Automated needle localisation for electric field computation during an electroporation ablation

Eloïse Inacio, Luc Lafitte, Olivier Sutter, Olivier Seror, Baudouin Denis de Senneville, Clair Poignard

Automated needle localisation for electric field computation during an electroporation ablation

1st Eloïse Inacio
Project team MONC
Univ. Bordeaux, UMR CNRS 5251, INRIA
Talence, France
eloise.inacio@inria.fr

2nd Luc Lafitte
Project team MONC
Univ. Bordeaux, UMR CNRS 5251, INRIA
Talence, France
luc.lafitte@inria.fr

3rd Olivier Sutter
Interv. Radiol. Unit
Univ. Hosp. Paris Seine Saint Denis
Avicenne Hosp., APHP, Univ. Paris 13
Bobigny, France,
olivier.sutter@aphp.fr

4th Olivier Seror
Interv. Radiol. Unit
Univ. Hosp. Paris Seine Saint Denis
Avicenne Hosp., APHP, Univ. Paris 13
Bobigny, France,
olivier.seror@aphp.fr

5th Baudouin Denis de Senneville
Project team MONC
Univ. Bordeaux, UMR CNRS 5251, INRIA
Talence, France
bdenisde@math.u-bordeaux.fr

6th Clair Poignard
Project team MONC
Univ. Bordeaux, UMR CNRS 5251, INRIA
Talence, France
clair.poignard@inria.fr

***Abstract*—The objective of this paper is to provide a step-forward towards the per procedural visualisation of the electric field distribution during a clinical irreversible electroporation (IRE) procedure. To this end, an automated workflow is needed to compute the electric field distribution on a single Cone Beam Computed Tomography (CBCT) scan.**

The aim of the current paper is to propose a deep learning strategy for the automatic segmentation of the needles. In particular, a novel coarse-to-fine approach is proposed to extract relevant needle information from the CBCT scan, despite inherent artefacts generated during capture. The obtained needle information is subsequently fed into a standard static linear model for the electric field computation. Since the set-up is performed in the medical image framework, the electric field distribution and the region of interest are visible to provide to the radiologist a visual evaluation of the treatment.

The segmentation results are evaluated on 8 of the 16 patients of the dataset using the Dice coefficient to compare the predicted segmentation with the ground truth. The proposed segmentation method is fast (around 2 minutes are needed with a commodity hardware), allowing its use in a clinical setting.

***Index Terms*—Deep Neural Network, Fine-object Segmentation, CBCT, Electric field distribution.**

I. Introduction

Irreversible electroporation (IRE) is a minimally invasive ablation technique notably used in the ablation of deep-seated tumour. It consists in delivering localised high electric pulses through conductive needles to permeabilise cell membranes and lead to apoptosis. It is mostly non-thermal, thus suitable for treating tumours near major blood vessels and organs as it will not affect nearby vital structures. However, due to the complexity of the technique, thorough treatment planning and assessment based on medical imaging are required.

In [1] the clinical workflow to perform IRE liver ablation was described: the target tumour is localised on a high resolution image, then the needles are inserted percutaneously and a low resolution Cone Beam Computed Tomography (CBCT) scan is acquired to verify the positioning. Finally the IRE pulses are delivered. Following this clinical workflow, a numerical workflow was described in [1] to provide a numerical assessment of the IRE procedure. The current paper proposes here an automated needle extraction method to automatically perform the treatment assessment.

Indeed one of the most important steps for the computation of the electric field during an IRE procedure lies in the localisation of the percutaneously inserted needles, as previously shown in [2]. The captured information is then used to determine the method (needle placement, etc) and efficacy of IRE for a patient.

CBCT scan is an image modality which is particularly beneficial intra-operatively since it is fast, restricted to the targeted organ and with lower radiation exposure than conventional Computed Tomography (CT) scans. However, the CBCT scans are subjected in turns to low contrast, low signal-to-noise ratio, and artifacts especially when needles are inserted. These drawbacks hamper the automated needle extraction task and require the development of novel deep learning approaches.

