

Titre: Automatic multiclass intramedullary spinal cord tumor
Title: segmentation on MRI with deep learning

Auteurs: Andréanne Lemay, Charley Gros, Zhizheng Zhuo, Jie Zhang, Yunyun
Authors: Duan, Julien Cohen-Adad, & Yaou Liu

Date: 2021

Type: Article de revue / Article

Référence: Lemay, A., Gros, C., Zhuo, Z., Zhang, J., Duan, Y., Cohen-Adad, J., & Liu, Y. (2021).
Citation: Automatic multiclass intramedullary spinal cord tumor segmentation on MRI with
deep learning. NeuroImage - Clinical, 31, 9 pages.
<https://doi.org/10.1016/j.nicl.2021.102766>

Document en libre accès dans PolyPublie

Open Access document in PolyPublie

URL de PolyPublie:
PolyPublie URL: <https://publications.polymtl.ca/9276/>

Version: Version officielle de l'éditeur / Published version
Révisé par les pairs / Refereed

Conditions d'utilisation:
Terms of Use: Creative Commons Attribution 4.0 International (CC BY)

Document publié chez l'éditeur officiel

Document issued by the official publisher

Titre de la revue:
Journal Title: NeuroImage - Clinical (vol. 31)

Maison d'édition:
Publisher: Elsevier

URL officiel:
Official URL: <https://doi.org/10.1016/j.nicl.2021.102766>

Mention légale:
Legal notice:



Automatic multiclass intramedullary spinal cord tumor segmentation on MRI with deep learning

Andreanne Lemay^{a,b}, Charley Gros^{a,b}, Zhizheng Zhuo^c, Jie Zhang^c, Yunyun Duan^c, Julien Cohen-Adad^{a,b,d,*}, Yaou Liu^{c,*}

^a NeuroPoly Lab, Institute of Biomedical Engineering, Polytechnique Montreal, Montreal, QC, Canada

^b Mila, Quebec AI Institute, Canada

^c Department of Radiology, Beijing Tiantan Hospital, Capital Medical University, Beijing, China

^d Functional Neuroimaging Unit, CRIUGM, Université de Montréal, Montreal, QC, Canada

ARTICLE INFO

Keywords:

Deep learning
Automatic segmentation
Spinal cord tumor
MRI
Multiclass
CNN

ABSTRACT

Spinal cord tumors lead to neurological morbidity and mortality. Being able to obtain morphometric quantification (size, location, growth rate) of the tumor, edema, and cavity can result in improved monitoring and treatment planning. Such quantification requires the segmentation of these structures into three separate classes. However, manual segmentation of three-dimensional structures is time consuming, tedious and prone to intra- and inter-rater variability, motivating the development of automated methods. Here, we tailor a model adapted to the spinal cord tumor segmentation task. Data were obtained from 343 patients using gadolinium-enhanced T1-weighted and T2-weighted MRI scans with cervical, thoracic, and/or lumbar coverage. The dataset includes the three most common intramedullary spinal cord tumor types: astrocytomas, ependymomas, and hemangioblastomas. The proposed approach is a cascaded architecture with U-Net-based models that segments tumors in a two-stage process: locate and label. The model first finds the spinal cord and generates bounding box coordinates. The images are cropped according to this output, leading to a reduced field of view, which mitigates class imbalance. The tumor is then segmented. The segmentation of the tumor, cavity, and edema (as a single class) reached $76.7 \pm 1.5\%$ of Dice score and the segmentation of tumors alone reached $61.8 \pm 4.0\%$ Dice score. The true positive detection rate was above 87% for tumor, edema, and cavity. To the best of our knowledge, this is the first fully automatic deep learning model for spinal cord tumor segmentation. The multiclass segmentation pipeline is available in the Spinal Cord Toolbox (<https://spinalcordtoolbox.com/>). It can be run with custom data on a regular computer within seconds.

1. Introduction

Intramedullary spinal cord tumors (IMSCT) represent 2 to 5% of all central nervous system tumors (Das et al., 2020). This relatively low prevalence contributes to the difficulty in understanding this malignant pathology (Claus et al., 2010). The first step to a better understanding of the disease is an improved characterization of the tumor. Segmentation informs the healthcare specialists on the tumor's position, size, and growth rate leading to quantitative monitoring of the tumor's progression. In addition, characterizing the edema and cavity (i.e., syrinx) associated with the tumor is clinically relevant (Balériaux, 1999; Kim

et al., 2014; Das et al., 2020; Huntoon et al., 2016). Segmentation of these three components is insightful as it reflects the characteristics of the tumor and symptoms, and it also helps for treatment and surgical planning. For instance, the cyst (i.e., cavities) growth rate is correlated with the development of new symptoms, which can lead to the resection of the tumor (Huntoon et al., 2016). As stated by (Balériaux, 1999), differentiating tumor infiltration from associated edema (which is achieved by multi-class segmentation methods) gives insights about the tumor's malignancy and is therefore relevant information for the neuro-radiologist. Manual labeling is tedious for clinicians and prone to intra- and inter-rater variability. Fully automatic segmentation models

Abbreviations: CNN, convolutional neural network; IMSCT, intramedullary spinal cord tumor; T1w, T1-weighted; T2w, T2-weighted.

* Corresponding authors at: Department of Radiology, Beijing Tiantan Hospital, Capital Medical University, No.119 Southern 4th Ring Road West, Fengtai District, Beijing 100070, China. Dept. Genie Electrique, L5610, Ecole Polytechnique, 2900 Edouard-Montpetit Bld, Montreal, QC H3T 1J4, Canada.

E-mail addresses: jcohen@polymtl.ca (J. Cohen-Adad), liuyaou@bjth.org (Y. Liu).

<https://doi.org/10.1016/j.nicl.2021.102766>

Received 30 March 2021; Received in revised form 22 June 2021; Accepted 17 July 2021

Available online 22 July 2021

2213-1582/© 2021 The Authors. Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

overcome these issues. Although automatic methods to segment brain tumors are numerous, there is currently, to the best of our knowledge, no automatic model to segment IMSCT. Heterogeneity in tumor size, intensity, location, in addition to images varying in resolution, dimensions, and fields of view represent a challenge for the segmentation task.

