



Research article

Neural networks for classification of strokes in electrical impedance tomography on a 3D head model[†]

Valentina Candiani^{1,*} and Matteo Santacesaria²

¹ Department of Mathematics and Systems Analysis, Aalto University, P.O. Box 11100, FI-00076 Aalto, Espoo, Finland

² MaLGa Center, Department of Mathematics, University of Genoa, Via Dodecaneso 35, 16146 Genova, Italy

[†] **This contribution is part of the Special Issue:** Mathematical aspects of machine learning
Guest Editors: Ernesto De Vito; Lorenzo Rosasco; Silvia Villa
Link: www.aimspress.com/mine/article/6026/special-articles

* **Correspondence:** Email: valentina.candiani@aalto.fi.

Abstract: We consider the problem of the detection of brain hemorrhages from three-dimensional (3D) electrical impedance tomography (EIT) measurements. This is a condition requiring urgent treatment for which EIT might provide a portable and quick diagnosis. We employ two neural network architectures – a fully connected and a convolutional one – for the classification of hemorrhagic and ischemic strokes. The networks are trained on a dataset with 40 000 samples of synthetic electrode measurements generated with the complete electrode model on realistic heads with a 3-layer structure. We consider changes in head anatomy and layers, electrode position, measurement noise and conductivity values. We then test the networks on several datasets of unseen EIT data, with more complex stroke modeling (different shapes and volumes), higher levels of noise and different amounts of electrode misplacement. On most test datasets we achieve $\geq 90\%$ average accuracy with fully connected neural networks, while the convolutional ones display an average accuracy $\geq 80\%$. Despite the use of simple neural network architectures, the results obtained are very promising and motivate the applications of EIT-based classification methods on real phantoms and ultimately on human patients.

Keywords: electrical impedance tomography; classification of brain strokes; fully connected neural networks; convolutional neural networks; computational head model

1. Introduction

Electrical impedance tomography (EIT) is a noninvasive imaging modality for recovering information about the electrical conductivity inside a physical body from boundary measurements of current and potential. In practice, a set of contact electrodes is employed to drive current patterns into the object and the resulting electric potential is measured at (some of) the electrodes. The reconstruction process of EIT requires the solution of a highly nonlinear inverse problem on noisy data. This problem is typically ill-conditioned [1–3] and solution algorithms need either simplifying assumptions or regularization strategies based on a priori knowledge.

In recent years, *machine learning* has arisen as a data-driven alternative that has shown tremendous improvements upon the ill-posedness of several inverse problems [4–6]. It has been already successfully applied in EIT imaging [7–10]. The purpose of this work is to apply machine learning to the problem of classification of brain strokes from EIT data.

Stroke, a serious and acute cerebrovascular disease, and a leading cause of death, can be of two types: *hemorrhagic*, caused by blood bleeding into the brain tissue through a ruptured intracranial vessel, and *ischemic*, caused by vascular occlusion in the brain due to a blood clot (thrombosis). Visible symptoms are precisely the same in both cases, which makes it extremely difficult to differentiate them without advanced imaging modalities. Ischemic stroke can be treated through the use of thrombolytic (or clot-dissolving) agents within the first few hours [11], since human nervous tissue is rapidly lost as stroke progresses. On the other hand, thrombolytics are harmful, or even potentially fatal, to patients suffering from hemorrhagic stroke, thus it is essential to distinguish between the two types. A rapid and accurate diagnosis is crucial to triage patients [12, 13] to speed up clinical decisions concerning treatments and to improve the recovery of patients suffering from acute stroke [14].

In this work we study how to accelerate the diagnosis of acute stroke using EIT. Currently, stroke can be classified only by using expensive hospital equipment such as X-ray CT. On the contrary, an EIT device is cheap, compact and could be carried in an ambulance (even though measurements would need to be taken while the patient is not moving). The main challenge, for emergency use, is that data are collected at a single time frame: this excludes time-difference imaging [15] and leaves absolute and frequency-difference [16] imaging as the only options. Another important application, where measurements at different times are available, is bedside real-time monitoring of patients after the acute stage of stroke. In both scenarios, getting a full image reconstruction with the existing inversion algorithms is computationally heavy and time-consuming, and thus machine learning techniques can be used to expedite the process. In this work we focus on the case of absolute imaging, i.e., where EIT measurements are available at a single time frame and at a single frequency, which can be considered the most challenging scenario.

Although EIT for brain imaging has been studied for decades [16–23], there are only few recent results that employ machine learning algorithms for stroke classification. The work [24] proposes the use of both Support Vector Machines (SVM) and Neural Networks (NN) for detecting brain hemorrhages using EIT measurements in a 2-layer model of the head. The main weakness of the model, however, is that it does not take into account the highly resistive skull layer, which is known to have a shielding effect when it comes to EIT measurements. Also, only a finite set of head shapes is considered and the model lacks the ability to generalize to new sample heads. The more recent work [25] considers a 4-layer model for the heads and uses data from 18 human patients [26] that are

classified using SVM. The main difference with our method is that our classification is made directly from raw electrode data, while in [25] a preprocessing step involving a precise knowledge of the anatomy of the patient's head is required. Moreover, only strokes of size 20 ml or 50 ml are considered in four specific locations, while our datasets include strokes with volume as small as 1.5 ml located anywhere within the brain tissue. Another methodology is shown in [27], where first Virtual Hybrid Edge Detection (VHED) [28] is used to extract specific features of the conductivity, then neural networks are trained to identify the stroke type. This approach is very promising but currently limited to a 2D model. Applications of deep neural networks to EIT have been also considered in [7–9, 29–31]. One could also use some machine learning techniques to form a model for the head shapes: see [10] for an approach that could potentially be applicable to head imaging.

In this work we consider two different types of NN, a fully connected and a convolutional NN, that we feed with absolute EIT measurements and produce a binary output. These measurements, which form the training and test datasets, are simulated by using the so-called *complete electrode model* (CEM) [32, 33] on a computational 3-layer head model, where each layer corresponds to a different head anatomical region: scalp, skull and brain tissue.

The training and test datasets are made of pairs of simulated electrode measurements at a single time frame and a label which indicates whether the data are associated with a hemorrhagic stroke or not. More precisely, label 1 is meant to indicate a hemorrhagic stroke, while label 0 stands for either an ischemic stroke or no stroke. This is motivated by the fact that detecting the presence or absence of hemorrhage may be sufficient to initiate appropriate treatment. More precisely, the datasets contain EIT measurements in the following proportions: 50% hemorrhages, 25% ischemic stroke and 25% healthy brains. We chose to include a large number of healthy patients in order to cover a broader range of potential applications and not restrict ourselves only to the emergency setting. The measurements in the training and test datasets are generated by varying the conductivity distribution, the electrode positions and noise, the shape of the scalp, the skull and the brain tissue. We model a hemorrhage as a volume of the brain with higher conductivity values with respect to the brain tissue, and an ischemic stroke with lower values, based on the available medical literature [34, 35]. In the training dataset the strokes are modeled as a single ball inclusion of higher or lower conductivity, while in the test datasets we consider different shapes of multiple inclusions. Concerning the variations in the geometry, a joint principal component model for the variations in the anatomy of the human head [36] is considered, so that we are able to generate realistic EIT datasets for brain stroke classification on a 3D finite element (FE) head model. The training dataset is made of 40 000 pairs of electrode data and labels, while every test dataset is made of approximately 5 000 samples. No validation was used in the training of the fully connected network, while for the convolutional one the training set was randomly split (83% training, 17% validation). These test datasets take into account a variety of possible errors in the measurement setup: slight variations in the background conductivity and in the contact impedances are considered, along with misplacement of electrodes and mismodeling of the head shape. The functionality of the chosen methods is demonstrated via the measures of accuracy, sensitivity and specificity of the networks trained with noisy EIT data.

Our numerical tests show that the probability of detecting hemorrhagic strokes is reasonably high, even when the electrodes are misplaced with respect to their intended location and the geometric model for the head is inaccurate. We find that both fully connected neural networks and convolutional neural networks are efficient tools for the described classification. More precisely, in our experiments we

observe that a shallow fully connected neural network generalizes better to the test datasets than a convolutional one.

This paper is organized as follows. In Section 2 we recall the CEM and the parametrized head model with the workflow for mesh generation. Neural networks and their specifications are introduced in Section 3. Section 4 presents the experiment settings, while numerical results are described in Section 5. Finally, Section 6 lists the concluding remarks.

