



HAL
open science

Anisotropic Hybrid Networks for liver tumor segmentation with uncertainty quantification

Benjamin Lambert, Pauline Roca, Florence Forbes, Senan Doyle, Michel Dojat

► **To cite this version:**

Benjamin Lambert, Pauline Roca, Florence Forbes, Senan Doyle, Michel Dojat. Anisotropic Hybrid Networks for liver tumor segmentation with uncertainty quantification. REMIA 2023 - 2nd Workshop on Resource-Efficient Medical Image Analysis, Oct 2023, Vancouver (BC), Canada. 10.48550/arXiv.2308.11969 . hal-04436251

HAL Id: hal-04436251

<https://hal.science/hal-04436251>

Submitted on 5 Feb 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Anisotropic Hybrid Networks for liver tumor segmentation with uncertainty quantification

Benjamin Lambert^{1,2}, Pauline Roca², Florence Forbes³, Senan Doyle², and Michel Dojat¹

¹ Univ. Grenoble Alpes, Inserm, U1216, Grenoble Institut Neurosciences, 38000, FR

² Pixyl, Research and Development Laboratory, 38000 Grenoble, FR

³ Univ. Grenoble Alpes, Inria, CNRS, Grenoble INP, LJK, 38000 Grenoble, FR

Abstract. The burden of liver tumors is important, ranking as the fourth leading cause of cancer mortality. In case of hepatocellular carcinoma (HCC), the delineation of liver and tumor on contrast-enhanced magnetic resonance imaging (CE-MRI) is performed to guide the treatment strategy. As this task is time-consuming, needs high expertise and could be subject to inter-observer variability there is a strong need for automatic tools. However, challenges arise from the lack of available training data, as well as the high variability in terms of image resolution and MRI sequence. In this work we propose to compare two different pipelines based on anisotropic models to obtain the segmentation of the liver and tumors. The first pipeline corresponds to a baseline multi-class model that performs the simultaneous segmentation of the liver and tumor classes. In the second approach, we train two distinct binary models, one segmenting the liver only and the other the tumors. Our results show that both pipelines exhibit different strengths and weaknesses. Moreover we propose an uncertainty quantification strategy allowing the identification of potential false positive tumor lesions. Both solutions were submitted to the MICCAI 2023 Atlas challenge regarding liver and tumor segmentation.

Keywords: Hepatocellular carcinoma · Liver · Segmentation · Uncertainty

1 Introduction

Liver cancer is the sixth most common cancer worldwide, the fourth most common cause of cancer-related death and hepatocellular carcinoma (HCC) in particular is the most common type of primary liver cancer in adults. HCC is a cancer with large therapeutic choices depending on the tumor staging. In case of unresectable HCC, transarterial radioembolisation (TARE) is one of the recommended treatments: it consists of the delivering of 20–60 μM -sized yttrium-90 (Y90) microspheres to the tumor arterial supply, causing tumor necrosis through radiation-induced DNA damage and eventual cell death [13]. To calculate dosimetry and plan the intervention, the spatial extent of the tumor needs

to be determined [14]. In clinical routine, the liver and tumor are delineated using contrast-enhanced magnetic resonance imaging with four phases (pre-contrast, arterial, portal venous, delayed phases) allowing to differentiate the tumor from the liver parenchyma. These tasks are time-consuming and required high medical expertise.

In this context, automatic tools could reduce intra- and inter-observer variability and help radiologists to save time. Several studies [1,16] proposed CE-MRI based deep learning methods to segment liver and HCC tumors using private datasets (of 174 and 190 patients respectively). However, open-source MRI datasets are scarce, and existing ones [3] do not contain manual annotations, inhibiting the development of automated algorithms. In this context, the "A Tumor and Liver Automatic Segmentation" (ATLAS) challenge and dataset have been proposed. The aim of the challenge is to produce automatic segmentation maps of liver and the tumor(s) using CE-MRI of patients with unresectable HCC based on a small annotated training dataset (60 patients) [12]. This challenge is organised in collaboration with the Medical Image Computing and Computer Assisted Intervention (MICCAI) Workshop on Resource-Efficient Medical Image Analysis (REMIA) which will take place during the MICCAI 2023 edition in Vancouver, Canada. This paper describes our contribution to the challenge, including two different pipelines to segment liver and HCC tumors in CE-MRI. The first pipeline correspond to a multiclass model that simultaneously segment the liver and tumors. In the second pipeline, we used two different binary models, one focusing on the liver and the other on the tumors. Additionally, we propose to complement the segmentations with the estimation of an uncertainty score for each identified tumor lesion, potentially guiding the reviewing of the automated result by the user.

