

Electrical Impedance Tomography Shape Classification Using Deep Neural Networks

Yigermal Bassie, Jefferson Martins, Ricardo Suyama, Fernando Silva de Moura, Aline Neves, André Kazuo Takahata

Abstract—Electrical impedance tomography (EIT) is often formulated as an inverse problem of reconstructing internal conductivity distribution from boundary-voltage measurements. In contrast to conventional image reconstruction approaches, this work addresses a supervised task classification of the shape of a single inclusion directly from experimentally acquired EIT voltage data. The dataset consists of measurements obtained from a circular saline-filled phantom instrumented with 16 boundary electrodes and containing one polyethylene terephthalate glycol (PETG) inclusion with circular, square, or triangular geometry. A protocol-ordered voltage feature was used as input to a fully connected deep neural network classifier. The proposed model achieved approximately 97% test accuracy under the adopted experimental split, while a random forest baseline produced comparable performance. These findings show that measured EIT voltage measurements signals contain discriminative information for object-shape classification. However, they also indicate that the dataset is strongly separable and that future work may include additional phantom configurations, and comparisons with reconstruction-based and convolutional approaches.

Keywords—Shape classification, Electrical Impedance Tomography, Deep Neural Networks, phantom measurements, voltage-domain classification.

I. INTRODUCTION

Electrical Impedance Tomography (EIT) is a non-invasive imaging technique that estimates information about the internal conductivity distribution of a domain from electrical currents applied through boundary electrodes and the resulting boundary electrical voltage measurements [1], [2]. The basic EIT forward model is governed by the generalized Laplace equation, which is given by:

$$\nabla \cdot (\sigma(\mathbf{r}) \nabla V(\mathbf{r})) = 0, \quad (1)$$

where $\sigma(\mathbf{r})$, $V(\mathbf{r})$, and $\mathbf{r}$ denotes the conductivity distribution, electrical potential, and spatial coordinates, respectively. Estimating $\sigma(\mathbf{r})$ from boundary measurements is difficult because the inverse problem is nonlinear and ill-posed; small measurement errors or modeling inaccuracies can produce unstable reconstructions. For this reason, EIT image reconstruction methods often require regularization, accurate electrode modeling, and robust noise-handling strategies.

Due to its non-invasive nature, low cost, portability, and real-time imaging capability, EIT has attracted considerable

attention in biomedical imaging [1], industrial process monitoring [3], and geophysical exploration [4]. In biomedical applications, EIT has been explored for lung monitoring, breast imaging, and brain imaging because it does not expose patients to ionizing radiation and can be implemented using compact instrumentation [5]. Nevertheless, the computational and mathematical difficulties of EIT reconstruction continue to motivate alternative data-driven approaches for tasks such as classification, detection, and decision support.

Machine learning and deep learning methods have attracted increasing attention in EIT-related studies because they can learn nonlinear relationships directly from measured or simulated data [6], [7]. Particularly, Deep Neural Networks (DNNs) are able to extract nonlinear features through multiple trainable layers and have shown strong performance in high-dimensional signal-processing problems [6], [7], [8], [9]. In EIT, such models have been investigated for reconstruction support, anomaly detection, and classification tasks [10], [11], [12], [13], [14]. For instance, previous studies have used machine learning models for impedance-based target classification, stroke-type classification, and other EIT-related diagnostic problems [12], [13], [14]. However, it is important to distinguish image reconstruction from the task studied in this paper. The present work does not reconstruct an EIT image. Instead, it classifies the shape of a single inclusion directly from boundary-voltage measurements.

Convolutional Neural Networks (CNNs) are widely used when the input has an explicit spatial structure, such as images or reconstructed conductivity maps [15], [16]. In this work, however, the primary input is not an image. It is a fixed-length voltage vector of protocol-ordered EIT voltage measurements. Therefore, a fully connected DNN is adopted as a direct voltage-domain classifier. CNN-based approaches using reconstructed images or reshaped voltage maps remain relevant alternatives, and they are identified as future comparisons rather than presented as the main method of this study.