Artificial neural networks have been increasingly used in image processing tasks such as segmentation, classification or object tracking. In the past decade, it has come to outperform more traditional techniques thanks to great improvements in Graphics Processing Unit (GPU) technology, allowing more complex architectures, and to the growing number of datasets,

This study has been partly granted by the Plan Cancer Projet MECI, n°21CM119-00.
The numerical experiments presented in this paper were carried out using the PlaFRIM experimental testbed, supported by Inria, CNRS (LABRI and IMB), Université de Bordeaux, Bordeaux INP and Conseil Régional d'Aquitaine (see https://www.plafrim.fr/).

required for supervised training. In particular for the segmentation task, most state-of-the-art convolutional network approaches are based on an architecture designed in 2015 in [3]: the U-Net. The output combines contextual information as well as detailed information, achieving high quality segmentation maps. It was first applied to biomedical data, specifically neuronal structure and cell segmentation. The simple yet powerful architecture has then been declined to fit specific data (*i.e.,* liver and liver tumor [4], liver lesion [5], etc). One advantage of this kind of approach is that they rely on features learned by the network during training and not hand-crafted features whose quality and suitability limit performance. However, data presented to the network in the training phase should be chosen carefully and fine-tuning of the model might be necessary if used on a very different dataset.

The aim of this paper is a strategy to obtain accurately and fastly the position of the needles from medical imaging. This automated segmentation is the first step towards an on-line visualization of the electric field distribution as delivered by the clinical IRE procedure described in [1]. The automated workflow is designed to enable the visualization of the electric field distribution directly on a single CBCT scan. The proposed strategy for automated needle segmentation tackles the following successive tasks:

1) The needles are coarsely segmented in the CBCT scan using a deep convolutional network. Challenges arise from the nature of the scan (*i.e.,* CBCT low contrast, low signal-to-noise ration and needle artefacts) and from the label imbalance in the dataset. We address this issue by the use of a suited loss function along a patch-based optimization strategy.
2) A fine representation of the tip of the electrode is subsequently obtained. However, incomplete coarse segmentation makes the precise detection of the needles more difficult. The proposed solution consists in completing the missing data by using an analytical needle representation based on the Hough transform [6].

The distribution of the electric field is then numerically computed using the standard static non linear model as presented in [2] directly in the CBCT framework.

II. Method

A. Coarse needle segmentation

1) Implemented deep neural network: A U-Net model is adapted to the segmentation of the needles used for the IRE of liver tumor. The architecture is presented in Fig. 1.

A contracting path encodes the input into a space of lower dimensionality via down-sampling operations, here max-pooling. An expanding path then decodes the smaller vectors into the required output via up-sampling operations that increase the output resolution, here up-sampling, faster than the original deconvolution layer. At each level of the encoder and the decoder, two $3 \times 3 \times 3$ zero-padded convolutions are applied, followed by a ReLU activation function and a Batch

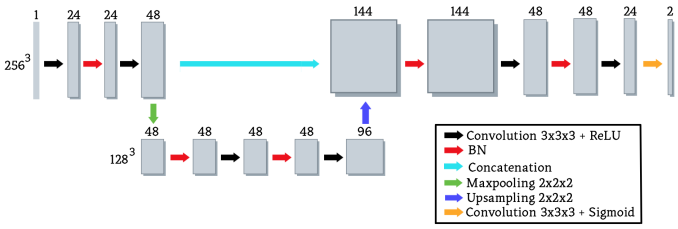

Fig. 1. Proposed 3D U-Net architecture for the calculation of the coarse needle segmentation map: Max pooling is the down-sampling operation, ReLU and Sigmoid are activation functions and BN is batch normalisation. The number of input channels (NC) as well as the number of $3 \times 3 \times 3$ filters are indicated on the top of each block and the size of the feature maps are on the left.

Normalisation (BN) layer while in the original U-Net, unpadded convolutions were used. The feature maps thus had to be cropped before the concatenation allowing to combine contextual information from the encoder and detailed information from the decoder. Padding the input prior to convolution allows to avoid an additional operation and to have a quicker training and inference. Only two levels are used in the adapted U-Net as it is sufficient for the patch size considered and again allows a less time consuming processing of the data. Also, more importance is given to detailed information by concatenating twice more decoder feature maps than there are encoder feature maps as very fine objects are being segmented.

