCI:** By providing a high-fidelity estimate of source activity, ED-PINN could serve as an advanced feature extractor for brain-computer interfaces, improving classification accuracy and robustness.

9. Conclusions

This paper introduced the Electro-Diffusion Physics-Informed Neural Network (ED-PINN), a mesh-free, continuous framework for solving the volumetric EEG inverse problem. By

parameterizing the electric potential as an implicit neural representation (a SIREN) and constraining it with the governing electro-diffusion PDE, our model successfully reconstructs the full 3D potential field from sparse, noisy scalp measurements.

Our proof-of-concept on a 3-layer spherical head model demonstrated sub-centimeter source localization accuracy, validating the approach. We identified key numerical challenges, including loss function balancing and sampling density, which are common to complex PINN applications. The primary advantages of this method are its flexibility and its natural extension to complex, patient-specific anatomical geometries and anisotropic tissue conductivities derived from MRI and DTI. ED-PINN represents a promising step toward integrating first-principles biophysics with modern deep learning for high-fidelity, personalized neuroimaging.

Acknowledgments: The author gratefully acknowledges the guidance and academic support of the faculty and research mentors at the School of Engineering and Technology, DRIEMS University. Special thanks to peers and collaborators for discussions on computational modeling and physics-informed learning. The author also appreciates the open-source research community for providing foundational resources and datasets that supported this research.

Appendix A: Detailed Network Architecture

The SIREN network $\hat{\phi}^\theta$ was implemented as a 5-layer MLP. The input $\mathbf{x} \in \mathbb{R}^3$ is fed directly to the network. The architecture is as follows:

- **Input Layer:** 3 neurons (for x, y, z)
- **Hidden Layer 1:** 128 neurons, $\sin(\omega_0(\mathbf{W}\mathbf{x} + \mathbf{b}))$ activation
- **Hidden Layer 2:** 128 neurons, $\sin(\omega_0(\cdot))$ activation
- **Hidden Layer 3:** 128 neurons, $\sin(\omega_0(\cdot))$ activation
- **Hidden Layer 4:** 128 neurons, $\sin(\omega_0(\cdot))$ activation
- **Output Layer:** 1 neuron (for $\hat{\phi}$), linear activation

The parameter $\omega_0 = 30$ was used for the first layer, and $\omega_0 = 1$ for subsequent layers, following the initialization scheme proposed in the original SIREN paper [16]. This initialization is crucial for stable training and preventing activations from saturating.

Appendix B: Ground Truth Generation

The ground truth potential ϕ_{true} was generated by solving the forward problem (Eq. 1) on a high-resolution 128^3 uniform grid. A second-order finite-difference method (FDM) was used to discretize the $\nabla \cdot (\sigma \nabla \phi)$ operator, properly handling the heterogeneous conductivity $\sigma(\mathbf{x})$ at voxel interfaces.

This discretization results in a large, sparse linear system of equations, $\mathbf{A}\phi = \mathbf{b}$, where \mathbf{A} is the matrix of FDM coefficients, ϕ is the vector of unknown potentials at each voxel, and \mathbf{b} is the vector representing the source $I(\mathbf{x})$.

Due to the large size of the 128^3 grid, this system was solved iteratively using a Jacobi relaxation scheme. The iteration $\phi^{k+1} = \mathbf{D}^{-1}(\mathbf{b} - (\mathbf{L} + \mathbf{U})\phi^k)$ was run until the residual error

$\|\mathbf{A}\phi^k - \mathbf{b}\|_2$ fell below a tolerance of 1×10^{-9} . This $\phi^{k_{final}}$ solution serves as our high-fidelity

ϕ_{true} .

Appendix C: Notes on Hyperparameter Tuning

The loss weights λ_{pde} and λ_{bc} are critical. If they are too small, the network only fits the data and produces a physically nonsensical field in the interior. If they are too large, the optimizer forces the PDE residual to zero *at the expense of* fitting the data, resulting in a poor solution that does not match the measurements.

