

# Journal of Medical Imaging

MedicalImaging.SPIEDigitalLibrary.org

## **Anatomically consistent CNN-based segmentation of organs-at-risk in cranial radiotherapy**

Pawel Mlynarski  
Hervé Delingette  
Hamza Alghamdi  
Pierre-Yves Bondiau  
Nicholas Ayache

**SPIE.**

Pawel Mlynarski, Hervé Delingette, Hamza Alghamdi, Pierre-Yves Bondiau, Nicholas Ayache, "Anatomically consistent CNN-based segmentation of organs-at-risk in cranial radiotherapy," *J. Med. Imag.* 7(1), 014502 (2020), doi: 10.1117/1.JMI.7.1.014502

# Anatomically consistent CNN-based segmentation of organs-at-risk in cranial radiotherapy

Pawel Mlynarski,<sup>a,\*</sup> Hervé Delingette,<sup>a</sup> Hamza Alghamdi,<sup>b</sup>  
Pierre-Yves Bondiau,<sup>b</sup> and Nicholas Ayache<sup>a</sup>

<sup>a</sup>Université Côte d'Azur, Inria, Epione Research Team, Nice, France

<sup>b</sup>Université Côte d'Azur, Centre Antoine Lacassagne, Nice, France

**Abstract.** Planning of radiotherapy involves accurate segmentation of a large number of organs at risk (OAR), i.e., organs for which irradiation doses should be minimized to avoid important side effects of the therapy. We propose a deep learning method for segmentation of OAR inside the head, from magnetic resonance images (MRIs). Our system performs segmentation of eight structures: eye, lens, optic nerve, optic chiasm, pituitary gland, hippocampus, brainstem, and brain. We propose an efficient algorithm to train neural networks for an end-to-end segmentation of multiple and nonexclusive classes, addressing problems related to computational costs and missing ground truth segmentations for a subset of classes. We enforce anatomical consistency of the result in a postprocessing step. In particular, we introduce a graph-based algorithm for segmentation of the optic nerves, enforcing the connectivity between the eyes and the optic chiasm. We report cross-validated quantitative results on a database of 44 contrast-enhanced T1-weighted MRIs with provided segmentations of the considered OAR, which were originally used for radiotherapy planning. In addition, the segmentations produced by our model on an independent test set of 50 MRIs were evaluated by an experienced radiotherapist in order to qualitatively assess their accuracy. The mean distances between produced segmentations and the ground truth ranged from 0.1 to 0.7 mm across different organs. A vast majority (96%) of the produced segmentations were found acceptable for radiotherapy planning. © 2020 Society of Photo-Optical Instrumentation Engineers (SPIE) [DOI: [10.1117/1.JMI.7.1.014502](https://doi.org/10.1117/1.JMI.7.1.014502)]

**Keywords:** segmentation; organs at risk; radiotherapy; convolutional neural networks; magnetic resonance image.

Paper 19183R received Jul. 22, 2019; accepted for publication Jan. 17, 2020; published online Feb. 13, 2020.

## 1 Introduction and Related Work

Malignant tumors of the central nervous system cause more than 200,000 deaths per year worldwide.<sup>1</sup> Many brain cancers are treated with radiotherapy, often combined with other types of treatment, in particular, surgery and chemotherapy. Radiotherapy planning requires segmentation of target volumes (visible tumor mass and regions likely to contain tumor cells) and anatomical structures that are susceptible to damage by ionizing radiation exposure during treatment. The segmented volumes are used for computation of optimal irradiation doses, with the objective of maximizing irradiation of cancer cells while minimizing damage of neighboring healthy structures, called organs at risk (OAR). Magnetic resonance images (MRIs)<sup>2</sup> are commonly used for imaging of tumors and organs in the head. In this work, we address the challenging problem of multiclass segmentation of organs in MRI of the brain.

Delineation of OAR is today manually performed by experienced clinicians. Due to a large number of structures to be accurately segmented, the segmentation process takes usually several hours per patient. Manual segmentation represents, therefore, a very high cost and eventually delays the beginning of the therapy. Moreover, a high intraobserver and interobserver variability is observed.<sup>3</sup> Automatic methods for segmentation of OAR are therefore of particular interest. We can distinguish two main types of approaches proposed in the literature.

---


\*Address all correspondence to Pawel Mlynarski, E-mail: [pawel.mlynarski@inria.fr](mailto:pawel.mlynarski@inria.fr)

The first type of methods corresponds to atlas-based approaches.<sup>4-6</sup> The input image is typically registered to one<sup>7,8</sup> or several<sup>9,10</sup> annotated images, from which the segmentation is extrapolated. When multiple atlases are used, the candidate segmentations may be combined, for instance, by voting strategies<sup>10</sup> or by the STAPLE algorithm.<sup>11</sup> An important advantage of atlas-based methods is to produce anatomically consistent results. However, their main drawback is their limited generalization capacity. The important variability between cases results not only from the natural anatomical differences between patients but also from pathological factors. In particular, healthy organs are deformed by growing tumors, which may appear at different locations and which are typically not present in atlases. Some organs may even be missing because of surgeries undergone previously by the patient.

The second group of approaches is based on a discriminative classification of voxels with machine learning models such as random forests<sup>12-14</sup> or convolutional neural networks (CNNs).<sup>15</sup> These discriminative methods are less constrained than atlas-based approaches and may, therefore, better adapt to the diversity of cases. However, in general, voxelwise classifiers may produce results that are inconsistent in terms of shapes and locations of organs.

OAR in the head have complex shapes and are surrounded by other structures sharing similar voxel intensities in MRI. Moreover, there are large differences related to acquisition of MRI, especially when images come from different medical centers. In order to segment organs from MRI, a complex and abstract information has, therefore, to be extracted. CNNs are suitable for this task, as they have the ability to automatically learn complex and relevant image features. In this work, we propose a system based on CNNs for multiclass segmentation of OAR in brain MRI.

In this work, we assume nonexclusive classes, i.e., that one voxel may belong to zero or several classes (Fig. 1). This is in contrast with the majority of segmentation models, which assign one unique label to each voxel following the format of public segmentation challenges such as the BRATS.<sup>16</sup> However, some works addressing OAR segmentation consider nonexclusive classes<sup>17-20</sup> similar to our work. An important difficulty to train machine learning models for multiclass OAR segmentation is the varying availability of ground truth segmentations of different classes among patients, depending on clinical needs. Although some organs, such as the optic nerve, are systematically segmented during radiotherapy planning, annotation of other structures may be available for only a subset of patients. One solution to this problem is to independently train one model per class, as was proposed in some recent deep learning works.<sup>19-21</sup> A limitation of this approach is, however, the need to perform time-consuming trainings for every class while the number of classes of interest may be large. In this work, we propose a loss function and an algorithm to train neural networks for an end-to-end multiclass segmentation, taking into account



**Fig. 1** Segmentation of OAR in radiotherapy planning. (a) T1-weighted MRI acquired after injection of a gadolinium-based contrast agent. (b) Manual annotations of several OAR. In contrast to standard segmentations problems, one voxel may belong to zero or several classes (for instance, the eye and the lens).

the problem of missing annotations. To the best of our knowledge, the only deep learning method for end-to-end multiclass OAR segmentation that addresses this issue is the one proposed<sup>22</sup> for the segmentation of head and neck OAR in CT scans.

The network architecture used in our work is a modified version of 2-D U-Net.<sup>23</sup> As computation of gradients of the loss function by backpropagation represents high computational costs and GPU memory load, large segmentation CNNs often cannot be trained on entire MRIs or CT scans (representing several millions of voxels). Input images are often downsampled<sup>18,24</sup> to allow large 3-D CNNs to capture information from distant regions, i.e., to have a long-range 3-D receptive field. In this work, we have chosen a 2-D architecture to limit the memory load while capturing a long-range spatial context without the need to downsample input images. In particular, downsampling of inputs may affect segmentation of small structures such as the optic nerve or the lens.

