

Generation of synthetic training data for SEEG electrodes segmentation

Anja Pantovic¹, Xiaoxi Ren¹, Cédric Wemmert¹, Irène Ollivier² and Caroline Essert¹

¹ICube Laboratory, Université de Strasbourg.

²Department of Neurosurgery, Strasbourg University Hospital.

Abstract

Purpose: Stereoelectroencephalography (SEEG) is a minimally invasive surgical procedure, used to locate epileptogenic zones. An accurate identification of the metallic contacts recording the SEEG signal is crucial to ensure effectiveness of the upcoming treatment. However, due to the presence of metal, postoperative CT scans contain strong streak artefacts that interfere with deep learning segmentation algorithms and require a lot of training data to distinguish from actual contacts. We propose a method to generate synthetic data and use them to train a neural network to precisely locate SEEG electrode contacts.

Methods: Random electrodes were generated following manufacturer's specifications and dimensions, and placed in acceptable regions inside metal-free CT images. Metal artefacts were simulated in the generated dataset using radon transform, beam hardening, and filtered back projection. A UNet neural network was trained for the contacts segmentation task using various training setups combining real data, basic augmented data and synthetic data. The results were compared.

Results: We reported a higher accuracy when including synthetic data during the network training, while training only on real and basic augmented data more often led to misclassified artefacts or missed contacts. The network segments post-operative CT slices in less than 2 seconds using 4 GeForce RTX2080 Ti GPUs and in under a minute using a standard PC with GeForce GTX1060.

Conclusion: Using synthetic data to train the network significantly improves contact detection and segmentation accuracy.

Keywords: Stereoelectroencephalography, Epilepsy, Data Augmentation, Segmentation, Sinogram, Radon transform

1 Introduction

Epilepsy is one of the most frequent serious neurological disorders. Current antiepileptic drugs are effective in 60-70% of individuals. For patients with pharmacoresistent focal epilepsy of structural etiology [1], surgical resection may be an appropriate treatment option. In such cases, a successful resection relies on the accurate identification of the epileptogenic zone which is usually achieved through Stereoelectroencephalography (SEEG)[2]. Deeply located areas of the brain are explored owing to the surgical implantation of multiple depth electrodes, recording neuronal activity through metallic contacts evenly spaced along their body. The recording of SEEG signals and spatial association with the contact location allows for the detection of the spatial and temporal organization of the epileptic seizure [3], its starting point, and decide of the zone to resect. Identifying the exact location of up to 250 metallic contacts is an extremely time consuming task that is usually performed manually and is highly prone to human error and inaccuracy. In order to save valuable time and potentially increase the effectiveness of the upcoming surgical procedure, several groups have been working on automatic SEEG electrode localisation [4–8]. In a previous work, we proposed an approach to automatically segment the electrode contacts using 2D and 3D UNet [8], where post-implantation CT scans showing the SEEG electrodes were used to train a neural network.

However, due to the presence of metal, artefacts appear on the CT scans as bright streaks around the metal and dark streaks between two metallic contacts, and impair their visibility and detection. Models trained with this kind of data in insufficient number often result in false positives. Usual data augmentation provides better results, but does not bring more variety in the training samples and continues to mark artefacts as false positives. To bring more variety to the training dataset, we propose to generate synthetic data with a sufficient level of realism in terms of electrode and contact locations and appearance of the final image. The three objectives are to easily increase the size of the training dataset while controlling the variety of the configurations and automatically generating the ground truth during the process to avoid a tedious annotation task.

This paper introduces an approach to generate synthetic volumes mimicking post-implantation CT images. Section 2 details the selection of locations for the generated contacts, and the generation of the final image with artefact simulation. In Section 3, we present the UNet architecture and the experimental setup. Section 4 shows and discusses the obtained results, after which a conclusion is drawn.

2 Methods

The synthetic volumes were built from representative CT volumes without metal implant. Four plain CT head volumes were selected as base volumes from the CQ500 dataset publicly available on Qure.ai [9]. From that base, contacts were added at random but relevant locations and the artefacts were simulated.