2 Materials and Methods

2.1 Dataset

The ATLAS challenge dataset is composed of an open-source training fold (60 subjects), allowing the development of segmentation models, as well as an hidden test set (30 subjects) for evaluation [12]. For each subject, a unique CE-MRI sequence is available, with one of the following phase: arterial, delayed, post-contrast, non-contrast or unknown. A ground truth segmentation mask is associated to each image, containing 3-classes: background, healthy liver and tumor lesions. Training images exhibit variable resolution, ranging from $0.684mm$ to $1.4mm$ in the XY plane, and from $2mm$ to $4.6mm$ in the Z-axis (see Figure 1). To develop our approach, we adopt a 5-fold cross validation scheme. In each fold, 40 subjects are used for training, 8 for validation and 12 for testing. The distribution of data in each fold is resumed in Table 1.

Fold	Training (N=40)					Validation (N=8)					Test (N=12)				
	A	D	NC	P	U	A	D	NC	P	U	A	D	NC	P	U
1	23	6	1	5	5	3	0	1	3	1	7	2	0	2	1
2	24	6	1	4	5	2	1	0	4	1	7	1	1	2	1
3	24	5	1	4	6	2	1	1	4	0	7	2	0	2	1
4	25	5	1	5	4	2	1	1	3	1	6	2	0	2	2
5	24	6	0	6	4	3	1	1	2	1	6	1	1	2	2

Table 1: Data distribution in each fold. A: arterial, D: delayed, NC: non-contrast, P: portal, U: unknown

2.2 Segmentation pipelines

In this section, we start by describing the design choices that are common to both pipelines, namely the neural network architecture and loss function. Then, we detail the particularity of both approaches (see Figure 2 for an illustration).

Joint design choices The first challenge to tackle is the anisotropy of the data, with the Z resolution being up to 4 times that in the plane. Instead of preprocessing the data to obtain a uniform spatial resolution, we instead chose to use a segmentation neural network able to handle the intrinsic anisotropy of the raw data. To this end, our two pipelines revolve around the use of the Anisotropic Hybrid U-Net (AHUNet) [8]. This architecture is composed of a pre-trained 2D convolutional encoder ignoring between-slice information, and a 3D convolutional decoder that reconstitute the 3D context. Second, although not the scope of the ATLAS challenge, we aim at enhancing the segmentation of the




Fig. 1: Distribution of the spatial resolution in the training dataset. Colors indicate the corresponding CE-MRI phase. To facilitate visualization, the plot is generated using jitter.

models by an uncertainty quantification module allowing to estimate a confidence score for each tumor lesion. Thus, the calibration of the output probability is of primary importance, which is usually poor with modern networks [5]. To circumvent this limitation, we resort to using the recently proposed Margin-based Label Smoothing loss [9] which has been proposed to improve the calibration of medical image segmentation models.