The proposed training strategy is designed to counterbalance the highly skewed data, as background is more represented than needles in the labels. The network is trained using supervised training. Input is fed to the network whose output is compared with ground truth data using a loss function that measures similarity according to the objective. The loss function (L) is a combination of binary cross entropy (L_{BCE} [7]) and Dice loss (L_D [8]), which is based on the Dice coefficient [9] and demonstrated to be well suited for imbalanced structure segmentation [10]. L is mathematically expressed as follows:

$$L = L_{BCE} + L_D \tag{1}$$

with

$$L_{BCE} = -\frac{1}{|\Omega|} \sum_{\vec{r} \in \Omega} \left(y_{\vec{r}} \log(p_{\vec{r}}) + (1 - y_{\vec{r}}) \log(1 - p_{\vec{r}}) \right) \tag{2}$$

$$L_D = 1 - \frac{2 \sum_{\vec{r} \in \Omega} y_{\vec{r}} p_{\vec{r}} + \varepsilon}{\sum_{\vec{r} \in \Omega} y_{\vec{r}} + \sum_{\vec{r} \in \Omega} p_{\vec{r}} + \varepsilon} \tag{3}$$

where Ω is the image domain, $\vec{r}$ the 3D voxel coordinate, y the true segmentation map, ε a smoothing factor (we took $\varepsilon = 1$) and p the predicted probability map (a value $p_{\vec{r}}$ close to 1 means that a high probability for the presence of a needle is predicted at the voxel coordinate $\vec{r}$).

The loss function is minimized with the gradient descent method: weights of network are updated in the direction opposite the gradient, proportionally to a learning rate. These steps are repeated until a convergence criterion is fulfilled: in this case, a number of iteration.

The following parameters were used for training of the 3D

U-Net: batch size = 1 and optimizer = Adam with a fixed learning rate of $1e^{-3}$. The number of epochs and the resolution of the input depend on the experiment.

2) Proposed patch-based approach: Though patches are used as in the original paper in order to reduce the memory requirements while maintaining good resolution, the patch-overlap strategy [3] is not.

CBCT scans captured during operation with needles in place for electroporation constitute highly imbalanced data label-wise. Indeed, the needles only take a small part of the image. Thus, a convolutional network trained on these data will have more difficulty learning the needles than the background.

One way to give more importance to needles is to select patches of the image where at least one voxel belongs to a needle. In this way, the ratio of "needle" voxels to "background" voxels may be increased and the label imbalance is reduced. The "needle" patches are selected using the true background segmentation maps.

An equilibrium between "needle" information and "background" information is reached when the training set consists of 50% of patches containing part of a needle and 50% of empty patches. This optimised patch set is referred to as the balanced set.

"Background" patches used to complete the set come from a single scan in order to ensure that all background structures are represented and learned: ribs, liver, etc.

3) Inference threshold optimization: The label imbalance, explained in the section above, is noticeable in the probability maps $p_{\vec{r}}$, then thresholded to obtain binary segmentation maps (noted $\hat{y}_{\vec{r}}$ throughout the rest of the manuscript): the probabilities when predicting a "background" voxel are much higher than the probabilities predicting a "needle" voxel. To improve the segmentation quality, we propose to vary the inference threshold θ in order to artificially give more weight to the needles when deciding on a label.

B. Fine needle segmentation of the tip of the electrodes

Difficulties arise when parts of the needles are not segmented on the mask. One can note that only the tips of the electrodes are useful for the computation of the electric field distribution as it is where the pulses are delivered. The above-mentioned coarse segmentation masks are processed to extract a finer representation of the tips of the electrodes. Practically, a Hough transform [6] is used to detect needles represented by lines despite the missing data as follows.

The coarse binary mask is first converted to a point cloud in order to apply the Hough transform. It consists in parameterizing the curve to detect, and fill an accumulator array whose dimensions correspond to each parameter of the curve. The scores of the cell in the accumulator then correspond to how represented the curves are in the image, *i.e.,* the number of points belonging to a given curve. In this implementation, Roberts' optimal line representation [12] is used and additional parameters are taken into account: the number of lines to be detected (chosen as the number of needles N multiplied by a coefficient), the maximum distance allowed between the detected line and the points belonging to the line (*i.e.,* the radius of the needles, 8 to 9.5 gauge for NanoKnife probes [13]), the minimum score in the accumulator array for a line to be detected and the maximum distance between two points of a line.

