

Interactive Segmentation for COVID-19 Infection Quantification on Longitudinal CT scans

Michelle Xiao-Lin Foo, Seong Tae Kim, Magdalini Paschali, Leili Goli, Egon Burian, Marcus Makowski, Rickmer Braren, Nassir Navab, Thomas Wendler

Abstract—Consistent segmentation of COVID-19 patients' CT scans across multiple time points is essential to assess disease progression and response to therapy accurately. Existing automatic and interactive segmentation models for medical images only use data from a single time point (static). However, valuable segmentation information from previous time points is often not used to aid the segmentation of a patient's follow-up scans. Also, fully automatic segmentation techniques frequently produce results that would need further editing for clinical use. In this work, we propose a new single network model for interactive segmentation that fully utilizes all available past information to refine the segmentation of follow-up scans. In the first segmentation round, our model takes 3D volumes of medical images from two-time points (target and reference) as concatenated slices with the additional reference time point segmentation as a guide to segment the target scan. In subsequent segmentation refinement rounds, user feedback in the form of scribbles that correct the segmentation and the target's previous segmentation results are additionally fed into the model. This ensures that the segmentation information from previous refinement rounds is retained. Experimental results on our in-house multiclass longitudinal COVID-19 dataset show that the proposed model outperforms its static version and can assist in localizing COVID-19 infections in patient's follow-up scans.

Index Terms—COVID-19, CT, Interactive Segmentation, Deep Learning

I. INTRODUCTION

OVER the last few years, there has been immense progress in the medical image analysis field due to the introduction of deep learning networks. Specifically, such automatic segmentation methods are nowadays widely adopted in data annotation software.

This work was supported in part by the Bavarian Research Foundation (BFS) under grant agreement AZ-1429-20C, the Bavarian Ministry of Economics, Development and Energy under grant agreement DIK-2003-0003/DIK0127/02, the Korean Government (MSIT) under IITP Grant 2021-0-02068 and NRF Grant 2021R1G1A1094990. (Corresponding author: Seong Tae Kim)

M.X.-L. Foo, M. Paschali, N. Navab, and T. Wendler are with Chair for Computer Aided Medical Procedures and Augmented Reality, Technical University of Munich, Germany.

S.T. Kim is with the Department of Computer Science and Engineering, Kyung Hee University, South Korea (email: st.kim@khu.ac.kr)

L. Goli is with Department of Electrical and Computer Engineering, University of Toronto, Canada

E. Burian, M. Makowski, and R. Braren are with Institute for Diagnostic and Interventional Neuroradiology, Technical University of Munich, Germany

In December 2019, the first cases of a new coronavirus disease, COVID-19, a severe acute respiratory illness, emerged in Wuhan, China [28]. This highly infectious respiratory virus rapidly spread worldwide and threw the world into a global pandemic. According to the numbers on the COVID-19 monitoring site by Johns Hopkins University¹, as of August 30th, 2021, more than 216 million people have been infected and 4.498 million people have succumbed due to the complications caused by the virus.

During the initial outbreak of COVID-19, there was an urgent need for fast annotation of medical scans in order to further understand the disease. Numerous researchers utilized deep learning-based methods for this task [7], [20], [23]. Computed tomography (CT) scans provide crucial diagnostic information in the assessment and treatment of COVID-19 patients [1], [13]. However, medical reports show observations that the imaging features of COVID-19 are mixed and diverse among patients, and the change in radiological patterns during the course of the disease is inconsistent [15], [21], [29]. The subtle anatomical boundaries and variations in size, density, location, and texture of the disease pose a challenge for automatic segmentation techniques.

Human interactions coupled with deep learning models can offer a way to overcome the challenges faced by automatic segmentation models and to improve segmentation results as shown by [18], [22], [27]. However, previous works on interactive segmentation only used single time point data for segmentation. The readily available segmentation information from previous time points has not been exploited to segment a patient's follow-up scans.

To address the limitations of automatic static segmentation models, in this paper, we propose an interactive segmentation method that segments COVID-19 infection on longitudinal CT scans. The proposed method is designed to leverage available longitudinal information and user feedback to improve segmentation quality. The main contributions of this paper can be summarized as follows:

- We propose a new segmentation approach that utilizes information from previous time point, past segmentation refinement rounds, and user feedback. To the best of our knowledge, it is the first work to design an interactive segmentation on longitudinal CT scans. Our method is also the first interactive segmentation of COVID-19

¹<https://coronavirus.jhu.edu/map.html>

infection using longitudinal COVID-19 CT scans.

- Our method can be used to extend existing static models for longitudinal interactive segmentation with minimal effort.
- We conduct an extensive study using an in-house longitudinal COVID-19 dataset to showcase the improved performance of the longitudinal interactive segmentation model over the static interactive segmentation model.

II. RELATED WORKS

A. COVID-19 Infection CT Segmentation

After the outbreak of COVID-19, in order to understand the disease and support the radiologists several studies were presented. For the task of learning from noisy labels to segment COVID-19 pneumonia lesions from lung CT scans, Wang et al. [23] proposed a noise-robust training method. Shan et al. [20] presented a modified 3D convolutional neural network called VB-Net, which combines V-Net [16], a fully convolutional neural network for volumetric medical image segmentation with a bottle-neck deep residual learning framework for quantitative COVID-19 infection assessment. Fan et al. [7] proposed a semi-supervised learning approach for segmenting diverse radiological patterns such as ground-glass opacity and consolidation from lung CT scans. However, due to the subtle anatomical boundaries, pleural-based location, and high variations in infection characteristics, it is still challenging to automatically identify and quantify CT image findings related to COVID-19 [21], [29].