The experimental dataset used in this work was provided by Krukewitt and Thönes (2023), and contains of EIT voltage measurements acquired from a circular saline-filled phantom with 16 boundary electrodes [17]. Each measurement corresponds to a phantom containing one PETG inclusion whose shape is circle, square, or triangle. The classification objective is to map the measured voltage pattern to the corresponding inclusion shape. To the best of our knowledge, this is the first work employing this dataset for object shape classification using deep learning techniques. The proposed approach utilizes a fully connected DNN model trained on protocol-ordered voltage measurements to classify the embedded object shape.

Yigermal Bassie, Jefferson Martins, Ricardo Suyama, Fernando Silva de Moura, Aline Neves, André Kazuo Takahata are with the Engineering, Modeling and Applied Social Sciences Center, Federal University of ABC, SP, Brazil, e-mail: yigermal.bassie, jefferson.martins, ricardo.suyama, fernando.moura, aline.neves, andre.t@ufabc.edu.br. This work was financed by the São Paulo Research Foundation (FAPESP)-Process numbers: 2023/13729-0, 2022/03243-0, 2013/07559-3, 2021/11905-0, CAPES-Finance Code 001, and BIOS (FAPESP Grant # 2020/09838-0)

Unlike conventional EIT studies that focus on reconstructing conductivity images, the present work addresses a supervised classification problem directly from voltage-domain data. The results demonstrate that the measured voltage signals contain discriminative information for shape classification task.

The paper is organized as follows. In section II, the neural network architecture is explained, and the proposed sequential neural network architecture is also described. The descriptions of the the dataset, preprocessing steps, and training procedure baselines, and ablation protocol is presented in section III. In section IV, the discussion of the results is explained. Finally, we conclude by summarize the key findings in section V.

II. DEEP NEURAL NETWORK MODEL

Deep neural networks (DNNs) are machine learning models composed of multiple layers of trainable nonlinear transformations. Each layer maps its input to a new representation through weighted connections and activation functions, allowing the network to learn complex relationships directly from data [6]. DNNs have been applied to many high-dimensional signal-processing problems, including image classification [15], speech recognition [18], and EIT-related learning problems [19].

In this work, the objective is to classify the shape of a single inclusion directly from measured EIT voltage data. Each sample is represented by a feature vector $\mathbf{x} \in \mathbb{R}^{196}$, composed of 192 valid voltage measurements and four summary voltage features such as mean, standard deviation, maximum value, and minimum value. The classification problem is defined over the class set $\mathcal{C} = \{\text{circle, square, triangle}\}$. The hidden layers learn intermediate feature representations from the voltage vector, and the final layer maps these representations to the probability of each shape class.

The proposed architecture is shown in Fig. 1. It is a fully connected sequential DNN composed of dense layers with rectified linear unit (ReLU) activation, batch normalization, dropout, and L2 weight regularization. This architecture was selected because the input is a vector of voltage measurements rather than an image. Therefore, the network is designed as a practical, moderate-capacity configuration rather than as a theoretically optimal network. The first 256-unit block can form nonlinear combinations across all voltage channels, and the subsequent 128- and 64-unit blocks progressively compress the learned representation before the three-class output. ReLU activations provide nonlinearity, batch normalization supports stable optimization [20], and dropout and L2 penalties provide regularization. The output layer uses the softmax activation function to produce class probabilities for the three possible inclusion shapes. For a given sample, the output layer produces the probability assigned to class c as

$$\hat{y}_c = \frac{\exp(z_c)}{\sum_{j=1}^C \exp(z_j)}, \quad c = 1, \dots, C, \quad (2)$$

where z_c denotes the logit associated with class c , $\hat{y}_c$ is the predicted probability for that class, and $C = 3$ is the total number of classes. The predicted label is then selected as the class with the highest softmax probability. This architecture was




Fig. 1: Sequential fully connected DNN architecture used for object-shape classification from EIT voltage features.

selected because the input is a vector of voltage measurements rather than an image