After the given number of lines is selected, a voting procedure further selects the curves representing the needles. Several criteria are taken into account: the inverse of the number of points belonging to the line, the distance between the extremities of the segment as represented in the point cloud, *i.e.,* the length of the needle segmented by the 3D U-Net, the mean distance from the points to their projection on the line and the distance between the extremity of the segment and the centre of the image. The latter constraint comes from the hypothesis that the tip of the needle where the pulses are delivered is the furthest extremity from the border of the scan. The 4 scores are normalised across the set of lines detected by the Hough algorithm and summed to produce the final score. Then, the N lines with the lowest final scores are selected as the finer representation of the needles extremities.

C. Assessment of the proposed needle extraction approach

1) Datasets: During IRE of liver tumours, a CBCT scan of the abdominal cavity is captured after needle insertion and before pulse delivery. Low signal to noise ratio, low contrast and artefacts typical to CBCT present an additional challenge to the segmentation task.

The patient dataset is divided into a training set and a testing set as follows. It consists of a total of 16 patients: 8 patients are randomly selected for training and the last 8 are used to assess the segmentation accuracy of the algorithm.

Ground truth, consisting of binary masks, is generated for all patients using the segmentation tool available in ITK-SNAP [16].

The dimension of the scans are 512 in the Right-Left and Anterior-Posterior directions and between 195 and 512 in the Superior-Inferior directions. The voxel sizes are 0.45 mm in all directions. Every scan is set to common dimensions before being presented to the convolutional network.

An intensity-scale is applied on CBCT scans prior to being processed by the neural network, during learning and testing phases. As intensity values are absolute, the same intensity scaling may be applied to the whole dataset (Fig. 2).

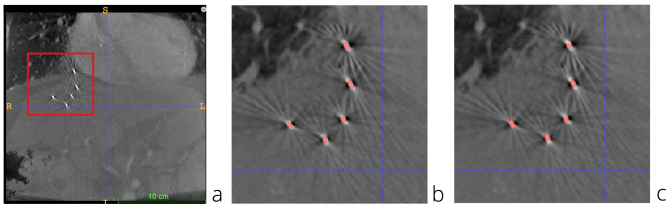

Fig. 2. Typical example of an intensity-scaled scan and the corresponding segmentation maps: (a) is a slice of the scan to be segmented with intensity (min, max) = (-999, 1500), (b) is the region framed in a with manual segmentation in red and (c) is the region framed in a with automatic segmentation in red.

The data is highly imbalanced label-wise. Indeed, the needle only constitutes a small portion of the CBCT scans: in the case of our dataset, there is 1 "needle" voxel for 10 000 "background" voxels in average.

2) Performance assessment: On the one hand, coarse segmentation masks are individually assessed prior to the Hough transform. To do so, the ground truth data generated for the test set ($y_{\vec{r}}$) are used, and compared to the predicted labels ($\hat{y}_{\vec{r}}$) using the Dice score (noted DSC):

$$DSC = \frac{2\sum_{\vec{r}\in\Omega} y_{\vec{r}}\hat{y}_{\vec{r}}}{\sum_{\vec{r}\in\Omega} y_{\vec{r}} + \sum_{\vec{r}\in\Omega} \hat{y}_{\vec{r}}} \tag{4}$$

On the other hand, fine segmentation is assessed by considering the coordinate of the extremity of the needle. It is compared to the manual coordinates using the Euclidean norm:

$$D = \sqrt{(x_t - x_p)^2 + (y_t - y_p)^2 + (z_t - z_p)^2} \tag{5}$$

where x_t, y_t, z_t are the true coordinates and x_p, y_p, z_p are the predicted coordinates.