Even if most of the proposed deep learning methods for OAR segmentation do not apply anatomical constraints on the output of neural networks, some approaches include shape priors in models. For instance, Tong et al.<sup>25</sup> proposed to learn latent representations of shapes of organs by a stacked autoencoder and to use these learned representations in the loss function of a segmentation network, in order to compare the shape of the output with the shape of the ground truth. References 26 and 27 proposed to adapt triangulated meshes representing organ boundaries to medical images and to use neural networks for regression of distances between centers of triangles and organ boundaries. This type of approach may, therefore, be seen as atlas-based with the use of deep learning for boundary detection.

Spatial relations between anatomical structures in the head could be explicitly modeled using, for instance, graph-based representations as it was proposed in model-based methods.<sup>28,29</sup> However, inclusion of constraints related to connectivity and relative positions of organs in loss functions of CNNs is not trivial due to considerable computational costs. In order to apply such constraints, a neural network would have to segment large regions of the input images during the training phase. Segmentation of large 3-D volumes containing different anatomical structures requires a considerable amount of the GPU memory, as outputs of all layers of the network are stored in the GPU during computation of gradients of the loss function. Moreover, penalization of anatomical inconsistencies during the training does not guarantee anatomically consistent results at the test phase. To the best of our knowledge, none of the proposed deep learning methods explicitly enforces consistency of OAR segmentation in terms of relative positions of organs. However, some methods define regions of interest of organs, for instance by registering the image to a set of atlases.<sup>19</sup>

In our work, we enforce some anatomical constraints in a postprocessing stage, starting from the segmentation produced by majority voting of 2-D CNNs processing the image by axial, coronal, and sagittal slices. In particular, we propose an anatomically consistent segmentation of the optic nerves, with an approach based on the search of the shortest path in a graph, using outputs of neural networks to define weights of edges in the graph. Application of anatomical constraints in postprocessing modules rather than in the deep learning model is mainly motivated by computational costs of CNNs but also by their “black box” aspect.

We consider eight classes of interest, corresponding to anatomical structures systematically segmented during radiotherapy planning for brain cancers: eye, lens, optic nerve, optic chiasm, pituitary gland, hippocampus, brainstem, and brain (including cerebrum, cerebellum, and brainstem). The anatomical structures composed of left and right components (eye, lens, optic nerve, and hippocampus) are seen as one entity by the neural network but are separated in the post-processing step.

Most of the proposed deep learning methods for segmentation of OAR were applied on CT scans in the context of head and neck cancers,<sup>30</sup> i.e., cancers of the upper parts of respiratory and digestive systems (mouth, larynx, and throat). To the best of our knowledge, the only deep learning method for segmentation of OAR in MRIs of the brain is the one proposed in Ref. 27 (MRI T1 and T2).

Our method is tested on a set of contrast-enhanced T1-weighted MRIs acquired in the Centre Antoine Lacassagne in Nice (France). First, our method is quantitatively evaluated on a set of 44 MRIs with provided segmentation of different anatomical structures. Segmentation performances are measured by three different metrics: Dice score, Hausdorff distance, and the mean

distance between the output and the ground truth. Then the segmentations produced by our method on a different set of 50 MRIs are qualitatively evaluated by an experienced radiotherapist. Our system was able to produce segmentations with an accuracy level that was found acceptable for radiotherapy planning in a large majority of cases (96%). The mean distances between the output segmentation and the ground truth for different organs were between 0.1 and 0.7 mm.

## 2 Methods

### 2.1 Deep Learning Model


#### 2.1.1 Network architecture

The architecture used in our work is a modified version of 2-D U-Net,<sup>23</sup> which is composed of an encoding part and a decoding part. The encoding part is a sequence of convolutional and max-pooling layers. The number of feature maps is doubled after each pooling, taking advantage of their reduced dimensions. The decoding part is composed of convolutional and upsampling layers. Feature maps of the encoding part are concatenated in the decoding part in order to combine low-level and high-level features and to ease the flow of gradients during the optimization process. The final convolutional layer (the segmentation layer) of the standard U-Net has two feature maps, representing pixelwise classification scores of the class 0 (“background”) and the class 1. During training, these two final feature maps are normalized by the softmax function.

We adapt this architecture to our problem of multiclass segmentation with nonexclusive classes, where each pixel may belong to zero or several classes. In the following,  $C$  denotes the number of classes (in our experiments,  $C = 8$ ) and the classes are numbered from 1 to  $C$ . In our model, each class  $c$  has its dedicated binary segmentation layer (Fig. 2), composed of two feature maps corresponding to pixelwise scores of the class and of the background. Each segmentation layer takes as input the second to last convolutional layer of U-Net. We use batch normalization<sup>31</sup> in all convolutional layers of the network, except segmentation layers.

#### 2.1.2 Training of the model

Our loss function and training scheme were designed to deal with class imbalance and the problem of missing annotations (for a given image, the ground truth is available only for a subset of classes). In a given training image  $i$ , each pixel  $(x, y)$  has three possible labels for the class  $c$ : 0



**Fig. 2** Architecture of our model. The rectangles represent layers and their height represents the sampling factor (increasing with max-poolings, decreasing with upsamplings). The numbers of features maps is specified below layers. The proposed model is a modified version of U-Net, having one segmentation layer per class in order to perform an end-to-end multiclass segmentation with nonexclusive classes.

(negative), 1 (positive), or  $-1$  (unknown). If the ground truth segmentation of the class  $c$  is unavailable for the image  $i$ , all pixels are labeled as unknown for the class  $c$  by default. However, missing annotations may be partially reconstructed from segmentations of other classes. For example, if the segmentation of the “lens” class is not available but the “eye” class is segmented, all pixels outside the eye may be labeled as negative for the lens.

Given a training batch of  $M$  images and the estimated parameters  $\theta$  of the network, the segmentation layer of the class  $c$  is penalized by the following loss function, which can be seen as pixelwise cross-entropy with adaptative weights. Let us note  $N_0^c$ ,  $N_1^c$ , and  $N_{-1}^c$  the numbers of pixels labeled, respectively, 0, 1, and  $-1$  for the class  $c$  in the training batch. The weight  $w_{(x,y)}^i$  of the pixel  $(x, y)$  of the image  $i$  has three possible values, according to the label of the pixel. If the label is unknown, then  $w_{(x,y)}^i = 0$ . If its label is 1, then  $w_{(x,y)}^i = t_c/N_1^c$ , where  $0 < t_c < 1$  is a fixed hyperparameter, which we call the “target weight.” If the pixel is labeled 0, then  $w_{(x,y)}^i = (1 - t_c)/N_0^c$ . The introduced hyperparameter  $t_c$  controls, therefore, the relative weight of positive and negative pixels of the class  $c$  (positive pixels have the total weight of  $t_c$  and negative pixels have the total weight of  $1 - t_c$ ). This type of weighting strategy has been used in our previous work<sup>32</sup> to counter the problem of class imbalance. The loss function of the segmentation layer of the class  $c$  is defined by  $\text{loss}_c(\theta) = -\sum_{i=1}^M \sum_{(x,y)} w_{(x,y)}^i \log[p_{i,(x,y)}^l(\theta)]$ , where  $p_{i,(x,y)}^l$  is the softmax score given by the network for the ground truth label  $l$  of the pixel.

The loss function of the model is a convex combination of losses of all segmentation layers:  $\text{loss}(\theta) = (1/C) \sum_{c=1}^C \text{loss}_c(\theta)$ .

