



CT Evaluation of 2D and 3D Holistic Deep Learning Methods for the Volumetric Segmentation of Airway Lesions

Amel Imene Hadj Bouzid, Baudouin Denis de Senneville, Fabien Baldacci, Pascal Desbarats, Patrick Berger, Ilyes Benlala, Gaël Dournes

► To cite this version:

Amel Imene Hadj Bouzid, Baudouin Denis de Senneville, Fabien Baldacci, Pascal Desbarats, Patrick Berger, et al.. CT Evaluation of 2D and 3D Holistic Deep Learning Methods for the Volumetric Segmentation of Airway Lesions. ISBI 2024 - 21st IEEE International Symposium on Biomedical Imaging, May 2024, Athènes, Greece. pp.1-5, 10.1109/ISBI56570.2024.10635201 . hal-04793701

HAL Id: hal-04793701

<https://hal.science/hal-04793701v1>

Submitted on 21 Nov 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.

CT EVALUATION OF 2D AND 3D HOLISTIC DEEP LEARNING METHODS FOR THE VOLUMETRIC SEGMENTATION OF AIRWAY LESIONS

*Amel Imene Hadj Bouzid*¹ *Baudouin Denis de Senneville*² *Fabien Baldacci*³
*Pascal Desbarats*³ *Patrick Berger*¹ *Ilyes Benlala*⁴ *Gaël Dournes*⁴

¹Univ. Bordeaux, INSERM, CRCTB, U 1045, F-33000 Bordeaux, France

²Univ. Bordeaux, CNRS, Bordeaux INP, IMB, UMR 5251, F-33400 Talence, France

³Univ. Bordeaux, CNRS, Bordeaux INP, LaBRI, UMR 5800, F-33400 Talence, France

⁴Service d'Imagerie Médicale Radiologie Diagnostique et Thérapeutique, CHU de Bordeaux, France

ABSTRACT

This research embarked on a comparative exploration of the holistic segmentation capabilities of Convolutional Neural Networks (CNNs) in both 2D and 3D formats, focusing on cystic fibrosis (CF) lesions. The study utilized data from two CF reference centers, covering five major CF structural changes. Initially, it compared the 2D and 3D models, highlighting the 3D model's superior capability in capturing complex features like mucus plugs and consolidations. To improve the 2D model's performance, a loss adapted to fine structures segmentation was implemented and evaluated, significantly enhancing its accuracy, though not surpassing the 3D model's performance. The models underwent further validation through external evaluation against pulmonary function tests (PFTs), confirming the robustness of the findings. Moreover, this study went beyond comparing metrics; it also included comprehensive assessments of the models' interpretability and reliability, providing valuable insights for their clinical application.

Index Terms— Segmentation, CT, cystic fibrosis, loss function

1. INTRODUCTION

Deep learning algorithms are increasingly used in semantic segmentation of medical imaging, especially in identifying and segmenting multi-class disease-related abnormalities in CT scans [1]. This progress is crucial in airway imaging, where non-invasive methods are key for quantifying and characterizing lung abnormalities, essential in tracking the progression and severity of diseases like cystic fibrosis (CF) [2].

Cystic fibrosis, characterized by excessive mucus production, results in serious respiratory complications. CT scans play a vital role in evaluating lung disease severity in CF. However, the commonly used visual scoring systems are subjective, with interpretations varying among evaluators [3]. Moreover, these systems typically categorize lung damage in broad terms,

potentially lacking the sensitivity required for detailed, longitudinal monitoring of the disease.

As CF treatments advance, the need for non-invasive, quantitative biomarkers to assess lung disease has become more pressing [4]. While 2D CNNs have been effective in detecting CF markers in CT scans [5], the full potential of 3D volumetric analysis to accurately depict the diverse forms of CF lesions remains largely unexplored. These lesions, varying from tubular bronchiectasis and peribronchial thickening to centrilobular mucus with "tree-in-bud" patterns, and dense consolidations, suggest that a 3D approach could offer more comprehensive insights into these complex structures.


In addressing the segmentation challenges of CF lesions, we employed the nnU-Net architecture in both its 2D and 3D variants for semantic segmentation [6]. This approach facilitated an in-depth comparison between these methodologies in segmenting CF lesions, which are characterized by a variety of shapes, spatial distributions, and textures. We also refined the nnU-Net's loss function, moving beyond the traditional Dice coefficient to better address the segmentation of both larger and smaller, more intricate lesions [7].