3) Hardware and implementation: The adapted 3D U-Net was implemented with the Tensorflow 1.4 [17] and Keras 2.2.4 [18] packages. The platform used for the training of the CNN is an Intel Xeon E5-2683 2.4 GHz (2 Hexadeca-core) with 256 GB of RAM equipped by a GPU Nvidia Tesla P100 with 16 GB of memory. The commodity hardware used to test the segmentation is an Intel 2.5 GHz i7 workstation (8 cores) with 32 GB of RAM.

D. Preliminary assessment of an IRE procedure

To compute the electric field in tissues, different models of tissue electroporation can be chosen, and there is currently no definitive answer neither on the best choice of the model (static, nonlinear, dynamical, biphasic...) nor on the value of the parameters. One can refer to [19] for a comparison of different models. However, it is well-known that the linear static model under evaluate the region of electroporation, compared with other non linear models. Therefore, from a medical point of view, it is sufficient, as a first approximation, to compute the electric field as the gradient of the linear static potential φ satisfying:

$$-\nabla\cdot(\sigma\nabla\varphi) = 0, \text{in the region of computation } \Omega, \tag{6}$$

complemented with Dirichlet boundary conditions on the active needles, a floating potential condition [15] on the passive needle and a Fourier-Robin condition on the boundary of the simulation box. The liver conductivity was set to $\sigma = 0.5$ S/m while the tumour conductivity was set to $\sigma = 0.7$ S/m following [20]. The isolines 400 V/cm and 650 V/cm and the manually segmented tumour are then superimposed on the CBCT with needles, in the same way as [1].

III. Results

A. Assessment of the proposed patch-based approach

Three sets are created with different proportions of "needle" and "background" voxels to compare with the balanced set, designed to counter-balance the skewed data. The first one has 100% "needle" patches, the second has approximately 66% "needle" patches, and the balanced set has nearly 50% "needle" patches. The segmentation results, of the network when all patches are used, are also considered for comparison. We observe increasing segmentation quality as a balance is found between "needle" information and "background" information (Table I).

The first dataset is composed of 100% "needle" voxel. 38 patches across the training set were detected for common dimension = $256\times256\times256$ voxels and patch size = $64\times64\times64$ voxels. The segmentation quality is slightly over a half of what was obtained when all patches were used ($DSC = 0.2552$). It is due to the fact that the network does not have enough data to fully learn what the background looks like. This is further shown by the lack of efficiency of decreasing the inference threshold θ, as seen in Table I. Indeed, since the network does not learn the background as much as before, it cannot predict with its previous confidence the "background" labels.

The second dataset is composed of the 38 patches previously mentioned and 23 "background" patches, all extracted from the same image in order to ensure that every structures present in the background are included. That is, there are 62% of "needle" patches for 38% of "background" patches. The segmentation quality is doubled as compared to the previous dataset. The performance obtained using every patch is exceeded while the training time is lowered and the network learned better the "needle" class. Furthermore, decreasing the inference threshold θ allows again to improve the segmentation: the network learns the background better than previously.

The balanced set is composed of the 61 patches included in the second patch set and the last 35 "background" patches in the randomly chosen scan used previously. That is, there are 44% of "needle" patches for 56% of "background" patches. We notice a smaller increase of segmentation quality. Also, decreasing the inference threshold θ still improves performance.

B. Inference threshold optimization

In order to counter-balance the effect of a skewed dataset on learning, different inference thresholds are considered.

We observe that the regular threshold, θ = 0.5, may not be optimal on the used dataset (Fig. 3c).

On the one hand, for difficult images (*i.e.,* low contrast, weak signal, etc), needles are hardly recognizable. There is an almost linear relationship between the segmentation quality

TABLE I
PERFORMANCE OF THE COARSE NEEDLE SEGMENTATION APPROACH USING THE DIFFERENT PATCH SETS

Set	***Mean Dice coefficient***	***Standard deviation***	***Time/epoch (s)***
All Patch	0.4532	0.0272	6749
Patch Set 100-0%	0.1310	0.0075	847
Patch Set 67-33%	0.5179	0.0227	1701
Balanced Set%	0.5536	0.0163	2698