We propose a sampling strategy to construct training batches so that there are positive and negative pixels for each of the  $C$  classes in each training batch.

For each image of the training database, we precompute bounding boxes of all classes with provided segmentations. For bilateral classes such as the eyes, there are generally two bounding boxes per image corresponding to left and right components, unless one of the components is missing (e.g., an organ removed by surgery). The precomputed bounding boxes are used during the training in order to sample patches containing positive pixels of different classes.

At the beginning of the training, for each class  $c$ , we construct a list  $I_c$  of training images with provided ground truth segmentation of the class  $c$ . To sample a 2-D patch which is likely to contain positive pixels of the class  $c$ , we randomly choose an image  $i$  from  $I_c$  and a random point  $(x, y, z)$  from the bounding box (or two bounding boxes if the class has left and right components) of the class  $c$  in the chosen image. Once the point is chosen, a 2-D patch centered on this point is extracted from the image  $i$  and segmentations of all available classes are read. In the following, we refer to this procedure as extracting a patch centered on the class  $c$ .

We assume that the number of images in each training batch ( $M$ ) is larger than the number of classes  $C$ , in order to be able to sample at least one image/patch centered on each of the classes. Each training batch is constructed as follows. The first  $C$  images of the batch are centered, respectively, on each of the  $C$  classes. At this stage, the batch is likely to contain positive and negative pixels of each class. The remaining  $M - C$  images may be chosen randomly or be centered on larger classes. In our case,  $C = 8$ ,  $M = 10$ , and the last images are centered on the largest class we segment the brain, whose bounding box occupies almost an entire volume of the head.

As the model is trained for multiclass segmentation with nonexclusive classes, several binary segmentation maps have to be read for each input image in each iteration of the training. If the ground truth segmentations are not optimally stored in the memory, these reading operations may considerably slow down the training. The ground truth label of a given pixel can be represented by one bit (0 or 1). However, to store binary segmentation masks in commonly used formats such as HDF5,<sup>33</sup> each label would have to be represented by at least one byte. We propose, therefore, to store multiclass segmentations in a specifically encoded format, where every bit represents a label of a given class  $c$ . A binary segmentation mask of the class  $c$  is retrieved by the “bitwise and operation” between the encoded multiclass segmentation and the code of the class, corresponding to a power of 2.

The size of extracted 2-D patches should be chosen according to the capacities of the GPU. In our experiments, the training batches were composed of 10 patches of size  $230 \times 230$ . Given that



in our network we use unpadded operations (convolutions, max-poolings, etc.), the dimensions of the outputs of segmentation layers are considerably smaller.

The model is trained with a variant of stochastic gradient descent with momentum presented in our previous work.<sup>32</sup> The main characteristics of this algorithm are that gradients are computed over several batches in each iteration of the training, in order to use many training examples despite GPU memory limitations.


## 2.2 Postprocessing and Enforcing Anatomical Consistency

Fully convolutional neural networks such as our model produce segmentations by individually classifying every voxel based on intensities of voxels within the corresponding receptive field. Such classification is performed by extracting powerful and automatically learned image features. However, as this classification is performed on a voxel by voxel basis, there is no guarantee of obtaining an anatomically consistent result, especially when the number of training images is limited. In particular, CNNs do not explicitly take into account aspects such as relative positions of different structures or adjacency of voxels belonging to the same structure. Including constraints related to these aspects in loss functions of neural networks or conceiving architectures that produce anatomically consistent results is difficult, in particular because of computational costs (need to simultaneously segment large 3-D regions of input images). We propose, therefore, to improve consistency of segmentations in a postprocessing step. We also separate left and right components of classes such as the eye, as these components are considered separately for radiotherapy planning.

We combine, by majority voting, segmentations produced by three networks trained, respectively, on axial, coronal, and sagittal slices. The goal of this combination is to take into account the three dimensions and to improve the robustness of the method. We subsequently apply a few rules described in the following, in order to correct some observed inconsistencies.

### 2.2.1 Segmentation of the brain

Brain (including the cerebrum, the cerebellum, and the brainstem) is the largest class to be segmented. For various reasons, some voxels within this structure may be inconsistently classified as negative by networks, which appears as “holes” in the segmentation or unrealistically sharp borders. We propose, therefore, a procedure that we call triplanar hole-filling (Fig. 3). For each axial, coronal, and sagittal plan of the 3-D segmentation, we compute connected components of the background (negative voxels) and we remove components (changing their label from 0 to 1), which are not connected to the border of the plan. The reason of applying this procedure in 2-D is that some holes may easily be connected to the outside of the class in 3-D.



**Fig. 3** (a) Example of “holes” in the original output segmentation (left image) on a test example. (b) Segmentation obtained after our postprocessing (triplanar hole-filling).

The bounding box of the segmentation of the brain is subsequently used to separate left and right components of bilateral classes. Note that the head of the patient may appear at different locations of the image, depending on acquisition conditions and performed preprocessings. For a given class expected to have left and right components (eye, lens, optic nerve, and hippocampus), barycenter of each connected component is computed. In order to decide which side corresponds to a connected component, the coordinate  $x$  (right–left) of its barycenter is compared to min and max coordinates  $x$  of the bounding box of the brain.

### 2.2.2 Segmentation of the visual system


We propose an anatomically consistent segmentation of the visual system (eyes, lenses, optic nerves, and chiasm), starting from the segmentations predicted by neural networks.

The eye is probably the less challenging organ for automatic segmentation as it has a simple spherical shape. However, some false positives are possible, especially in cases where an eye has been removed by surgery, resulting in false positives within the orbit. We propose, therefore, to remove connected components of eye segmentation whose volume is below an expected minimum value, which is set to  $4 \text{ cm}^3$ .

We constraint segmentation of the lenses to be inside the eyes, i.e., we assign the 0 label to all voxels outside the predicted masks of the eyes. Segmentation of the optic chiasm is obtained by taking the largest connected component of the segmentation predicted by the networks. We distinguish left and right sides of the chiasm in order to compute landmarks for segmentation of the two optic nerves as described in the following.

Segmentation of the optic nerve in MR images is particularly challenging as the nerve is thin and may have an appearance similar to neighboring structures at some locations. However, it has a rather regular shape, which can be seen as a tube connecting an eye and the optic chiasm. The nerve is generally very visible at some locations, particularly close to the eye. A human expert is able to track the trajectory of the nerve to distinguish it from neighboring structures at more difficult locations. Based on this observation, we propose a graph-based algorithm for segmentation of the optic nerves in order to guarantee connectivity between the eyes and the optic chiasm and to decrease the number of false positives. The algorithm is based on the search of the shortest path between two nodes in a graph. Outputs of neural networks are used to define weights of the edges in the graph. The different steps of the algorithm (applied separately for left and right nerves) are described below.

First, we detect landmarks corresponding to the two endpoints of an optic nerve based on the initial segmentation of the visual system produced by neural networks (Fig. 4). The first



**Fig. 4** Approximate position of the optic nerve landmarks (displayed on the same axial slice) found by the system on a test example. For each of the optic nerves, the graph-based algorithm ensures the connectivity between the two landmarks.




landmark of the left optic nerve is the barycenter of  $P$  points initially predicted as the left optic nerve and which are closest to the left eye. The second landmark is similarly computed but searching  $P$  points of the left side of the optic chiasm, which are the closest to the initial prediction of the left optic nerve. We take the barycenter of several points (in our experiments  $P = 30$ ) in order to obtain a point that is more likely to be close to the centerline of the nerve. If the detected chiasm landmarks for the two optic nerves are abnormally close, the procedure is applied only for one nerve, connecting the landmark with the closest eye.