1.1. Heterogeneity in spinal cord tumor characteristics

This study covers the three main tumor types (95% of IMSCT (Balériaux, 1999)): astrocytoma, ependymoma, and hemangioblastoma, with cervical, thoracic, and lumbar coverage (Fig. 1). On T2-weighted (T2w) scans, all tumor types display isointense to hyperintense signals (Fig. 1A), and on T1w-weighted (T1w) scans from isointense to hypointense signals, with well or ill-defined boundaries (Kim et al., 2014; Baker et al., 2000). In the case of isointense tumors or ill-defined margins, it is challenging to segment the lesion properly. Gadolinium-enhanced T1w MRI yields a hyperintense signal from the tumor, but the enhancement patterns are different depending on the tumor type. For example, astrocytomas usually present partial, moderate (Fig. 1E), or no enhancement (Fig. 1D) (Balériaux, 1999), while hemangioblastomas generally yield intense enhancement (Fig. 1F) (Baker et al., 2000). All tumor types do not display the same characteristic on MRI scans which is a challenge when developing a robust deep learning model. Tumors can be associated with tumoral or non-tumoral cystic components (i.e., liquid-filled cavities in the spinal cord) (Fig. 1C and F) or extensive edema, but these cannot be systematically seen with the tumor (Balériaux, 1999; Kim et al., 2014; Baker et al., 2000). Finally, tumors can be present as a single (Fig. 1A, 1B, 1D, 1E) or multiple lesions (Fig. 1C and 1F) (Balériaux, 1999) and measure from less than 10 mm (Chu et al., 2001) up to 19 vertebral bodies long (Balériaux, 1999). This heterogeneity in IMSCT can be an obstacle to precisely delineate each component related to the tumor.

1.2. Previous work

1.2.1. Brain tumor models

The annual BraTS challenge contributes to the progress of deep learning models for brain tumor segmentation. The top-performing models on BraTS are mostly convolutional neural networks (CNNs)

(Havaei et al., 2017; Isensee et al., 2017, 2019; Kamboj et al., 2018; Naceur et al., 2018). Although they are both part of the central nervous system, brain tumors and spinal tumors need to be processed differently. A critical difference between the brain and the spinal cord is the size resulting in different fields of view. The spinal cord has an average diameter of 13.3 ± 2.2 mm at its largest point (transverse diameter in C5) (Frostell et al., 2016), while the length can range from 420 to 450 mm (Boonpirak and Apinhasmit, 1994). In contrast, the brain has more isotropic dimensions. This discrepancy leads to challenging decisions regarding cropping and patch sizes. In comparison with the brain, spinal cord imaging is also more hampered by respiratory and cardiac motion artifacts (Stroman et al., 2014). Another possible added difficulty of spinal tumor/edema/cavity segmentation is that, in comparison with the brain, these pathological presentations generally interfaces with the cerebrospinal fluid, making it sometimes difficult to separate them, especially on T2w scans, where the fluid appears as a hyperintense signal.

1.2.2. Spinal tumor models

Previous studies introduced models for segmenting tumors located in the spine (i.e., present in bone vs. spinal cord) (Hille et al., 2020; Reza et al., 2019; Wang, 2017). In 2020, Hill et al. developed a U-Net-based model to segment spine metastases (Hille et al., 2020). They reached a Dice score of 77.6% on average. Reza et al., 2019 presented a cascaded architecture for spine chordoma tumor segmentation (i.e., rare tumor type usually in the bone near the spinal cord and the skull base) (Reza et al., 2019). Their approach yielded a median absolute difference of 48%. Even if spine tumors and IMSCT are both present in the spine area, they exhibit different intensities and sizes and are juxtaposing with different tissue types (i.e., bone tissue vs. neuronal tissue), making spine segmentation models not applicable to the spinal cord. Also, previous models for spine tumors are single class (i.e., segmentation of tumor only), while IMSCT is associated with multiple components that should be segmented separately. Moreover, the trained models are not publicly available hence cannot directly benefit the medical community wanting to apply the model on new data.

1.3. Study scope

In this work, we present a multi-contrast (gadolinium-enhanced T1w




Fig. 1. IMSCT heterogeneity in intensities, size, and longitudinal location. The first row (A, B, C) presents T2w scans, while the second row (D, E, F) are the gadolinium-enhanced T1w scans. The first column (A, D) displays an example of astrocytoma, the second column (B, E) an example of ependymoma, and the last column (C, F) an example of hemangioblastoma. The white arrows point to solid tumor components, the green arrows point to a liquid-filled cavity, and the pink arrow to edema. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

and T2w) segmentation model for IMSCT. We chose a two-stage cascaded architecture composed of two U-Nets (Ronneberger et al., 2015) (i.e., semantic-wise CNNs). The first step consists of localizing the region of interest, i.e., the spinal cord. The second step labels the structures of IMSCT. Both models are based on the modified 3D U-Net implemented by Isensee et al. for the 2017 BraTS challenge (Isensee et al., 2017). The framework segments the three main components associated with IMSCT: the enhanced and non-enhanced tumor component, liquid-filled cavities, and edema. The choice of a multi-contrast design was motivated by the relevance of both contrasts (gadolinium-enhanced T1w and T2w) for delineating all components of the IMSCT (Baker et al., 2000). In clinical routine, both contrasts are generally acquired (Balériaux, 1999). Our pipeline is robust to varying resolution, fields of view, and tumor sizes. The cascaded models were trained on the three most common IMSCT types, which also happen to present very different image characteristics (shape, contrast, presence of edema/cavity). We integrated the model in the SCT open-source software (De Leener et al., 2017). Hence, the pipeline can be easily applied to custom data within seconds in a single command-line or via the graphical user interface (GUI) available in SCT (v5.0 and higher).

2. Material and methods

2.1. Dataset

The data used for this experiment includes 343 MRI scans acquired for preoperative examination (before application of any treatment) from Beijing Tiantan Hospital, Capital Medical University from October 2012 to September 2018, with heterogeneous vertebral coverage (cervical, thoracic, and lumbar), from patients diagnosed with spinal cord tumor: astrocytoma ($n=101$), ependymoma ($n=122$), and hemangioblastoma ($n=120$). T2-weighted (T2w) and Gadolinium-enhanced T1-weighted (T1w) images were available for each patient, as well as the manual segmentation for the tumor, edema, and cavity. 36 astrocytoma cases were high-grade while the rest of the cases were low-grade. The segmentation was performed by a neuro-radiologist with five years of experience, quality controlled and corrected accordingly. Table 1 presents the demographic data. The ranges of native resolution, in mm, of the sagittal scans were [0.34, 1.33] for the in-plane resolution and [1.5, 5.2] for the slice thickness. Table 2 includes the volume information for each component by type.

The ground truths for the spinal cord detection model were generated in two steps using SCT (De Leener et al., 2017). The centerline of the spinal cord for the images was manually identified. The current version of the spinal cord detection algorithm (Gros et al., 2018) not robust in the presence of tumors, hence could not be leveraged as an automatic algorithm for initial localization or as ground truth. From the centerline, a mask of 30 mm diameter was automatically generated. This value was chosen in accordance with the average spinal cord diameter with an extra buffer to ensure full coverage of the spinal cord on the right-left and anterior-posterior axes.

2.2. Data preparation

To maximize the segmentation performance, image registration was performed using SCT (De Leener et al., 2017). T1w images were

Table 2

Tumor component volume by type (MEAN \pm STD). Patients without cavity or edema were excluded from the average.