2. OVERVIEW


In this study, we utilized 2D and 3D nnU-Net networks with customized loss functions to segment CF lesions from 75 clinical exams. For the training phase, we selected 50 CT scans, ensuring a balanced representation of both healthy individuals and patients with a range of lesion counts. For the network input, two modalities were prepared: one including the lung envelope and another excluding it.

The internal evaluation involved cross-validation on the 50 CT using metrics specifically designed for detecting lesions of varying sizes. The best-performing models in 2D and 3D were then externally evaluated on a separate cohort of 25 patients, correlating their performance with FEV1% scores. Additionally, these models underwent model-agnostic analysis with Gradient-weighted Class Activation Mapping (Grad-CAM)

[8] and data-agnostic uncertainty calculation to provide interpretability and enable a comprehensive comparison between 2D and 3D models. The methodological approach is illustrated in the Figure 1.



(a) Example of patient labeling.



(b) Training and Testing Pipeline

Fig. 1. Study Method. CF lesions are color-coded as follows: Red for Bronchiectasis, Green for Peribronchial Thickening, Blue for Bronchial Mucus, Yellow for Bronchiolar Mucus, and Cyan for Consolidation.

3. EXPERIMENTAL SETUP

3.1. Data Description

75 CT scans were conducted using equipment from major manufacturers (GE Revolution® and Siemens Somatom Force®), with slice thicknesses ranging from 1 to 1.25 mm. Manual lesion segmentation was carried out using 3D Slicer 4.11 software by three experienced observers from a CF reference center. The CT images were displayed using parenchymal window settings (width at 1500 Hounsfield Units; level at -450 Hounsfield Units). For validation purposes, a minimum of 80% visual agreement among the observers was required.

3.2. Segmentation model

nnU-Net served as the cornerstone of our segmentation framework. Its encoder, comprising multiple blocks, each contained a sequence of layers. These layers included a convolutional layer, dropout, instance normalization, and Leaky ReLU activation, collaboratively working to extract and process features from the input data. The decoder in nnU-Net was divided into two main sections. The first section involved deconvolution, crucial for creating high-resolution feature maps, and was constituted by a single convolutional layer. The second section, known as Features Localization, consisted of three layers: a

convolutional layer, instance normalization, and Leaky ReLU activation. The convolution kernel sizes in the 2D version of nnU-Net were 3×3 with a 2×2 stride, while the 3D variant utilized $3 \times 3 \times 3$ kernels with a stride of $2 \times 2 \times 2$. Skip connections between the encoder and decoder established direct links between layers at various depths within the network, facilitating the propagation of information and contributing to enhanced segmentation accuracy. The learning rate (LR) schedule began at 10^{-2} and was updated every 30 epochs. In the optimization process, Stochastic Gradient Descent (SGD) with Nesterov momentum of 0.99 was employed, aiding in iterative parameter updates and model convergence. Notably, nnU-Net used patch sizes of 512×512 in 2D and $96 \times 160 \times 160$ in 3D, with batch sizes set to 12 in 2D and 2 in 3D.

3.3. Loss Function

We considered the **cross-entropy** (CE), which was defined as $CE = -\frac{1}{N} \sum_{i=1}^N [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$, where y_i was the true label of pixel i , $\hat{y}_i$ was the predicted probability for that pixel, and N represented the total number of pixels in the image. Additionally, the **Dice coefficient** was utilized, defined as $Dice = \frac{2 \times \sum_{i=1}^N y_i \hat{y}_i}{\sum_{i=1}^N y_i + \sum_{i=1}^N \hat{y}_i}$, where y_i and $\hat{y}_i$ were the true and predicted labels, respectively, for each pixel i .

The initial loss function was defined as $Loss = CE + Dice$, which had been effective in accurately segmenting larger lesions [9]. However, considering the complexity of tasks involving lesions of various sizes, we adopted a loss function that initially prioritized the Dice coefficient for the effective segmentation of larger lesions within the folds. As training progressed, the focus shifted towards improving the segmentation of the 50% of pixels that were less accurately predicted by the model. This loss function, which leveraged cross-entropy as its base and incorporated overlap measurements as weighted regularizers, demonstrated stability during training, particularly in adapting to smaller structures [10]. The Top50 variant of CE, focusing on the most challenging 50% of the pixels, was defined as $Top50 = -\frac{1}{N} \sum_{i \in \mathcal{P}_{50}} [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)]$, where $\mathcal{P}_{50}$ represented the set of pixels with the largest prediction errors.

The customized loss function was the Weighted Dice Top50, formulated as $WDiceTop50 = (1 - \alpha) \times Dice + \alpha \times Top50$, where α indicated the ratio between the current training epoch and the total number of epochs.