Before applying the graph-based algorithm, we refine the initial segmentation of the optic nerves based on voxel intensities (specific to each image). In fact, the optic nerves are surrounded by fat, which appears hyperintense on MR T1-weighted images and can be rather easily distinguished from the optic nerve. We compute an approximate range of intensities of voxels of the fat by computing the 98% quantile of a small volume surrounding the eye–nerve landmark. Voxels whose intensities are above 80% of this value are classified negative for the optic nerve in order to eliminate common false positives.


Given the two computed landmarks and the refined initial segmentation, we estimate the centerline of the optic nerve (Fig. 5) by computing the shortest path in an oriented graph. The nodes of the graph correspond to voxels within a region of interest (cuboids containing the two landmarks) and which are reachable from the starting point. The connectivity of nodes is defined by adjacency of voxels with increasing  $y$  coordinate, i.e., the child of the node  $(x, y, z)$  are nodes  $(x + d_x, y + 1, z + d_z)$  with  $d_x \in \{-1, 0, 1\}$  and  $d_z \in \{-1, 0, 1\}$ . We, therefore, assume strictly increasing  $y$  of the centerline toward the second landmark (from anterior to posterior).

Each node  $(x, y, z)$  of the graph has its associated cost based on three criteria (listed by decreasing importance):

- Label  $l_{(x,y,z)}$  initially assigned to the voxel  $(x, y, z)$ . A strong penalty is applied to voxels predicted as negative, in order to force the centerline to pass by points initially predicted as positive. The associated cost is  $c_{(x,y,z)}^{\text{label}} = 0$  if  $l_{(x,y,z)} = 1$  and  $c_{(x,y,z)}^{\text{label}} = C^l$  otherwise, where  $C^l$  is a fixed number controlling the importance of this cost (we set  $C^l = 100$ ).
- If the predicted label  $l_{(x,y,z)}$  is positive, distance  $d_{\text{border}}$  to the closest point classified as negative. The penalty is inversely proportional to this distance, to give priority to points that are far from predicted borders of the optic nerve (preference to central points). This cost is expressed by  $c_{(x,y,z)}^{\text{border}} = 0$ , if  $l_{(x,y,z)} = 0$  and  $c_{(x,y,z)}^{\text{border}} = R - d_{\text{border}}$  otherwise, where  $R$  is the radius of a search zone around the voxel  $(x, y, z)$ . As the visible nerve is larger close to the eye,  $R$  varies with the coordinate  $y$  (interpolation between  $R = 7$  and  $R = 3$ , expressed in number of voxels).
- Distance  $d_{\text{target}}$  to the target point (i.e., the nerve–chiasm landmark). The penalty is proportional to this distance in order to force the centerline to immediately go toward the target point if other criteria do not give priority to some points. In particular, when one part of the optic nerve has not been initially detected (negative voxels), the line should go in the



**Fig. 5** The centerlines of the optic nerves computed by our system on a test example (displayed on three different coronal slices). We assume one point of the centerline for each coronal slice between the two landmarks of the optic nerve.



**Fig. 6** Use of mathematical morphology for reduction of false positives. (a) A coronal patch centered on an optic nerve. (b) Result obtained by the system on a test example without using mathematical morphology. (c) Result obtained after application of morphological opening followed by taking the largest connected component.

direction of the target point. The associated cost is  $c_{(x,y,z)}^{\text{distance}} = C^l d_{\text{target}}$ , where  $C^l$  controls the importance of this cost. We fixed  $C^l = 0.001$ , to make it negligible compared to the previous criteria.

The cost of the node  $(x, y, z)$  is the sum of the three components:  $c_{(x,y,z)} = c_{(x,y,z)}^{\text{label}} + c_{(x,y,z)}^{\text{border}} + c_{(x,y,z)}^{\text{distance}}$ . The introduced cost determines the weights of edges in the graph. A directed edge between the point  $(x_1, y_1, z_1)$  and  $(x_2, y_2, z_2)$  has the weight of  $c_{(x_2,y_2,z_2)}$ . The shortest path between nodes corresponding to the two endpoints of the optic nerve is computed by Dijkstra’s algorithm.<sup>34,35</sup> The start point is the eye–nerve landmark as the optic nerve is generally well visible close to the eye. To the best of our knowledge, our approach is the first to combine deep learning with the search of the shortest path in a graph for segmentation of tubular anatomical structures. However, the idea of computing optimal distances for segmentation of tubular structures appears in interactive level-set methods.<sup>36–38</sup> The objective of these methods is to find a geodesic between two points in the image chosen by the user. The Eikonal equation is constructed based on voxel intensities and contrasts, and the problem is solved by fast marching,<sup>39</sup> similar to Dijkstra’s algorithm. Application of methods based only on image intensities may be difficult for segmentation of the optic nerves in MRI due, for instance, to the noise in images and local inhomogeneity of intensities within the optic nerve.

The final segmentation of the optic nerve is constructed from the centerline. As the optic nerve has a variable thickness, around each point  $(x, y, z)$  of the centerline, we consider two spherical volumes  $S_{(x,y,z)}^1$  and  $S_{(x,y,z)}^2$  with associated radii  $R_1 \leq R_2$ . All voxels within  $S_{(x,y,z)}^1$  are classified positive (optic nerve). Voxels of  $S_{(x,y,z)}^2$  that are not within  $S_{(x,y,z)}^1$  are classified positive only if they were positive in the original segmentation. We fixed  $R_1 = 2.5$  and  $R_2$  corresponds to the radius  $R$  defined previously (large close to the eye, smaller close to the optic chiasm).

Finally, we apply mathematical morphology<sup>40</sup> to reduce false positives corresponding to structures that are “attached” to the optic nerve and have a similar appearance. As these false positives are often connected to the correct segmentation by thin segments (Fig. 6), we apply the morphological opening with three 1-D structuring elements of size 2 in the three directions and we take the largest connected component.

### 3 Experiments

#### 3.1 Data and Preprocessing

We constructed a database of contrast-enhanced T1 MRIs acquired in the Centre Antoine Lacassagne (Nice, France), which is one of the three cancer centers in France equipped with proton therapy<sup>41</sup> systems.

**Table 1** Numbers of provided ground truth segmentations for different classes (in the database of 44 MRIs).

	Number of segmentations
Hippocampus	39
Brainstem	39
Eye	41
Lens	34
Optic nerves	40
Optic chiasm	41
Pituitary gland	29
Brain	37

The database contains 44 MRIs with provided segmentations of OAR and 50 nonannotated MRIs. The annotated images are used for training and cross-validated quantitative evaluation. For each scan, the ground truth segmentation was provided only for a subset of classes. The numbers of available segmentations for each class are reported in Table 1. The images without annotations are used for qualitative evaluation by a radiotherapist, as described in Sec. 3.4.

The images were originally provided in Dicom format<sup>42</sup> and were heterogeneous in terms of image intensities and geometrical properties such as size, spatial resolution, and the visible part of the head. The ground truth segmentations were the ones used for routine radiotherapy planning and were provided in Dicom RT-Struct files,<sup>43</sup> representing coordinates of polygons corresponding to contours of anatomical structures.

In order to use these images, we performed the following preprocessings. We used 3-D Slicer<sup>44</sup> and its extension SlicerRT<sup>45</sup> to generate 3-D volumes (in “nrrd” format) from Dicom slices and to generate binary label masks from RT-Struct files. All images were resampled to the same spatial resolution,  $0.7 \times 0.7 \times 0.9$  (isotropic in axial slices, 0.9 spacing between slices) and then resized to dimensions  $320 \times 365 \times 200$ . The 200 axial slices start from the top of the head, i.e., if an image originally has more than 200 axial slices, the bottom slices (close to the neck) are ignored. However, the input images had generally around 150 axial slices and the bottom slices were filled with zeros. To approximately normalize the image intensities, first we compute the maximum of an image, which is likely to be reached by a point on the fat or contrast-enhanced blood vessels. Then all voxel values are divided by the value of the maximum and multiplied by a fixed constant.