	Tumor	Volume (cm ³)	
		Edema	Cavity
Astrocytoma	8.3 \pm 5.9	7.2 \pm 7.8	34.4 \pm 31.8
Ependymoma	8.1 \pm 2.9	5.7 \pm 4.0	23.0 \pm 17.9
Hemangioblastoma	2.7 \pm 4.5	17.2 \pm 13.1	49.0 \pm 37.6

registered (Avants et al., 2009) onto the T2w scans using affine transformations with the cross-correlation metric on the SCT software. Images with different dimensions were resampled to a common resolution ($1 \times 1 \times 2$ mm³) before registration. Every registration has been manually verified and corrected if needed. Manual labels of vertebral discs were added and used to coregister the subjects when the first registration method failed.

2.3. Processing pipeline

The automatic pipeline includes preprocessing, a cascaded neural network and postprocessing. State-of-the-art pipelines for medical imaging tasks exploited cascaded architectures (Akkus et al., 2017; Gros et al., 2019; Christ et al., 2017; Hussain et al., 2017). A rationale of cascaded models is to isolate the region of interest with a first CNN and then segment the desired structure. Gros et al., 2019 benefits from this approach for spinal cord multiple sclerosis segmentation (Gros et al., 2019). The first CNN finds the centerline while the second performs segmentation. This helped to limit class imbalance and focus the task on the spinal cord.

Data preprocessing and model training was implemented with ivadomed (Gros et al., 2020). Data is preprocessed before training or inference. The preprocessing steps are included in the model's pipeline. The resolution of the sagittal images are set to 1 mm (superior-inferior), 1 mm (anterior-posterior), 2 mm (right-left). This choice of resolution was based on preliminary investigations and is a compromise between computational time and segmentation precision. The resampled images were cropped voxel-wise with dimensions of $512 \times 256 \times 32$, which corresponds to a bounding box of 51.2 cm \times 25.6 cm \times 6.4 cm applied at the center of the field of view. These dimensions are consistent with the adult spinal cord anatomy and allow for slight angulation in the right-left direction (e.g., in cases of scoliosis). In cases where the field of view is smaller in one or more axes, the image is zero-padded instead of cropped. The intensity of each scan was normalized by subtracting the average intensity over the standard deviation. The framework of the model is a cascaded architecture composed of two steps (Fig. 2). The first step aims to localize the spinal cord and crop the image around the spinal cord mask with a 3D bounding box (Fig. 2 step 1). Both Gd-e T1w and T2w images are preprocessed, cropped and concatenated before being used as input for the second step of the pipeline being the tumor segmentation task (Fig. 2 step 2). The first step reduces the field of view on the tumor and makes the model robust for MRI with varying fields of view or dimensions. Smaller images also lead to faster training and inference in addition to mitigate class imbalance. For the tumor segmentation model, the cropping is done with the spinal cord mask. Thus, the cropping size is different for every patient according to the field of view, dimensions of the image, and size of the spinal cord. The

Table 1

Demographic information of patients by tumor type. First row: number of subjects. Second row: sex distribution where M is male and F is female. Third row: median age in years. Fourth row: Age range in years (minimum–maximum). The last column presents the demographic information for all tumor types combined.

	Astrocytoma	Ependymoma	Hemangioblastoma	All
Subjects	101	122	120	343
Sex	60M:41F	70M:52F	68M:52F	198M:145F
Median age	30	42	35.5	37
Age range (min–max)	1–63	8–76	8–66	1–76




Fig. 2. Fully automatic spinal cord tumor segmentation framework. Step 1: The spinal cord is localized using a 3D U-Net and the image is cropped around the generated mask. Step 2: The spinal cord tumors are segmented.

segmentation yields 4 different labels: tumor core (enhanced and non-enhanced), cavity, edema, and the whole tumor composed of all structures together.

2.4. Model

A modified 3D U-Net (Isensee et al., 2017) architecture was used for both spinal cord localization and tumor segmentation. The main differences with the original 3D U-Net (Çiçek et al., 2016) reside in the addition of deep supervision modules (Kayalibay et al., 2017), the use of dropout layers (Srivastava et al., 2014), instance normalization (Ulyanov et al., 2016) (instead of batch normalization), and leaky ReLU (Maas et al., 2013) (instead of ReLU). Table 3 enumerates the main training parameters used for both models. Initial hyperparameter optimization was performed to find these values. Sigmoid was used as the final activation of the model, even for the multiclass model, to allow the prediction of a class composed of all structures.

2.4.1. Spinal cord localization

The spinal cord localization model was trained with a Dice loss (Milletari et al., 2016) using a single patch input of size $512 \times 256 \times 32$ on T2w MRI scans. Preliminary experiments with a multi-contrast model did not improve the performance: a single channel model was hence used for this task. Due to the large size of each 3D patch, the batch size was limited to 1. The training lasted for a maximum of 200 epochs with an early stopping of 50 epochs with an epsilon of 0.001 (i.e., minimal loss function difference to be considered as an improvement). The initial

learning rate was 0.001, and the cosine annealing scheduler was used. The number of downsampling layers (i.e., the depth) was set to 4, and the number of base filters (i.e., the number of filters used on the first convolution) to 8. Finally, random affine transformations are applied during the training: rotation (± 5 degrees applied on one of the three axes), scaling ($\pm 10\%$ in the sagittal plane), and translation ($\pm 3\%$ in the sagittal plane).

2.4.2. Tumor segmentation

The tumor segmentation is multi-contrast since both gadolinium-enhanced T1w and T2w MRI scans carry valuable information. Being a multiclass task, the model was trained with a multiclass Dice loss, which is the average of Dice loss for each class. Preliminary investigations favored this loss function in terms of global performance for all tumor components. Patches of size $128 \times 128 \times 32$ with a stride of $64 \times 64 \times 32$ and a batch size of 8 were used to train the model. At inference time, the patches were stitched together by averaging for voxels having more than one prediction value. The same parameters used for spinal cord localization for depth, early stopping, initial learning rate, and learning rate scheduler were selected. Multiclass segmentation of IMST is a more complex task than localizing the spinal cord and required a higher number of base filters to reach optimal performance, which was set to 16.

2.4.3. Dataset split

For both models, 60% of the subjects were used for training, 20% for validation, and 20% for testing. Since only 45 subjects had tumors located in the lumbar region, all subjects were included in the training and validation sets to maximize the exposure of the model to IMST located in the lumbar section of the spinal cord. Hence, the model has not been validated on lumbar tumors.

2.5. Postprocessing

Postprocessing steps are applied to the model's prediction. The steps include: binarization with a threshold of 0.5, filling holes, and removing tumor prediction smaller than 0.2 cm^3 , as well as edema and cavity predictions smaller than 0.5 cm^3 to limit false positives and noise. These postprocessing steps can be customized when applying the model with SCT.

2.6. Evaluation

2.6.1. Segmentation pipeline

The Dice score was used to evaluate the segmentation performance of both spinal cord localization and tumor segmentation models. To assess the localization task, the inclusion of all the IMST volumes into their respective bounding box, generated by the first step of the pipeline, was verified. Other metrics used to evaluate the multiclass segmentation

Table 3

Training parameters for spinal cord localization and tumor segmentation model. Abbreviation: Gd-e: Gadolinium-enhanced.