3.4. Internal Evaluation

To assess the accuracy of the CF lesion segmentation models, the Dice coefficient was primarily utilized. This measure was proficient in determining the degree of overlap between predicted and actual lesion areas. However, its performance was reduced for smaller structures [11]. Consequently, the Normalized Surface Distance (NSD) was also considered, owing to its

strength in identifying cases where predictions were close but not exactly coincident with the actual lesions [7]. The NSD assessed the similarity between the predicted segmentation and the ground truth, taking into account a margin around the object’s boundary. In our study, this margin was set at 3 pixels ($\approx 1.8\text{mm}$) to accommodate small variations or uncertainties at the segmentation’s edge. Sensitivity and Specificity were used to ensure the thorough detection of significant structures. Furthermore, the Area Under the Curve (AUC) reflected the model’s ability to effectively differentiate between various classes.

3.5. External Evaluation

In the external evaluation of our study, we specifically focused on the correlation between the volumes predicted by the 2D and 3D nnU-Net models and the FEV1% volumes measured in a cohort of 25 patients. FEV1% indicates the volume of air a patient can forcibly exhale in one second. To establish this relationship, we employed the Spearman’s rank correlation test, which was particularly apt for our analysis due to the involvement of non-linear relationships and non-normally distributed data. Our goal was to validate the results of the internal evaluation by confronting the models with clinical data.

3.6. Interpretability

3.6.1. Model-agnostic

To enhance our understanding of nnU-Net’s functionality, we used Grad-CAM [8], which produces heatmaps highlighting key image regions according to the model’s focus. This process involved isolating each label, inputting unseen CT-images into nnU-Net, calculating gradients for feature maps in the final convolutional layer, and determining the importance of each feature map. A final Grad-CAM heatmap was generated by applying ReLU activation to emphasize positively contributing features.

3.6.2. Data-agnostic

We modified the network into a Bayesian framework, treating weights as probabilistic distributions. This involved using Monte Carlo Dropout [12] with a 0.3 dropout rate in the convolutional block, creating five distinct network versions. Predictions were computed for each version, and variance among these predictions was calculated across cross-validation folds. This approach quantified the model’s uncertainty and provided insights into its prediction confidence.

4. RESULTS AND DISCUSSION

The internal evaluation, as shown in Table 1, revealed that both 2D and 3D nnU-Net models were effective in segmenting




Fig. 2. The comparison of CF lesion segmentations is presented. The first column displays the Ground Truth 3D segmentations, while the second and third columns show the predictions by the modified nnU-Net 3D and 2D models, respectively.




Fig. 3. The analysis showcases CF lesion feature maps: Ground Truth segmentations in the first row, followed by Grad-CAM feature maps of modified nnU-Net 3D and 2D models in the second and third rows, respectively. The maps’ intensity levels correlate with prediction probabilities, where higher intensity signifies higher probabilities.

bronchiectasis and peribronchial thickening, achieving Dice scores of 80% and 64%, respectively, and NSD scores of 80% for both. The 3D model, however, demonstrated superior performance in detecting smaller lesions such as mucus and consolidations. The custom loss function, designed to enhance the detection of smaller structures, significantly improved the 2D model’s performance. Specifically, bronchial and bronchiolar mucus detection improved by 6% and 4% ($p < 0.05$) in Dice scores, and 2% ($p < 0.05$) in NSD for both. While the Dice score for consolidation detection in the 2D model decreased by 2%, the NSD increased by 6%, indicating more accurate detection despite imperfect overlap in multiple slices. Despite these improvements, the 2D model could not match the 3D model’s performance, which saw a marginal increase in performance due to the change in loss function, with a 1% overall increase in Dice and NSD scores. Sensitivity and specificity analysis showed that all four models were specific, with the modified loss enhancing the sensitivity of the 2D model. However, the 2D model’s AUC of 77% was still lower than the 3D model’s 80%.