### 3.2 Metrics for Quantitative Evaluation

To quantitatively evaluate our system, we perform fivefold cross-validation on the set of 44 annotated MRIs. In each fold, 80% of the database is used for training and 20% is used for test. For each class of interest, two results are reported. First, we report results obtained with our model trained on axial slices (denoted “U-Net multiclass, axial” in the following), i.e., the raw output of the neural network, without postprocessing. Then we report results obtained after majority voting and postprocessing (denoted “final result” in the following).

The first metric we use is the Dice score, which measures the voxelwise overlap between the output and the ground truth segmentation. An important limitation of this metric is that it gives the same importance to very close and very distant mismatches. As the ground truth is often uncertain and noisy close to the boundaries of structures, the Dice scores are generally considerably lower for small structures. This is why, in addition to raw Dice scores, we also report results (Dice, sensitivity, and specificity) obtained when a margin of one voxel is allowed, i.e., ignoring mismatches on the borders of the ground truth. This assumption means that a false positive on a voxel  $(x, y, z)$  that is directly neighboring with the ground truth segmentation is

ignored, i.e., it is neither counted as false positive nor true positive. Similarly, a false negative (nondetection) on the border of the ground truth is ignored.

The second used metric is the undirected Hausdorff distance expressed in millimeters (the coordinates of points are expressed in real values). The Hausdorff distance measures the length of the farthest mismatch between the output and ground truth (false positive or false negative). It is, therefore, useful to assess the consistency of the result, i.e., presence of very distant mismatches. However, its limitation is that it only measures the value of the maximal distance, and therefore, one misclassified voxel is sufficient to considerably increase the Hausdorff distance.

Therefore, we also measure the mean distance between the output segmentation  $A$  and the ground truth  $B$ , defined as follows:

$$M(A, B) = \frac{1}{|A| + |B|} \left[ \sum_{a \in A} \inf_{b \in B} d(a, b) + \sum_{b \in B} \inf_{a \in A} d(b, a) \right], \quad (1)$$

where  $d$  is the Euclidean distance.

### 3.3 Quantitative Results

The mean distances between produced segmentations and the ground truth segmentation ranged from 0.08 (for the brain) to 0.69 mm (for the pituitary gland), as reported in Table 5. The results are variable across the different organs, according to their size, the number of ground truth segmentations available for training and the overall complexity of the segmentation task.

The Dice scores are usually higher for large anatomical structures such as the brain and the brainstem. In particular, the borders of the ground truth are usually very uncertain, which represents a problem for quantitative evaluation for smaller classes. In large classes, the border region is small compared to the entire volume of the class, and therefore, the mismatches on borders do not cause large drops of the metric. The highest Dice score was obtained for the brain (Dice score of 96.8). The lowest performances were obtained for the pituitary gland (mean Dice of 58, mean distance of 0.69 mm between the output and the ground truth). Segmentation of the pituitary gland is particularly challenging as it is small and difficult to be differentiated from surrounding structures. Moreover, the pituitary gland was the class with the lowest number of training examples (29 annotated cases, i.e., around 23 training cases in each of the five folds).

To take into account the uncertain borders of the ground truth, we also reported Dice scores, sensitivity, and specificity ignoring mismatches on the border of the ground truth, as described previously. As most mismatches between the outputs and the ground truth are on noisy borders of organs, there is a considerable difference between the raw Dice score (Table 2) and the Dice score with tolerance to one voxel (Table 3).

**Table 2** Mean Dice scores (fivefold cross-validation) obtained on a set of 44 MRIs. “Final result” denotes the result obtained after majority voting and postprocessing.

	U-Net multiclass, axial	Final result
Hippocampus	69.2	71.4
Brainstem	88.1	88.6
Eye	88.3	89.6
Lens	55.8	58.8
Optic nerves and chiasm	63.9	67.4
Pituitary gland	53.6	58.0
Brain	96.5	96.8

**Table 3** Mean Dice score (fivefold cross-validation), sensitivity, and specificity with tolerance to one voxel (ignoring mismatches on the borders due to the uncertainty of the ground truth).

	Dice score	Sensitivity	Specificity
Hippocampus	88.2	92.7	85.0
Brainstem	95.1	95.5	95.6
Eye	97.5	98.3	96.8
Lens	82.1	88.2	78.4
Optic nerves and chiasm	91.1	96.2	87.1
Pituitary gland	79.7	83.3	77.5
Brain	98.6	98.0	99.4

**Table 4** Hausdorff distances in millimeters (fivefold cross-validation). “Final result” denotes the result obtained after majority voting and postprocessing.

	U-Net multiclass, axial	Final result
Hippocampus	42.1	6.9
Brainstem	45.5	7.8
Eye	75.9	3.0
Lens	31.0	3.7
Optic nerves and chiasm	76.7	6.3
Pituitary gland	52.5	4.6
Brain	30.4	9.8

However, the measured Hausdorff distances (Table 4) are higher for large classes. The highest mean Hausdorff distance is observed for the brain, for which it is almost equal to 1 cm.

The combination of neural networks (trained, respectively, on axial, coronal, and sagittal slices) by majority voting improved almost all metrics. The improvements were particularly large for the Hausdorff distance (Table 4) and the mean distance (Table 5). We observe that the majority voting removes almost all distant false positives and yields more robust results than

**Table 5** Mean distances in millimeters (fivefold cross-validation).

	U-Net multiclass, axial	Final result
Hippocampus	0.97	0.66
Brainstem	0.26	0.26
Eye	0.35	0.11
Lens	1.29	0.63
Optic nerves and chiasm	1.09	0.48
Pituitary gland	2.45	0.69
Brain	0.07	0.08




a raw output of one neural network. The results were subsequently improved by additional postprocessings.

The postprocessing of the eyes consisted in setting a lower bound on the physical volume of the output segmentation. This simple procedure allowed to remove false positives and decreased the mean Hausdorff distance from 12.2 mm (result of the majority voting) to 3 mm.


The postprocessing of the optic nerve decreased the number of false positives and enforced connectivity between the eyes and the chiasm, as described in Sec. 2.2.2. False positives are removed when they are either too far from the centerline, hyperintense in T1-weighted MRI (fat surrounding eyes) or are disconnected from the main connected component after application of morphological opening removing thin segments. The Dice score with one-voxel tolerance increased from 89.6 (result of the majority voting) to 91.1 (after postprocessing) for the optic nerves and chiasm. The raw Dice score increased from 66.3 to 67.4.

The postprocessing of the brain consisted in taking the largest connected component and filling the “holes” of the segmentation in axial, coronal, and sagittal planes. As these holes are usually small compared to the whole volume of the class (occupying a large part of the image), the variation of the metrics is limited. The Dice score increased from 96.7 to 96.8 and the Hausdorff distance decreased from 10.2 to 9.8.

To the best of our knowledge, the only deep learning work for segmentation of OAR in MRI is the one proposed,<sup>27</sup> which reported cross-validated results (mean distances in mm) on a set of 16 MRIs. The authors used a model-based segmentation<sup>46</sup> combined with a neural network for detection of boundaries of anatomical structures. The results reported by the authors for the



**Fig. 7** Segmentation of the hippocampus produced by our system on a test example (three orthogonal slices passing by the same point). The output segmentation is represented by the green region, and the ground truth annotation is represented by the red contour.