Parameters	Model 1:	Model 2:
	Spinal cord localization	Tumor segmentation
Input	T2w	Gd-e T1w + T2w (multi-contrast)
Loss	Dice loss (Milletari et al., 2016)	Multiclass Dice loss
Batch size	1	8
Patch size	$512 \times 256 \times 32$	$128 \times 128 \times 32$
Stride	—	$64 \times 64 \times 32$
Early stopping	50	50
Learning rate	0.001	0.001
Learning rate scheduler	Cosine Annealing	Cosine Annealing
Depth	4	4
Number of base filters	8	16
Data augmentation		
Rotation	± 5 degrees	± 5 degrees
Scaling	$\pm 10\%$	$\pm 10\%$
Translation	$\pm 3\%$	$\pm 3\%$

model were: tumor true positive detection rate, false negative detection rate, precision, recall, as well as absolute and relative volume difference. A structure was considered detected if there was an overlap of at least 6 mm^3 between the ground truth and the prediction. This stands for true positive and negative detection rates. The relative volume difference was computed by subtracting the volume prediction to the volume of the ground truth and dividing by the ground truth volume. A negative relative volume difference signifies an over-segmentation from the model. The absolute volume difference is the absolute value of the relative volume difference. 12 random splittings were performed on the dataset in order to derive meaningful statistics on the results.

2.6.2. Comparison between cascaded architecture and single-step architecture

Training and inference time as well as Dice scores of the cascaded architecture were compared to the pipeline without the first step of the model. The same patch size and stride were used for both approaches. Since a single epoch takes several hours for the segmentation model without the initial cropping step and dozens of epochs are needed, only one training was performed without the first step of the pipeline (i.e., automatic cropping around the spinal cord). To ensure comparability of the two models, the same dataset split was used for both models. The inference time was computed on 100 images with average dimensions of $545 \times 545 \times 20$ on a regular computer (i.e., without GPU) with 16 GB of RAM.

2.7. Implementation

The implementation of the models and pipeline is publicly available and can be found in the deep learning toolbox ivadomed v2.6.1 (<http://ivadomed.org/>). The packaged and ready-to-use model can be applied to custom data via SCT's function `sct_deepseg`¹. The models were converted into the ONNX (Open Neural Network Exchange) format for fast prediction on CPUs.

3. Results

3.1. Spinal cord localization model

The spinal cord localization model reached a Dice score of 88.7% on the test set. When verifying if all tumor segmentations were contained in the bounding box generated by the first step of the pipeline, only two subjects (i.e., success rate of 99.4%) had part of their cavity located out of the bounding box. For both patients, the cavity rose up to the level of the brainstem, which explains why the model did not include this section (cavity extending out of the spinal cord region). Despite this partial truncation, the model correctly located the spinal cord area. Hence, the model successfully located all spine regions for all subjects.

3.2. Tumor segmentation

Fig. 3 illustrates qualitative results from four testing subjects. The first row represents a subject where the segmentation for every class was successful (Dice > 85% for all structures). The second and third rows illustrate two examples of false positive detection for cavity and edema, respectively. A hyperintense signal on the T2w scan (white arrows), similar to what can be seen for liquid-filled cavities, led to the false identification of cavity tissue. The false positive edema classification was caused by a moderate hyperintense signal on the T2w image (yellow arrows) that can be confused with edema. The last row shows an example where the model failed to correctly identify the different tumor's structure. As seen in the third row, a moderate hyperintense signal

on the T2w image generated a false positive edema prediction (pink arrows). The tumor was also misclassified, probably due to a strong hyperintense signal on T2w combined with a dim enhancement on the gadolinium-enhanced T1w (blue arrows). Cavities usually have a hyperintense signal with no enhancement on gadolinium-enhanced T1w.

Table 4 presents the principal metrics for IMSCT related structures. An average Dice score across all models for tumor, cavity, and edema as a single class reached $76.7 \pm 1.5\%$. When identifying each structure separately, the Dice score drops to 61.8%, 55.4%, and 31.0% for tumor, cavity, and edema, respectively. The lower scores associated with cavity and edema are partly due to false positive detection of the structure (e.g., the false detection rate of 46.3% for edema) as presented in the second and third rows of Fig. 3. False positive detections of structure absent in the subject are systematically associated with a Dice score of 0% even if the predicted volume is a few voxels. When keeping only the subjects with edema or cavity, the Dice scores improve to 47.4% for edema and 65.2% for cavities (i.e., increase of respectively 16.4% and 9.8%). However, there is a relatively high true positive detection rate for all structures, i.e., higher than 87% for each class, meaning that when a structure is present, the model usually detects it. As for the volume difference, when segmenting all structures as a single class, the relative volume difference is -3.4% (corresponding to a slight over-segmentation), and the absolute volume difference is 22.1%. The volume difference has also a performance drop when analyzing the structure separately (see Table 4). The model has a tendency to over-segment, i.e., negative relative volume difference, except for the edema, which is slightly under-segmented.

3.3. Advantages of a cascaded architecture

The cascaded architecture is associated with faster training time, inference time, and segmentation performance compared with the pipeline without the initial cropping step.

3.3.1. Training and inference time

Training a single epoch using the same patch and batch sizes without the first step of the cascaded architecture took $\sim 3.4 \text{ h}$, while one epoch lasted $\sim 14 \text{ min}$ for the proposed cascaded pipeline on a single NVIDIA Tesla P100 GPU.

The average inference time for the cascaded architecture took 22 s per image. In comparison, the model without the cropping step took 48 s. These times include the loading, resampling, preprocessing, spinal cord localization, multiclass tumor segmentation, postprocessing, and saving of the images. This time difference can become considerable when the model is sequentially applied to a large number of images.

3.3.2. Segmentation performance

In comparison with the single model architecture, the model using a cascaded architecture yielded a Dice score improvement of 5.0%, 30.0%, and 4.6% for the tumor, edema, and cavity respectively. The important class imbalance between the tumor and the background class paired with an imbalance of the edema class in the whole dataset led to consistent empty predictions for the edema class by the model using only the second step of the pipeline.