B. Interactive Segmentation

In interactive segmentation, feedback from users is used to enhance the predictions of machine learning models. This human-in-the-loop method shows potential in improving segmentation results [18]. However, interactive segmentation needs to be accurate and efficient in order to be helpful in a clinical setting. Xu et al. [25] introduced a deep interactive object selection method where user-provided clicks are transformed into Euclidean distance maps. However, Euclidean distance does not exploit image context information. In contrast, the geodesic distance transform by [5] further encodes spatial regularization and contrast-sensitivity information but suffers from sensitivity to noise. The deep learning-based interactive segmentation framework by Wang et al. [22] incorporated user-provided bounding boxes and scribbles (lines drawn on wrongly segmented areas). With the inclusion of this user feedback, they demonstrated an increase in segmentation performance. Zhou et al. [27] showed that with a small number of user interactions, segmentation accuracy can be substantially improved. Kitrungratsakul et al. [12] proposed a segmentation refinement module that can be appended to automatic segmentation networks and utilized a skip connection attention module to improve important features for both segmentation and refinement tasks. Withal, the methods introduced above are designed for single time point data.

C. Longitudinal Image Segmentation

To effectively study how a disease progresses, consistent segmentation of the affected regions based on scans from multiple time points can further provide important information. Birenbaum et al. [4] suggest the use of multiple longitudinal networks to process longitudinal patches from different views where the model concatenates the output of the encoder to produce Multiple Sclerosis (MS) lesion segmentation. Their method shows that the inclusion of information from multiple time points is beneficial to the model. To improve segmentation of MS lesions by taking advantage of the longitudinal information, [6] proposes a longitudinal network with an early fusion of two-time points scans to encode the structural differences implicitly. But due to the minute structural differences in MS lesions across different time points, it is still a challenge for the network to achieve a high accuracy. Kim et al. [11] present a framework that leverages spatio-temporal cues between longitudinal scans to improve quantitative assessment of the progression of COVID-19 infection in chest CT scans. However, the implementations above do not utilize the available previous time scan segmentation mask to segment a patient's follow-up scans.

III. PROPOSED METHOD

A. Interactive Segmentation Network

The proposed method extends the baseline longitudinal network by Denner et al. [6] to fully exploit longitudinal information and user feedback for interactive segmentation. The baseline longitudinal model is referred to as a 2.5D model as it retains the global context of the CT volume by combining the per-slice prediction (2D) from three anatomical planes (coronal, sagittal, and axial views) to produce the segmentation for one voxel. Each slice is processed by FC-DenseNet56 [10], which is a fully convolutional dense network for 2D segmentation. Previous work [2], [3], [8], [26] has shown that state-of-the-art results on various medical image segmentation problems are achievable through 2.5D approaches. This is due to the reasons that fully 3D approaches induce a high computational cost, and the patch-based 3D approaches lose the global structural information in the slice. Therefore, our baseline model is built as the 2.5D model, which uses two-time points consisting of stacked 2D slices from three anatomical views. This method can preserve global information along the two axes and local information from the third axis [26].

Let $\mathbf{V}_{t+1} \in \mathbb{R}^{h \times w \times s}$ and $\mathbf{V}_t \in \mathbb{R}^{h \times w \times s}$ denote the follow-up target CT volume and the reference previous volume, respectively. h and w are the height and width of the input, and s is the number of slices in the volume. \mathbf{X}_{t+1} and \mathbf{X}_t indicate the individual slices of the volumes. In this study, we assume that the segmentation masks for the reference previous volume $\mathbf{S}_t \in \{0, 1\}^{C \times h \times w}$ is available. C represents the number of foreground classes, and it is set to 2 in this study since we are targeting the segmentation of two foreground classes (i.e., ground-glass opacity and consolidation). In addition, let $\mathbf{E}_{t+1} \in \{-1, 1\}^{C \times h \times w}$ denote the editing masks on the target segmentation during the user feedback. Note that the

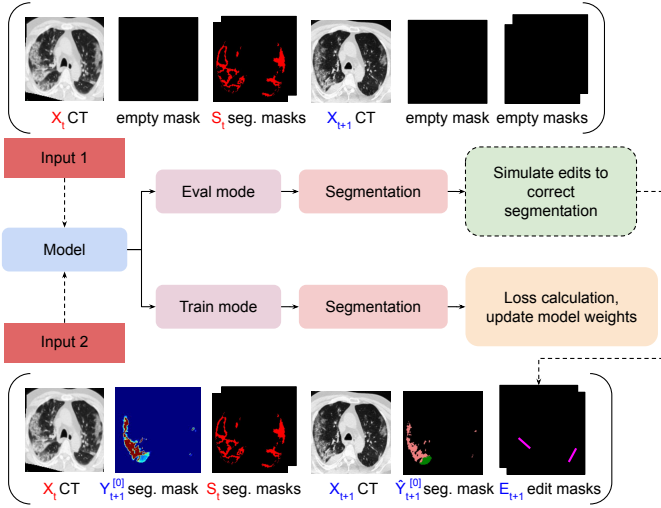


Fig. 1: Training flow of the proposed interactive longitudinal segmentation model. Training alternates between two training inputs that represent different scenarios: *Input 1* for initial segmentation round. *Input 2* for interactive segmentation rounds. Note that \mathbf{S}_t and \mathbf{E}_{t+1} two channels in our case as $C = 2$. Accordingly if \mathbf{E}_{t+1} is not available as it is in the scenario of *Input 1*, \mathbf{E}_{t+1} is a pair of empty masks.