III. METHODOLOGY

A. Experimental data and task definition

The dataset used in this study consists of experimental EIT measurements acquired from a saline-filled circular phantom equipped with 16 boundary electrodes [17]. A single PETG inclusion was placed inside the phantom, and the classification task was to identify whether the inclusion was a circle, a square, or a triangle. Figure 2 illustrates the physical setup considered in this work. It is important to note that this schematic is provided only to clarify the experimental setup; the proposed model receives measured boundary-voltage features as input rather than an image of the phantom. The measurements




Fig. 2: Schematic representation of the controlled phantom task. Only one PETG inclusion is present in each acquisition.

were obtained using an opposite current-injection protocol combined with adjacent voltage measurements. In this protocol, voltage values associated with current-injection electrodes are not used, resulting in protocol-induced missing entries. After excluding these entries and arranging the valid measurements in a consistent order, each sample contains 192 valid voltage values.

Let $V_j^{(k)}$ denote the voltage measured at channel j during repeated frame k , where $k = 1, \dots, 20$. The repeated frames are averaged to reduce the influence of random measurement noise. The mean voltage for channel j is computed as

$$\bar{V}_j = \frac{1}{20} \sum_{k=1}^{20} V_j^{(k)}, \quad j = 1, \dots, 192. \quad (3)$$

A channel denotes one retained adjacent voltage difference associated with one ordered current-injection pattern. The implemented 192-column representation can be organized as 16 injection blocks with 12 retained adjacent differences per block, giving $16 \times 12 = 192$ values. The four global descriptors are then appended, producing 196 input features. The dataset contains 3,293 samples for the circle class, 4,383 samples for the square class, and 4,383 samples for the triangle class. This distribution indicates moderate class imbalance, mainly due to the smaller number of circle samples. Figure 3 presents class-wise averages of the four descriptors. Figure 4 summarizes




Fig. 3: Class-wise averages of the per-sample voltage mean, standard deviation, maximum, and minimum. Each bar summarizes all samples of one class.

the protocol-ordered voltage measurements. The class means differ at several indices, but the bands also show substantial within-class variability. This observation supports the feasibility of supervised shape classification directly from EIT voltage measurements.




Fig. 4: Mean protocol-ordered voltage signatures for the three classes. Each solid curve is the mean across all samples in that class, and each translucent band represents the mean plus or minus one standard deviation. Thus, the figure summarizes multiple samples rather than a single acquisition. It is shown across both columns to make the overlapping mean curves and the large dispersion at selected indices easier to inspect.

B. Preprocessing

The preprocessing pipeline was designed to reduce data leakage and improve reproducibility. First, valid voltage measurements were concatenated into a 192-dimensional vector.

Second, four summary voltage features were added, giving a 196-dimensional feature vector. Third, the class labels were encoded for multiclass classification. Fourth, the dataset was split using stratified sampling: 15% was reserved for testing, and the remaining data were divided into training and validation sets, resulting in approximately 70% training, 15% validation, and 15% testing. Fifth, missing feature values, if present, were imputed using statistics estimated only from the training set. Sixth, feature standardization was fitted only on the training set and then applied to the validation and test sets. Finally, the Synthetic Minority Over-sampling Technique (SMOTE) was applied only to the training split to reduce the class imbalance [21].

C. Model training

The DNN was trained using supervised learning. The model parameters were optimized by minimizing categorical cross-entropy, which is suitable for multiclass classification with a softmax output layer [22]. For a single sample, the loss can be defined as [23]:

$$\mathcal{L}_{\text{CE}}(\mathbf{y}, \hat{\mathbf{y}}) = - \sum_{c=1}^C y_c \log(\hat{y}_c), \quad (4)$$

where $\mathbf{y}$ is the one-hot encoded ground-truth label, $\hat{\mathbf{y}}$ is the predicted probability vector, and $C = 3$. For a dataset containing M samples, the empirical cross-entropy loss is

$$\mathcal{L}_{\text{total}} = \frac{1}{M} \sum_{i=1}^M \mathcal{L}_{\text{CE}}(\mathbf{y}^{(i)}, \hat{\mathbf{y}}^{(i)}). \quad (5)$$