3.4. Inter-rater variability

Six randomly selected subjects (two subjects per tumor type) from the dataset were segmented by a second neuro-radiologist with 13 years of experience to assess the inter-rater variability for this task. The average Dice score between the two expert segmentations was 80.2% for the tumor, 59.2% for edema, and 0% for cavity. Within these six subjects, each expert identified only two patients with cavities, but disagreed on the segmentation of this component. While this is an extremely limited sample to quantify inter-rater variability, this

¹ https://spinalcordtoolbox.com/en/latest/user_section/command-line.html#sct-deepseg

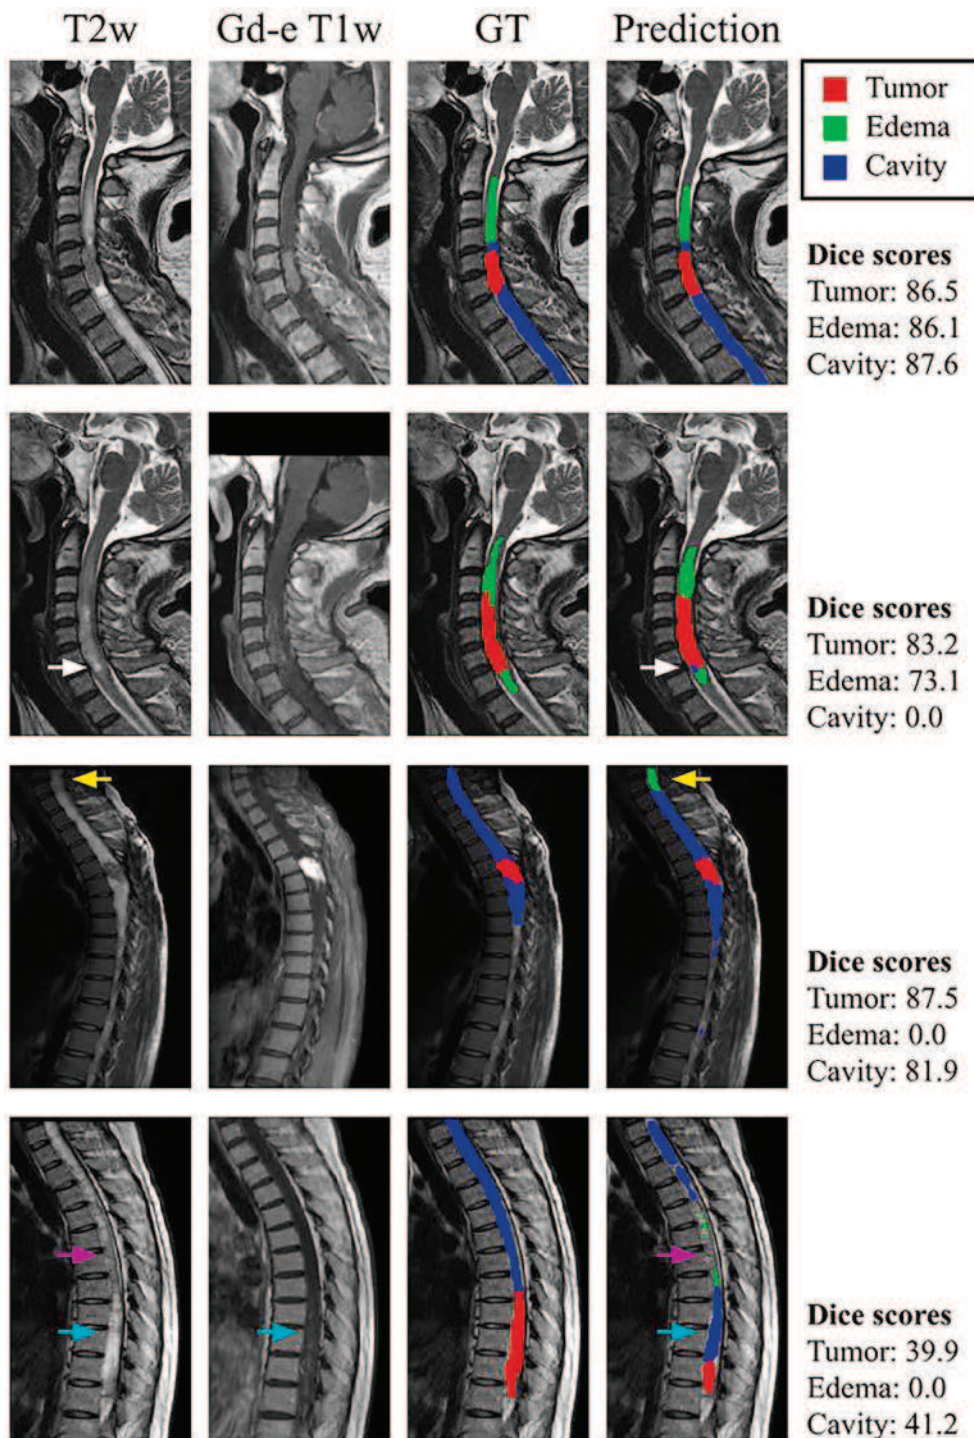


Fig. 3. Sagittal MRIs showing the model's prediction on a subject containing tumor, edema, and cavity. Column 1 presents T2w input image, column 2 Gadolinium-enhanced T1w input image, column 3 the manual annotation, and column 4 the model's prediction. The white arrows point to a false positive cavity segmentation due to a hyperintense signal on T2w. The yellow arrows show a false positive edema segmentation, a consequence of a moderate hyperintense signal on T2w. The pink arrows point towards a false negative cavity and false positive edema segmentations caused by a moderate hyperintense signal on the T2w scan. The blue arrows indicate a misclassification of tumor tissue due to a hyperintense signal on T2w, also characteristic of liquid-filled cavities, and a moderate enhancement on the Gd-e T1w scan. Abbreviations: Gd-e: Gadolinium-enhanced, GT: Ground truth. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

illustrates the difficulty of the task.

4. Discussion

We introduced an IMSCT segmentation model publicly available that can be applied to custom images within seconds with SCT. The cascaded architecture of the pipeline mitigates the class imbalance, speeds up the training and inference, and is associated with higher segmentation performance. The first step of the pipeline localizes the spinal cord and crops the image to isolate the region of interest. The robustness of the detection step is crucial since the segmentation model depends on this task (Gros et al. 2019). A Dice score of 88.7% on the test set was reached

for this step and the spinal cord was correctly located for all subjects. The segmentation of the tumor, edema, and cavity as a single class reached 76.7% of Dice score. Segmenting the tumor, cavity, and edema separately is more challenging for the model yielding Dice scores of 61.8%, 55.4%, and 31.0%, respectively. The performance is hampered by false positive detections due to class imbalance throughout the dataset (not all subjects have a cavity or edema). However, the model has a high true positive detection rate (>87%) for every class.

4.1. Tissue heterogeneity

As seen in Table 3, the model reaches good performance when

Table 4Multiclass IMSCT segmentation metrics on 12 dataset random splittings (MEAN \pm STD).