segmentation masks for the reference previous volume and the editing masks are concatenated to the input.

1) *Training:* To train the model to adapt to different input information combinations during inference, throughout the training process, the model is randomly trained with two different inputs as shown in Figure 1.

Input 1 represents the scenario where the user passes the data to the model at the beginning of interactive segmentation to produce the first segmentation for the target slice, whereas *input 2* represents the input data in the subsequent editing rounds. In both cases, scans from the two-time points with additional data are concatenated along the channel dimension so that structural changes that are evident between them are utilized by the model to improve its segmentation performance. Empty masks are used in place of information that is not available in the first segmentation round, such as user feedback and target prediction.

Let $\mathbf{I}_{t+1}^{[T]} \in \{0, 1\}^{h \times w \times 8}$ denote the input tensor for the segmentation round T . The input $\mathbf{I}_{t+1}^{[T]}$ consists of the reference previous CT slice \mathbf{X}_t , the segmentation masks on the reference previous CT slice \mathbf{S}_t , the target CT slice \mathbf{X}_{t+1} , the highest class probability per pixel on the target slice from the previous segmentation round $\mathbf{Y}_{t+1}^{[T-1]} \in \{0, 1\}^{h \times w}$, the predicted class on the target slice from the previous segmentation round $\hat{\mathbf{Y}}_{t+1}^{[T-1]} \in \{0, 1, 2\}^{h \times w}$ (0 for background, 1 for ground-glass opacity, and 2 for consolidation), the editing masks \mathbf{E}_{t+1} as shown in Figure 1. For *Input 2*, the empty mask is used for $\mathbf{Y}_{t+1}^{[T-1]}$, $\hat{\mathbf{Y}}_{t+1}^{[T-1]}$, \mathbf{E}_{t+1} . The interactive segmentation refinement network (ISR), f_{ISR} outputs the per-slice segmentation of the target image as follows:

$$\mathbf{Y}_{t+1}^{[T]} = f_{ISR}(\mathbf{I}_{t+1}^{[T]}) \in \{0, 1\}^{h \times w} \quad (1)$$

To produce the simulated edits for training with *input 2* the model is first set to evaluation mode (without model weights updates), and an initial segmentation of the target is generated. This is then used to produce the simulated edits. The edit simulation process will be further introduced in Section III-B.

During training, the segmentation of the slices is treated as a 2D segmentation problem. During inference with real user feedback, the prediction on slices from the three anatomical orientations are combined to produce the segmentation output for every voxel. The pseudocode of the training process of the proposed interactive longitudinal interactive segmentation model is presented in Algorithm 1.

Algorithm 1 Training of Proposed Interactive longitudinal Segmentation model, f_{ISR}

```

1: for  $1 \leq e \leq epochs$  do
2:   for  $i = 1$  to  $N$  do
3:     Get 1st predictions for input batch  $i$ 
4:     Generate random number,  $Z \in [0, 1]$ 
5:     if  $Z > 0.5$  then
6:       Generate simulated edits for the predictions
7:       Append simulated edits & outputs from
8:       1st prediction round to inputs
9:       Get 2nd predictions
10:      Calculate loss using 2nd predictions
11:    else
12:      Calculate loss using 1st predictions
13:    end if
14:    Backpropagate & update model weights
15:  end for
16: end for

```

B. Edit Simulation during Training

During training, simulated user edits for wrongly segmented regions are automatically generated. Incorrectly segmented regions are areas that are under- or over-segmented. The segmentation output from the model is compared with the ground truth to choose the slice region for simulating the user edits. Lines are automatically drawn on the selected regions as simulated feedback. As mentioned before, the edit information is concatenated to the CT scans as additional channels, one for each class, with foreground interaction having a value of 1 and background interaction -1. Because the model input is 2.5D instead of 3D, edits are simulated in the axial, coronal, and sagittal slices as opposed to only the axial slices as in [27]. Zhou et al. [27] simulate edits only for the most extensive 2D incorrectly segmented slice region. However, due to the scattered nature of the COVID-19 infections in our case, the top-5 largest wrongly segmented regions in each slice is used for edit simulation. The total number of generated edits, independent of the different classes are limited to prevent the model from overfitting and also to avoid considerable slowdown in the training process caused by the long edit simulation time needed when large numbers of incorrectly segmented regions are detected.