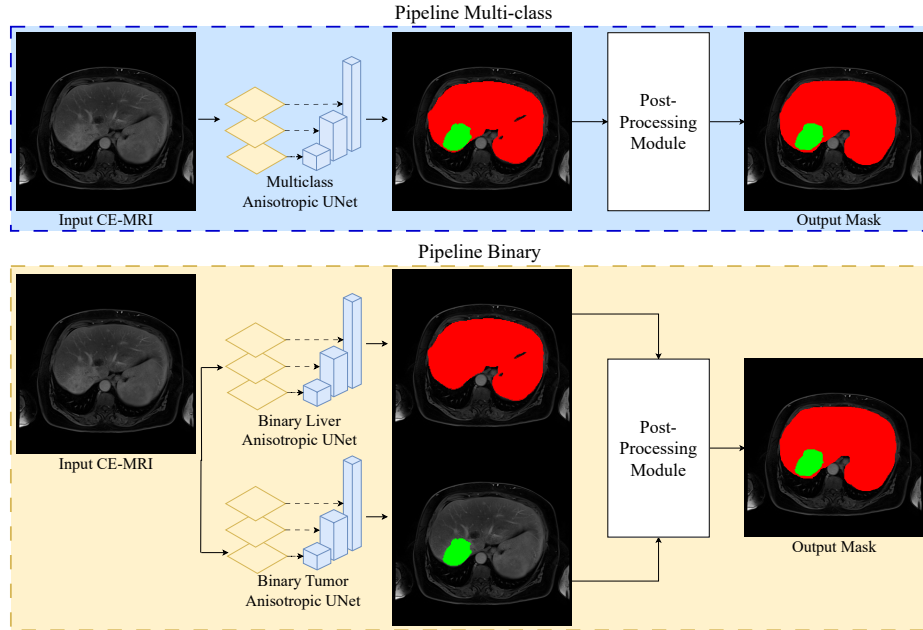


Fig. 2: Illustration of our two proposed pipelines, namely multiclass and binary.

2.3 Pipeline Multi-class

The most direct approach to obtain the segmentation of the liver and tumors is to train a multi-class AHUnet using the 3-classes ground truth masks. We adopt a patch-based approach, following which the input 3D MRI is divided into patches that are processed independently by the model. To guide the selection of the patches size, we computed bounding-boxes around each liver and tumor lesion in the training dataset using the ground truth masks (see Figure 3). This study indicates that defining a patch size of $256 \times 256 \times 64$ allows the encompassment of all tumor lesions, while also encompassing most of the livers.

2.4 Dual binary pipeline

Motivated by the observation that the segmentation of the liver was a much more easier task than tumor segmentation, we propose an alternative pipeline that disentangles both tasks. To do so, two distinct binary models are trained: i) a model that segments the *overall* liver (concatenation of the healthy liver class and the tumors) and ii) a model that segments the tumor lesions. From the bounding-box analysis (Figure 3), we kept the patch size of $256 \times 256 \times 64$ for the binary liver model. However, for the binary tumor model, we use a smaller patch size of $128 \times 128 \times 64$ as it allows to reduce the imbalance between background and tumor voxels, while still encompassing most of the nodules. At inference, the input MRI is processed by each model. The two resulting binary masks are then aggregated to reconstitute the final 3-class segmentation.

2.5 Post-processing Module

Post-processing has been proved as an important step for the automated segmentation of the liver in CT and MRI [4]. We adopt a 3-steps post-processing procedure, which is similar for both pipelines:

- Step 1:** Only the largest connected component for the liver is kept. A binary morphology *fill holes* operation [15] is performed on the resulting binary mask.
- Step 2:** Tumor lesions outside the liver are removed. This process differs slightly depending on the pipeline. For the multi-class approach, the binary mask of each identified tumor lesion is dilated. If this dilated mask has no intersection with liver voxels, the lesion is deleted. For the binary-approach, the process is more straight-forward, as we simply multiply the liver mask by the tumor mask. As an effect, tumor voxels outside the liver are filtered out.
- Step 3:** A binary closing morphological operation is applied to the remaining tumor lesions to improve the smoothness of the borders.

2.6 Lesion Uncertainty Quantification

Additionally to the segmentation of the liver and tumors, we propose to add a lesion uncertainty quantification module to our pipelines. This module complements the segmentation with an uncertainty score for each unique nodule. To achieve this, we collect the tumor probabilities $p_{i,tumor}$ ($i \in [1, N]$) for each of the N voxels of the lesion. The *lesion-wise* uncertainty score is then taken as:

$$L_{unc} = 1 - \frac{1}{N} \sum_{i=1}^N p_{i,tumor} \quad (1)$$




Fig. 3: Dimensions of bounding boxes encompassing the tumor lesions (left) and liver (right) on the Atlas dataset. In overlay we present the selected patch sizes of $256 \times 256 \times 64$ (red) and $128 \times 128 \times 64$ (blue).

2.7 Evaluation Protocol

The ATLAS challenge proposes to compare participants based on 9 metrics. The Dice, Surface Dice (SD), Hausdorff Distance (HD) and the Average Symmetric Surface Distance (ASD) are computed for both the *overall* liver (concatenation of the healthy liver and tumor voxels) and the tumors. Lastly, the Root Mean Square Error (RMSE) on tumor burden is computed, based on the computed volumes for the liver and tumors. To compare each pipeline, we train models on each of the 5 folds, and metrics are aggregated over the 5 test folds. Finally, qualitative assessment of the computed tumor uncertainty scores is assessed by monitoring the distribution of uncertainties for TP (non-zero intersection with the ground truth tumor mask) and FP (null intersection) lesions.