2.1 Generation of electrodes

2.1.1 Definition of the acceptable regions

The chosen contact locations must be relevant enough to avoid obviously meaningless images, for instance with electrodes parallel to the skull or entirely in the brain. In usual practice, 10–18 SEEG electrodes with 5–18 contacts each are implanted in the brain using a patient-specific scheme. The trajectory of each electrode is defined by an entry point and a tip point chosen preoperatively according to hypotheses based on the symptoms and preliminary non-invasive exams. For aesthetic reasons, the electrodes are usually inserted where the hair grows. The curvature of the skull may cause electrode deviation if the electrode is not perpendicular to the cranial vault. The frontal and occipital poles are little explored due to the difficulty of resection in that area, mostly for functional reasons. The tip point is located in a deep region of the brain. Most of the electrodes are usually implanted laterally in one hemisphere, surrounding the suspected volume. A few contralateral electrodes can also be implanted in the opposite hemisphere. It is also not rare to implant orthogonal (vertical) and oblique electrodes [10].

To simulate the above rules, two regions of interest (ROI) to contain respectively the entry point (*external region*) and tip point (*internal region*) of the generated electrodes were automatically defined on the base CT volumes. An observation of retrospective CT images of patients who underwent SEEG exploration showed that for most of the electrodes the first contact was at the edge of the skull and the last contact was located towards the centre of the brain. To build the external region, the brain and its contour were semi-automatically segmented on each base CT volume using 3DSlicer [11] and contour detection [12]. Since the electrodes cannot be inserted in the front and back of the head, the contours were filtered and the corresponding parts were removed based on their y and z coordinates.

In order to delineate the internal region corresponding to the tip of the electrode, a cylinder was defined in the centre of the head. The height of the cylinder was set to 40 voxels, the length of the large radius to 100 voxels and the length of the small radius to 80 voxels, to fit most electrode models and lengths. An example of the resulting internal and external regions can be seen on Fig.1. Considering that a majority of electrodes are implanted in the same hemisphere, each ROI was divided into two classes: ‘right’ and ‘left’. This allows to control the distribution of the electrodes between the hemispheres.

2.1.2 Generation of electrodes within the regions of interest

To generate a synthetic electrode, a candidate point was randomly selected within each of the two, and their direction vector was computed. A distance of at least 15mm must be maintained between all the electrodes. Therefore, the distances between the line defined by the two candidate points and all the previously generated electrodes were computed. If at least one distance was under the fixed threshold, the two points were discarded and two new points were




Fig. 1 Examples of external (green) and internal (yellow) regions of interest, and generated electrodes (blue), visualised as 3D models

selected. Once an acceptable couple of candidate points was found, an electrode model was chosen among available models provided by the manufacturer (Microdeep[®] depth electrodes, DIXI Medical, Besançon, France). The semi-rigid depth electrodes are 0.8mm in diameter, and may have various lengths (16–80.5mm) and number of contacts (5–18) according to the model. Based on the manufacturer’s specifications and dimensions evenly spaced cylinders representing the depth electrode contacts in our CT volumes were generated. As shown in Fig.1, all the synthetic electrodes start on the external region and aim at the internal region. The output of this process is a three-dimensional binary *contact mask* \mathcal{V}_m representing the cylinders.

2.2 Generation of synthetic CT images with metal artefact simulation

The simulation of CT metal artefact has been well studied in medical imaging applications [13–16], mostly to build test cases for metal artefact reduction algorithms, or in other fields such as simulation of threat images [17]. Metal streak artefacts are caused by multiple mechanisms, including beam hardening, scatter and Poisson noise [18]. Beam hardening and scatter both produce dark streaks between two metal objects. When the X-ray passes through an object with a high attenuation coefficient, they cause a strong attenuation and do not pass through metal under certain incidences, resulting in corrupt or missed data in the projection (*sinogram*). Consequently, the CT reconstruction process amplifies the missing data in the sinogram due to the logarithmic function and causes the light and dark streaks on the CT image [13, 14, 17].