**Fig. 8** Segmentation of the brainstem produced by our system on a test example (three orthogonal slices passing by the same point). The output segmentation is represented by the green region, and the ground truth annotation is represented by the red contour.


anatomical structures we also segment are: 0.608 mm for the brainstem, 0.563 mm for the eyes, 0.268 mm for the lenses, and 0.41 mm for the optic nerves and chiasm. Overall, the ranges of mean distances are, therefore, comparable to the ours.

Examples of the output segmentations (comparison to the ground truth) for the hippocampus, the brainstem and the optic nerve are displayed on Figs. 7–9.

### 3.4 Qualitative Evaluation by a Radiotherapist

The segmentations produced by our system on a set of 50 nonannotated MRIs are qualitatively evaluated by an experienced radiotherapist in order to assess their accuracy and utility for radiotherapy planning. For each of the 50 patients, the radiotherapist qualitatively evaluates the segmentations produced by our system for 12 anatomical structures (counting separately left and right components of bilateral classes), i.e., 600 segmentations are evaluated in total. The results are evaluated visually: no comparative dosimetric planning has been performed for corrected and uncorrected segmentations. The segmentations are displayed with 3-D Slicer.<sup>44</sup> Each of the 600 segmentations is assigned to one of the following categories.

- *Accept.* The radiotherapist would keep the segmentation for radiotherapy planning without any changes.
- *Accept, minor modifications.* The segmentation is still acceptable for radiotherapy planning, i.e., some minor errors are observed but leaving out the modifications should not affect the dose distribution after the addition of safety margins of target volumes and OAR: planning target volumes (PTV) and planning OAR volumes (PRV).



**Fig. 9** Segmentation of the optic nerves produced by our system on a test example (three orthogonal slices passing by the same point). The output segmentation is represented by the green region, and the ground truth annotation is represented by the red contour.

- *Accept, major modifications.* The segmentation has necessarily to be corrected, i.e., the resulting PTV and PRV differences before and after correction would have an important impact on the dose distribution (even if only few voxels are misclassified). The segmentation is, however, still good enough to be kept, i.e., it is less time-consuming to perform the necessary modifications than segmenting the structure from the beginning.
- *Reject.* The segmentation has failed and keeping it would not save time compared to manually segmenting the structure from the beginning.
- *Not assigned.* The structure is absent (e.g., organ removed by surgery) or invisible in the image because of a tumor.

The results are summarized in Table 6. 73% of the segmentations were assigned to the category “accept,” i.e., would be kept for radiotherapy planning without any modifications. Approximately 23% of the segmentations were assigned to the second category, i.e., acceptable for radiotherapy planning but with recommendation to perform some minor corrections, usually on extremities of organs. The system produced, therefore, satisfactory segmentations in a large majority of cases. It was able to correctly delineate organs despite the important difficulties such as the presence of tumors and the resulting mass effects, motion artifacts in MRI, different orientations of heads of patients, and anatomical modifications resulting from previous surgeries undergone by the patient (removed tissues).

Segmentations of the eyes had the highest rate of immediate acceptance: 93 out of 100 segmentations were assigned to the accept category. The only segmentation that required a major

**Table 6** Clinical evaluation by a radiotherapist on 50 test cases.

	Accept	Accept, minor corrections	Accept, major corrections	Reject	N/A
Hippocampus left	39/50	8	1	2	0
Hippocampus right	45/50	5	0	0	0
Brainstem	22/50	26	1	1	0
Eye left	48/50	2	0	0	0
Eye right	45/50	4	1	0	0
Lens left	39/50	7	4	0	0
Lens right	42/50	6	2	0	0
Optic nerve left	44/50	6	0	0	0
Optic nerve right	40/50	10	0	0	0
Optic chiasm	19/50	26	4	1	0
Pituitary gland	19/50	25	3	0	3
Brain	36/50	14	0	0	0
Total	438/600	139	16	4	3

modification was a case with a lesion inside the eye, possibly the polypoidal choroidal vasculopathy. The lesion was not classified by the system as part of the eye, and therefore, one part of the eye was not detected. The minor modifications recommended for other cases were generally to correct few nondetected voxels on the border of the eye (top or bottom axial slices) or few false positives on the anterior part of the orbit.

All segmentations of the optic nerves were found acceptable for radiotherapy planning: 84 out of 100 segmentations were assigned to the accept category and the remaining 16 cases required only minor corrections. Most of the minor errors were nondetections for few voxels on the extremity of the optic nerve close to the eye (e.g., on the top axial slice). There was also at least one case of false positives on the neighboring arteries, close to the optic chiasm.

Even if in the previous, quantitative evaluation, the metrics for the lenses were significantly lower than for other structures, most of their segmentations on the set of 50 MRIs were found satisfactory by the radiotherapist. Minor corrections were required in 13 out of 100 cases and major corrections were required in 6 cases. Most of the problems were nondetections, for instance, observed in cases where the patient looks to the side and the system does not detect one side of the lens. The lenses are very small structures, and their visibility is highly impacted by motion artifacts in MRI.

For the optic chiasm, corrections were more frequently required but were usually minor: 19 out of 50 cases were assigned to the accept category and 26 cases required minor corrections. The minor errors were often false positives on the hypothalamus (the same issue was observed in the ground truth used to train the model) and sometimes on arteries neighboring the chiasm. The major corrections (four cases) were mainly nondetections of a small subpart of the beginning of an optic nerve. In fact, even if only a small number of voxels is not detected (false negatives), the corrections are necessary as an excessive irradiation of one part of the optic nerve could make the entire nerve dysfunctional.<sup>47</sup> One segmentation was rejected due to nondetection of one part of the chiasm. This error appeared in a challenging case where the anatomy of the patient was modified by an important mass effect caused by a tumor.

Similar performances were obtained for the pituitary gland, located below the optic chiasm. Most of the minor (26 cases) and major (3 cases) required corrections correspond to nondetections, typically on the 1–2 lowermost slices. In at least 2 cases, few false negatives were observed

on the pituitary stalk (also observed in some ground truth segmentations used for training), which is the connection between the pituitary gland and the hypothalamus.

Even if segmentation of the hippocampus is difficult (low contrast with neighboring structures), in our evaluation it had one of the highest acceptance rates, with 84 segmentations in the accept category. However, it is also the only structure for which more than one segmentation was rejected. The two rejected segmentations correspond to cases where a large tumoral mass has grown near to the hippocampus, causing an edema having a similar intensity in T1-weighted MRI. Moreover, the tumors had a large necrotic core, which may be confused with a ventricle by the system. In other cases, the required corrections (mostly minor) correspond usually to false positives (in particular on the amygdales, neighboring hippocampi, and having a similar intensity in MRI T1) or some nondetections on the extremities of the hippocampus.

For the brainstem, 48 out of 50 segmentations were found acceptable for radiotherapy planning but required minor modifications in approximately half cases. The required corrections (false positives or nondetections) were almost exclusively on the uppermost axial slices (typically on 2 slices) which correspond to the top extremity of the brainstem. The only rejected segmentation corresponds to a case with a tumor adjacent to the brainstem and which was mistakenly included in the segmentation (false positives).

Finally, all segmentations of the brain (occupying a large part of the head) were found acceptable for radiotherapy planning even if they required minor corrections in almost one third of cases. The recommended corrections include, for instance, nondetections close to the cribriform plate (between the eyes) and false positives on bones.

In particular, we observe that the only two structures for which all segmentations were found acceptable for radiotherapy planning (without any major correction) are the ones for which a specific postprocessing was performed, i.e., the optic nerves and the brain.

## 4 Conclusion and Future Work

In this work, we proposed a CNN-based method for segmentation of OAR from MR images in the context of neuro-oncology. The method was evaluated on clinical data.