L2 regularization was included in the training objective:

$$\mathcal{L}_{\text{reg}} = \mathcal{L}_{\text{total}} + \frac{\lambda}{2} \sum_l \|\mathbf{W}^{(l)}\|_2^2, \quad (6)$$

where $\mathbf{W}^{(l)}$ is the weight matrix of layer l and $\lambda = 10^{-4}$ is the regularization coefficient. This term penalizes large weights and is commonly used to reduce overfitting. Its empirical effect in this study is evaluated in the ablation protocol rather than assumed a priori. Besides, the model parameters were optimized using the Adam optimizer [24]. The updated adaptive learning rates for each parameter can be computed by using

$$\theta_{t+1} = \theta_t - \eta \frac{\hat{m}_t}{\sqrt{\hat{v}_t + \epsilon}}, \quad (7)$$

where η is the learning rate, $\hat{m}_t$ and $\hat{v}_t$ are bias-corrected estimates of the first and second moments of the gradient, and ϵ is a small constant for numerical stability. The initial learning rate was set to 5×10^{-4} . Early stopping was used to stop training when the validation loss did not improve for 10 consecutive epochs, and the best model weights were restored.

D. Baselines and ablation protocol

To address whether the models can obtain comparable performance, three baseline classifiers were trained using the same feature representation and data split: multinomial logistic regression, linear SVM, and random forest. These baselines help clarify whether the high accuracy is specific to the DNN

architecture or whether the voltage features are already strongly class-discriminative.

An ablation study was also performed to examine the influence of selected training components. Starting from the full DNN configuration, three variants were evaluated: (i) DNN without SMOTE, (ii) DNN without dropout, and (iii) DNN without L2 regularization. This analysis supports a more careful interpretation of the regularization and balancing techniques used in the model.

IV. RESULTS AND DISCUSSION

A. Classification performance

The DNN model achieved approximately 97% test accuracy under the adopted sample-level stratified split. Table I summarizes the class-wise precision, recall, and F_1 -score. The results show high performance for all three classes, with the circle class being slightly easier to identify than the two polygonal classes. The results indicate that the voltage vectors contain

TABLE I: Class-wise DNN performance on the test set.

Class	Precision	Recall	F_1 -score
Circle	0.99	0.98	0.98
Square	0.96	0.97	0.96
Triangle	0.97	0.97	0.97

shape-dependent information that can be learned by supervised classifiers. Misclassifications occur mainly between square and triangle samples, which is physically plausible because both polygonal inclusions may produce more similar boundary-voltage perturbations than the circular inclusion. However, because the experiment uses a controlled phantom with one object at a time, these results should be interpreted as evidence of feasibility rather than as a general solution for EIT image interpretation.

B. Baseline comparison

Table II reports the results obtained by the baseline classifiers. The random forest achieved performance comparable to the DNN result. This result is important because it shows that the dataset is highly discriminative under the adopted split. Therefore, the main conclusion that voltage-domain classification is feasible and that both deep and classical machine learning models can exploit the information present in the measured voltage features.

TABLE II: Baseline models trained on the same voltage-domain feature representation.

Model	Accuracy	Macro F_1
Logistic regression	0.93	0.94
Linear SVM	0.94	0.94
Random forest	0.97	0.97
DNN	0.97	0.98

C. Ablation results

Table III presents the ablation results for selected DNN training components. Removing dropout reduced the accuracy more clearly than removing SMOTE or L2 regularization, suggesting that dropout contributed more strongly to generalization in this split. The no-L2 variant obtained performance similar to the full model, indicating that L2 regularization should not be claimed as a major source of improvement without further repeated experiments. Similarly, SMOTE produced a modest improvement, which is reasonable because the class imbalance is moderate rather than severe. These ablation results

TABLE III: Ablation study of the DNN training components.

DNN variant	Accuracy	Macro F_1
Full DNN	0.97	0.97
Without SMOTE	0.96	0.96
Without dropout	0.95	0.95
Without L2 regularization	0.96	0.96

support a more balanced interpretation of the model design. Dropout appears useful for this dataset, while SMOTE and L2 regularization have smaller effects. Because neural-network training can vary with random initialization and data splitting, stronger claims would require repeated runs and grouped evaluation protocols.