Our method to simulate metal artefacts is a simplified monochromatic version inspired by [15], where the authors used synthetic artefacts generation to test artefact reduction algorithms. The process starts by mapping the *metal-free* CT volume \mathcal{V}_{CT} and the *metal-only* 3D contact mask \mathcal{V}_m to the projection domain using radon transform using the ASTRA module [19], to obtain respectively sinograms \mathcal{S}_{CT} and \mathcal{S}_m . The artefacts are then simulated by normalising \mathcal{S}_m in the interval $[0, 1]$, and hardening it to obtain a \mathcal{S}'_m , as follows:

$$\mathcal{S}'_m(x) = \begin{cases} 1, & \mathcal{S}_m(x) > 0 \\ 0, & \mathcal{S}_m(x) = 0 \end{cases} \quad (1)$$




Fig. 2 Flow chart of the metal artefact simulation process

$S'_m(x)$ is then combined with S_{CT} to form S_c as: $S_c = S_{CT} + \omega S_m$, where ω is a parameter defining how strong the artefacts will be. Poisson random noise is applied to S_c based on the X-ray spectrum to simulate realistic quantum noise in the projection, and form a noisy projection S_n . Finally, the metal-inserted volume V_f is reconstructed from S_n using three-dimensional filtered back projection. The overall process is illustrated on Fig.2.

3 Validation and experimental setup

In order to evaluate the impact of the synthetic data on the segmentation of SEEG electrode contacts, UNet neural network architecture [20] has been trained in five different setups described in the subsection below.

3.1 Datasets and training setups

Three types of data were used, that will later be referred to as *real*, *augmented* and *synthetic* data/volumes. The *real* volumes consist of post-implantation CT scans from 18 patients who underwent SEEG at Strasbourg University Hospital. Each CT slice is of resolution 512×512 , with the number of slices in each scan varying between 166 and 259, accounting for a total of 4139 slices. Annotation of the electrode contacts was done manually using 3D Slicer[11]. To enlarge the dataset, the *augmented* volumes were built from the real data using some basic augmentation techniques - rotation and flipping. The annotations were transformed accordingly. The *synthetic* volumes were generated using the approach described in Section 2. The artefact strength parameter was set to $\omega = 0.08$, chosen to simulate artefacts similar to those that appear in the real dataset. Metal-inserted CT volumes were generated with corresponding binary masks of electrode contacts as a ground truth. Each generated case contains between 217 and 253 slices of resolution 512×512 , and an average of 185 metallic contacts. Examples of real and synthetic data are shown on Fig.3.

Real, augmented and synthetic data were used to train UNet for the electrode segmentation task. In order to compare impact of the different training datasets on the segmentation results, network was trained using five combinations of the different datasets, that constitute five experimental setups:




Fig. 3 Examples of the real (a,b) and synthetic (c,d) images

- **Real:** Training dataset composed of real data only, which includes 4139 slices from 18 different patients.
- **Augmented:** Training dataset composed of half real data and half augmented data, for a total of 36 volumes (18 volumes of each). Each real data volume was used to create one augmented volume, for a total of 8278 slices.
- **Synthetic:** Training dataset composed of half real data and half synthetic data, for a total of 36 volumes (18 volumes of each). The 18 synthetic volumes together contain 4121 slices, so this dataset consists of 8260 images.
- **Augmented + Synthetic:** Training dataset composed of 18 real volumes, 18 augmented volumes and 18 synthetic volumes, for a total of 12399 images.
- **Large Augmented:** Training dataset composed of 18 real volumes and 54 augmented volumes, *i.e.* 16556 images. The objective of this setup is to assess whether a possible increase in accuracy when using our synthetic data would be due to a higher amount of data used for training or rather because the synthetic data has a specific impact on what the network learns.

3.2 Network Configuration

To segment SEEG electrode contacts from the post-implantation CT scans, we have implemented the encoder-decoder UNet architecture [20]. The left side of the “U” shape represents a general convolutional process, consisting of the repeated application of two 3×3 convolutions. Each convolutional layer is followed by a ReLU activation and batch normalisation. Then, spatial dimensions are reduced by applying a 2×2 max pooling operation. The right side of the “U” shape is built by up-sampling of the feature map followed by a 2×2 transpose convolution. At the final layer, a 1×1 convolution is used after which each pixel is classified either as an electrode contact or a background pixel.