First, we proposed a deep learning model and a training algorithm for segmentation of multiple and nonexclusive anatomical structures. The proposed methodology addresses problems related to computational costs and the variable availability of ground truth segmentations of the different anatomical structures (unsegmented classes). The neural network used in our method is a modified version of U-Net. The network is trained separately for segmentation in axial, coronal, and sagittal slices. The three versions of the network are combined by majority voting.

Second, we proposed procedures to enforce anatomical consistency of the result in a post-processing stage. In particular, we proposed a graph-based algorithm for segmentation of the optic nerves, which are among the most difficult anatomical structures for automatic segmentation. The proposed postprocessings have shown their efficiency particularly in the qualitative evaluation by a radiotherapist. In particular, all segmentations of the optic nerves were found acceptable for radiotherapy planning.

The method was evaluated quantitatively on a set of 44 annotated MRIs, with fivefold cross-validation and using several metrics. The segmentations produced by our system on a set of 50 nonannotated MRIs were qualitatively evaluated by an experienced radiotherapist. Despite the limited size of the training database (44 annotated MRIs) and the different challenges of the segmentation tasks (in particular, presence of tumors), a large majority of the output segmentations were found sufficiently accurate to be used for computation of irradiation doses in radiotherapy.

An important step of the future work is to adapt the method to multimodal data. Often, several types of images are acquired during radiotherapy planning for one patient, including CT scans and different MR sequences (T1, T2, FLAIR). Inclusion of different imaging modalities could improve segmentation of several structures but it comes also with new challenges related, for instance, to intermodality registration and training of models on cases with missing modalities.

In our work, we used a variant of 2-D U-Net to limit the GPU memory load. 2-D CNNs have, however, the drawback of ignoring one spatial dimension. An interesting alternative



would be to use 3-D CNNs with anisotropic receptive fields, such as the model proposed in Ref. 48.

As discussed in Sec. 3, using the voxelwise overlap as the metric to compare the output segmentation with the ground truth has its limits, in particular due to uncertain borders of the ground truth. An interesting direction for future work would be to use a loss function penalizing distances between the output and the ground truth. However, computation of mean distances may represent important computational costs as it implies computation of distances for a large number of pairs of coordinates.

Even if the proposed postprocessing modules (hole-filling, thresholding on the size of connected components, inclusion of one class in another, graph-based algorithm to find centerlines) have been applied to specific anatomical structures, they can be adapted for segmentation of other structures. In particular, our graph-based algorithm could be used to compute centerlines of different tubular structures, such as blood vessels. The proposed postprocessings are not specific to MRI, except the intensity-based refinement of the segmentation of the optic nerve (hyperintensity in MRI T1 of the fat surrounding the optic nerve).

As for other segmentation tasks in medical imaging, availability of annotated training data is an important problem. Methods able to exploit weaker forms of annotations (bounding boxes and slice-level labels) for training of segmentation models are therefore of interest. In particular, methods combining weakly annotated and fully annotated training images were recently proposed in Refs. 49 and 50. As our system was able to produce accurate segmentations in a large majority of cases and the rare observed errors were mainly on boundaries of organs, the system could be used for generation of bounding boxes (subsequently verified by a human), which could be used to train segmentation models that are able to exploit this type of annotations.

Another important direction of the future work is to combine segmentation of OAR and segmentation of radiotherapy target volumes. In particular, a large variability of methods for tumor segmentation<sup>32,48,51–53</sup> were proposed in recent years. Deep learning could also be used for computation of irradiation doses<sup>54</sup> in radiotherapy planning.

## Disclosures

The authors have no conflicts of interest to disclose.

## Acknowledgments

Pawel Mlynarski was funded by the Microsoft Research-Inria Joint Centre, France. This work was supported by the Inria Sophia Antipolis—Méditerranée, “NEF” computation cluster.

## References

1. T. Vos et al., “Global, regional, and national incidence, prevalence, and years lived with disability for 310 diseases and injuries, 1990–2015: a systematic analysis for the Global Burden of Disease Study 2015,” *Lancet* **388**(10053), 1545–1602 (2016).
2. S. Bauer et al., “A survey of MRI-based medical image analysis for brain tumor studies,” *Phys. Med. Biol.* **58**(13), R97 (2013).
3. C. L. Brouwer et al., “3D variation in delineation of head and neck organs at risk,” *Radiat. Oncol.* **7**(1), 32 (2012).
4. D. Ciardo et al., “Atlas-based segmentation in breast cancer radiotherapy: evaluation of specific and generic-purpose atlases,” *Breast* **32**, 44–52 (2017).
5. P.-Y. Bondiau et al., “Atlas-based automatic segmentation of MR images: validation study on the brainstem in radiotherapy context,” *Int. J. Radiat. Oncol. Biol. Phys.* **61**(1), 289–298 (2005).
6. S. Alchatzidis, A. Sotiras, and N. Paragios, “Local atlas selection for discrete multi-atlas segmentation,” in *IEEE 12th Int. Symp. Biomed. Imaging (ISBI)*, IEEE, pp. 363–367 (2015).
7. O. Commowick, V. Grégoire, and G. Malandain, “Atlas-based delineation of lymph node levels in head and neck computed tomography images,” *Radiother. Oncol.* **87**(2), 281–289 (2008).

8. O. Commowick, S. K. Warfield, and G. Malandain, "Using Frankenstein's creature paradigm to build a patient specific atlas," *Lect. Notes Comput. Sci.* **5762**, 993–1000 (2009).
9. L. Ramus and G. Malandain, "Assessing selection methods in the context of multi-atlas based segmentation," in *IEEE Int. Symp. Biomed. Imaging: From Nano to Macro*, IEEE, pp. 1321–1324 (2010).
10. L. Ramus and G. Malandain, "Multi-atlas based segmentation: application to the head and neck region for radiotherapy planning," in *MICCAI Workshop Med. Image Anal. Clin.-A Grand Challenge*, pp. 281–288 (2010).
11. S. K. Warfield, K. H. Zou, and W. M. Wells, "Simultaneous truth and performance level estimation (STAPLE): an algorithm for the validation of image segmentation," *IEEE Trans. Med. Imaging* **23**(7), 903–921 (2004).
12. A. Criminisi et al., "Regression forests for efficient anatomy detection and localization in CT studies," *Lect. Notes Comput. Sci.* **6533**, 106–117 (2010).
13. A. Criminisi et al., "Regression forests for efficient anatomy detection and localization in computed tomography scans," *Med. Image Anal.* **17**(8), 1293–1303 (2013).
14. R. Gauriau et al., "Multi-organ localization with cascaded global-to-local regression and shape prior," *Med. Image Anal.* **23**(1), 70–83 (2015).
15. Y. LeCun and Y. Bengio, "Convolutional networks for images, speech, and time series," *Handb. Brain Theory Neural Networks* **3361**(10), 1995 (1995).
16. B. H. Menze et al., "The multimodal brain tumor image segmentation benchmark (BRATS)," *IEEE Trans. Med. Imaging* **34**(10), 1993–2024 (2015).
17. S. Nikolov et al., "Deep learning to achieve clinically applicable segmentation of head and neck anatomy for radiotherapy," arXiv:1809.04430 (2018).
18. Y. Wang et al., "Organ at risk segmentation in head and neck ct images using a two-stage segmentation framework based on 3D U-Net," *IEEE Access* **7**, 144591–144602 (2019).
19. M. Larsson, Y. Zhang, and F. Kahl, "Robust abdominal organ segmentation using regional convolutional neural networks," *Appl. Soft Comput.* **70**, 465–471 (2018).
20. B. Ibragimov and L. Xing, "Segmentation of organs-at-risks in head and neck CT images using convolutional neural networks," *Med. Phys.* **44**(2), 547–557 (2017).
21. K. Men et al., "More accurate and efficient segmentation of organs-at-risk in radiotherapy with convolutional neural networks cascades," *Med. Phys.* **46**(1), 286–292 (2019).
22. W. Zhu et al., "AnatomyNet: deep learning for fast and fully automated whole-volume segmentation of head and neck anatomy," *Med. Phys.* **46**(2), 576–589 (2019).
23. O. Ronneberger, P. Fischer, and T. Brox, "U-Net: convolutional networks for biomedical image segmentation," *Lect. Notes Comput. Sci.* **9351**, 234–241, (2015).
24. H. R. Roth et al., "Hierarchical 3D fully convolutional networks for multi-organ segmentation," arXiv:1704.06382 (2017).
25. N. Tong et al., "Fully automatic multi-organ segmentation for head and neck cancer radiotherapy using shape representation model constrained fully convolutional neural networks," *Med. Phys.* **45**(10), 4558–4567 (2018).
26. T. Brosch et al., "Deep learning-based boundary detection for model-based segmentation with application to MR prostate segmentation," *Lect. Notes Comput. Sci.* **11073**, 515–522 (2018).
27. E. Orasanu et al., "Organ-at-risk segmentation in brain MRI using model-based segmentation: benefits of deep learning-based boundary detectors," in *Int. Workshop on Shape in Med. Imaging*, Springer, pp. 291–299 (2018).
28. I. Bloch et al., "Fusion of spatial relationships for guiding recognition, example of brain structure recognition in 3D MRI," *Pattern Recognit. Lett.* **26**(4), 449–457 (2005).
29. G. Fouquier, J. Atif, and I. Bloch, "Sequential model-based segmentation and recognition of image structures driven by visual features and spatial relations," *Comput. Vision Image Understanding* **116**(1), 146–165 (2012).
30. A. Argiris et al., "Head and neck cancer," *Lancet* **371**(9625), 1695–1709 (2008).
31. S. Ioffe and C. Szegedy, "Batch normalization: accelerating deep network training by reducing internal covariate shift," arXiv:1502.03167 (2015).
32. P. Mlynarski et al., "3D convolutional neural networks for tumor segmentation using long-range 2D context," *Computer. Med. Imaging Graphics* **73**, 60–72 (2019).