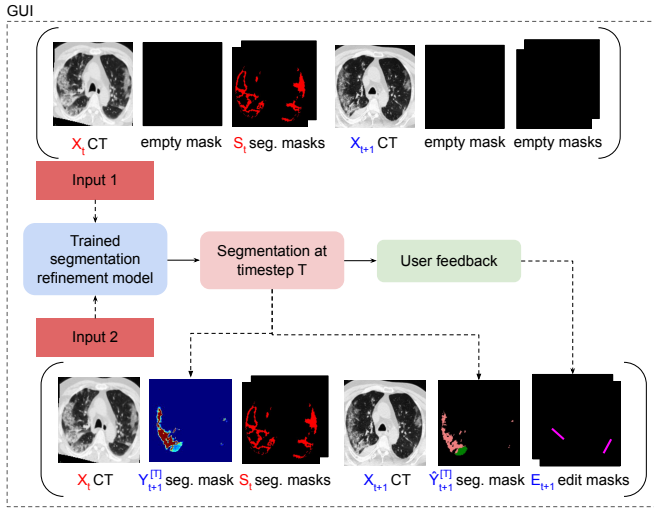


Fig. 2: Overview of the interactive segmentation flow with the GUI. Input to the segmentation model prior to the first user interaction resembles *Input 1* that is used during training. Accordingly, input of subsequent editing rounds resembles *Input 2* as in Figure 1 which enables assisted segmentation refinement via the GUI.

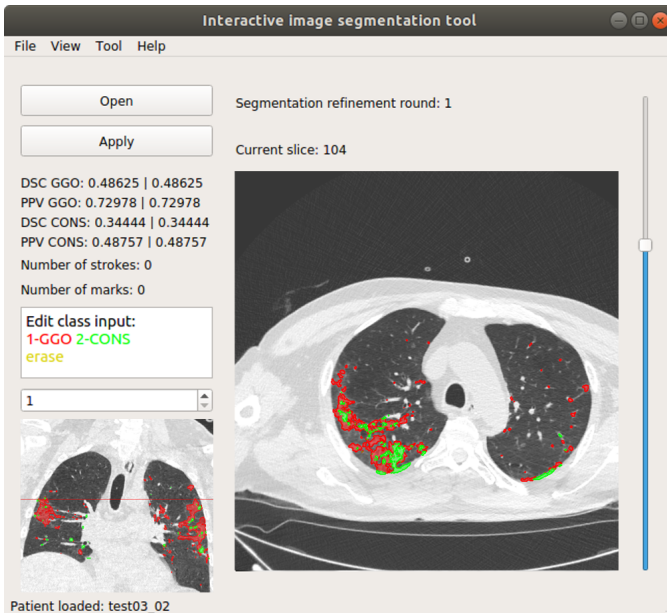


Fig. 3: Our GUI for COVID-19 lung infection interactive segmentation. The spin box in the side bar shows the class for the current brush input, here 1, i.e., the brush for ground-glass opacity (GGO).

C. Inference with GUI for User Feedback

Figure 2 shows the interactive segmentation refinement stage. An editing graphical user interface (GUI) as shown in Figure 3 is implemented using Qt Designer². The GUI automatically loads the trained segmentation refinement model during start-up. Then, the user can load the CT volume to be segmented. After the initial segmentation round, the predicted

segmentation will be overlaid on the scans for inspection. The user can then use the brush to edit wrongly segmented areas and run the data through the model again. This process can be repeated as many times as necessary.

For each segmentation refinement round, the user feedback on incorrectly segmented regions from current and previous rounds is summed up in each slice’s edit mask. The current round user feedback has higher priority, and so it is multiplied by two before being added to the previous user edits. The values of the mask are then clipped to $[-1, 1]$. This is done so that the previous edit information is not lost. The mask is then concatenated to its corresponding slice image and fed into the segmentation refinement model.

D. Implementation

The baseline longitudinal segmentation model from [6] is modified and used in this study. It is an end-to-end 2.5D segmentation network based on FC-DenseNet56 [10] and implemented in PyTorch 1.4 [17]. Mean Squared Error (MSE) loss, Adam optimizer with AMSGrad [19] and a learning rate of 0.0001 are used for training. The inference time for processing a COVID-19 patient’s 2.5D data with a size of $3 \times (150 \times 150 \times 150)$ takes 15 seconds on an NVIDIA GeForce RTX 2080 Ti with 11GB GPU.

IV. EXPERIMENTS

A. COVID-19 Segmentation Dataset and Preprocessing

An in-house clinical dataset collected from the Radiology Department of Technical University of Munich during the COVID-19 first wave (March-June 2020) is used for training and evaluation. It consists of 30 longitudinal low-dose native CT scans from patients age between 46 and 82 years old with a positive polymerase chain reaction (PCR) test for COVID-19. The time gap between the follow-up scans and the previous scans is 17 ± 10 days; the scans were taken during admission and hospitalization (33 ± 21 days, 0–71 days). The scans were performed using two different CT imaging devices (IQon Spectral CT and iCT 256, Philips Healthcare, Best, the Netherlands) with the same parameters (X-ray current 140-210 mA, voltage 120kV peak, slice thickness 0.9mm) and covered the entire lung. The data was collected with the approval of the institutional review board of TUM (ethics approval 111/20 S-KH).