The network was trained in the five above setups. In order to reduce memory usage, all slices were cropped to the dimension of 400×400 pixels, removing only background pixels and preserving the whole skull within the frame. Each training was done in the identical hyper-parameter setup. Adam optimiser was used with the initial learning rate $1 * 10^{-5}$, $\beta_1 = 0.9$, $\beta_2 = 0.999$, $\epsilon = 1 * 10^{-8}$ and the *learning decay* of $1.99 * 10^{-7}$ was used to help fine-tune the network. The loss function was binary crossentropy. The number of epochs was determined empirically and set to 150. Networks were cross-validated using the

Table 1 Training time for 150 epochs and average results of the UNet model after leave-one-out cross validation performed in five different setups, evaluated on seven metrics.

<i>Setup</i>	Training time	Dice coef.	IoU	TP	FP	FN	PPV	TPR
<i>Real</i>	45'	0.807	0.704	180.6	9.6	4.4	0.950	0.976
<i>Augmented</i>	1h31'	0.850	0.748	182.2	7.6	2.8	0.960	0.985
<i>Synthetic</i>	1h29'	0.848	0.742	183.6	3.4	1.4	0.981	0.992
<i>Augm.+Synth.</i>	5h45'	0.862	0.764	184.3	2.6	0.7	0.986	0.996
<i>Large Augm.</i>	6h38'	0.857	0.756	183.0	5.3	2.0	0.972	0.989

‘leave-one-out’ approach: for each training phase one real volume was left out for testing. In all five setups, testing was done on the set of slices belonging to a single patient from the real dataset, unseen during the training. The data that were created from the current test case were also left out during training.

3.3 Experimental conditions and metrics

Training was performed on 4 NVIDIA GeForce RTX2080 Ti GPUs, each with 11GB RAM. Networks were trained for 150 epochs for each cross-validation fold and evaluated based on the Dice coefficient and Intersection over Union (IoU). Furthermore, an algorithm based on the connected component count was implemented to compute the number of true positive (TP), false positive (FP) and false negative (FN) segmented contacts for each test case, which were further used to calculate the *Positive Predictive Value*, $PPV = \frac{TP}{TP+FP}$ and the *True Positive Rate*, $TPR = \frac{TP}{TP+FN}$. Predictions for the left-out case were made in less than 2 sec. using this configuration. For reference, on a standard PC with one Nvidia GeForce GTX1060, prediction time is under one minute.

4 Results and Discussion

Fig.4 shows the results on 4 examples with different levels of complexity. For each example, the original CT image is shown on row 1, and the predictions from the five different setups on rows 2-5. TP (contacts properly found) are represented in green, FP (artefacts detected as contacts) in red, and FN (contacts not found) in yellow. In all five setups most of the contacts are properly found. The training times and metrics are summarised in Table 1. Given values are averages across all validation folds.

As we can observe on Table 1, all of the data augmentation setups outperformed the *Real* setup. With this setup, the network segments well the ‘clear’ electrodes when artefacts are not too strong (Fig.4, row 2, column 1). When the artefacts are very bright, the network often misclassifies them as contacts (row 2, col. 2-3, FP in red). On the contrary, the network fails to segment contacts that coincide with the skull or the screw in the CT scan (row 2, col. 4, FN in yellow). IoU and Dice coefficient are relatively low, indicating that the segmentation is imprecise. This observation confirms the need to augment the training datasets to obtain better results.