2. Generation of electrode data

2.1. Forward model

We start by recalling the CEM [33] of EIT. Let $\Omega \subset \mathbb{R}^3$ denote a bounded Lipschitz domain and assume there is a set of $M \in \mathbb{N} \setminus \{1\}$ contact electrodes E_1, \dots, E_M attached to its boundary $\partial\Omega$. When a single measurement by an EIT device is performed, net currents $I_m \in \mathbb{R}$, $m = 1, \dots, M$ are driven through each E_m and the resulting constant electrode potentials $U_m \in \mathbb{R}$, $m = 1, \dots, M$, are measured. Due to conservation of electric charge, any applicable current pattern $I = [I_1, \dots, I_M]^T$ belongs to the mean-free subspace

$$\mathbb{R}_\diamond^M := \left\{ J \in \mathbb{R}^M \mid \sum_{m=1}^M J_m = 0 \right\}.$$

The contact impedances at the electrode-object interfaces are modeled by a vector $z = [z_1, \dots, z_M]^T \in \mathbb{R}_+^M$. The electrode patches are identified with the nonempty, connected and open subsets of $\partial\Omega$ that they cover and assumed to be well separated, i.e., $\overline{E}_m \cap \overline{E}_l = \emptyset$ if $m \neq l$. We denote $E = \cup E_m$. The electromagnetic potential u inside Ω and the piecewise constant potentials on the electrodes U weakly satisfy

$$\begin{aligned} \nabla \cdot (\sigma \nabla u) &= 0 && \text{in } \Omega, \\ \nu \cdot \sigma \nabla u &= 0 && \text{on } \partial\Omega \setminus \overline{E}, \\ u + z_m \nu \cdot \sigma \nabla u &= U_m && \text{on } E_m, \quad m = 1, \dots, M, \\ \int_{E_m} \nu \cdot \sigma \nabla u \, dS &= I_m, && m = 1, \dots, M, \end{aligned} \tag{2.1}$$

where $\nu \in L^\infty(\partial\Omega, \mathbb{R}^3)$ denotes the exterior unit normal of $\partial\Omega$. Moreover, the isotropic conductivity distribution σ describing the electric properties of Ω is assumed to belong to

$$L_+^\infty(\Omega) := \{ \zeta \in L^\infty(\Omega) \mid \text{ess inf } \zeta > 0 \}. \tag{2.2}$$

A physical justification of (2.1) can be found in [32]. Given an input current pattern $I \in \mathbb{R}_\diamond^M$, a conductivity σ and contact impedances z with the properties described above, the pair $(u, U) \in H^1(\Omega) \oplus \mathbb{R}_\diamond^M$ is the unique solution of the elliptic boundary value problem (2.1) according to [32, 33]. Note that the use of \mathbb{R}_\diamond^M corresponds to systematically choosing the ground level of potential so that the mean of the electrode potentials is zero. The measurement, or current-to-voltage map of CEM is defined as the mapping $I \mapsto U$, from \mathbb{R}_\diamond^M to \mathbb{R}_\diamond^M .

2.2. Head model

The head model used in this work follows the same approach as in [36], though slightly modifying and upgrading the setting to a three-layers model. We define a *layer* for each one of the anatomical structures that we are considering for this particular head model. There are $L = 3$ different layers: the scalp layer, i.e., the outer one corresponding to the skin, the resistive skull layer and the interior brain layer (see Figure 1).

For each layer, the library of $n = 50$ heads from [37] is used to build the model for the variations in the shape and size of the human head. We can represent the crown of the l th layer in the j th head, for $l = 1, 2, 3$ and $j = 1, \dots, n$, as the graph of a function

$$S_j^l : \begin{cases} \mathbb{S}_+ \rightarrow \mathbb{R}^3, \\ \hat{x} \mapsto r_j^l(\hat{x}) \hat{x}, \end{cases} \quad (2.3)$$

where \mathbb{S}_+ is the upper unit hemisphere, i.e.,

$$\mathbb{S}_+ = \{x \in \mathbb{R}^3 \mid \|x\|_2 = 1 \text{ and } x_3 > 0\},$$

and $r_j^l : \mathbb{S}_+ \rightarrow \mathbb{R}_+$ gives the distance from the origin to the surface of the l th layer of the j th head as a function of the direction $\hat{x} \in \mathbb{S}_+$, where the origin is set at approximately the center of mass of each bottom face of the heads (see Figure 1).

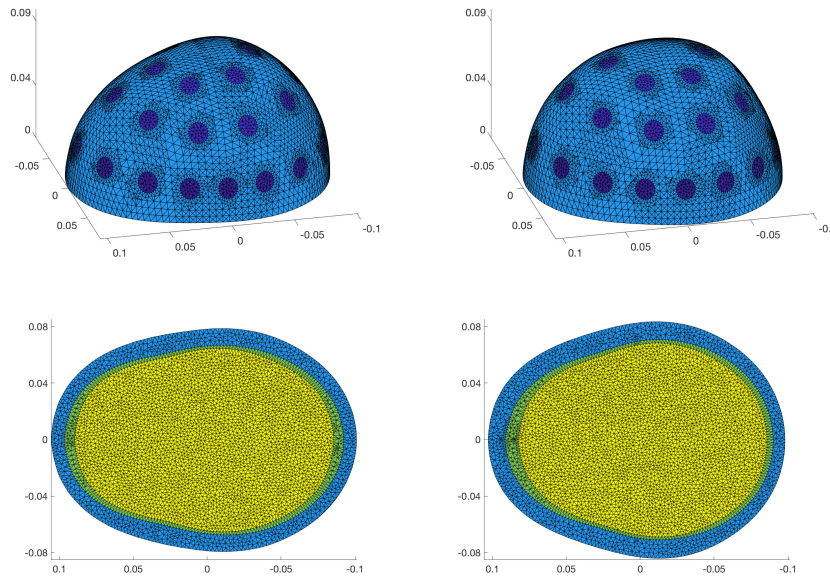


Figure 1. Top row: two different head models, oriented with the forehead on the left and the back of the head on the right. The 32 electrodes are at their intended positions, with the FE mesh associated to the head appropriately refined around them. Bottom row: the corresponding bottom face of each head model, where the three layers associated to scalp, skull and brain tissues are visible. Note that the first model has a more flattened forehead (top row) and it is narrower in the coronal plane (bottom row, y direction), corresponding to a shorter distance between the ears. Also, the thickness of the scalp layer is clearly different. The unit of length is meter.

Then, for each layer l , we introduce the average head and perturbations

$$\bar{r}^l = \frac{1}{n} \sum_{j=1}^n r_j^l \quad \text{and} \quad \rho_j^l = r_j^l - \bar{r}^l, \quad j = 1, \dots, n, \quad l = 1, 2, 3, \quad (2.4)$$

where \bar{r}^l describes the l th layer of the average head and $\rho_1^l, \dots, \rho_n^l$ are the corresponding perturbations that define the employed library of heads. We assume the functions $\rho_1^l, \dots, \rho_n^l$ belong to $H^1(\mathbb{S}_+)$ and are linearly independent for every l .

Mimicking the formulation in [36], we introduce a joint principal component model involving all the three layers. A single head in the library defines an object in the space $[H^1(\mathbb{S}_+)]^3$, that is, a three dimensional vector whose components are in the space $H^1(\mathbb{S}_+)$ and define the three layers. The reason for choosing $H^1(\mathbb{S}_+)$ is two-fold. First of all, according to the numerical tests in [36], the simplest option $s = 0$ in $H^s(\mathbb{S}_+)$ leads to undesired cusps in the optimal basis $\hat{\rho}_1, \dots, \hat{\rho}_n$ defined below. On the other hand, $s \in \mathbb{N} \setminus \{1\}$ would require the use of higher order conforming finite elements and the implementation of the needed inner products on \mathbb{S}_+ for $s \in \mathbb{R}_+ \setminus \mathbb{N}$ would lead to unnecessary technical considerations. The construction of the new principal component model for the head library is performed as in [36], with the exception that the inner product between two elements of $[H^1(\mathbb{S}_+)]^3$, say, v and w , is defined as

$$(v, w)_{[H^1(\mathbb{S}_+)]^3} := \sum_{i=1}^3 (v_i, w_i)_{H^1(\mathbb{S}_+)}. \quad (2.5)$$

Our aim is to look for an \tilde{n} -dimensional subspace $V_{\tilde{n}} \subset [H^1(\mathbb{S}_+)]^3$, $1 \leq \tilde{n} \leq n$, that satisfies

$$\sum_{j=1}^n \min_{\eta \in V_{\tilde{n}}} \|\rho_j - \eta\|_{[H^1(\mathbb{S}_+)]^3}^2 \leq \sum_{j=1}^n \min_{\eta \in W} \|\rho_j - \eta\|_{[H^1(\mathbb{S}_+)]^3}^2, \quad (2.6)$$

for all \tilde{n} -dimensional subspaces W of the Sobolev space $[H^1(\mathbb{S}_+)]^3$. The purpose is to find a low dimensional subspace that on average contains the best approximations for the perturbations $\{\rho_j = (\rho_j^1, \rho_j^2, \rho_j^3)\}_{j=1}^n$, where the quality of the fit is measured by the squared norm of $[H^1(\mathbb{S}_+)]^3$.