An expert rater (radiologist with four years of experience) annotated the dataset at voxel-level with the ImFusion Labels software (ImFusion GmbH, Munich, Germany³). Lung masks (lung parenchyma vs. other tissues) and pathology masks for four classes: healthy lung (HL), ground-glass opacity (GGO), consolidation (CONS), and pleural effusion (PLEFF) are generated. Due to the large variations in intensity range, size, and alignment, the raw CT volumes have to be preprocessed before they are used for training. The volumes are cropped to the lung regions using manually annotated lung masks. Intensity values outside the range $(-1024, 600)$ are clipped, and min-max normalization is performed on the volumes before they

²<https://doc.qt.io/qt-5/qt designer-manual.html>

³ImFusion, <https://www.imfusion.com/>




Fig. 4: Deformable registration example of reference scan to target scan. *Left*: Reference scan *Right*: Target scan.




Fig. 5: Examples of GGO and CONS on axial slices of COVID-19 patients' lung CTs. Arrows point to the infection areas. *Left*: GGO *Right*: CONS.

are resized to $150 \times 150 \times 150$ pixels. Slices that have a voxel-value variation smaller than 0.001% between their min and max values are considered empty and removed.

Similar to [11], we also use the deformable registration algorithm by [14], where the image is deformed through a B-Spline Transform that uses a sparse set of grid points overlaid onto the fixed domain of the image, to register the reference scan to the follow-up scan and to resolve the misalignment error between scans. Registration is performed on the lung masks to avoid registration errors that may arise due to the pathological changes in the lung parenchyma. An example of aligned CT scans from different time points are shown in Figure 4.

According to [9], GGO is the most common findings in COVID-19 patients' CT scans, followed by CONS. Figure 5 shows examples of GGO and CONS from our dataset. For our experiments, we only segment GGO and CONS, due to the low occurrences of PLEFF in the patient cohort of the dataset.

For the 30 patients, the reference and follow-up CT scans of each patient after registration have an average structural similarity index (SSIM) [24] of 29.71%. This shows that the scans taken at different time points broadly differ from one another perceptually. Besides that, the average change in the percentage of GGO and CONS in the patients' lung CTs from different time points is 13.68% and 6.59%, respectively. This indicates the noticeable difference in the disease progression over time in the dataset. Table I shows the percentage of GGO and CONS in the lungs of the patients at each timestep.

The training set is made up of 16 patients (37 volumes), with training (n=12) validation (n=4) split. Our model is tested on an independent test set consisting of 14 patients (28 volumes).

TABLE I: Percentage of GGO and CONS in the lungs of patients. T refers to the timestep of the scans. n is the number of patient volumes at each timestep

Radiomic	Average		Std. Dev.		Average		Std. Dev.	
	T-1 (n=30)	T-2 (n=30)	T-1 (n=30)	T-2 (n=30)	T-3 (n=5)	T-3 (n=5)	T-3 (n=5)	
GGO	15.23%	14.76%	20.17%	18.07%	15.02%	11.91%		
CONS	6.52%	7.07%	8.15%	11.51%	11.28%	11.07%		

B. Experimental Settings

1) *Evaluation Metrics*: The segmentation performance of the models is evaluated using the following metrics.

- **Dice Similarity Coefficient (DSC)** is a statistical measure of the similarity between two segmentations.

$$DSC = \frac{2TP}{2TP + FP + FN} \quad (2)$$

- **Positive Predictive Value (PPV, or precision)** displays the fraction of correctly segmented regions over all predicted segmentations.

$$PPV = \frac{TP}{TP + FP} \quad (3)$$

- **True Positive Rate (TPR, recall or sensitivity)** shows the proportion of correct segmentation outputs with respect to the ground truth.

$$TPR = \frac{TP}{TP + FN} \quad (4)$$

- **Volume Difference (VD)** is calculated as the absolute difference in the predicted lesion segmentation volume and ground truth lesion segmentation volume over the ground truth lesion segmentation volume.

$$VD = 100 \times \frac{|\text{Lesion_volume}_{pred} - \text{Lesion_volume}_{gt}|}{\text{Lesion_volume}_{gt}} \quad (5)$$

C. Ablation Study

In order to study how the additional information concatenated to the inputs influences the model's segmentation performance, an ablation study was carried out using our longitudinal COVID-19 dataset. In addition to the longitudinal baseline network [6], we also implemented a static version of the network for comparison. `long_edit+ref_seg` is the baseline `long.+ref_seg` extended for interactive segmentation. The proposed model additionally incorporates past prediction outputs of the target as additional information to guide the following segmentation, and `static_edit` is the version of the proposed model without reference information. Table II summarizes all tested models and their inputs. The reference manual segmentation in the table refers to ground truth masks of the reference images, whereas the edit masks contain the user feedback on the target segmentation.

TABLE II: Input for different models tested in ablation study

Model name	Model input				
	Target image	Reference image	Ref. manual segmentation	Edit mask	Target previous segmentation
Baseline static network	✓				
Baseline long. network	✓	✓			
Baseline long.+ref_seg	✓	✓	✓		
static_edit	✓			✓	✓
long_edit+ref_seg	✓	✓	✓	✓	
Proposed	✓	✓	✓	✓	✓

As it can be inferred from Table III, the baseline models performed better than the interactive segmentation models in the first segmentation round without edits and the baseline longitudinal models have higher Dice scores compared to the baseline static model. Among the longitudinal baseline models, concatenating reference segmentation to the input CTs can further improve its GGO Dice by 1.44% and CONS Dice by 0.76%. Comparing the interactive segmentation models, the longitudinal interactive segmentation models output better initial segmentation than the static model. The proposed model's GGO Dice is 8.45% higher than the static model, whereas the long_edit+ref_seg model has an improvement of 15.7% in its CONS Dice over the static model.