* Using an inference threshold of 0.2

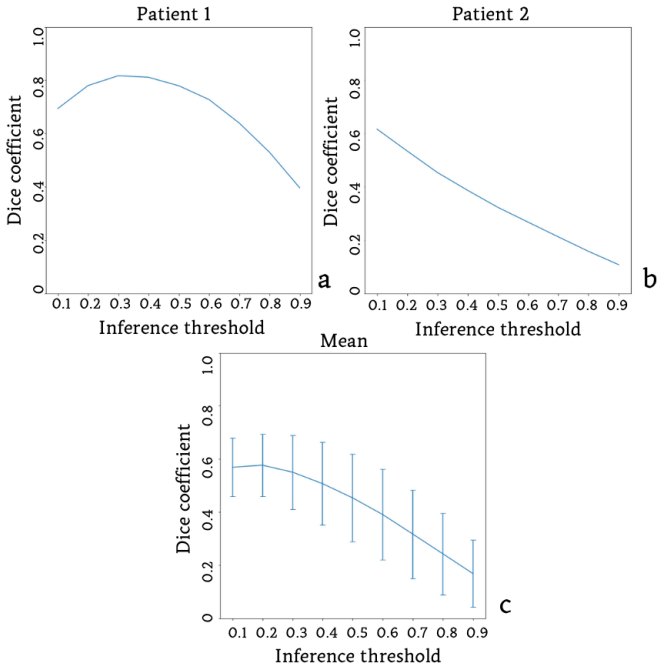

Fig. 3. Segmentation quality for varying inference threshold: (a) is the plot of the Dice coefficient against the inference threshold for a high quality scan, (b) is the plot of the Dice coefficient against the inference threshold for a low quality scan and (c) is the plot of the mean Dice coefficient computed across the test set against the inference threshold (error bar = standard deviation). The network was trained on $64 \times 64 \times 64$ patches extracted from CBCT scans re-sampled to a common dimension of $128 \times 128 \times 128$ voxels and using 30 epochs.

and the inference threshold (Fig. 3b). We suppose it is due to the difficulty to discern needles from background. That is, the more positive labels the threshold allows, the more the needles have a chance to be segmented.

On the other hand, for higher quality images, the relationship is parabolic with a peak between $\theta = 0.2$ and $\theta = 0.3$ (Fig. 3a). We believe that this is due to the fact that the network was able to learn what a needle is and recognizes the object on the scans with good signal. It can thus predict with a higher precision which voxels correspond to a needle.

The best results are obtained for $\theta = 0.2$ in average (Fig. 3c). This value is used in the all experiments presented in this paper.

C. Assessment of the fine segmentation

From the fine representation of the needles extremities, the coordinates of the probe end is extracted and compared to the manually obtained coordinates to evaluate the quality of segmentation. The mean distances between predicted and true coordinates across all needles detected in a given scan are shown in Table II for the 8 patients of the test set. We compare different coarse segmentation methods: the proposed 3D U-Net and a simpler intensity thresholding method (the scan is first thresholded with a high threshold of 500 to determine the best needle, then with a low threshold of 100 around the best needle to detect the other curves).

TABLE II
PERFORMANCE OF THE FINE NEEDLE SEGMENTATION WITH A SIMPLE INTENSITY THRESHOLDING METHOD AND WITH THE PROPOSED 3D U-NET

Patient	***Average distance* (mm) with thresholding***	***Average distance* (mm) with 3D U-Net***
1	76.98	1.88
2	0.97	1.44
3	1.25	0.72
4	1.11	1.17
5	7.69	5.80
6	4.99	6.07
7	3.66	2.18
8	18.13	1.56

* Over all needles in the given CBCT scan

Using the commodity hardware, the coarse segmentation requires about 2 minutes and the fine segmentation algorithms less than 10 seconds to compute the representation of the needles extremities. The speed of computation could be further enhanced by GPU acceleration during the coarse segmentation but is already suitable for clinical settings.

Including the 3D U-Net in the coarse-to-fine approach allows a significant improvement in the needle localisation in many cases. Indeed, in average, the distance between prediction and ground truth is of 2.61 mm against 14.35 mm without the deep learning method (Table II). Also, the statistical range is smaller with the 3D U-Net, indicating that the proposed method is more robust. Finally, the algorithm always outputs a needle representation whereas, without the deep learning based segmentation, the method may crash.