Labels	Tumor + Cavity + Edema	Tumor	Cavity	Edema
Dice score [%]	76.7 \pm 1.5	61.8 \pm 4.0	55.4 \pm 4.4	31.0 \pm 4.6
Optimal value: 100				
True positive detection rate [%]	98.7 \pm 0.9	90.6 \pm 4.1	93.8 \pm 2.9	87.3 \pm 6.4
Optimal value: 100				
False positive detection rate [%]	9.9 \pm 3.5	17.2 \pm 5.1	23.3 \pm 4.3	46.3 \pm 7.2
Optimal value: 0				
Precision [%]	78.7 \pm 2.6	74.3 \pm 4.4	60.5 \pm 3.7	41.4 \pm 5.8
Optimal value: 100				
Recall [%]	77.5 \pm 2.7	64.3 \pm 4.0	67.5 \pm 3.8	46.9 \pm 7.5
Optimal value: 100				
Absolute volume difference [%]	22.1 \pm 6.0	75.3 \pm 39.5	61.8 \pm 12.2	67.0 \pm 31.7
Optimal value: 0				
Relative volume difference [%]	-3.4 \pm 9.9	-26.0 \pm 45.9	-34.4 \pm 33.3	4.5 \pm 27.6
Optimal value: 0				

segmenting all structures associated with IMSCT together, but the metrics drop when segmenting each structure separately. The varying intensity patterns and the ambiguous delimitation of each structure can partly explain this behavior. On T2w scans, tumors can present hyperintense signals which can be confused with the hyperintense signal associated with liquid-filled cavities. A hyperintense signal on gadolinium-enhanced T1w modality helps to distinguish a tumor from a cavity, but tumors don't always exhibit an enhancement (see Fig. 1A and 1D). The varying intensity patterns for tumors cause misclassification. The same goes for distinguishing edema from tumors or cavities from edema. For example, a cavity presenting a low voxel intensity can be mistaken for edema (see Fig. 3 third row). Another challenging aspect of differentiating each component of IMSCT is the ill-defined boundaries between each structure. Fig. 4 illustrates this phenomenon. The tumor displays edema intensity patterns on the T2w scan and has poor enhancement on gadolinium-enhanced T1w scan. The boundaries of the tumor are ambiguous which is challenging for clinicians, hence for the deep learning model.

4.2. Class imbalance for edema and cavity

Another challenge that explains the lower performances for separate

segmentation of tumor, edema, and cavity is the class imbalance of edema and cavity across the dataset. 67% of subjects have a cavity and only 51% of subjects have edema. The lower number of examples presented combined with the small size of edema and/or cavity compared to the background class (i.e., background represents 99.5% of voxels) hinder the model's ability to properly generalize leading to a high number of false positives. False positive detections of a structure are associated with a Dice score of 0% which highly impacts the metrics. This false positive rate is coupled with a high true positive rate meaning that when a structure is present the model usually detects it which is generally preferred in clinical settings.

4.3. Usage in clinical settings

The medical field can benefit from the integration of deep learning models into clinical settings (Perone and Cohen-Adad, 2019; Valliani et al., 2019). Even if imperfect, segmentation models can save hours of work to clinicians requiring only modifications to the automatic prediction. An automated segmentation pipeline could increase efficiency and consistency when integrated into clinical settings. Longitudinal follow-up to monitor volume growth or tumor delineation before radiotherapy, treatment, or biopsy are two typical scenarios where our framework could help. Also, the automated segmentation opens the door to investigate the clinical significance of spinal cord tumor characteristics (location, size, composition), predicting survival, assessing clinical disability (sensorimotor deficits), correlating to molecular status (e.g., H3 K27M mutation), evaluation of treatment effect, and predicting the prognosis. For instance, the work of Eden et al. quantified the impact of multiple sclerosis lesion distribution across various spinal white matter tracts (Eden et al., 2019), using a spinal atlas-based analysis (Lévy et al., 2015). Similar investigations could be done with spinal cord tumors.

4.4. Limitation

Tumor segmentation is prone to inter-rater disagreement as seen in section 3.4. An important limitation of this study is that the ground truth was performed by a single expert. Poorly delineated borders for the tumor components susceptible to high inter-rater variability might partly explain the poor performance of the model on edema segmentation. Another limitation of the study is that the model was only trained on intramedullary tumors, hence knowledge of this tumor category is assumed before applying the model. The performance of the model on other tumor types was not assessed.

4.5. Perspectives

4.5.1. Continuous learning

Continuous learning is increasingly popular in the medical field (Pianyhk et al., 2020; Gonzalez et al., 2020) since it facilitates multiple




Fig. 4. Subject with an astrocytoma with unclear tumor boundaries.

center data training while addressing privacy issues related to data sharing (Gonzalez et al., 2020). Integrated tools for continual training would benefit the IMSCCT model. Continuous learning allows the model to learn from new examples presented to the model, and become more and more robust through time. A challenge associated with continual training is catastrophic forgetting which occurs when the model gives disproportionate weight to new learning examples, hence forgetting the initial training data (Kirkpatrick et al., 2017; Gonzalez et al., 2020). Implementing techniques properly addressing this issue would allow to develop a robust continuous learning pipeline. This opens the possibility of having a cooperative model trained on data from different centers without direct data sharing that would improve through time.

4.5.2. Missing modalities

In this work, a multi-contrast model was chosen to improve the robustness of the segmentation. The presented model requires both T2w and gadolinium-enhanced T1w contrasts to generate a prediction. Users missing one of the contrasts would be unable to use the model, but solutions exist to mitigate this problem. HeMIS (Havaei et al., 2016) is a deep learning approach addressing this issue. A model robust to missing modalities would benefit the users since less images are necessary to have a prediction. Future work could focus on a HeMIS U-Net version of the model presented here.

4.5.3. SoftSeg

Recent work praises the benefits of leaning towards a soft training approach (Gros et al., 2021; Müller et al., 2019). SoftSeg (Gros et al., 2021) is a training pipeline yielding and propagating soft values to address partial volume effect and certainty calibration of the model for segmentation tasks. Gros et al. reports higher Dice scores when using SoftSeg compared to a conventional segmentation pipeline. SoftSeg could be generalized to multiclass predictions and be applied to IMSCCT segmentation. This method takes into account partial volume effect, hence could benefit the volume measurement of unhealthy tissues. Also, the soft predictions would give more insight on the model's certainty especially towards ill-defined boundaries, which should be associated with higher uncertainty. The broader range of prediction values yields more information and allows a more enlightened postprocessing.

4.5.4. Use of additional data for improved segmentation

A recent study has shown that conditioning the model with the tumor type helps to improve the overall segmentation performance (Lemay et al., 2021). A possible avenue to enhance the model's performance would be to input relevant clinical information to the model such as the tumor type, the grade, demographic information, etc. Metadata information is often available in clinical settings and could be fed to the model with the images to improve the final output.

5. Conclusion

In this work, we presented the first model for multiclass segmentation for IMSCCT with a two-stage U-Net based cascaded architecture. The choice of a cascaded architecture is associated with faster training and inference time, as well as higher dice scores. An average Dice score across 12 random splittings of 76.7% was reached to segment all IMSCCT structures as a single class. Varying intensity patterns, ill-defined boundaries, and high class imbalance was a challenge to the separate labeling of tumors, cavities, and edemas. The model is associated with true positive detection rates above 87% for all components (tumor, edema, and cavity). Future work could focus on the implementation of continual learning tools, techniques to address missing modalities, or a SoftSeg version of the model. The segmentation pipeline is available in the SCT software and can be applied to custom data through a single command-line in a few seconds.