2.8 Ablation Study

To demonstrate the positive impact of using AHUnet over standard 3D segmentation architecture, we reproduced the multi-class pipeline by replacing the AHUnet by a Dynamic Unet 3D (DynUnet) which has been successful in various medical image segmentation challenges [6]. We also run the evaluation procedure without post-processing to demonstrate its gain on performance.

2.9 Implementation Details

Our framework is implemented in Pytorch [10] and experiments are carried on a Nvidia RTX A5000 GPU. AHUnet and DynUnet are implemented using the open-source MONAI library [2]. Models are trained using the ADAM optimizer [7] with a learning rate of $2e-4$ and a batch size of 2, until the loss on validation

data stagnates over 30 epochs. During training we use Data Augmentation implemented using the TorchIO library [11] comprising spatial (rotation, flips, affine transforms) as well as intensity (contrast, gamma, bias, noise) augmentations.

3 Results and Discussion

Table 2: Performance on the 60 subjects aggregated over the 5 test folds. Mean results are presented along the standard deviation.

	Post-processing	Multi-class AHUnet		Dual binary AHUnet		DynUnet
		✓	×	✓	×	✓
Liver	ASD ↓	2.06 ± 1.78	3.31 ± 3.53	1.99±2.35	5.03 ± 5.01	2.57 ± 2.70
	Dice ↑	0.94±0.04	0.93 ± 0.04	0.94±0.06	0.92 ± 0.07	0.93 ± 0.06
	HD (×10 ¹) ↓	3.22 ± 1.78	11.27 ± 6.03	2.85±1.69	12.79 ± 6.76	3.60 ± 2.60
	SD ↑	0.82 ± 0.10	0.80 ± 0.11	0.84±0.12	0.81 ± 0.12	0.81 ± 0.13
Tumor	ASD (×10 ¹) ↓	1.46±2.5	1.62 ± 2.50	1.53±2.31	3.51 ± 4.21	1.90 ± 2.25
	Dice ↑	0.57±0.28	0.57±0.28	0.54 ± 0.29	0.49 ± 0.30	0.49 ± 0.29
	HD (×10 ¹) ↓	6.73±3.84	8.87 ± 5.74	7.77 ± 4.13	18.0 ± 7.89	8.49 ± 4.34
	SD ↑	0.40±0.21	0.40±0.21	0.41±0.22	33.9 ± 0.21	0.33 ± 0.20
	RMSE (×10 ⁻²) ↓	0.73 ± 1.38	0.75 ± 0.01	0.72±1.30	1.17 ± 2.10	1.04 ± 2.19

Segmentation performance of each pipeline is presented in Table 2. Qualitative assessment of lesion uncertainty is presented in Figure 4.

For both pipelines, the liver is far better segmented than tumors. On the multi-class AHUnet pipeline for example a mean Dice score of 0.94 is reached for liver segmentation whereas for tumor we reached a value of 0.57. This can be explained by the small volume of tumors as compared to the liver, and by the important variability in appearance depending on the contrast phase. Moreover, these results are similar to Bousabarah study [1] that obtained Dice scores per case of 0.91 and 0.48 for liver and tumor segmentation respectively using a deep-learning method based on the four phases of CE-MRI exams. Our post-processing module enabled an important gain on the 9 metrics, for both pipelines. The dual binary pipeline obtained better metrics on the liver than the multi-class one. However, for tumor, the multi-class AHUnet is slightly superior. Finally, the multi-class pipeline relying on the DynUnet (ablation study) is worst than the multi-class AHUnet for all metrics, showing the relevance of using a neural network specifically designed for anisotropic data.

To assess if the differences between both proposed pipelines are statistically significant, we performed a significance analysis, presented in Table 3. We first test the hypothesis that the obtained scores are drawn from a normal distribution (Shapiro’s test). If the hypothesis is valid, one-sided student T-test are

performed, else we opt for Wilcoxon’s tests. Results show that the Liver metrics are significantly better for the Dual binary pipeline than for the multiclass pipeline. On the contrary, the Tumor HD is significantly lower for the multiclass pipeline than for the binary one. Other metrics are not statistically different between both pipelines (Tumor ASD, Dice, HD and RMSE).