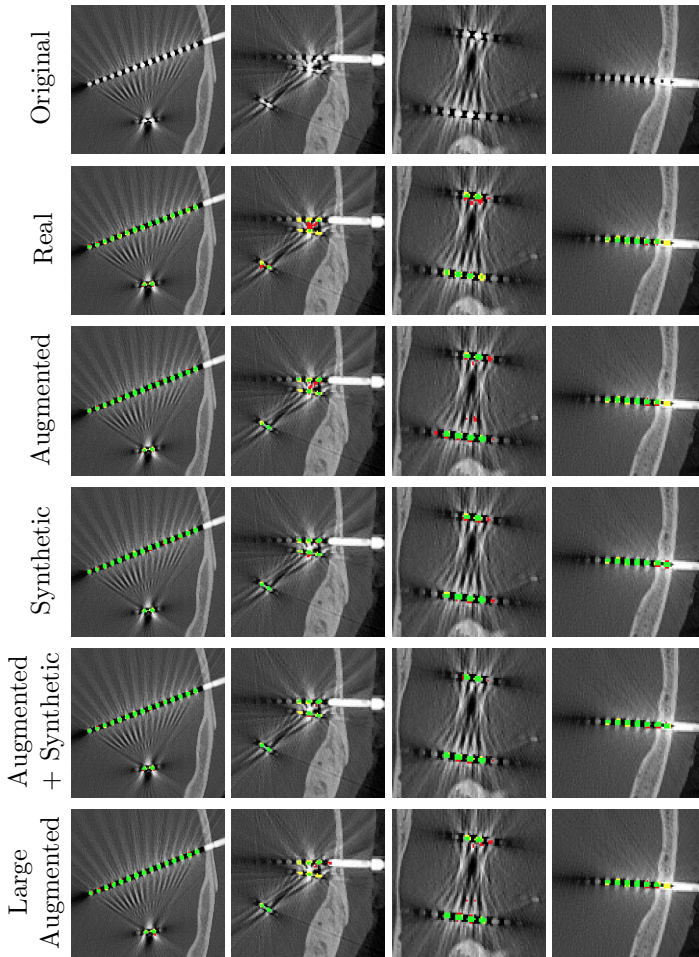


Fig. 4 Original images and predictions obtained from the five different setups. Columns represent different situations: simple/general case (column 1), strong artefacts (columns 2-3), contact close to structures of similar brightness such as screw and skull (column 4).

When comparing training with the *Augmented* and *Synthetic* setups, that both doubled the number of training data, we observe that the *Augmented* setup led to higher Dice coefficient and IoU, but also higher numbers of FP and FN. This indicates that the *Augmented* setup was better at delineating the identified contacts, but that more contacts were missed (Fig.4, row 3, col. 4) and more artefacts were misclassified as contacts (row 3, col. 2-3). In the *Augmented* setup, the network did not overcome the issue of strong artefacts nor the contacts coinciding with the skull. The *Synthetic* setup seems less sensitive to artefacts (row 4, col. 2-3), and found more often the skull contacts, (row 4, col. 4). Slightly lower Dice and IoU scores obtained in the *Synthetic* than in the *Augmented* setup may be explained by the fact that real/augmented data

were manually annotated, while the labels of the synthetic data were automatically saved during electrode generation. Training with synthetic data may have introduced discrepancies of annotations between real/synthetic labels.

The highest validation scores have been reached after training the network on the real, augmented and synthetic data all together (*Augmented + Synthetic* setup). This is consistent with the visualised results (5th row of Fig.4). The number of FP and FN contacts was reduced while the delineation of each contact seems more accurate. As we can see, the artefacts were avoided and the contacts in the skull area were well segmented. This setup seems to take advantage of both additional datasets for a refined prediction compensating the discrepancy. One could object that the results were better because the amount of training data was tripled. As mentioned in Section 3.1, the *Large Augmented* setup quadrupled the training data but used only basic augmentation. We can observe on Table 1 that the *Augmented + Synthetic* setup performed better than the *Large Augmented* setup for a comparable (even a little lower) number of data. This may be due to a better variety in the electrode configurations and examples of artefacts. Due to a lack of variety, the *Large Augmented* setup remained sensitive to artefacts and failed to segment some skull contacts.

We conclude that the generated synthetic data have a positive impact on the accuracy of the segmentation, which is supported by the statistical analysis of the results, performed over the PPV and the TPR. A paired two-sample Student's *t-test* was performed, testing the *Augmented + Synthetic* setup against each of the other setups without synthetic data. The difference was found statistically significant (*p-values* < 0.05) over both metrics. The configurations of electrodes, as well as the amount and strength of the artefacts in the synthetic images can be fine-tuned, which helps the network learn to better distinguish between the artefacts and actual electrode contacts. In the synthetic images, electrode contacts closest to the screws are always placed on the edge of the skull, which improves the training on this particular configuration. However, in rare cases even with the synthetic data such contacts are sometimes missed, when they are visually completely connected to the screw on the CT image.