D. Learning curves

Figures 5 and 6 show the training and validation accuracy and loss curves for full DNNs model. The training accuracy increases steadily, while the validation accuracy stabilizes after the initial epochs. The validation curves are used only for model selection and early stopping; the model parameters are not updated using validation data. This distinction is important to avoid confusing validation monitoring with training adaptation.




Fig. 5: Training and validation accuracy across epochs.




Fig. 6: Training and validation loss across epochs.

The learning curves indicate stable convergence, although the gap between training and validation performance shows that overfitting remains a relevant concern. This justifies the use

of early stopping and regularization. It also motivates future evaluation under stricter data partitions, such as grouped splits by object position or measurement session.

E. Limitations

The present study has four main limitations. First, the dataset contains a controlled phantom configuration with only one inclusion at a time. Second, only three simple object shapes are considered. Third, the evaluation uses a sample-level stratified split; if repeated measurements from similar object positions appear in both training and test sets, performance may be optimistic. Fourth, the proposed model operates directly on voltage-domain features and is not compared with reconstruction-based classifiers or CNNs trained on reconstructed conductivity images. These limitations do not invalidate the feasibility demonstrated by the results, but they restrict the generality of the conclusions.

V. SUMMARY AND CONCLUSION

This paper investigated object-shape classification from experimental EIT boundary-voltage measurements using a fully connected DNN. Unlike conventional EIT image-reconstruction studies, the proposed approach directly classifies the shape of a single PETG inclusion from voltage-domain features. The dataset contains measurements from a circular saline-filled phantom with 16 boundary electrodes and three inclusion geometries, such as circle, square, and triangle.

The final feature representation consisted of 192 protocol-ordered voltage measurements augmented with four summary statistics. The DNN combined dense layers, ReLU activations, batch normalization, dropout, L2 regularization, and a softmax output layer. Under the adopted stratified split, the DNN achieved approximately 97% test accuracy. Baseline results showed that random forest performance was comparable, indicating that the voltage features are strongly discriminative for this controlled task. The ablation study further showed that dropout had the clearest positive effect among the evaluated training components, while SMOTE and L2 regularization produced smaller changes.

Overall, the results support the feasibility of voltage-domain EIT object-shape classification. However, they should not be interpreted as evidence that the DNN is generally superior to simpler classifiers or that the approach is immediately generalizable to more complex EIT scenarios. Future work should evaluate grouped splits by object position, include additional shapes and multiple inclusions, compare with CNNs trained on reconstructed or reshaped inputs, test noise robustness, and investigate physics-informed methods to improve interpretability and generalization.

ACKNOWLEDGEMENTS

This work was financed by the São Paulo Research Foundation (FAPESP)-Process numbers: 2023/13729-0, 2022/03243-0, 2013/07559-3, 2021/11905-0, CAPES-Finance Code 001, and BIOS (FAPESP Grant # 2020/09838-0)