Finally, the qualitative assessment of our lesion uncertainty module (Figure 4, left) shows that FP lesions tend to be attributed to higher uncertainty scores than TP lesions). The proposed score could thus be used to draw the user’s attention to uncertain lesions, that are more likely to be FP. The proposed lesion uncertainty score is highly correlated with the lesion volume (Figure 4, right), reaching a Spearman correlation coefficient of -0.82 , which is intuitive as FP lesions are more frequent among small lesions. To complement this analysis, we provide two examples of the proposed lesion quantification in Supplementary Material.

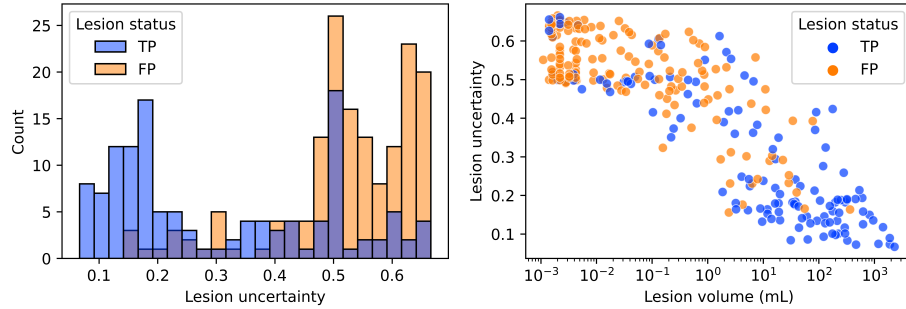


Fig. 4: Qualitative evaluation of the tumor lesion uncertainty obtained with the multi-class pipeline. Left: histogram of uncertainty estimates with respect to the lesion status (TP or FP). Right: lesion uncertainty with respect to the lesion volume (in log-scale).

4 Conclusion

In this work, we addressed the problem of segmenting the liver and tumors in CE-MRI. Using the AHU-net architecture as cornerstone, we propose two different pipelines, multi-class and dual binary, demonstrating a gain on tumor and liver segmentation respectively. Additionally, we propose to quantify the uncertainty of identified nodules, which can be used to improve the interpretability of the automated predictions. For the final ATLAS 2023 challenge submissions, models trained on the 5 folds are aggregated in an ensemble, for each pipeline separately. Both pipelines (binary and multiclass) are submitted as separate algorithms.

Table 3: Significance analysis between the two pipelines. B: binary pipeline; M: multi-class pipeline.

		Shapiro p-value	Hypothesis	Test	p-value	Significance ($p < 0.05$)
Liver	ASD ↓	4.9×10^{-13}	B < M	Wilcoxon	6.2×10^{-4}	✓
	Dice ↑	7.1×10^{-14}	B > M	Wilcoxon	7.1×10^{-14}	✓
	HD ↓	3.0×10^{-4}	B < M	Wilcoxon	4.6×10^{-2}	✓
	SD ↑	1.9×10^{-9}	B > M	Wilcoxon	1.1×10^{-7}	✓
Tumor	ASD ↓	3.0×10^{-7}	M < B	Wilcoxon	2.0×10^{-1}	×
	Dice ↑	2.3×10^{-4}	M > B	Wilcoxon	2.3×10^{-1}	×
	HD ↓	4.1×10^{-4}	M < B	Wilcoxon	1.4×10^{-3}	✓
	SD ↑	3.8×10^{-1}	B > M	T-test	7.6×10^{-1}	×
	RMSE ↓	2.7×10^{-8}	B < M	Wilcoxon	7.7×10^{-1}	×

5 Supplementary Material




Fig. 5: Case A: patient with a FP lesion (dashed arrow) associated with a high uncertainty value (= 0.47); Case B: patient presenting one TP lesion (arrow) associated with a low lesion uncertainty (0.13) and a second FP lesion (dashed arrows) associated with a low uncertainty (0.17) in hypersignal on arterial phase. Complementary information (precontrast, portal and delayed CE-MRI T1 phases) are required to assess if this nodule is a CHC.