Following the approach in [36], we define the matrix R that takes into account the variations in every layer:

$$R_{ij} = \sum_{l=1}^3 (\rho_i^l, \rho_j^l)_{H^1(\mathbb{S}_+)} = (\rho_i, \rho_j)_{[H^1(\mathbb{S}_+)]^3}, \quad i, j = 1, \dots, n.$$

By applying Lemma 3.1 from [36] with minor modifications, we obtain the following set of orthonormal basis functions for $V_{\tilde{n}}$:

$$\hat{\rho}_k := (\hat{\rho}_k^1, \hat{\rho}_k^2, \hat{\rho}_k^3) \quad \text{with} \quad \hat{\rho}_k^l := w_k^T \rho^l \quad k = 1, \dots, \tilde{n}, \quad (2.7)$$

where

$$w_k = \frac{1}{\sqrt{\lambda_k}} v_k, \quad k = 1, \dots, \tilde{n}, \quad (2.8)$$

λ_k, v_k are eigenvalues and orthonormal eigenvectors of $R \in \mathbb{R}^{n \times n}$ and we have defined $\rho^l = [\rho_1^l, \dots, \rho_n^l]^T : \mathbb{S}_+ \rightarrow \mathbb{R}^n$, for $l = 1, 2, 3$. The positive eigenvalues λ_k are listed in descending order, and the corresponding l -dependent eigenvectors are employed in the definition of $\hat{\rho}_k^l$ for all $l = 1, 2, 3$.

The parametrization for the l th layer in our head model can then be written as

$$S^l(\hat{x}; \alpha) = \left(\bar{r}^l(\hat{x}) + \sum_{k=1}^{\tilde{n}} \alpha_k \hat{\rho}_k^l(\hat{x}) \right) \hat{x}, \quad \hat{x} \in \mathbb{S}_+, \quad l = 1, 2, 3, \quad (2.9)$$

where $\hat{\rho}_k^l$ are defined as in (2.7), α_k are free shape coefficients and $1 \leq \tilde{n} \leq n$ is chosen appropriately (cf. [36]). When generating random head structures for one numerical experiment, the vector of shape coefficients $\alpha \in \mathbb{R}^{\tilde{n}}$ is drawn from a Gaussian distribution $\mathcal{N}(0, \Gamma_\alpha)$, with the covariance matrix:

$$(\Gamma_\alpha)_{kh} = \frac{1}{\tilde{n} - 1} \sum_{j=1}^{\tilde{n}} (\rho_j, \hat{\rho}_k)_{[H^1(\mathbb{S}_+)]^3} (\rho_j, \hat{\rho}_h)_{[H^1(\mathbb{S}_+)]^3}; \quad (2.10)$$

see once again [36] for further details.

2.3. Generating the FEM mesh

Our workflow for generating a tetrahedral mesh for the head model consists of three steps: generation of an initial surface mesh, insertion of electrodes and tetrahedral mesh generation. The initial surface mesh is constructed by subdividing k times a coarse surface partition consisting of four triangles, where $k \in \mathbb{N}$ can be chosen by the operator of the algorithm, then M electrodes are inserted in the surface mesh following the process described in [36]. A (dense) mesh \mathcal{T}_m for the resulting polygonal domain is generated using the Triangle software [38]. After inserting all the electrodes, the process is completed by generating a tetrahedral partition for the whole volume by TetGen [39] starting from the formed surface mesh.

With the current head library, the average head size is approximately 20×16 cm in the axial plane, while its height is about 9 cm. The head shapes obtained by changing the shape parameters in (2.9) are variations from this average, up to a maximum of approximately 2 cm difference in each dimension. The thickness of the scalp varies within the range 10 – 20 mm, while the skull is about 2 – 15 mm thick. Throughout all the experiments, the number of electrodes is chosen to be $M = 32$ and we select $\tilde{n} = 10$.

Remark 2.1. *Despite being an upgrade with respect to the computational head model used in [36], this three-layer model is clearly still a simplified version of the true head anatomy. In particular, it does not take into account the shunting effect of the highly conductive cerebrospinal fluid (CSF) layer inside the skull. The CSF is known to represent a major challenge in EIT brain imaging, for it is extremely difficult to distinguish it from a bleed.*

3. Neural networks

There is a wide range of machine learning classification algorithms available in the literature. We chose to use neural networks (NN) since they consistently outperformed kernel methods for this specific kind of nonlinear dataset in our preliminary numerical tests.

We consider a 2-layer fully connected network and a convolutional neural network for the classification of brain strokes from electrode data. In both cases, as detailed in Section 4, the input is a vector of size $M(M - 1)$, which represents a single set of electrode measurements, where M is the

number of electrodes. The output layer is a single scalar value between 0 and 1. A rounding is then applied to the output in order to obtain a binary value for the classification. Since in our experiments $M = 32$, the input layer is composed of 992 neurons.

3.1. Fully connected neural networks

We consider a fully connected neural network (FCNN) that takes as input electrode data generated as discussed in Section 2 and gives a binary output: 1 for hemorrhage, 0 for no hemorrhage.

Our FCNN has two layers with weights, as shown in Figure 2. The input layer has 992 neurons, while the second and the third layer have 7 and 1 neurons, respectively. The size of the hidden layer has been chosen on the basis of the results obtained in preliminary tests on smaller datasets. The network can be represented by the following real-valued function:

$$f(x, \theta) = g(W_2(g(W_1x + b_1)) + b_2), \quad (3.1)$$

where we denoted by $\theta = \{W_1, W_2, b_1, b_2\}$ the set of weights and biases and

- $x \in \mathbb{R}^{992}$ is the input electrode data,
- $W_1 : \mathbb{R}^{992} \rightarrow \mathbb{R}^7$, $b_1 \in \mathbb{R}^7$ are the weights and the bias of the first layer,
- $W_2 : \mathbb{R}^7 \rightarrow \mathbb{R}$, $b_2 \in \mathbb{R}$ are the weights and the bias of the second layer,
- $g : \mathbb{R} \rightarrow \mathbb{R}$ is the sigmoid function $g(t) = 1/(1 + e^{-t})$ applied component-wise.

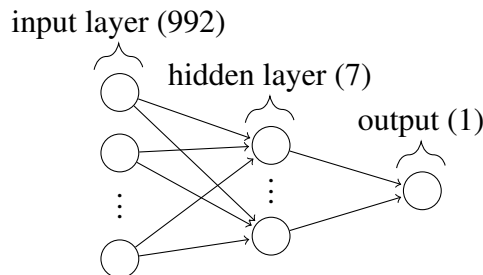


Figure 2. An illustration of the architecture of our FCNN, with one input layer with 992 nodes, one hidden layer with 7 nodes and one output layer with a single node for the binary classification.

The network $f(x, \theta)$ is trained by minimizing the cross-entropy loss

$$L(\theta) = - \sum_j \left(y_j \log(f(x_j, \theta)) + (1 - y_j) \log(1 - f(x_j, \theta)) \right), \quad (3.2)$$

where the sum is over the samples x_j in the training set, with y_j being the corresponding true binary label.

3.2. Convolutional neural networks

We also consider a convolutional neural network (CNN) for our classification task. We refer to [40] for more details on the architecture of a CNN.

As depicted in Figure 3, our network has six layers with weights. The first two are convolutional and the last four are fully-connected. As with the FCNN, our CNN is trained by minimizing the cross-entropy loss (3.2). The architecture was motivated by similar CNNs used in image classification [40], and the hyperparameters have been chosen after preliminary tests on smaller datasets.

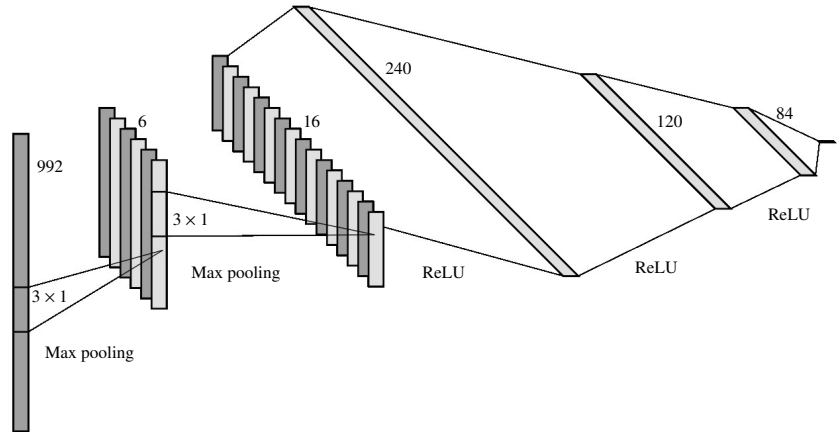


Figure 3. A visual illustration of the architecture of our CNN, with two convolutional layers and four fully connected layers.