The evaluations results in Table III further revealed that the CONS TPR for the static models are considerably lower than the longitudinal models, with the non interactive methods having a difference of 8.58% between the baseline static network and baseline long. network. For the interactive methods,

the static_edit model CONS TPR is 10.34% less than the longitudinal interactive model with lower TPR. These results proves that longitudinal models can improve segmentation of a more complicated class such as CONS.

D. Segmentation Results and Discussions

The following evaluations are carried out using the GUI to obtain real user feedback on the segmentation output.

1) *Quantitative Results:* The aim of the interactive segmentation model is to assist and reduce the user's workload during the segmentation of new data. Thus, its desired function is to take in rough user interactions and improve the segmentation output. To further determine whether the segmentation refinement model serves this purpose, the initial segmentation that is manually corrected and fed into the segmentation refinement model is compared with the model refined segmentation output. The results are shown in Table IV.

As can be seen from the results in Table IV, the interactive segmentation refinement model was able to use the previous segmentation results and user feedback to improve the segmentation of the target scans. Out of the three models, the proposed model has the highest Dice scores for GGO and CONS after one round of segmentation refinement. However, by comparing the change in Dice of the edited initial segmentation by the user and refined segmentation by the model, the static interactive segmentation model showed the most considerable improvement in its segmentation output, with an increase of 12.73% in GGO Dice and 31.15% in CONS Dice, whereas

TABLE III: Evaluation results on the test set before user interactions. Values displayed are the mean and standard errors. Bold values represent the best results for each metric. Empty masks are used in place of the edit mask and target image previous segmentation mask for static_edit, long_edit+ref_seg and the proposed model.

Model	Dice (%)		PPV (%)		TPR (%)		VD (%)	
	GGO	CONS	GGO	CONS	GGO	CONS	GGO	CONS
Non interactive methods								
Baseline static network	44.15 ± 3.33	19.75 ± 4.89	62.12 ± 2.89	42.86 ± 8.53	38.98 ± 5.09	14.96 ± 3.60	49.98 ± 6.07	102.0 ± 28.84
Baseline long. network	45.42 ± 3.03	27.63 ± 5.61	67.34 ± 3.03	41.35 ± 8.01	37.96 ± 4.46	23.54 ± 4.32	48.56 ± 6.09	93.18 ± 38.04
Baseline long.+ref_seg	46.86 ± 3.12	28.39 ± 6.39	62.61 ± 3.86	42.12 ± 8.53	41.63 ± 4.16	24.36 ± 5.61	43.31 ± 6.53	90.36 ± 39.43
Interactive methods								
static_edit	35.97 ± 4.20	13.99 ± 3.51	66.82 ± 3.01	41.20 ± 8.45	29.03 ± 4.99	9.28 ± 2.27	59.12 ± 7.20	98.11 ± 22.21
long_edit+ref_seg	37.70 ± 3.89	29.69 ± 5.86	64.66 ± 3.03	37.85 ± 7.60	30.39 ± 4.46	26.80 ± 5.35	55.46 ± 6.75	64.88 ± 27.62
Proposed	44.42 ± 3.61	23.94 ± 6.12	62.65 ± 3.49	41.89 ± 8.86	40.38 ± 5.42	19.62 ± 5.92	51.84 ± 7.40	49.71 ± 7.94

TABLE IV: Evaluation results of the interactive methods on the test set with real user interactions. Values displayed are the mean and standard errors. Bold values represent the best results for each metric.

Model	Dice (%)		PPV (%)		TPR (%)		VD (%)	
	GGO	CONS	GGO	CONS	GGO	CONS	GGO	CONS
Initial segmentation with user edits								
static_edit	40.86 ± 3.16	24.33 ± 3.14	67.80 ± 3.01	54.37 ± 7.45	32.77 ± 4.33	17.16 ± 1.97	54.15 ± 6.31	87.99 ± 20.73
long_edit+ref_seg	44.34 ± 2.34	36.71 ± 5.57	67.03 ± 3.30	44.38 ± 6.84	35.72 ± 3.39	33.91 ± 5.05	48.76 ± 5.41	54.67 ± 20.78
Proposed	49.22 ± 2.58	36.33 ± 5.19	64.34 ± 3.88	50.50 ± 7.10	44.64 ± 4.50	30.91 ± 5.13	46.63 ± 6.54	40.08 ± 6.03
Output segmentation after one round of segmentation refinement by model								
static_edit	53.59 ± 2.33	55.48 ± 4.48	71.63 ± 4.12	71.34 ± 4.34	44.62 ± 2.98	47.01 ± 4.68	40.21 ± 3.76	33.68 ± 5.66
long_edit+ref_seg	54.79 ± 1.95	54.01 ± 5.03	66.70 ± 3.62	51.14 ± 5.16	48.54 ± 2.54	64.44 ± 5.67	34.74 ± 4.55	74.43 ± 36.10
Proposed	59.86 ± 2.57	58.81 ± 4.36	62.06 ± 4.48	58.18 ± 5.15	61.86 ± 3.00	62.39 ± 4.06	36.39 ± 11.28	34.11 ± 12.75