Another remaining issue is the segmentation of the vertically implanted electrode. Vertical electrodes appear on the axial plane as a single contact (see Fig.3b) and on most slices as an uninterrupted shape, with only slight brightness variations and no clear space between neighbouring contacts. In 3D, most of their contacts appear connected. In Fig.5, we show 3D views of ground truth (left) and prediction masks obtained from the *Augmented + Synthetic* setup (middle and right) for two *real* datasets. The left and middle columns show the connected components in different colour which allows us to easily detect merged contacts. In the right column, TP are shown in grey while FP (contacts detected instead of a screw) are highlighted in red.

Further work needs to be done to overcome the few remaining issues. The merged contacts on vertical electrodes could be separated using dilation and erosion operations. The remaining FP or FN could be identified when grouping the contacts into electrodes and correlating with electrode specifications.




Fig. 5 3D views of ground truth (left) and prediction masks of the *Augmented + Synthetic* (middle and right) for two *real* datasets. The left and middle column show the connected components labelled in different colours. In the left column, FP are highlighted in red.

5 Conclusion

In this paper we proposed an approach to automatically generate synthetic volumes mimicking SEEG post-implantation CT images, to augment the training dataset and improve automatic contact segmentation. This approach offers a possibility to easily build large training datasets and better control their variety, while automatically generating the ground truth. To validate the impact of the synthetic data on the training, a UNet neural network was trained with five different data configurations. The results showed that basic augmentation methods were not sufficient for a network to learn to separate contacts from the metallic artefacts, whereas adding synthetic images had a significantly positive impact allowing to overcome this issue. Some limitations persist however even with the synthetic data, with vertically implanted electrodes where the segmented contacts may be merged, or with contacts close to structures of similar brightness that may remain misclassified.

Declarations

Funding: This work was supported by ArtIC “Artificial Intelligence for Care” grant (ANR-20-THIA-0006-01), co-funded by Région Grand Est, Inria Nancy - Grand Est, IHU Strasbourg, University of Strasbourg and University of Haute-Alsace, France.

Conflict of interest: The authors have no conflict of interest to declare.

Ethics approval This research study was conducted retrospectively from anonymised data, in accordance with the ethical standards of our institution and with the 1964 Helsinki declaration and its later amendments or comparable ethical standards.

Informed consent: Informed consent was obtained from all individual participants included in the study.