Regarding the two convolutional layers, we chose to use 1D kernels. This choice might not be optimal since we are losing some geometric information about the electrode configuration. On the other hand, even though a single set of electrode measurements can be represented as a matrix of size 32×31 , there is no obvious advantage in considering it as an image.

In the first layer of our CNN, the 992×1 input vector is filtered with 6 kernels of size 3×1 with stride 1 and zero padding. Then a max-pooling layer, with kernel and stride of size 2 and zero padding, is applied to each output channel of the first layer. This is then filtered, in the second convolutional layer, with 16 kernels of size 3×1 (with stride 1 and zero padding). Then another max-pooling layer, with kernel and stride of size 2 and zero padding, is applied to the output. The third, fourth, fifth and sixth layers are fully connected and have sizes 240, 120, 84 and 1, respectively (see Section 3.1 for more details). The rectified linear unit (ReLU) activation function $g(x) = \max(0, x)$, is applied to the output of each fully connected layer, except the last one.

4. Experimental settings

4.1. Simulation of measurement data

We assume to be able to drive $M - 1$ linearly independent current patterns $I^{(1)}, \dots, I^{(M-1)} \in \mathbb{R}_{\diamond}^M$ through the M electrodes and measure the corresponding noisy electrode potentials $V^{(1)}, \dots, V^{(M-1)} \in \mathbb{R}^M$. The employed current patterns are of the form $I^{(j)} = e_p - e_j$, $j = 1, \dots, p-1, p+1, \dots, M$, where $p \in \{1, \dots, M\}$ is the label of the so-called current-feeding electrode and e_k denotes a standard basis vector. Such current patterns have been used in [36] and with real-world data in [41] and [42]. In our tests, the current-feeding electrode is always the frontal one on the top belt of electrodes (cf. Figure 4).

The potential measurements are stacked into a single vector

$$\mathcal{V} := [(V^{(1)})^T, \dots, (V^{(M-1)})^T]^T \in \mathbb{R}^{M(M-1)}, \quad (4.1)$$

and we analogously introduce the stacked forward map

$$\mathcal{U} : \mathbb{R}_+^N \times \mathbb{R}_+^M \times \mathbb{R}^{\bar{n}} \times (0, \pi/2)^M \times [0, 2\pi)^M \rightarrow \mathbb{R}^{M(M-1)}$$

via

$$\mathcal{U}(\sigma, z, \alpha, \theta, \phi) = [U(\sigma, z, \alpha, \theta, \phi; I^{(1)})^T, \dots, U(\sigma, z, \alpha, \theta, \phi; I^{(M-1)})^T]^T.$$

Here, the conductivity $\sigma \in \mathbb{R}_+^N$ is identified with the $N \in \mathbb{N}$ degrees of freedom used to parametrize it, i.e., the number of nodes in the mesh, the contact impedances are identified with the vector $z = [z_1, \dots, z_M] \in \mathbb{R}_+^M$, $\alpha \in \mathbb{R}^{\bar{n}}$ is the parameter vector in (2.9) determining the shape of the computational head model, and $\theta \in (0, \pi/2)^M$ and $\phi \in [0, 2\pi)^M$ define the polar and azimuthal angles of the electrode center points, respectively.

For each forward measurement, the parameters $\alpha \in \mathbb{R}^{\bar{n}}$ defining the shape of the head are drawn from the distribution $\mathcal{N}(0, \Gamma_\alpha)$, where $\Gamma_\alpha \in \mathbb{R}^{\bar{n} \times \bar{n}}$ is the diagonal covariance matrix defined componentwise by (2.10); for a motivation of this choice, as well as for a proof of Γ_α being diagonal, see the principal component construction in [36, Section 3.2].

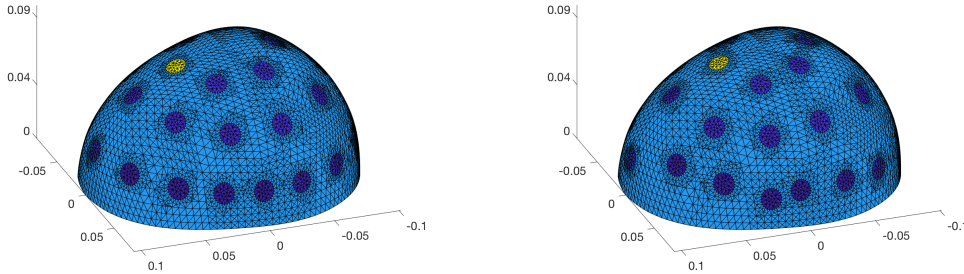


Figure 4. Left: a head model with electrodes at their intended positions. Right: the same head model with misplaced electrodes, where $\varsigma_\theta = \varsigma_\phi = 0.03$ in (4.2). Counting upwards from the bottom belt, there are altogether $M = 16 + 10 + 6 = 32$ electrodes of radius $R = 0.75$ cm. The current-feeding electrode is $p = 27$, i.e., the frontal one on the top belt of electrodes, highlighted in yellow. The FE mesh is refined appropriately around the electrodes. The unit of length is meter.

Every set of measurements is performed with $M = 32$ electrodes of radius $R = 0.75$ cm, organized in three belts around the head (see Figure 4). The expected values for the polar and azimuthal angles of the electrode centers, $\bar{\theta}$ and $\bar{\phi}$, correspond to the correct angular positions of the electrodes, i.e., the positions where one originally aims to attach the electrode patches. The actual central angles of the target electrodes, i.e., θ and ϕ , are then drawn from the distributions $\mathcal{N}(\bar{\theta}, \Gamma_\theta)$ and $\mathcal{N}(\bar{\phi}, \Gamma_\phi)$, where

$$\Gamma_\theta = \varsigma_\theta^2 \mathbb{I} \quad \text{and} \quad \Gamma_\phi = \varsigma_\phi^2 \mathbb{I}. \quad (4.2)$$

Here, $\mathbb{I} \in \mathbb{R}^{M \times M}$ is the identity matrix and $\varsigma_\theta, \varsigma_\phi > 0$ determine the standard deviations in the two angular directions. Notice that ς_θ and ς_ϕ must be chosen so small that the electrodes are not at a risk to overlap or move outside the crown of the computational head.

The relative contact impedances $z_m \in \mathbb{R}_+$, $m = 1, \dots, M$ are independently drawn from $\mathcal{N}(\bar{z}, \varsigma_z^2)$, where $\bar{z} > 0$ is chosen so much larger than $\varsigma_z > 0$ that negative contact impedances never occur in practice.

Finally, the conductivity distributions are defined as follows. For each layer (scalp, skull and brain) and stroke type (hemorrhagic or ischemic) we draw parameters $\sigma_{scalp}, \sigma_{skull}, \sigma_{brain}$ and σ_h or σ_i from Gaussian distributions of the form $\mathcal{N}(\bar{\sigma}, \varsigma_\sigma^2)$ (see Table 1). With these parameters we construct a conductivity that is constant on each layer and in the stroke region (if present). The final conductivity $\sigma \in \mathbb{R}_+^N$ corresponds to a piecewise linear parametrization on a dense FE mesh with $N \approx 20\,000$ nodes and about 85 000 tetrahedra associated to the head defined by the shape parameters α and refined appropriately around the electrodes, whose positions are determined by θ and ϕ . The number of principal components for the shape parameters $\alpha_1, \dots, \alpha_{\tilde{n}}$ (see Section 2.2) is chosen to be $\tilde{n} = 10$ in all the experiments.

Table 1. Parameter values related to the training set. These same values are also used in the simulation of the test sets if not explicitly mentioned otherwise.

Parameter values for the training set	
conductivity of the scalp σ_{scalp}	$\sim \mathcal{N}(0.4, 0.006)$ (S/m)
conductivity of the skull σ_{skull}	$\sim \mathcal{N}(0.06, 0.004)$ (S/m)
conductivity of the brain σ_{brain}	$\sim \mathcal{N}(0.2, 0.016)$ (S/m)
conductivity of hemorrhagic stroke σ_h	$\sim \mathcal{N}(1.8, 0.012)$ (S/m)
conductivity of ischemic stroke σ_i	$\sim \mathcal{N}(0.12, 0.008)$ (S/m)
radius of the inclusion r	$\sim \mathcal{U}(0.7, 2.3)$ (cm)
height of the cylindrical inclusion h	$\sim \mathcal{U}(1, 3)$ (cm)
standard deviation for electrode positions $\varsigma_\theta, \varsigma_\phi$	$\in \{0.01, 0.02, 0.03\}$ (rad)
contact impedances values z_m	$\sim \mathcal{N}(0.01, 0.001)$ (S/m ²)
shape parameters α	$\sim \mathcal{N}(0, \Gamma_\alpha)$
relative noise level ς_η	10^{-3}
stroke location	uniformly random within the brain tissue

We approximate the ideal data $\mathcal{U}(\sigma, z, \alpha, \theta, \phi)$ by FEM with piecewise linear basis functions and denote the resulting (almost) noiseless data by $\mathcal{U} \in \mathbb{R}^{M(M-1)}$. The actual noisy data is then formed as

$$\mathcal{V} = \mathcal{U} + \eta, \quad (4.3)$$

where $\eta \in \mathbb{R}^{M(M-1)}$ is a realization of a zero-mean Gaussian with the diagonal covariance matrix

$$\Gamma_\eta = \varsigma_\eta^2 \left(\max_j (\mathcal{U})_j - \min_j (\mathcal{U})_j \right)^2 \mathbb{I}, \quad (4.4)$$

with $\mathbb{I} \in \mathbb{R}^{M(M-1) \times M(M-1)}$ the identity matrix. The free parameter $\varsigma_\eta > 0$ can be tuned to set the relative noise level. Such a noise model has been used with real-world data, e.g., in [41].