CRedit authorship contribution statement

Andreanne Lemay: Conceptualization, Investigation, Methodology, Validation, Formal analysis, Software, Writing - original draft. **Charley Gros:** Conceptualization, Investigation, Methodology, Software, Writing - review & editing. **Zhizheng Zhuo:** Resources, Data curation. **Jie Zhang:** Resources, Data curation. **Yunyun Duan:** Resources, Data curation. **Julien Cohen-Adad:** Conceptualization, Methodology, Supervision, Writing - review & editing. **Yaou Liu:** Resources, Data curation, Supervision.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The authors want to thank Olivier Vincent, Lucas Rouhier, Anthime Bucquet, Joseph Paul Cohen, and Sara Paperi for insightful discussions during the development of the model.

Funding

This work was supported by IVADO [EX-2018-4], Canada Research Chair in Quantitative Magnetic Resonance Imaging [950-230815], the Canadian Institute of Health Research [CIHR FDN-143263], the Canada Foundation for Innovation [32454, 34824], the Fonds de Recherche du Québec - Santé [28826], the Fonds de Recherche du Québec - Nature et Technologies [2015-PR-182754], the Natural Sciences and Engineering Research Council of Canada [RGPIN-2019-07244], the Canada First Research Excellence Fund (IVADO and TransMedTech), the Courtois NeuroMod project and the Quebec BioImaging Network [5886, 35450], Spinal Research and Wings for Life (INSPIRED project), the National Science Foundation of China [81870958, 81571631], the Beijing Municipal Natural Science Foundation for Distinguished Young Scholars [JQ20035], and the Special Fund of the Pediatric Medical Coordinated Development Center of Beijing Hospitals Authority [XTYB201831]. A.L. has a fellowship from NSERC and FRQNT. The authors thank the NVIDIA Corporation for the donation of a Titan X GPU.

References

- Akkus, Z., Galimzianova, A., Hoogi, A., Rubin, D.L., Erickson, B.J., 2017. Deep Learning for Brain MRI Segmentation: State of the Art and Future Directions. *J. Digit. Imaging* 30 (4), 449–459.
- Avants, B.B., Tustison, N., Song, G., 2009. Advanced Normalization Tools (ANTs). *The Insight Journal* 2 (365), 1–35.
- Baker, K.B., Moran, C.J., Wippold, F.J., Smirniotopoulos, J.G., Rodriguez, F.J., Meyers, S.P., Siegal, T.L., 2000. MR Imaging of Spinal Hemangioblastoma. *AJR Am. J. Roentgenol.* 174 (2), 377–382.
- Balériaux, D.L.F., 1999. Spinal Cord Tumors. *Eur. Radiol.* 9 (7), 1252–1258.
- Boonpirak, N., Apinhasmit, W., 1994. Length and Caudal Level of Termination of the Spinal Cord in Thai Adults. *Acta Anat.* 149 (1), 74–78.
- Christ, P.F., Ettlinger, F., Grün, F., Elshaera, M.E.A., Lipkova, J., Schlecht, S., Ahmaddy, F., et al., 2017. Automatic Liver and Tumor Segmentation of CT and MRI Volumes Using Cascaded Fully Convolutional Neural Networks. *arXiv [cs.CV]*. [arXiv: https://arxiv.org/abs/1702.05970](https://arxiv.org/abs/1702.05970).
- Chu, B.C., Terae, S., Hida, K., Furukawa, M., Abe, S., Miyasaka, K., 2001. MR Findings in Spinal Hemangioblastoma: Correlation with Symptoms and with Angiographic and Surgical Findings. *AJNR Am. J. Neuroradiol.* 22 (1), 206–217.
- Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O., 2016. 3D U-Net: Learning Dense Volumetric Segmentation from Sparse Annotation. In: *International conference on medical image computing and computer-assisted intervention*. Springer, Cham, pp. 424–432.
- Claus, E.B., Abdel-Wahab, M., Burger, P.C., Engelhard, H.H., Ellison, D.W., Gaiano, N., Gutmann, D.H., et al., 2010. Defining Future Directions in Spinal Cord Tumor Research: Proceedings from the National Institutes of Health Workshop. *Journal of Neurosurgery: Spine* 12 (2), 117–121.
- Das, J.M., Hoang, S., Mesfin, F.B., 2020. *Cancer, Intramedullary Spinal Cord Tumors*. StatPearls Publishing.