The values $\lambda_{pde} = 200$ and $\lambda_{bc} = 50$ were found through a simple empirical grid search.

1. We first set $\lambda_{pde} = 1$, $\lambda_{bc} = 1$ and observed the magnitudes of the loss gradients.

2. We observed that the L_{pde} gradients were several orders of magnitude larger than the L_{data} gradients.

3. We manually adjusted the λ values to bring the *initial* gradient magnitudes into a similar range (e.g., within 1-2 orders of magnitude of each other).

This manual process is a known limitation and a primary target for improvement. More sophisticated methods, such as those described by Wang et al. [22], would automate this balancing act and likely lead to faster convergence and a lower final error.

References

1. M Hämäläinen, R Hari, RJ Ilmoniemi, J Knuutila, and OV Lounasmaa. Magnetoencephalography – theory, instrumentation, and applications to noninvasive studies of the working human brain. *Reviews of Modern Physics*, 65(2):413–497, 1993.
2. PL Nunez and R Srinivasan. *Electric Fields of the Brain: The Neurophysics of EEG*. Oxford University Press, 2006.
3. S Baillet, JC Mosher, and RM Leahy. Electromagnetic brain mapping. *IEEE Signal Processing Magazine*, 18(6):14–30, 2001.
4. R Grech, T Cassar, J Muscat, KP Camilleri, SG Fabri, M Zervakis, P Xanthopoulos, V Sakkalis, and B Vanrumste. Review on solving the inverse problem in eeg source analysis. *Journal of NeuroEngineering and Rehabilitation*, 5:25, 2008.
5. M Fuchs, M Wagner, and J Kastner. Source reconstruction and forward models for eeg and meg. *Brain Topography*, 16(3):145–158, 2002.
6. A Tarantola. *Inverse Problem Theory and Methods for Model Parameter Estimation*. SIAM, 2005.
7. O Hauk and M Stenroos. A comparison of eeg/meg source localization methods. *NeuroImage*, 147:14–28, 2017.
8. F Lucka, S Pursiainen, M Burger, and C Wolters. Bayesian inference for eeg and meg inverse problems. *Inverse Problems*, 28(5):055012, 2012.
9. M Dannhauer, B Lanfer, C Wolters, and TR Knösche. Head modeling strategies for eeg source localization. *Biomedical Engineering Online*, 10(1):83, 2011.
10. JC Mosher, RM Leahy, and PS Lewis. Eeg and meg forward solutions for inverse methods. *IEEE Transactions on Biomedical Engineering*, 46(3):245–259, 1999.
11. M Raissi, P Perdikaris, and GE Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.
12. GE Karniadakis, IG Kevrekidis, L Lu, P Perdikaris, S Wang, and L Yang. Physics-informed machine learning. *Nature Reviews Physics*, 3:422–440, 2021.
13. S Cai, Z Mao, Z Wang, M Yin, and GE Karniadakis. Physics-informed neural networks for forward and inverse problems in stochastic differential equations. *Journal of Computational Physics*, 425:109913, 2021.
14. Y Bai, Z Meng, and GE Karniadakis. Physics-informed neural networks for scientific computing: A comprehensive review. *Computer Methods in Applied Mechanics and Engineering*, 403:115671, 2023.
15. Y Yao, X Zhou, and GE Karniadakis. Pinns for anisotropic electrostatics and heterogeneous media. *Journal of Computational Physics*, 476:111896, 2023.
16. V Sitzmann, J Martel, A Bergman, D Lindell, and G Wetzstein. Implicit neural representations with periodic activation functions. *Advances in Neural Information Processing Systems*, 33:7462–7473, 2020.
17. F Tadel, S Baillet, JC Mosher, D Pantazis, and RM Leahy. Brainstorm: A user-friendly application for meg/eeg analysis. *Computational Intelligence and Neuroscience*, 2011:879716, 2011.
18. Y Liu, O Sourina, and MK Nguyen. Deep learning in eeg: Recent advances and new perspectives. *IEEE Transactions on Neural Systems and Rehabilitation Engineering*, 29:1–13, 2021.
19. H Yang, Y Sun, and M Zhou. Brain source imaging with physics-informed neural networks: methodology and challenges. *Frontiers in Neuroscience*, 16, 2022.
20. G Pang, L Lu, and GE Karniadakis. Physics-informed neural networks for electromagnetics. *IEEE Transactions on Neural Networks and Learning Systems*, 2021.