For external evaluation, as detailed in Table 2, models with the modified loss were selected based on their superior perfor-

Metrics	Configuration	Bronchiectasis	Peribronchial Thickening	Bronchial mucus	Bronchiolar mucus	Consolidation	Avg
Dice	2D <i>DiceCE</i>	0.77	0.63	0.55	0.28	0.61	0.57
	2D <i>WDiceTop50</i>	0.80	0.64	0.61	0.32	0.59	0.59
	3D <i>DiceCE</i>	0.80	0.64	0.64	0.39	0.61	0.61
	3D <i>WDiceTop50</i>	0.80	0.65	0.64	0.39	0.60	0.62
NSD	2D <i>DiceCE</i>	0.79	0.81	0.60	0.44	0.53	0.63
	2D <i>WDiceTop50</i>	0.78	0.79	0.62	0.46	0.59	0.65
	3D <i>DiceCE</i>	0.80	0.80	0.67	0.51	0.67	0.69
	3D <i>WDiceTop50</i>	0.80	0.81	0.68	0.51	0.65	0.70
Sensibility	2D <i>DiceCE</i>	0.74	0.63	0.47	0.22	0.65	0.54
	2D <i>WDiceTop50</i>	0.76	0.64	0.52	0.43	0.62	0.59
	3D <i>DiceCE</i>	0.77	0.63	0.58	0.34	0.64	0.59
	3D <i>WDiceTop50</i>	0.77	0.64	0.57	0.36	0.61	0.60
Specificity	2D <i>DiceCE</i>	1.00	1.00	1.00	1.00	0.99	0.99
	2D <i>WDiceTop50</i>	1.00	1.00	1.00	0.80	0.99	0.96
	3D <i>DiceCE</i>	1.00	1.00	1.00	1.00	0.99	0.99
	3D <i>WDiceTop50</i>	1.00	1.00	1.00	1.00	0.99	0.99
AUC	2D <i>DiceCE</i>	0.87	0.81	0.73	0.61	0.82	0.77
	2D <i>WDiceTop50</i>	0.88	0.82	0.76	0.61	0.80	0.77
	3D <i>DiceCE</i>	0.88	0.81	0.79	0.67	0.82	0.79
	3D <i>WDiceTop50</i>	0.88	0.82	0.78	0.68	0.80	0.80

Table 1. Internal Evaluation of the nnU-Nets predictions. Best results are reported with bold characters.

n=25		2D FEV1%	3D FEV1%
Bronchiectasis	ρ	-0.46	-0.60
	p-value	0.02	<0.001
Peribronchial Thickening	ρ	-0.45	-0.59
	p-value	0.02	<0.001
Bronchial mucus	ρ	-0.34	-0.50
	p-value	0.09	0.005
Bronchiolar mucus	ρ	-0.30	-0.49
	p-value	0.10	0.006
Consolidation	ρ	-0.24	-0.39
	p-value	0.18	0.04

Table 2. External Evaluation on FEV1% of the modified 2D and 3D nnU-Nets. ρ = Spearman correlation coefficient.

mance in internal validation. Correlation analysis with FEV1% indicated significant correlations for bronchiectasis and peribronchial thickening volumes predicted by the modified 2D and 3D nnU-Net models. However, correlations for mucus and consolidation were not significant for the 2D model, whereas they were significant for the 3D model, supporting the findings of the internal evaluation. Figure 2 presents an example in 3D of CF predictions, where mucus and consolidation are less detected in the modified nnU-Net 2D compared to its 3D counterpart.

Interpretability, as demonstrated in Figure 3 with Grad-CAM feature maps for the same patient, showed that the 2D model exhibited more hesitation, particularly in identifying mucus, compared to the 3D model. Additionally, uncertainty estimation, calculated through variances across the five versions of nnU-Net in both 2D and 3D, indicated only a 10^{-4}

deviation, suggesting comparable robustness in network predictions between the 2D and 3D models.

Our study challenged the prevalent belief in the literature [13] about the superiority of 2D models, showing that 3D models perform better in segmenting diverse CF lesion forms, notably tubular bronchiectasis and centrilobular mucus. The 3D models’ performance is attributed to their efficient use of spatial information, a crucial factor even with limited training data. The importance of this discovery is amplified by our dataset’s challenges, dominated by small mucus plugs and bronchiolar impactions, as depicted in grads 1 and 2.

Exploring the use of separate networks for each lesion type and analyzing metrics individually for each class [14] could provide more in-depth insights. Furthermore, incorporating attention mechanisms [15] might enhance decoder guidance through spatial attention maps. This study contributes to the ongoing discourse on the role of Dice loss in semantic segmentation [16], highlighting the necessity for a variety of metrics and approaches in the design and evaluation of networks, particularly in the complex context of multi-class segmentation.

5. CONCLUSION

The 2D and 3D models demonstrate comparable effectiveness in detecting bronchiectasis and peribronchial thickening. However, the 3D model is better in identifying more complex lesions like mucus and consolidations, benefiting from its ability to analyze spatial relationships across sequential slices. Our study also highlights the potential advantages of alternative loss functions over the traditional Dice loss in semantic segmentation, particularly for nuanced structures. The results of this study could pave the way for new, fully automated,

volumetric scoring systems to quantify CF severity in future clinical applications.