- De Leener, B., Dupont, S.M., Fonov, V.S., Nikola Stikov, D., Collins, L., Callot, V., Cohen-Adad, J., 2017. SCT: Spinal Cord Toolbox, an Open-Source Software for Processing Spinal Cord MRI Data. *NeuroImage* 145 (Pt A), 24–43.
- Eden, D., Gros, C., Badji, A., Dupont, S.M., De Leener, B., Maranzano, J., Zhuoquiong, R., Liu, Y., Granberg, T., Ouellette, R., Stawiarz, L., 2019. Spatial distribution of multiple sclerosis lesions in the cervical spinal cord. *Brain* 142 (3), 633–646.
- Frostell, A., Hakim, R., Thelin, E.P., Mattsson, P., Svensson, M., 2016. A Review of the Segmental Diameter of the Healthy Human Spinal Cord. *Front. Neurol.* 7 (December), 238.
- Gonzalez, C., Sakas, G. and Mukhopadhyay, A., 2020. What Is Wrong with Continual Learning in Medical Image Segmentation? arXiv [cs.CV]. arXiv. <http://arxiv.org/abs/2010.11008>.
- Gros, C., De Leener, B., Badji, A., Maranzano, J., Eden, D., Dupont, S.M., Talbott, J., Zhuoquiong, R., Liu, Y., Granberg, T., Ouellette, R., Tachibana, Y., Hori, M., Kamiya, K., Chougar, L., Stawiarz, L., Hillert, J., Bannier, E., Kerbrat, A., Edan, G., Labauge, P., Callot, V., Pelletier, J., Audoin, B., Rasosanandrianina, H., Brisset, J.-C., Valsasina, P., Rocca, M.A., Filippi, M., Bakshi, R., Tauhid, S., Prados, F., Yiannakas, M., Kearney, H., Ciccarelli, O., Smith, S., Treaba, C.A., Mainero, C., Lefeuvre, J., Reich, D.S., Nair, G., Auclair, V., McLaren, D.G., Martin, A.R., Fehlings, M.G., Vahdat, S., Khatibi, A., Doyon, J., Shepherd, T., Charlson, E., Narayanan, S., Cohen-Adad, J., 2019. Automatic Segmentation of the Spinal Cord and Intramedullary Multiple Sclerosis Lesions with Convolutional Neural Networks. *NeuroImage* 184, 901–915.
- Gros, C., De Leener, B., Dupont, S.M., Martin, A.R., Fehlings, M.G., Bakshi, R., Tummala, S., et al., 2018. Automatic Spinal Cord Localization, Robust to MRI Contrasts Using Global Curve Optimization. *Med. Image Anal.* 44, 215–227.
- Gros, C., Lemay, A., Vincent, O., Rouhier, L., Bucquet, A., Cohen, J.P., Cohen-Adad, J., 2020. ivadomed: A Medical Imaging Deep Learning Toolbox. *Journal of Open Source Software* 6 (58), 2868.
- Gros, C., Lemay, A., Cohen-Adad, J., 2021. SoftSeg: Advantages of soft versus binary training for image segmentation. *Medical Image Analysis* 71, 102038.
- Havaei, M., Davy, A., Warde-Farley, D., Biard, A., Courville, A., Bengio, Y., Pal, C., Jodoin, P.-M., Larochelle, H., 2017. Brain Tumor Segmentation with Deep Neural Networks. *Med. Image Anal.* 35, 18–31.
- Havaei, M., Guizard, N., Chapados, N., Bengio, Y., 2016. “HeMIS: Hetero-Modal Image Segmentation”. In *Medical Image Computing and Computer-Assisted Intervention* – Springer International Publishing.
- Hille, G., Steffen, J., Dünnwald, M., Becker, M., Saalfeld, S. and Tönnies, K., 2020. Spinal Metastases Segmentation in MR Imaging Using Deep Convolutional Neural Networks. arXiv [eess.IV]. arXiv. <http://arxiv.org/abs/2001.05834>.
- Huntoon, K., Wu, T., Elder, J.B., Butman, J.A., Chew, E.Y., Linehan, W.M., Oldfield, E.H., Lonser, R.R., 2016. Biological and Clinical Impact of Hemangioblastoma-Associated Peritumoral Cysts in von Hippel-Lindau Disease. *J. Neurosurg.* 124 (4), 971–976.
- Hussain, S., Anwar, S.M., Majid, M., 2017. Brain Tumor Segmentation Using Cascaded Deep Convolutional Neural Network. *Conference Proceedings: Annual International Conference of the IEEE Engineering in Medicine and Biology Society. IEEE Engineering in Medicine and Biology Society. Conference 2017 (July): 1998–2001*.
- Isensee, F., Kickingereder, P., Wick, W., Bendszus, M., Maier-Hein, K.H., 2017. Brain Tumor Segmentation and Radiomics Survival Prediction: Contribution to the BRATS 2017 Challenge. In: *International MICCAI Brainlesion Workshop*. Springer, Cham, pp. 287–297.
- Isensee, F., Kickingereder, P., Wick, W., Bendszus, M., Maier-Hein, K.H., 2019. No New-Net. In: *Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries*. Springer International Publishing, pp. 234–244.
- Kamboj, A., Rani, R., Chaudhary, J., 2018. Deep Learning Approaches for Brain Tumor Segmentation: A Review. In: *2018 First International Conference on Secure Cyber Computing and Communication (ICSCCC)*, pp. 599–603.
- Kayalibay, Baris, Grady Jensen, and Patrick van der Smagt. 2017. “CNN-Based Segmentation of Medical Imaging Data.” arXiv [cs.CV]. arXiv. <http://arxiv.org/abs/1701.03056>.
- Kim, D.H., Kim, J.-H., Choi, S.H., Sohn, C.-H., Yun, T.J., Kim, C.H., Chang, K.-H., 2014. Differentiation between Intramedullary Spinal Ependymoma and Astrocytoma: Comparative MRI Analysis. *Clin. Radiol.* 69 (1), 29–35.
- Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A.A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., Hassabis, D., 2017. Overcoming catastrophic forgetting in neural networks. *Proceedings of the national academy of sciences* 114 (13), 3521–3526.
- Lemay, A., Gros, C., Vincent, O., Liu, Y., Cohen, J.P., Cohen-Adad, J., 2021. Benefits of Linear Conditioning with Metadata for Image Segmentation. In *Medical Imaging with Deep Learning*.
- Lévy, S., Benhamou, M., Naaman, C., Rainville, P., Callot, V., Cohen-Adad, J., 2015. White Matter Atlas of the Human Spinal Cord with Estimation of Partial Volume Effect. *NeuroImage* 119, 262–271. <https://doi.org/10.1016/j.neuroimage.2015.06.040>.
- Maas, A.L., Hannun, A.Y., Andrew, Y.N., 2013. Rectifier Nonlinearities Improve Neural Network Acoustic Models. *Proc. Icm1* 30, 3.
- Milletari, F., Navab, N., Ahmadi, S.A., 2016. V-net: Fully convolutional neural networks for volumetric medical image segmentation. In: *2016 fourth international conference on 3D vision (3DV)*. IEEE, pp. 565–571.
- Müller, R., Kornblith, S., Hinton, G.E., 2019. In: “When Does Label Smoothing Help?” In *Advances in Neural Information Processing Systems 32*. Curran Associates Inc., pp. 4694–4703.
- Naceur, M.B., Saouli, R., Akil, M., Kachouri, R., 2018. Fully Automatic Brain Tumor Segmentation Using End-To-End Incremental Deep Neural Networks in MRI Images. *Comput. Methods Programs Biomed.* 166 (November), 39–49.
- Perone, C.S., Cohen-Adad, J., 2019. Promises and limitations of deep learning for medical image segmentation. *Journal of Medical Artificial Intelligence* 2.
- Pianyk, O.S., Langs, G., Dewey, M., Enzmann, D.R., Herold, C.J., Schoenberg, S.O., Brink, J.A., 2020. Continuous Learning AI in Radiology: Implementation Principles and Early Applications. *Radiology* 297 (1), 6–14.
- Reza, S.M., Roy, S., Park, D.M., Pham, D.L., Butman, J.A., 2019. Cascaded convolutional neural networks for spine chordoma tumor segmentation from MRI. In *Medical Imaging 2019: Biomedical Applications in Molecular, Structural, and Functional Imaging*. International Society for Optics and Photonics 10953, 1095325.
- Ronneberger, O., Fischer, P., Brox, T., 2015, October. U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical image computing and computer-assisted intervention*. Springer, Cham, pp. 234–241.
- Srivastava, N., Hinton, G., Krizhevsky, A., Sutskever, I., Salakhutdinov, R., 2014. Dropout: A Simple Way to Prevent Neural Networks from Overfitting. *Journal of Machine Learning Research: JMLR* 15, 1929–1958.
- Stroman, P.W., Wheeler-Kingshott, C., Bacon, M., Schwab, J.M., Bosma, R., Brooks, J., Cadotte, D., Carlstedt, T., Ciccarelli, O., Cohen-Adad, J., Curt, A., Evangelou, N., Fehlings, M.G., Filippi, M., Kelley, B.J., Kollias, S., Mackay, A., Porro, C.A., Smith, S., Strittmatter, S.M., Summers, P., Tracey, I., 2014. The Current State-of-the-Art of Spinal Cord Imaging: Methods. *NeuroImage* 84, 1070–1081.
- Ulyanov, D., Vedaldi, A. and Lempitsky, V., 2016. Instance normalization: The missing ingredient for fast stylization. arXiv [cs.CV]. arXiv <http://arxiv.org/abs/1607.08022>.
- Valliani, A.-A., Ranti, D., Oermann, E.K., 2019. Deep Learning and Neurology: A Systematic Review. *Neurology and Therapy* 8 (2), 351–365.
- Wang, X., 2017. Development of Segmentation Algorithms and Machine Learning Classification Methods for Characterization of Breast and Spine Tumors on MRI. University of California, Irvine.