21. H Yang, Y Zhou, and S Li. Brainpinn: A physics-informed neural network for brain conductivity estimation from eeg data. *IEEE Access*, 9:132942–132954, 2021.
22. S Wang, H Teng, and P Perdikaris. Understanding and mitigating gradient pathologies in physics-informed neural networks. *SIAM Journal on Scientific Computing*, 44(6):A3647– A3670, 2022.
23. Y Huang, LC Parra, and S Haufe. Numerical modeling of electrical fields for transcranial brain stimulation – a review. *NeuroImage*, 140:1–15, 2016.
24. A Gramfort, M Luessi, E Larson, D Engemann, D Strohmeier, C Brodbeck, L Parkko- nen, and M H"am"al"ainen. Meg and eeg data analysis with mne-python. *Frontiers in Neuroscience*, 7:267, 2013.
25. B Moseley, A Markham, and T Nissen-Meyer. Finite basis physics-informed neural networks (fbpinns): A scalable domain decomposition approach for solving differential equations. *Advances in Neural Information Processing Systems*, 33:8471–8481, 2020.
26. CM Bishop. *Pattern Recognition and Machine Learning*. Springer, 2006.
27. I Goodfellow, Y Bengio, and A Courville. *Deep Learning*. MIT Press, 2016.
28. L Ruthotto and E Haber. Deep neural networks motivated by partial differential equations.
29. *Journal of Mathematical Imaging and Vision*, 62:352–364, 2020.
30. L Zhang, H Huang, and S Yao. Material modeling with physics-informed neural networks.
31. *npj Computational Materials*, 9:41, 2023.
32. J Carrasquilla and RG Melko. Machine learning phases of matter. *Nature Physics*, 13:431–434, 2017.
33. D Srivastava and P Kalra. Advances in inverse eeg source localization. *Clinical Neuro- physiology Practice*, 3:13–24, 2018.
34. J Mairal, F Bach, J Ponce, and G Sapiro. Sparse modeling for image and vision process- ing. *Foundations and Trends in Computer Graphics and Vision*, 8(2-3):85–283, 2014.
35. L Lu, X Meng, Z Mao, and GE Karniadakis. Deeponet: Learning nonlinear operators.
36. *Nature Machine Intelligence*, 3:218–229, 2021.
37. MD McKay, RJ Beckman, and WJ Conover. Comparison of three methods for selecting values of input variables in the analysis of output from a computer code. *Technometrics*, 21(2):239–245, 1979.
38. RM Neal. *Bayesian learning for neural networks*. Springer, 1996.
39. L Greengard and JY Lee. Accelerating the nonuniform fast fourier transform. *SIAM Review*, 46:443– 454, 2004.
40. A Grebet et al. Deep learning for eeg signal classification: A review. *Neural Computing and Applications*, 2019.
41. Q Liu et al. Deep learning for eeg and microstate analysis. *IEEE Access*, 2017.
42. Z Wang et al. A review of physics-informed neural networks and their applications.
43. *Applied Sciences*, 2019.
44. K Luo et al. Scaling pinns for inverse problems. *Journal of Computational Physics*, 2020.
45. Y Guo et al. Applications of pinns in medical imaging. *Medical Image Analysis*, 2022.
46. L Lu, X Meng, Z Mao, and GE Karniadakis. Deepxde: A deep learning library for solving differential equations. *SIAM Review*, 63(1):208–228, 2021.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.