33. M. Folk et al., “An overview of the HDF5 technology suite and its applications,” in *Proc. EDBT/ICDT 2011 Workshop Array Databases*, ACM, pp. 36–47 (2011).
34. T. H. Cormen et al., *Introduction to Algorithms*, MIT Press (2009).
35. F. B. Zhan and C. E. Noon, “Shortest path algorithms: an evaluation using real road networks,” *Transp. Sci.* **32**(1), 65–73 (1998).
36. T. Deschamps and L. D. Cohen, “Fast extraction of minimal paths in 3D images and applications to virtual endoscopy,” *Med. Image Anal.* **5**(4), 281–299 (2001).
37. L. D. Cohen and R. Kimmel, “Global minimum for active contour models: a minimal path approach,” *Int. J. Comput. Vision* **24**(1), 57–78 (1997).
38. F. Benmansour and L. D. Cohen, “Tubular structure segmentation based on minimal path method and anisotropic enhancement,” *Int. J. Comput. Vision* **92**(2), 192–210 (2011).
39. J. A. Sethian, “Fast marching methods,” *SIAM Rev.* **41**(2), 199–235 (1999).
40. F. Zana and J.-C. Klein, “Segmentation of vessel-like patterns using mathematical morphology and curvature evaluation,” *IEEE Trans. Image Process.* **10**(7), 1010–1019 (2001).
41. W. Levin et al., “Proton beam therapy,” *Br. J. Cancer* **93**(8), 849–854 (2005).
42. P. Mildenerger, M. Eichelberg, and E. Martin, “Introduction to the DICOM standard,” *Eur. Radiol.* **12**(4), 920–927 (2002).
43. M. Y. Law and B. Liu, “DICOM-RT and its utilization in radiation therapy,” *Radiographics* **29**(3), 655–667 (2009).
44. A. Fedorov et al., “3D Slicer as an image computing platform for the quantitative imaging network,” *Magn. Reson. Imaging* **30**(9), 1323–1341 (2012).
45. C. Pinter et al., “SlicerRT: radiation therapy research toolkit for 3D Slicer,” *Med. Phys.* **39**(10), 6332–6338 (2012).
46. O. Ecabert et al., “Automatic model-based segmentation of the heart in CT images,” *IEEE Trans. Med. Imaging* **27**(9), 1189–1201 (2008).
47. P. Källman, A. Ågren, and A. Brahme, “Tumour and normal tissue responses to fractionated non-uniform dose delivery,” *Int. J. Radiat. Biol.* **62**(2), 249–262 (1992).
48. G. Wang et al., “Automatic brain tumor segmentation using cascaded anisotropic convolutional neural networks,” *Lect. Notes Comput. Sci.* **10670**, 178–190 (2017).
49. P. Mlynarski et al., “Deep learning with mixed supervision for brain tumor segmentation,” *J. Med. Imaging* **6**(3), 034002 (2019).
50. M. P. Shah, S. Merchant, and S. P. Awate, “MS-Net: mixed-supervision fully-convolutional networks for full-resolution segmentation,” *Lect. Notes Comput. Sci.* **11073**, 379–387 (2018).
51. A. Myronenko, “3D MRI brain tumor segmentation using autoencoder regularization,” *Lect. Notes Comput. Sci.* **11384**, 311–320 (2018).
52. K. Kamnitsas et al., “Efficient multi-scale 3D CNN with fully connected CRF for accurate brain lesion segmentation,” *Med. Image Anal.* **36**, 61–78 (2017).
53. S. Parisot et al., “Concurrent tumor segmentation and registration with uncertainty-based sparse non-uniform graphs,” *Med. Image Anal.* **18**(4), 647–659 (2014).
54. E. A. Andres et al., “PO-1002 pseudo computed tomography generation using 3D deep learning-application to brain radiotherapy,” *Radiother. Oncol.* **133**, S553 (2019).

**Pawel Mlynarski** is a PhD student at Inria Sophia Antipolis, France. His research work addresses problems of automatic analysis of medical images in oncology. His thesis focuses on segmentation of brain tumors and organs at risk in the context of radiotherapy planning. He obtained his master’s degree of applied mathematics at the École Normale Supérieure (Cachan, France).

**Hervé Delingette** is a research director at Inria, director of an academy of excellence at Université Côte d’Azur, and member of the board of the MICCAI Society. He also holds a chair at the new AI institute 3IA Côte d’Azur. He received his engineering and PhD degrees from the École Centrale Paris. His research focuses on various aspects of artificial intelligence in medical image analysis, computational physiology, and surgery simulation.

**Hamza Alghamdi** is a radiation oncologist at the Centre Antoine Lacassagne (Nice, France).

**Pierre-Yves Bondiau** presented his thesis as medical doctor in 1995 about a automatic eye reconstruction software. His scientific thesis was made at the Inria Sophia Antipolis, France, about automatic segmentation for radiotherapy treatment. Then he implemented the first radio-surgery robot, and he started the high energy proton device project at the Antoine Lacassagne Anticancer Center, Nice, France. Presently he is involved in the application of AI in cancer treatments.

**Nicholas Ayache** is research director at Inria, head of the Epione research team dedicated to e-patients for e-medicine, and scientific director of the new AI institute 3IA Côte d'Azur. He is a member of the French Academy of Sciences and the Academy of Surgery. His current research focuses on AI methods to improve diagnosis, prognosis, and therapy from medical images and clinical, biological, behavioral, and environmental data available on the patient.