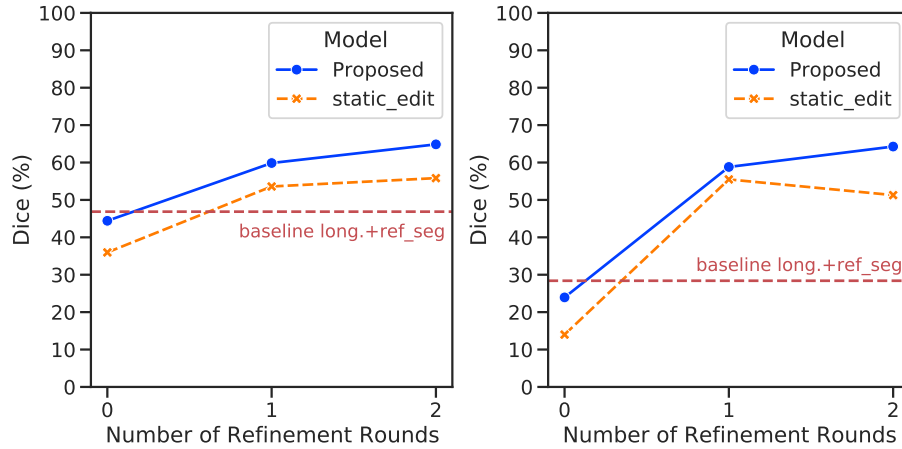


Fig. 6: Average change in Dice vs Number of Refinement Rounds for 14 test patients. *Left*: class GGO and *Right*: class CONS.

the proposed model showed an increase of 10.64% in GGO Dice and 22.48% in CONS Dice.

The proposed model and static_edit model are further evaluated with another round of segmentation refinement. Figure 6 shows how the Dice of 14 patients change after each segmentation refinement rounds with real user edits. The best baseline model: baseline long.+ref_seg is added for comparison.

Experimental results showed that the proposed model improves the Dice by a significant amount just after two rounds of segmentation refinement with real user feedback. The total average increase of the segmentation Dice after two rounds of segmentation refinement is 20.44% for GGO and 40.33% for CONS, with an average increase of 4.99% for GGO and 5.46% for CONS between the first and second segmentation refinement rounds. As for the static_edit model, after two rounds of segmentation refinement, the GGO and CONS Dice is 9.03% and 12.98% lower than the proposed model.

From the plots in Figure 6, it is visible that the proposed model is superior in terms of segmentation refinement performance compared to the static interactive segmentation model. The static_edit model is shown to be unable to further refine the segmentation correctly at the second refinement step; the GGO segmentation Dice improved, whereas the CONS segmentation Dice dropped. Closer inspection of the segmentation output shows that the static_edit model is more likely to wrongly classify regions, which leads to the Dice decrease. As for the proposed model, the inclusion of reference scan information improved its segmentation classification.

Since the degree of infection severity varies among the test patients, which influences the output segmentation, the amount of rough user edit strokes needed to edit a patient's initial segmentation range from 1 to 10 per edited slice for the proposed model. In most cases, the increase in Dice is larger for CONS after refinement with the model. Qualitative results of the initial segmentation showed that CONS is frequently under segmented due to its high similarity to the background class (blood vessels and walls of airways have a very similar Hounsfield unit histogram) making it difficult for the model to segment the region correctly without user guidance.

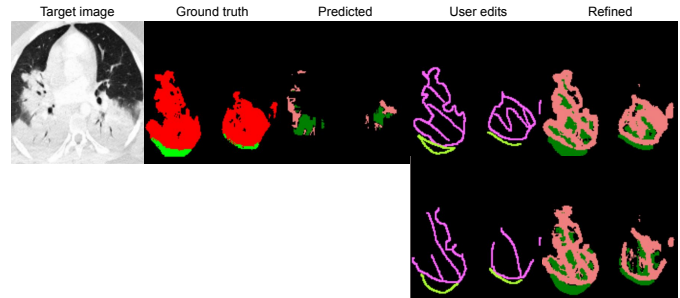


Fig. 7: Example of how the refined segmentation differs when different types of user edits (top and lower rows) are drawn to refine the segmentation. (Red: GGO ground truth, Green: CONS ground truth, Pink: predicted GGO, Dark green: predicted CONS, Magneta: foreground edit for GGO, Neon green: foreground edit for CONS)

2) *Qualitative Results*: Qualitative results from the proposed model are presented in this section.

Comparison of Edits

Figure 7 shows the different segmentation refinement output given different types of user edits. As can be observed in the lower segmentation output, it is sufficient to draw the outer outline of the infected regions to separate them from the background. The model is able to further segment the unedited areas, but it is often incorrectly classified. However, as shown in Figure 8, the wrong segmentation can be corrected in the following refinement round.

A possible explanation for the wrong segmentation of GGO as CONS is due to their feature similarity in the CT scans. One of the main characteristics that differentiate GGO from CONS is the location of the segmentation. CONS is often located at the bottom of the axial lung CT. But in severe cases, it can be found higher up and its borders are often GGO as shown in the second row of Figure 8.