4.2. Training dataset

The training dataset consists of pairs of input vectors and the corresponding output binary values. An input vector of length $M(M - 1) = 992$ contains the noisy electrode potentials related to a random choice of the parameter values for the head model presented in Section 2.2. The output is a categorical vector $\{0, 1\}$, with label = 1 indicating the presence of a hemorrhage and label = 0 no hemorrhage (there is no stroke or there is an ischemic one). This choice has been motivated by practitioners, since ruling out the presence of a hemorrhage allows them to start treating the patient immediately with blood-thinning medications.

The simulated stroke is represented as a single ball with varying location inside the brain tissue, varying volume and varying conductivity levels drawn from Gaussiann distributions. The chosen parameters $\bar{\sigma}$ and ζ_{σ} for all conductivity values are in line with conductivity levels reported in the medical literature [34, 35, 43–45] and are displayed in detail in Table 1. The radius of the ball defining the stroke is drawn from a uniform distribution $r \sim \mathcal{U}(r_{\min}, r_{\max})$, with $r_{\min} = 0.7$ cm and $r_{\max} = 2.3$ cm, which corresponds to volumes ranging from 1.5 ml to 50 ml. The inclusion center (x_c, y_c, z_c) is chosen randomly under the condition that the whole inclusion is contained in the brain tissue.

For the azimuthal and polar angles, we have drawn different standard deviations ζ_{θ} and ζ_{ϕ} (cf. Eq (4.2)) uniformly from the set $\{0.01, 0.02, 0.03\}$ (radians) for each forward computation, in order to take into account different levels of electrode movements. These might depend on, e.g., the initial misplacement of the electrode helmet on a patient’s head, the differences between the geometry of the patients head and of the helmet, as well as the overall movement of the patient during the examination. For a better understanding on how the selected standard deviation affects the electrode positions, see Figure 4, where both ζ_{θ} and ζ_{ϕ} are chosen to be 0.03. In particular, this value is the highest standard deviation one could use in our computational head model before the electrode patches start overlapping, especially on the lower belt, where they are closer.

The relative noise level in (4.4) is set to $\zeta_{\eta} = 10^{-3}$ (cf. [41], where such a noise level has been used with real-world data). A complete summary of the parameter values used in the random generation of the training data is reported in Table 1.

The training dataset contains around 40 000 samples, approximately split into 50% conductive inclusions (hemorrhage), 25% resistive inclusions (ischemia), 25% no inclusions (healthy). Let us emphasize that all samples, including those associated with no inclusions (strokes) correspond to different head and inner layer shapes, electrode positions and measurement noise realizations. This same statement applies also to all test datasets employed in assessing the performance of the neural networks.

The computations presented were performed with a MATLAB implementation using computer resources within the Aalto University School of Science “Science-IT” project. Measurements were computed in parallel over 200 different nodes on Triton [46], the Aalto University high-performance computing cluster, and the overall computation time for generating the training data did not exceed two hours.

4.3. Test datasets

While training is performed on the large dataset introduced above, testing the accuracy of the classifiers is realized on independent test sets, based on three different models for the stroke

geometries. More precisely, we constructed 14 test sets that originate from the parameter ranges of Tables 1 and 4 different variations, over three families of geometric models for the strokes.

The three models chosen for the test strokes are the following (cf. Figure 5).

- (1). A single ball: a test conductivity sample has one ball-shaped inclusion or no inclusion. The corresponding label is equal to 1 if the inclusion corresponds to a hemorrhagic stroke, 0 otherwise. We approximately have 50% conductive inclusions, 25% resistive inclusions, 25% no inclusion.
- (2). A single ball or cylinder: a test conductivity sample still exhibits one or no inclusions, but different shapes are considered; a single ball or a cylinder-shaped stroke. The height of the cylinders is uniformly drawn from the distribution $h \sim \mathcal{U}(h_{\min}, h_{\max})$, with $h_{\min} = 1$ cm and $h_{\max} = 3$ cm, while the radius is the same as for the corresponding balls, with the values reported in Table 1. This again corresponds to volumes ranging from 1.5 ml to 50 ml. The same labels as in the first case are used. As before, we have 50% conductive inclusions, 25% resistive inclusions and 25% healthy cases; the inclusion shape can be either a ball or a cylinder with a 50% chance.
- (3). Balls or cylinders: we consider cases with one, two or zero inclusions of different shapes (balls and cylinders). The label is equal to 1, if there is *at least* one hemorrhagic inclusion, and 0 if the nature of the inclusion(s) is resistive or the brain is healthy. Again, approximately half of the cases have at least one conductive inclusion, 25% of cases have one or two resistive inclusions, 25% have no inclusion. In particular, it is possible that label 1 corresponds to one hemorrhagic and one resistive inclusion (cf. Figure 5), which is a case not encountered in the training data.

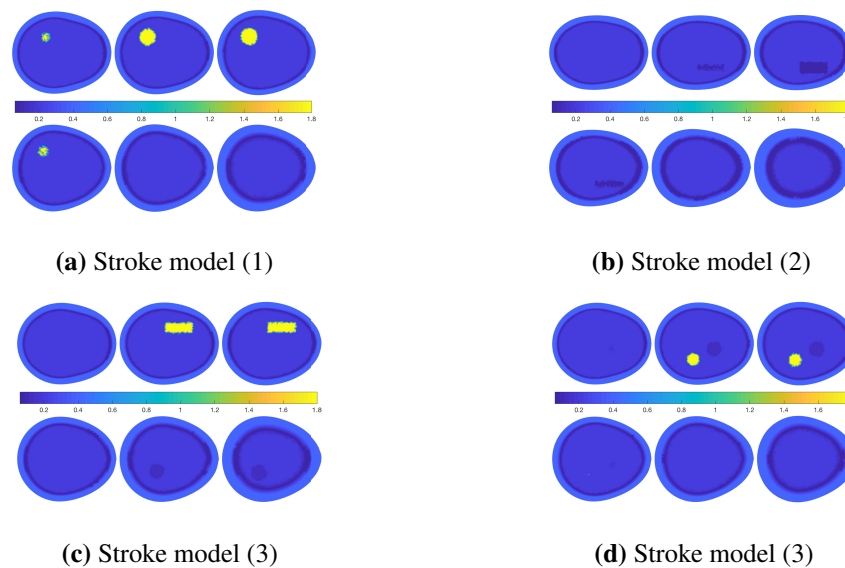


Figure 5. Four examples of six horizontal cross-sections of conductivities used in simulations of data samples in different datasets. The slices are at levels 2, 3, 4, 5, 6 and 7 cm. Top left: a conductivity sample from the training dataset, with one conductive ball (hemorrhagic stroke). Top right: a conductivity sample corresponding to the second geometric stroke model, with one resistive cylinder (ischemic stroke). Bottom row: two examples on the third geometric stroke model, the left panel shows one cylindrical conductive inclusion and one ball-shaped resistive inclusion, while the right panel displays two balls, one conductive and one resistive.

For each stroke model (1)–(3), we constructed 5 datasets for which we considered:

- standard parameters, i.e., the ones listed in Table 1,
- two different ranges for the expected inclusion radius,
- lower level of electrode misplacement,
- a higher amount of relative noise added to the data.

See Section 5 for more details. These alterations are to be understood in comparison to the random parameter models listed in Table 1 for the generation of the training data. This results in 14 different test sets with 5 000 samples each.

Remark 4.1. *In this setting, strokes are represented by a well defined ball or cylinder of a constant conductivity value embedded in a homogeneous background, whereas cases of nested inclusions are not considered. In fact, the presence of a penumbra or hypodense tissue would have a huge impact on the performance metrics. A penumbra is a region of normal to high blood volume that surrounds an ischemic stroke as the brain tries to balance the net blood pressure and flow. On the other hand, the hypodense tissue is formed around a hemorrhage when there is shortage of blood corresponding to a lower conductivity. Both situations are critical and detecting them is of vital importance, but we leave their investigation for future studies.*

5. Results

We start by first reviewing the computational details about the training and testing process. Next, we present and discuss the classification performance on each dataset.