References

1. Bousabarah, K., Letzen, B., Tefera, J., Savic, L., Schobert, I., Schlachter, T., Staib, L.H., Kocher, M., Chapiro, J., Lin, M.: Automated detection and delineation of hepatocellular carcinoma on multiphasic contrast-enhanced MRI using deep learning **46**(1), 216–225, publisher: Springer
2. Cardoso, M.J., Li, W., Brown, R., Ma, N., Kerfoot, E., Wang, Y., Murrey, B., Myronenko, A., Zhao, C., Yang, D., et al.: Monai: An open-source framework for deep learning in healthcare. arXiv preprint arXiv:2211.02701 (2022)

3. Clark, K., Vendt, B., Smith, K., Freymann, J., Kirby, J., Koppel, P., Moore, S., Phillips, S., Maffitt, D., Pringle, M., et al.: The cancer imaging archive (tcia): maintaining and operating a public information repository. *Journal of digital imaging* **26**, 1045–1057 (2013)
4. Furtado, P.: Loss, post-processing and standard architecture improvements of liver deep learning segmentation from computed tomography and magnetic resonance. *Informatics in Medicine Unlocked* **24**, 100585 (2021)
5. Guo, C., Pleiss, G., Sun, Y., Weinberger, K.Q.: On calibration of modern neural networks. In: *International conference on machine learning*. pp. 1321–1330. PMLR (2017)
6. Isensee, F., Petersen, J., Klein, A., Zimmerer, D., Jaeger, P.F., Kohl, S., Wasserthal, J., Koehler, G., Norajitra, T., Wirkert, S., et al.: nnu-net: Self-adapting framework for u-net-based medical image segmentation. *arXiv preprint arXiv:1809.10486* (2018)
7. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *3rd International Conference on Learning Representations, ICLR 2015* (2015)
8. Liu, S., Xu, D., Zhou, S.K., Pauly, O., Grbic, S., Mertelmeier, T., Wicklein, J., Jerebko, A., Cai, W., Comaniciu, D.: 3d anisotropic hybrid network: Transferring convolutional features from 2d images to 3d anisotropic volumes. In: *Medical Image Computing and Computer Assisted Intervention–MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part II* 11. pp. 851–858. Springer (2018)
9. Murugesan, B., Liu, B., Galdran, A., Ayed, I.B., Dolz, J.: Calibrating segmentation networks with margin-based label smoothing. *Medical Image Analysis* **87**, 102826 (2023)
10. Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., Killeen, T., Lin, Z., Gimelshein, N., Antiga, L., et al.: Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems* **32** (2019)
11. Pérez-García, F., Sparks, R., Ourselin, S.: Torchio: a python library for efficient loading, preprocessing, augmentation and patch-based sampling of medical images in deep learning. *Computer Methods and Programs in Biomedicine* **208**, 106236 (2021)
12. Quinton, F., Popoff, R., Presles, B., Leclerc, S., Meriaudeau, F., Nodari, G., Lopez, O., Pellegrinelli, J., Chevallier, O., Ginhac, D., et al.: A tumour and liver automatic segmentation (atlas) dataset on contrast-enhanced magnetic resonance imaging for hepatocellular carcinoma. *Data* **8**(5), 79 (2023)
13. Reig, M., Forner, A., Rimola, J., Ferrer-Fàbrega, J., Burrel, M., Garcia-Criado, A., Kelley, R.K., Galle, P.R., Mazzaferro, V., Salem, R., Sangro, B., Singal, A.G., Vogel, A., Fuster, J., Ayuso, C., Bruix, J.: BCLC strategy for prognosis prediction and treatment recommendation: The 2022 update **76**(3), 681–693. <https://doi.org/10.1016/j.jhep.2021.11.018>
14. Smits, M.L., Elschot, M., Sze, D.Y., Kao, Y.H., Nijsen, J.F., Iagaru, A.H., de Jong, H.W., van den Bosch, M.A., Lam, M.G.: Radioembolization dosimetry: the road ahead. *Cardiovascular and interventional radiology* **38**, 261–269 (2015)
15. Virtanen, P., Gommers, R., Oliphant, T.E., Haberland, M., et al.: SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods* **17**, 261–272 (2020)
16. Zheng, R., Wang, Q., Lv, S., Li, C., Wang, C., Chen, W., Wang, H.: Automatic liver tumor segmentation on dynamic contrast enhanced mri using 4d information: deep

learning model based on 3d convolution and convolutional lstm. IEEE Transactions on Medical Imaging **41**(10), 2965–2976 (2022)