D. Numerical assessment of the IRE procedure

The fine representation of needle extremities is then entered in the standard static non linear model of the electric field to produce visuals such as Fig. 4.

Two isolines corresponding to 400 V/cm and 650 V/Cm are represented on the CBCT with needles. These thresholds are supposed to be sufficient for reversible and irreversible electroporation respectively in the liver. The tumor, manually segmented by the surgeon, is added to the image so that a quick visual assessment of the IRE procedure is possible.

IV. DISCUSSION AND CONCLUSION

The proposed method greatly enhances localisation of needles utilized in the IRE process. Improving the quality of needle detection with a coarse-to-fine approach and a convolutional network leads to more precise simulation of the electric field distribution.

The U-Net architecture and learning strategy were adapted to highly skewed 3D data and to the clinical settings in which it will be used. Firstly, the Dice loss associated with binary cross-entropy constitutes the loss function instead of a weighted cross-entropy. The Dice coefficient is also used for the evaluation of the segmentation quality. Experiments using the weighted Dice coefficient, a modified Dice coefficient that includes weights for false negatives and for false

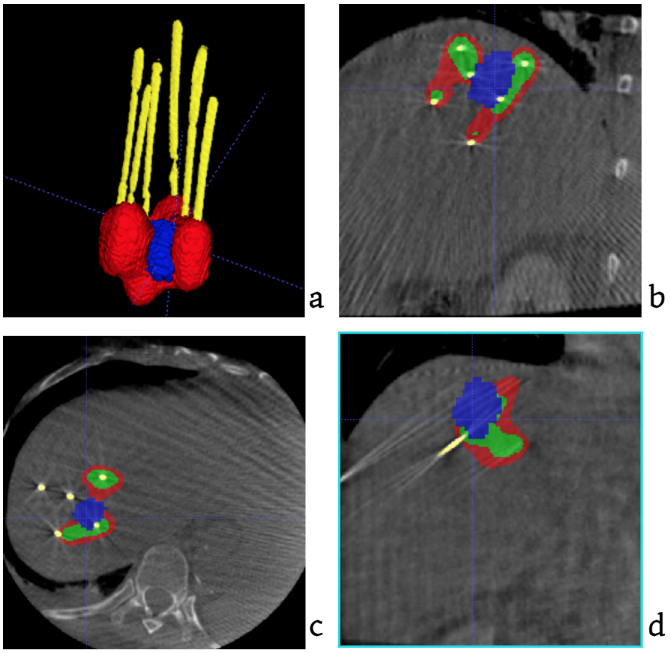

Fig. 4. Computation of the delivered electric field on the CBCT scan with needles in 3D (with slices): the tumor is in blue, the 400 V/cm isoline in red, the 650 V/cm isoline in green and the needles, as detected by the 3D U-Net, are represented in yellow. We note that the coverage of the tumor in this case is not perfect: some of the red sphere, representing the tumor, is not within the isoline in white.

positives in the denominator, to evaluate a network trained with the weighted Dice loss function confirmed that the Dice coefficient does not favor networks trained with the Dice loss function. It is thus suitable for the evaluation of the segmentation quality, including when comparing different loss functions. Secondly, the original patch-overlapping strategy [3] is not used. Preliminary experiments showed that increasing patch-overlapping does not improve segmentation quality but nearly triples processing time. Thirdly, changing the inference threshold enables to give more importance to the needle information, less represented in the data. Lastly, selecting patches considering the needle information they contain allows to focus the network learning on the needle information or on the background information depending on which set is used. However, this last option makes the training more unstable. Using a deeper network stabilises slightly the training but leads to a decrease in segmentation quality, probably due to over-fitting.

To conclude, our contribution consists in a learning strategy allowing good counter-balancing of highly skewed data by using an appropriate loss function, by carefully selecting images to be fed to the network during learning based on the class information they provide and by modifying the inference threshold to take advantage of the different amount of class information learned. Indeed, the main control parameters in this case are identified as the loss function, the inference threshold and the ratio of “needle” information to “background” information during training. The gain in segmentation quality leads to a more precise modelling of the IRE process.