REFERENCES

- [1] D. S. Holder, *Electrical impedance tomography: methods, history and applications*. CRC press, 2004.
- [2] A. Adler and A. Boyle, "Electrical impedance tomography: Tissue properties to image measures," *IEEE Transactions on Biomedical Engineering*, vol. 64, no. 11, pp. 2494–2504, 2017.
- [3] D. Constales, G. S. Yablonsky, and G. B. Marin, "Predicting kinetic dependences and closing the balance: Wei and prater revisited," *Chemical Engineering Science*, vol. 123, pp. 328–333, 2015.
- [4] W. Daily and A. Ramirez, "Electrical resistance tomography during in-situ trichloroethylene remediation at the savannah river site," *Journal of Applied Geophysics*, vol. 33, no. 4, pp. 239–249, 1995.
- [5] R. Bayford and A. Tizzard, "Bioimpedance imaging: an overview of potential clinical applications," *Analyst*, vol. 137, no. 20, pp. 4635–4643, 2012.
- [6] I. Goodfellow, Y. Bengio, A. Courville, and Y. Bengio, *Deep learning*. MIT press Cambridge, 2016, vol. 1, no. 2.
- [7] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, pp. 436–444, 2015.
- [8] M. E. Rayed, S. S. Islam, S. I. Niha, J. R. Jim, M. M. Kabir, and M. Mridha, "Deep learning for medical image segmentation: State-of-the-art advancements and challenges," *Informatics in Medicine Unlocked*, p. 101504, 2024.
- [9] M. Muzammul and X. Li, "Comprehensive review of deep learning-based tiny object detection: challenges, strategies, and future directions," *Knowledge and Information Systems*, pp. 1–89, 2025.
- [10] G. Wang, J. C. Ye, K. Mueller, and J. A. Fessler, "Image reconstruction is a new frontier of machine learning," *IEEE transactions on medical imaging*, vol. 37, no. 6, pp. 1289–1296, 2018.
- [11] T. Zhang, X. Tian, X. Liu, J. Ye, F. Fu, X. Shi, R. Liu, and C. Xu, "Advances of deep learning in electrical impedance tomography image reconstruction," *Frontiers in Bioengineering and Biotechnology*, vol. 10, p. 1019531, 2022.
- [12] J. Chen, S. Wang, K. Wang, P. Abiri, Z.-Y. Huang, J. Yin, A. M. Jabalera, B. Arianpour, M. Roustaei, E. Zhu *et al.*, "Machine learning-directed electrical impedance tomography to predict metabolically vulnerable plaques," *Bioengineering & translational medicine*, vol. 9, no. 1, p. e10616, 2024.
- [13] Y. Li, N. Wang, L.-F. Fan, P.-F. Zhao, J.-H. Li, L. Huang, and Z.-Y. Wang, "Robust electrical impedance tomography for biological application: A mini review," *Heliyon*, vol. 9, no. 4, 2023.
- [14] V. Candiani and M. Santacesaria, "Neural networks for classification of strokes in electrical impedance tomography on a 3d head model," *arXiv preprint arXiv:2011.02852*, 2020.
- [15] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," *Advances in neural information processing systems*, vol. 25, 2012.
- [16] A. Dhillon and G. K. Verma, "Convolutional neural network: a review of models, methodologies and applications to object detection," *Progress in Artificial Intelligence*, vol. 9, no. 2, pp. 85–112, 2020.
- [17] L. Krukewitt and J. P. Thönes, "Electrical impedance tomography dataset using the spectra bioimpedance and eit complete kit," Zenodo, 2023, <https://doi.org/10.5281/zenodo.7192883>.
- [18] G. Hinton, L. Deng, D. Yu, G. E. Dahl, A.-r. Mohamed, N. Jaitly, A. Senior, V. Vanhoucke, P. Nguyen, T. N. Sainath *et al.*, "Deep neural networks for acoustic modeling in speech recognition: The shared views of four research groups," *IEEE Signal processing magazine*, vol. 29, no. 6, pp. 82–97, 2012.
- [19] S. J. Hamilton and A. Hauptmann, "Deep d-bar: Real-time electrical impedance tomography imaging with deep neural networks," *IEEE transactions on medical imaging*, vol. 37, no. 10, pp. 2367–2377, 2018.
- [20] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*. pmlr, 2015, pp. 448–456.
- [21] N. V. Chawla, K. W. Bowyer, L. O. Hall, and W. P. Kegelmeyer, "Smote: synthetic minority over-sampling technique," *Journal of artificial intelligence research*, vol. 16, pp. 321–357, 2002.
- [22] TensorFlow, "Sparsecategoricalcrossentropy," https://www.tensorflow.org/api_docs/python/tf/keras/losses/SparseCategoricalCrossentropy, 2024, accessed: 2025-05-23.
- [23] Y. Ho and S. Wookey, "The real-world-weight cross-entropy loss function: Modeling the costs of mislabeling," *IEEE access*, vol. 8, pp. 4806–4813, 2019.
- [24] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization. international conference on learning representations (2015)," *San Diego, California*, 2015.