For each classification learner we evaluate the results in terms of standard *performance metrics* such as sensitivity, specificity and accuracy on the test dataset. Sensitivity, or true positive rate, measures the proportion of actual positives that are correctly identified as such, while specificity, or true negative rate, quantifies the proportion of correctly classified actual negatives. Finally, accuracy is the sum of true positives and true negatives over the total test population, that is, the fraction of correctly classified cases.

5.1. Training the networks

The training dataset (see Section 4.2) was normalized by subtracting its mean and scaling by its standard deviation before the actual training. The mean and the standard deviation were stored and used to normalize the test datasets in the evaluation phase.

We trained our FCNN and CNN using the Adam optimizer [47] with batches of size 256 and learning rate 0.001. The FCNN was trained for 1500 epochs on the full training set with no validation. This choice was motivated by preliminary tests with validation that showed that our FCNN was not overfitting the training data. The CNN was trained for 120 epochs on a 83%/17% random training/validation splitting of the original dataset. In order to show the stability of the classification, the performance metrics for the FCNN and the CNN are reported as the average value computed over 10 different network trainings.

Training and tests were performed with a Python implementation on a laptop with 8GB RAM and an Intel CPU having clock speed 2.3 GHz. The overall training times were 30 and 40 minutes for FCNN and CNN respectively.

5.2. Results with fully connected neural network

Tables 2–4 summarize the results for the trained FCNN and each model for the stroke (1)–(3), respectively. In each table, the highest and lowest accuracy values are highlighted in bold and italics, respectively, while the values concerning the training set in Tables 2 and 5 are separated from test cases with a thicker line. Overall, the leading trend is that the network performs considerably well in the case with only one ball-like inclusion, while accuracy somewhat degrades when testing the network in the other two cases, especially for datasets with one or two inclusions of potentially different shapes and type.

The top row of Table 2 corresponds to testing the network on the training data and thus the resulting accuracy of 0.9673 is expected to give an upper limit for the performance of the FCNN. However, as shown below, we find that the network performs slightly better on some test datasets. On the top rows of Tables 3 and 4 we present the results for the FCNN on the test sets generated with the same parameters as the training set, but with different geometric models for the stroke. In every table, the results on the other rows are obtained by testing the trained network on four different test sets, where each time some parameters are altered from the values in Table 1.

The second rows of all three tables correspond to testing our classifier with strokes of larger volumes, where the lower bound for the ball radius is increased to $r_{\min} = 1$ cm, leading to the inclusion radius being drawn from $r \sim \mathcal{U}(1, 2.3)$ cm (4.20 ml – 50 ml for ball-shaped inclusions). In the case of a cylindrical inclusion, its radius is modified accordingly, while its height remains in the range [1, 3] cm, as in Table 1 (9.40 ml – 50 ml for cylinder-shaped inclusions). This choice of parameters leads to a better performance of the network, due to the better average visibility of the test inclusions to EIT measurements. Accuracy ranges from 91% in the case (3) with one or more inclusions of different shapes up to 97% accuracy for the test (1) with only one ball inclusion, which is actually higher than the accuracy on the training set.

Conversely, as shown on the third rows of each table, when only smaller inclusions are considered, the overall performance metrics decrease, with the lowest accuracy of 87% reached for the last geometric stroke model (3). Again, the height of the possible cylinder-shaped inclusions remains as in Table 1. It should also be noted that in Table 2, as well as in all our other tests, test datasets with larger and smaller inclusions are not simply subsets of the training set, but they have instead been simulated with random choices for the remaining parameters (cf. Table 1), providing hence new and unseen sets of data for the neural network.

In the experiments documented on the fourth rows, we considered a lower degree of electrode misplacement, that is, a lower level of inaccuracy in the electrode position when a helmet of electrodes is placed on a patient’s head. This translates to selecting a lower standard deviation in formula (4.2) when simulating the test data, with $\varsigma_{\theta}, \varsigma_{\phi}$ being drawn from {0.005, 0.01, 0.02} radians. Results in this case are notably better than those obtained with the baseline dataset for each of the three random stroke models (1)–(3), with the accuracy ranging from 91% to almost 97%. This confirms that electrode positioning does indeed significantly affect the classification accuracy.

Finally, to test the network in a more realistic setup, we increased the relative noise level in the measurements in (4.4) from $\varsigma_{\eta} = 10^{-3}$ to $\varsigma_{\eta} = 10^{-2}$ in the test data. The resulting performance indicators are listed on the last rows of Tables 2–4. Despite the (inevitable) decrease in performance, the accuracy was still over 90% for all three models for test stroke generation (1)–(3).

Table 2. Means and standard deviations for the performance metrics obtained after training our FCNN 10 times on the same 40 000-sample training dataset and testing it on different test sets of 5 000 samples with a single ball-shaped inclusion. Note that the first row corresponds to testing the FCNN on the training dataset.

FCNN: Datasets with a single ball-shaped stroke						
Datasets	Accuracy		Sensitivity		Specificity	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Training dataset	0.9673	0.0034	0.9753	0.0116	0.9600	0.0122
Larger radius $r \in [1, 2.3]$ cm	0.9740	0.0040	0.9660	0.0107	0.9825	0.0054
Smaller radius $r \in [0.7, 1.5]$ cm	0.9190	0.0081	0.9500	0.0181	0.8933	0.0243
Less el. uncert. $\varsigma_{\theta}, \varsigma_{\phi} \in \{0.5, 1, 2\} \cdot 10^{-2}$ rad	0.9684	0.0039	0.9769	0.0121	0.9609	0.0107
Increased relative noise $\varsigma_{\eta} = 10^{-2}$	0.9527	0.0033	0.9617	0.0129	0.9445	0.0124

Table 3. Performance metrics after testing the FCNN on test sets with a single ball-shaped or cylinder-shaped inclusion. Compared to the values in Table 2, accuracy slightly drops but the trend remains the same, with tests on dataset corresponding to inclusions with larger radii being the most accurate and the ones with smaller radii displaying only a 89% accuracy.

FCNN: Datasets with a single ball or cylinder-shaped stroke						
Datasets	Accuracy		Sensitivity		Specificity	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Standard parameters	0.9347	0.0074	0.9460	0.0093	0.9243	0.0148
Larger radius $r \in [1, 2.3]$ cm	0.9523	0.0062	0.9439	0.0101	0.9609	0.0116
Smaller radius $r \in [0.7, 1.5]$ cm	0.8988	0.0128	0.9398	0.0110	0.8644	0.0240
Less el. uncert. $\varsigma_{\theta}, \varsigma_{\phi} \in \{0.5, 1, 2\} \cdot 10^{-2}$ rad	0.9412	0.0066	0.9508	0.0097	0.9323	0.0142
Increased relative noise $\varsigma_{\eta} = 10^{-2}$	0.9335	0.0081	0.9454	0.0100	0.9231	0.0155

Table 4. Performance metrics after testing the FCNN on test sets with ball-shaped or cylinder-shaped inclusions. Levels of accuracy are lower compared to Tables 2 and 3 but, once again, the overall trend remains unchanged.

FCNN: Datasets with balls or cylinders-shaped strokes						
Datasets	Accuracy		Sensitivity		Specificity	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Standard parameters	0.9038	0.0088	0.8954	0.0210	0.9142	0.0182
Larger radius $r \in [1, 2.3]$ cm	0.9135	0.0093	0.8945	0.0204	0.9351	0.0151
Smaller radius $r \in [0.7, 1.5]$ cm	0.8793	0.0108	0.8867	0.0237	0.8740	0.0205
Less el. uncert. $\varsigma_{\theta}, \varsigma_{\phi} \in \{0.5, 1, 2\} \cdot 10^{-2}$ rad	0.9166	0.0102	0.9055	0.0235	0.9297	0.0145
Increased relative noise $\varsigma_{\eta} = 10^{-2}$	0.9020	0.0094	0.8933	0.0200	0.9127	0.0171

5.3. Results with convolutional neural network

Following the same workflow as in Section 5.2, Tables 5–7 present the analogous results for the CNN trained on the training dataset and tested in the three cases related to the different geometric models for the inclusions. It can clearly be seen that the overall accuracy is significantly lower than in the case of FCNN. This is arguably due to the fact that our CNNs tend to overfit the data, resulting in a high accuracy for the training set (first row of Table 5), which significantly degrades in the other test cases. With the exception of the case of the smaller strokes, the overall accuracy remains steadily over 80% and roughly follows the same trends as for the FCNN. Tests in the case of larger inclusion volumes have by far the best accuracy level, ranging from 84% to almost 89%, while the least accurate classification is indeed obtained for the strokes with smaller volumes. The mean accuracy with a lower level of electrode uncertainty and increased relative noise is between 81% and 86%, depending on the model for generating the test strokes.