References

- [1] Scheffer, I.E., Berkovic, S., Capovilla, G., Connolly, M.B., French, J., Guilhoto, L., Hirsch, E., Jain, S., Mathern, G.W., Moshé, S.L., Nordli, D.R., Perucca, E., Tomson, T., Wiebe, S., Zhang, Y.-H., Zuberi, S.M.: ILAE classification of the epilepsies: position paper of the ILAE Commission for Classification and Terminology. *Epilepsia* **58**(4), 512–521 (2017)
- [2] Talairach, J., Bancaud, J.: Lesion, “irritative” zone and epileptogenic focus. *Stereotactic and Functional Neurosurgery* **27**(1-3), 91–94 (1966)
- [3] Minotti, L., Montavont, A., Scholly, J., Tyvaert, L., Taussig, D.: Indications and limits of stereoelectroencephalography (SEEG). *Neurophysiologie Clinique* **48**(1), 15–24 (2018)
- [4] Meesters, S., Ossenblok, P., Colon, A., Schijns, O., Florack, L., Boon, P., Wagner, L., Fuster, A.: Automated identification of intracranial depth electrodes in computed tomography data. In: *IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pp. 976–979 (2015)
- [5] Narizzano, M., Arnulfo, G., Ricci, S., Toselli, B., Tisdall, M., Canessa, A., Fato, M., Cardinale, F.: SEEG assistant: A 3DSlicer extension to support epilepsy surgery. *BMC Bioinformatics* **18**, 124 (2017)
- [6] Granados, A., Vakharia, V., Rodionov, R., Schweiger, M., Vos, S., O’Keeffe, A., Li, K., Wu, C., Miserocchi, A., Mcevoy, A., Clarkson, M., Duncan, J., Sparks, R., Ourselin, S.: Automatic segmentation of stereoelectroencephalography (SEEG) electrodes post-implantation considering bending. *International Journal of Computer Assisted Radiology and Surgery* **13**, 935–946 (2018)
- [7] Benadi, S., Ollivier, I., Essert, C.: Comparison of interactive and automatic segmentation of stereoelectroencephalography electrodes on computed tomography post-operative images: preliminary results. *Healthcare technology letters* **5**(5), 215–220 (2018)
- [8] Pantovic, A., Ollivier, I., Essert, C.: 2D and 3D-UNet for segmentation of SEEG electrode contacts on post-operative CT scans. In: *Proc. of Medical Imaging: Visualization and Image-Guided Procedures* (2022). to appear
- [9] Chilamkurthy, S., Ghosh, R., Tanamala, S., Biviji, M., Campeau, N.G., Venugopal, V.K., Mahajan, V., Rao, P., Warier, P.: Deep learning algorithms for detection of critical findings in head CT scans: a retrospective study. *The Lancet* **392**(10162), 2388–2396 (2018)
- [10] Mullin, J.P., Shriver, M., Alomar, S., Najm, I., Bulacio, J., Chauvel, P.,

- Gonzalez-Martinez, J.: Is SEEG safe? A systematic review and meta-analysis of stereoelectro-encephalography-related complications. *Epilepsia* **57**(3), 386–401 (2016)
- [11] Kikinis, R., Pieper, S.D., Vosburgh, K.G.: 3D Slicer: A platform for subject-specific image analysis, visualization, and clinical support. In: *Intraoperative Imaging and Image-guided Therapy*, pp. 277–289. Springer, New York (2014)
- [12] Suzuki, S., be, K.: Topological structural analysis of digitized binary images by border following. *Computer Vision, Graphics, and Image Processing* **30**(1), 32–46 (1985)
- [13] Boas, F.E., Fleischmann, D.: CT artifacts: causes and reduction techniques. *Imaging in Medicine* **4**(2), 229–240 (2012)
- [14] Mehranian, A., Ay, M., Rahmim, A., Zaidi, H.: Sparsity constrained sinogram inpainting for metal artifact reduction in x-ray computed tomography. In: *IEEE Symposium on Nuclear Science*, pp. 3694–3699 (2011)
- [15] Zhang, Y., Yu, H.: Convolutional Neural Network Based Metal Artifact Reduction in X-Ray Computed Tomography. *IEEE Transactions on Medical Imaging* **37**(6), 1370–1381 (2018)
- [16] Yu, L., Zhang, Z., Li, X., Xing, L.: Deep sinogram completion with image prior for metal artifact reduction in CT images. *IEEE Transactions on Medical Imaging* **40**(1), 228–238 (2020)
- [17] Megherbi, N., Breckon, T.P., Flitton, G.T., Mouton, A.: Radon transform based automatic metal artefacts generation for 3D threat image projection. In: *Optics and Photonics for Counterterrorism, Crime Fighting and Defence IX; and Optical Materials and Biomaterials in Security and Defence Systems Technology X*, vol. 8901, pp. 94–102 (2013). International Society for Optics and Photonics, SPIE
- [18] De Man, B., Nuyts, J., Dupont, P., Marchal, G., Suetens, P.: Metal streak artifacts in x-ray computed tomography: a simulation study. *IEEE Transactions on Nuclear Science* **46**(3), 691–696 (1999)
- [19] Palenstijn, W.J., Batenburg, K.J., Sijbers, J.: The astra tomography toolbox. In: *13th International Conference on Computational and Mathematical Methods in Science and Engineering*, pp. 1139–1145 (2013)
- [20] Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *Medical Image Computing and Computer-Assisted Intervention*, pp. 234–241 (2015). Springer