6. COMPLIANCE WITH ETHICAL STANDARDS

All patients were provided written informed consent for reusing data from their medical records according to French Laws and after Institutional Review Board approval (study registration number CHUBX2020RE0267).

7. ACKNOWLEDGMENTS

This study was conducted in the framework of the University of Bordeaux's IdEx "Investments for the Future" program RRI "IMPACT" which received financial support from the French government.

8. REFERENCES

- [1] T. Boeken, J. Feydy, A. Lecler, P. Soyer, A. Feydy, M. Barat, and al, "Artificial intelligence in diagnostic and interventional radiology: Where are we now?," *Diagnostic and Interventional Imaging*, vol. 104, pp. 1–5, 2022.
- [2] M. Kirby and BM. Smith, "Quantitative ct scan imaging of the airways for diagnosis and management of lung disease," *CHEST*, vol. 164, pp. 1150–1158, 2023.
- [3] AS. Brody, MR. Kosorok, Z. Li, LS. Broderick, JL. Foster, A. Laxova, and al, "Reproducibility of a scoring system for computed tomography scanning in cystic fibrosis," *Journal of Thoracic Imaging*, vol. 21, pp. 14–21, 2006.
- [4] KJ. Bayfield, TA. Douglas, T. Rosenow, JC. Davies, SJ. Elborn, M. Mall, and al, "Time to get serious about the detection and monitoring of early lung disease in cystic fibrosis," *Thorax*, vol. 76, pp. 1255–1265, 2021.
- [5] G. Dournes, CS. Hall, MM. Willmering, AS. Brody, J. Macey, S. Bui, B. Denis de Senneville, P. Berger, F. Laurent, I. Benlala, and JC. Woods, "Artificial intelligence in computed tomography for quantifying lung changes in the era of cftr modulators," *European Respiratory Journal*, vol. 59, 2022.
- [6] F. Isensee, PF. Jaeger, SAA. Kohl, J. Petersen, and KH. Maier-Hein, "nnu-net: a self-configuring method for deep learning-based biomedical image segmentation," *Nature Methods*, vol. 18, pp. 203–211, 2021.
- [7] A. Reinke, MD. Tizabi, CH. Sudre, M. Eisenmann, T. Rädtsch, M. Baumgartner, and al, "Common limitations of image processing metrics: A picture story," *ArXivOrg*, 2023.

- [8] RR. Selvaraju, M. Cogswell, A. Das, R. Vedantam, D. Parikh, and D. Batra, “Grad-cam: Visual explanations from deep networks via gradient-based localization,” *IEEE International Conference on Computer Vision (ICCV)*, pp. 618–626, 2017.
- [9] J. Ma, J. Chen, M. Ng, R. Huang, Y. Li, C. Li, and al, “Loss odyssey in medical image segmentation,” *Medical Image Analysis*, vol. 71, 2021.
- [10] SA. Taghanaki, K. Abhishek, JP. Cohen, J. Cohen-Adad, G. Hamarneh, M. Baumgartner, and al, “Deep semantic segmentation of natural and medical images: A review,” *ArXivOrg*, 2020.
- [11] L. Maier-Hein, E. Al, B. Menze, and al, “Metrics reloaded: Pitfalls and recommendations for image analysis validation,” *ArXivOrg*, 2022.
- [12] Y. Gal and Z. Ghahramani, “Dropout as a bayesian approximation: representing model uncertainty in deep learning,” *Proceedings of the 33rd International Conference on International Conference on Machine Learning*, vol. 48, pp. 1050–1059, 2016.
- [13] L. Crespi, D. Loiacono, P. Sartori, and R. Kobbi, “Are 3d better than 2d convolutional neural networks for medical imaging semantic segmentation?,” *International Joint Conference on Neural Networks (IJCNN)*, pp. 1–8, 2022.
- [14] D. Müller, I. Soto-Rey, and F. Kramer, “Towards a guideline for evaluation metrics in medical image segmentation,” *BMC Research Notes*, vol. 15, 2022.
- [15] X. Zhao, P. Zhang, F. Song, C. Ma, G. Fan, Y. Sun, and al, “Prior attention network for multi-lesion segmentation in medical images,” *IEEE Transactions on Medical Imaging*, vol. 41, pp. 3812–3823, 2022.
- [16] B. Liu, J. Dolz, A. Galdran, R. Kobbi, and I. Ben Ayed, “Do we really need dice? the hidden region-size biases of segmentation losses,” *Medical Image Analysis*, vol. 91, 2024.