Table 5. Means and standard deviations for the performance metrics obtained after training our CNN 10 times on the same 40 000-sample training dataset and testing it on different test sets of 5 000 samples with a single ball-shaped inclusion. Note that the first row corresponds to testing the CNN on the training dataset.

CNN: Datasets with a single ball-shaped stroke						
Datasets	Accuracy		Sensitivity		Specificity	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Training dataset	0.9593	0.0148	0.9752	0.0215	0.9479	0.0383
Larger radius $r \in [1, 2.3]$ cm	0.8892	0.0096	0.8747	0.0399	0.9110	0.0278
Smaller radius $r \in [0.7, 1.5]$ cm	0.7589	0.0237	0.8435	0.0401	0.7152	0.0452
Less el. uncert. $\varsigma_{\theta}, \varsigma_{\phi} \in \{0.5, 1, 2\} \cdot 10^{-2}$ rad	0.8611	0.0040	0.8823	0.0398	0.8489	0.0329
Increased relative noise $\varsigma_{\eta} = 10^{-2}$	0.8448	0.0041	0.8664	0.0403	0.8324	0.0338

Table 6. Performance metrics after testing the CNN on test sets with a single ball-shaped or cylinder-shaped inclusion. Compared to the values in Table 5, accuracy remains somewhat stable, with tests on the dataset corresponding to inclusions with larger radii being the most accurate and the ones with smaller radii displaying only a 76% accuracy.

CNN: Datasets with a single ball or cylinder-shaped stroke						
Datasets	Accuracy		Sensitivity		Specificity	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Standard parameters	0.8470	0.0057	0.8677	0.0299	0.8328	0.0294
Larger radius $r \in [1, 2.3]$ cm	0.8756	0.0045	0.8660	0.0322	0.8902	0.0300
Smaller radius $r \in [0.7, 1.5]$ cm	0.7608	0.0234	0.8465	0.0275	0.7120	0.0405
Less el. uncert. $\varsigma_{\theta}, \varsigma_{\phi} \in \{0.5, 1, 2\} \cdot 10^{-2}$ rad	0.8508	0.0049	0.8776	0.0280	0.8312	0.0282
Increased relative noise $\varsigma_{\eta} = 10^{-2}$	0.8477	0.0034	0.8646	0.0322	0.8378	0.0289

Table 7. Performance metrics after testing the CNN on test sets with ball-shaped or cylinder-shaped inclusions. Levels of accuracy are lower compared to Tables 5 and 6 but, with the exception of the case of smaller strokes, the overall accuracy remains steadily over 80%.

CNN: Datasets with balls or cylinders-shaped strokes						
Datasets	Accuracy		Sensitivity		Specificity	
	Mean	Std Dev	Mean	Std Dev	Mean	Std Dev
Standard parameters	0.8181	0.0247	0.7826	0.0504	0.8750	0.0203
Larger radius $r \in [1, 2.3]$ cm	0.8446	0.0304	0.7916	0.0514	0.9292	0.0163
Smaller radius $r \in [0.7, 1.5]$ cm	<i>0.7510</i>	0.0103	0.7548	0.0486	0.7590	0.0310
Less el. uncert. $\varsigma_{\theta}, \varsigma_{\phi} \in \{0.5, 1, 2\} \cdot 10^{-2}$ rad	0.8362	0.0238	0.7991	0.0483	0.8918	0.0178
Increased relative noise $\varsigma_{\eta} = 10^{-2}$	0.8144	0.0261	0.7840	0.0523	0.8646	0.0192

6. Conclusions

This work applies neural networks to the detection of brain hemorrhages from simulated absolute EIT data on a 3D head model. We developed large datasets based on realistic anatomies that we used to train and test a fully connected neural network and a convolutional neural network. Our classification tests show encouraging results for further development of these techniques, with fully connected neural networks achieving an average accuracy higher than 90% on most test datasets. Since the datasets included a large proportion of healthy patents, this work could motivate the development of a new, cheap, non-invasive screening test for brain hemorrhages.

The results demonstrate that the classification performance is affected by many factors, including the size of the stroke, the mismodeling of the electrode position and the noise in the data. Among these, the size of the stroke is the one most significantly altering the performance, which indicates that the method is probably unable to detect very small hemorrhages. Note that in the dataset generation we considered strokes with volumes of as small as 1.5 ml. Further, the FCNN, trained on single ball-shaped strokes, was able to generalize to data generated from multiple strokes of different shapes with little loss in the accuracy. We also want to stress that in every dataset the position of the stroke was chosen uniformly at random within the brain tissue. This means that on the one hand we did not restrict or assume to know in which brain hemisphere the stroke was located. On the other hand we considered strokes taking place potentially very close to the skull layer or very deep inside the brain tissue, a situation which is very challenging to analyze from EIT data.

The simulated datasets have several limitations nonetheless. The three-layer head model is a simplified version of the true head anatomy. In particular, it does not take into account the cerebrospinal fluid and other subtle anatomical tissues. Besides, the considered conductivity distributions have an average skull-to-brain ratio of 3 : 10, which has been shown to be a difficult quantity to estimate [44, 45] and appears to affect the performance metrics. A brain stroke is a complex phenomenon which evolves with time, and it is so far unclear how to precisely model its conductivity values, both in the hemorrhagic and in the ischemic case.

This work leaves many open directions for future studies. A natural one is to consider more realistic datasets. To our knowledge, the only publicly available dataset with real EIT data on human

patients has been released by University College London [26]. The dataset comprises data from only 18 patients, which is far from the size of our training set (40 000 samples), but could be used as a test set after training on synthetic data. Datasets from phantom data would also be an important intermediate step to validate the generation of synthetic measurements. On a related note, cases of nested inclusions should also be taken into account.

Another important direction is to implement more refined machine learning algorithms. We found that a simple 2-layer FCNN outperforms a more sophisticated CNN, though we did not investigate in depth the problem of designing an optimal network architecture for our EIT data. This certainly leaves room for improvements in the performance.

Finally, bedside real-time monitoring of stroke patients could become another valuable application of the presented techniques. The aim would be to predict whether an hemorrhage is increasing in volume over time or it is stable. This could be studied with a similar approach based on machine learning algorithms trained on datasets with several measurements taken at different times.

Acknowledgments

This work has been partially carried out at the Machine Learning Genoa (MaLGa) center, Università di Genova (IT). This material is based upon work supported by the Air Force Office of Scientific Research under award number FA8655-20-1-7027. MS is member of the “Gruppo Nazionale per l’Analisi Matematica, la Probabilità e le loro Applicazioni” (GNAMPA), of the “Istituto Nazionale per l’Alta Matematica” (INdAM). This work was also supported by the Academy of Finland (decision 312124) and the Finnish Cultural Foundation (grant number 00200214). The authors thank Nuutti Hyvönen (Aalto University) for useful discussions, Antti Hannukainen (Aalto University) for help with his FE solver and all the members of the project “Stroke classification and monitoring using Electrical Impedance Tomography”, funded by Jane and Aatos Erkkö Foundation, that motivated this work.

Conflict of interest

All authors declare no conflicts of interest in this paper.

References

1. L. Borcea, Electrical impedance tomography, *Inverse Probl.*, **18** (2002), R99–R136.
2. M. Cheney, D. Isaacson, J. C. Newell, Electrical impedance tomography, *SIAM Rev.*, **41** (1999), 85–101.
3. G. Uhlmann, Electrical impedance tomography and Calderón’s problem, *Inverse Probl.*, **25** (2009), 123011.
4. S. Arridge, P. Maass, O. Öktem, C.-B. Schönlieb, Solving inverse problems using data-driven models, *Acta Numer.*, **28** (2019), 1–174.
5. M. T. McCann, K. H. Jin, M. Unser, Convolutional neural networks for inverse problems in imaging: A review, *IEEE Signal Proc. Mag.*, **34** (2017), 85–95.