References

[1] O. Gallinato, B. Denis de Senneville, O. Seror, C. Poignard, Numerical workflow of irreversible electroporation for deep-seated tumor, Physics in Medicine and Biology, 2019

[2] O. Gallinato, B. Denis de Senneville, O. Seror, C. Poignard, Numerical Modelling Challenges For Clinical Electroporation Ablation Technique of Liver Tumors, Math. Model. Nat. Phenom., 2020

[3] O. Ronneberger, P. Fischer, T. Brox, U-Net: Convolutional Networks for Biomedical Image Segmentation, 2015

[4] X. Li, et al., H-denseunet: Hybrid densely connected unit for liver and liver tumor segmentation from CT volumes, 2017

[5] E. Vorontsof, et al., Liver lesion segmentation informed by joint liver segmentation, 2017

[6] L. Shapiro, G. Stockman, Computer Vision, Prentice-Hall, Inc., 2001

[7] M. Yi-de, L. Qing, Q. Zhi-bai, Automated image segmentation using improved PCNN model based on cross-entropy. Proc. 2004 Int. Symp. Intell. Multimed. Video Speech Process. 2004, https://doi.org/10.1109/ISIMP.2004.1434171.

[8] C. H. Sudre, W. Li, T. Vercauteren, S. Ourselin, M. J. Cardoso, Generalised Dice overlap as a deep learning loss function for highly unbalanced segmentations. ArXiv170703237 Cs 2017, 10553, 240–248. https://doi.org/10.1007/978-3-319-67558-9_28.

[9] L. R. Dice, Measures of the amount of Ecologic Association Between Species, Ecology, Vol. 26, No. 3, 1945

[10] T. Sugino, T. Kawase, S. Onogi, T. Kin, N. Saito, Y. Nakajima, Loss Weightings for Improving Imbalanced Brain Structure Segmentation Using Fully Convolutional Networks. Healthcare, 9, 938, 2021, https://doi.org/10.3390/healthcare9080938.

[11] W. Li, G. Wang, L. Fidon, S. Ourselin, M. J. Cardoso, and T. Vercauteren, On the Compactness, Efficiency, and Representation of 3D Convolutional Networks: Brain Parcellation as a Pretext Task, in Information Processing in Medical Imaging, Cham, pp. 348–360, 2017, https://doi.org/10.1007/978-3-319-59050-9_28.

[12] K. S. Roberts, A new representation for a line, Computer Vision and Pattern Recognition CVPR’88, 1988

[13] N. Jourabchi, K. Beroukhim, B. A. Tafti, S. T. Kee, E. W. Lee, Irreversible electroporation (NanoKnife) in cancer treatment, Gastrointestinal Intervention, Vol. 3, June 2014

[14] D. Sel, D. Cukjati, D. Batiuskaite, T. Slivnik, L.M. Mir and D. Miklavčič, Sequential finite element model of tissue electropermeabilization. Trans. Biomed. Eng. 2005

[15] A. Collin, S. Corridore, C. Poignard, Floating Potential Boundary Condition in Smooth Domains in an Electroporation Context. In: Suzuki T., Poignard C., Chaplain M., Quaranta V. (eds) Methods of Mathematical Oncology. MMDS 2020. Springer Proceedings in Mathematics Statistics, vol 370. Springer, Singapore, 2021

[16] P. A. Yushkevich, User-guided 3D active contour segmentation of anatomical structures: Significantly improved efficiency and reliability, NeuroImage, 2006

[17] M. Abadi, et al., TensorFlow: Large-scale machine learning on heterogeneous systems, 2015

[18] F. Chollet, et al., Keras, GitHub, 2015

[19] G. Jankowiak, C. Taing, C. Poignard, A. Collin, Comparison and calibration of different electroporation models. Application to rabbit livers experiments, ESAIM : ProcS, Vol.67. 2020.

[20] A. Peyman,B. Kos, M. Djokič, B.Trotovšek,C Limbaeck-Stokin, G. Serša, and D. Miklavčič, Variation in Dielectric Properties Due to Pathological Changes in Human Liver, Bioelectromagnetics, 36, 2015