6. A. Lucas, M. Iliadis, R. Molina, A. K. Katsaggelos, Using deep neural networks for inverse problems in imaging: beyond analytical methods, *IEEE Signal Proc. Mag.*, **35** (2018), 20–36.
7. S. J. Hamilton, A. Hauptmann, Deep D-bar: Real-time electrical impedance tomography imaging with deep neural networks, *IEEE Trans. Med. Imaging*, **37** (2018), 2367–2377.
8. S. J. Hamilton, A. Hänninen, A. Hauptmann, V Kolehmainen, Beltrami-Net: Domain-independent deep D-bar learning for absolute imaging with electrical impedance tomography (a-EIT), *Physiol. Meas.*, **40** (2019), 074002.
9. X. Y. Li, Y. Zhou, J. M. Wang, Q. Wang, Y. Lu, X. J. Duan, et al., A novel deep neural network method for electrical impedance tomography. *T. I. Meas. Control*, **41** (2019), 4035–4049.
10. J. K. Seo, K. C. Kim, A. Jargal, K. Lee, B. Harrach, A learning-based method for solving ill-posed nonlinear inverse problems: A simulation study of lung EIT, *SIAM J. Imaging Sci.*, **12** (2019), 1275–1295.
11. W. Hacke, M. Kaste, E. Bluhmki, M. Brozman, A. Dávalos, D. Guidetti, et al., Thrombolysis with alteplase 3 to 4.5 hours after acute ischemic stroke, *New Engl. J. Med.*, **359** (2008), 1317–1329.
12. T. Dowrick, C. Blochet, D. Holder, In vivo bioimpedance measurement of healthy and ischaemic rat brain: implications for stroke imaging using electrical impedance tomography, *Physiol. Meas.*, **36** (2015), 1273–1282.
13. L. Yang, W. B. Liu, R. Q. Chen, G. Zhang, W. C. Li, F. Fu, et al., In vivo bioimpedance spectroscopy characterization of healthy, hemorrhagic and ischemic rabbit brain within 10 Hz–1 MHz, *Sensors (Basel)*, **17** (2017), 791.
14. J. L. Saver, Time is brain—quantified, *Stroke*, **37** (2006), 263–266.
15. D. C. Barber, B. H. Brown, Applied potential tomography, *J. Phys. E: Sci. Instrum.*, **17** (1984), 723–733.
16. A. McEwan, A. Romsauerova, R. Yerworth, L. Horesh, R. Bayford, D. Holder, Design and calibration of a compact multi-frequency EIT system for acute stroke imaging, *Physiol. Meas.*, **27** (2006), S199–210.
17. D. Holder, Electrical impedance tomography (EIT) of brain function, *Brain Topogr.*, **5** (1992), 87–93.
18. L. Fabrizi, A. McEwan, T. Oh, E. J. Woo, D. S. Holder, An electrode addressing protocol for imaging brain function with electrical impedance tomography using a 16-channel semi-parallel system, *Physiol. Meas.*, **30** (2009), S85–101.
19. E. Malone, M. Jehl, S. Arridge, T. Betcke, D. Holder, Stroke type differentiation using spectrally constrained multifrequency EIT: evaluation of feasibility in a realistic head model, *Physiol. Meas.*, **35** (2014), 1051–1066.
20. A. Nissinen, J. P. Kaipio, M. Vauhkonen, V. Kolehmainen, Contrast enhancement in EIT imaging of the brain, *Physiol. Meas.*, **37** (2015), 1–24.
21. L. Yang, C. H. Xu, M. Dai, F. Fu, X. T. Shi, X. Z. Dong, A novel multi-frequency electrical impedance tomography spectral imaging algorithm for early stroke detection, *Physiol. Meas.*, **37** (2016), 2317–2335.

22. B. McDermott, M. O'Halloran, J. Avery, E. Porter, Bi-frequency symmetry difference EIT-feasibility and limitations of application to stroke diagnosis, *IEEE J. Biomed. Health Inform.*, **24** (2020), 2407–2419.
23. V. Kolehmainen, M. J. Ehrhardt, S. Arridge, Incorporating structural prior information and sparsity into EIT using parallel level sets, *Inverse Probl. Imag.*, **13** (2019), 285.
24. B. McDermott, M. O'Halloran, E. Porter, A. Santorelli, Brain haemorrhage detection using a SVM classifier with electrical impedance tomography measurement frames, *PLoS ONE*, **13** (2018), e0200469.
25. B. McDermott, A. Elahi, A. Santorelli, M. O'Halloran, J. Avery, E. Porter, Multi-frequency symmetry difference electrical impedance tomography with machine learning for human stroke diagnosis, *Physiol. Meas.*, **41** (2020), 075010.
26. N. Goren, J. Avery, T. Dowrick, E. Mackle, A. Witkowska-Wrobel, D. Werring, et al., Multi-frequency electrical impedance tomography and neuroimaging data in stroke patients, *Sci. Data*, **5** (2018), 180112.
27. J. P. Agnelli, A. Çöl, M. Lassas, R. Murthy, M. Santacesaria, S. Siltanen, Classification of stroke using neural networks in electrical impedance tomography, *Inverse Probl.*, **36** (2020), 115008.
28. A. Greenleaf, M. Lassas, M. Santacesaria, S. Siltanen, G. Uhlmann, Propagation and recovery of singularities in the inverse conductivity problem, *Anal. PDE*, **11** (2018), 1901–1943.
29. A. Adler, R. Guardo, A neural network image reconstruction technique for electrical impedance tomography, *IEEE Trans. Med. Imaging*, **13** (1994), 594–600.
30. M. Capps, J. L. Mueller, Reconstruction of organ boundaries With deep learning in the D-Bar method for electrical impedance tomography, *IEEE Trans. Biomed. Eng.*, **68** (2021), 826–833.
31. J. Lampinen, A. Vehtari, K. Leinonen, Using Bayesian neural network to solve the inverse problem in electrical impedance tomography, In: *Proceedings of 11th Scandinavian Conference on Image Analysis SCIA'99*, 1999, 87–93.
32. K. S. Cheng, D. Isaacson, J. S. Newell, D. G. Gisser, Electrode models for electric current computed tomography, *IEEE Trans. Biomed. Eng.*, **36** (1989), 918–924.
33. E. Somersalo, M. Cheney, D. Isaacson, Existence and uniqueness for electrode models for electric current computed tomography, *SIAM J. Appl. Math.*, **52** (1992), 1023–1040.
34. J. A. Latikka, J. A. Hyttinen, T. A. Kuurne, H. J. Eskola, J. A. Malmivuo, The conductivity of brain tissue: Comparison of results in vivo and in vitro measurement, In: *Proceedings of the 23rd Annual International Conference of the IEEE Engineering in Medicine and Biology Society*, IEEE, Istanbul, Turkey, 2001, 910–912.
35. H. McCann, G. Pisano, L. Beltrachini, Variation in reported human head tissue electrical conductivity values, *Brain Topogr.*, **32** (2019), 825–858.
36. V. Candiani, A. Hannukainen, N. Hyvönen, Computational framework for applying electrical impedance tomography to head imaging, *SIAM J. Sci. Comput.*, **41** (2019), B1034–B1060.
37. E. G. Lee, W. Duffy, R. L. Hadimani, M. Waris, W. Siddiqui, F. Islam, et al., Investigational effect of brain-scalp distance on the efficacy of transcranial magnetic stimulation treatment in depression, *IEEE T. Magn.*, **52** (2016), 1–4.

38. J. R. Shewchuk, Triangle: Engineering a 2D quality mesh generator and Delaunay triangulator, In: M. C. Lin, D. Manocha, Editors, *Applied computational geometry: towards geometric engineering*, Berlin: Springer-Verlag, 1996, 203–222.
39. S. Hang, TetGen, a Delaunay-based quality tetrahedral mesh generator, *ACM T. Math. Software*, **41** (2015), 1–36.
40. Y. LeCun, K. Kavukcuoglu, C. Farabet, Convolutional networks and applications in vision, In: *Proceedings of 2010 IEEE international symposium on circuits and systems*, IEEE, 2010, 253–256.
41. J. Dardé, N. Hyvönen, A. Seppänen, S. Staboulis, Simultaneous recovery of admittivity and body shape in electrical impedance tomography: an experimental evaluation, *Inverse Probl.*, **29** (2013), 085004.
42. J. Kourunen, T. Savolainen, A. Lehtikoinen, M. Vauhkonen, L. M. Heikkinen, Suitability of a PXI platform for an electrical impedance tomography system, *Meas. Sci. Technol.*, **20** (2008), 015503.
43. S. Gabriel, R. W. Lau, C. Gabriel, The dielectric properties of biological tissues: II. Measurements in the frequency range 10 Hz to 20 GHz, *Phys. Med. Biol.*, **41** (1996), 2251.
44. Y. Lai, W. Van Drongelen, L. Ding, K. E. Hecox, V. L. Towle, D. M. Frim, Estimation of in vivo human brain-to-skull conductivity ratio from simultaneous extra-and intra-cranial electrical potential recordings, *Clin. Neurophysiol.*, **116** (2005), 456–465.
45. T. F. Oostendorp, J. Delbeke, D. F. Stegeman, The conductivity of the human skull: results of in vivo and in vitro measurements, *IEEE Trans. Biomed. Eng.*, **47** (2000), 1487–1492.
46. Triton Aalto University School of Science “Science-IT” project. Available from: <https://scicomp.aalto.fi/triton/>.
47. D. P. Kingma, J. Ba, Adam: A method for stochastic optimization, In: Y. Bengio, Y. LeCun, Editors, *3rd International Conference on Learning Representations, ICLR 2015 - Workshop Track Proceedings*.



©2022 the Author(s), licensee AIMS Press. This is an open access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0>)