Fig. 8: Examples of two rounds of segmentation refinement, resulting in increasingly better results.




Fig. 9: Example of how edits on one slice are automatically propagated to other slices. After the initial prediction, edits are drawn on one slice. Red arrows point to false edit locations. Interestingly, the model learns that false edits should not be propagated further than one slice by the model.

Robustness Testing

A robustness test is carried out to evaluate how well the 2.5D model that uses stacked slices from three anatomical views as input is able to propagate edits on one slice to another.

Figure 9 displays the limited automatic propagation of edits from one slice to other slices. The further the slices are from the edited slice, the more segmentation are wrongly classified even though the correct regions are segmented and some regions are undersegmented. However, false edits are not propagated further than one slice away by the model. This shows the model's ability to detect edits that are incorrect by considering the features of the CT scans.

Influence of Reference Data on Target Segmentation

In some cases where a patient's reference scan is not correctly deformed to align with the follow-up scan, the initial segmentation of the follow-up scan for more difficult regions can

be negatively affected by the reference CT. Figure 10 shows some abnormalities in the initially predicted segmentation of the target image. The infected lower part of the target CT that is harder to discern from the background class is segmented according to the reference image. As observed in the target's initial predicted segmentation, its left outline looks similar to the reference image.

In another example, when the reference image is correctly aligned with the target image, as shown in Figure 11, the reference image can serve as a guide for the model to segment regions that are difficult for the model as previously mentioned, such as areas that look similar to the background class. In this example, the model produces a more accurate initial segmentation of the target scan.




Fig. 10: Example of how the target’s initial prediction is influenced by the reference image in a case where there is a misalignment between the reference and target image. The red arrows point to regions where the CT scans differ the most. The white arrows point to the interested regions.




Fig. 11: Example of a case where the reference image and target image are correctly aligned. The model segmented the difficult areas better compared to Figure 10. The red arrows point to regions where the CT scans differ the most. The white arrows point to the interested regions.

E. Limitations and Future Work

One of the limitations that we faced is the availability of a larger longitudinal COVID-19 dataset. Also, label noise is probably present in our data, since for severe cases, even expert radiologists struggle to separate in particular CONS from the background class and PLEFF. In future work, the potential of this method on improving segmentation results of other longitudinal medical images and in different clinical contexts can be further explored. The problem of wrongly classified segmentation in multiclass segmentation can be additionally examined through using an ensemble of binary interactive

segmentation models for each foreground class. Furthermore, it would be interesting to test the proposed method with other state of the art model architectures. Besides, this paper also reveals limitation of a 2.5D model in propagating edits from one slice to other slices further away, a 3D implementation of the interactive segmentation method which is not in the scope of this paper would perhaps be able to counter this issue at the cost of requiring bigger training datasets.

V. CONCLUSION

In this paper, an interactive segmentation method with 2.5D longitudinal network is proposed. Through concatenating the previous time-point reference segmentation mask, segmentation output of target image from past segmentation rounds and user interactions mask to target and reference images, all past information is fully utilized as input data to the model to improve segmentation results. Experiments on our in-house longitudinal COVID-19 dataset show that a large improvement in the Dice of both classes is obtained after one round of interactive segmentation refinement. Besides that, the proposed longitudinal interactive segmentation refinement model’s segmentation performance is superior compared to the static version of the interactive model. This concludes that, with the availability of longitudinal data, existing segmentation models can be easily adapted through our method and trained end-to-end for interactive segmentation refinement.

REFERENCES

- [1] Tao Ai, Zhenlu Yang, Hongyan Hou, Chenao Zhan, Chong Chen, Wenzhi Lv, Qian Tao, Ziyong Sun, and Liming Xia. Correlation of chest CT and RT-PCR testing for coronavirus disease 2019 (COVID-19) in china: A report of 1014 cases. *Radiology*, 296(2):E32–E40, Aug. 2020.
- [2] Ruba Alkadi, Ayman El-Baz, Dr. Fatma Taher, and Naoufel Werghi. *A 2.5D Deep Learning-Based Approach for Prostate Cancer Detection on T2-Weighted Magnetic Resonance Imaging: Munich, Germany, September 8-14, 2018, Proceedings, Part IV*, pages 734–739. Springer International Publishing, 01 2019.
- [3] Shahab Aslani, Michael Dayan, Loredana Storelli, Massimo Filippi, Vittorio Murino, Maria A. Rocca, and Diego Sona. Multi-branch convolutional neural network for multiple sclerosis lesion segmentation. *NeuroImage*, 196:1–15, Aug. 2019.
- [4] Ariel Birenbaum and Hayit Greenspan. Multi-view longitudinal cnn for multiple sclerosis lesion segmentation. *Eng. Appl. Artif. Intell.*, 65(C):111–118, Oct. 2017.
- [5] Antonio Criminisi, Toby Sharp, and Andrew Blake. GeoS: Geodesic image segmentation. In *Lecture Notes in Computer Science*, pages 99–112. Springer Berlin Heidelberg, 2008.
- [6] Stefan Denner, Ashkan Khakzar, Moiz Sajid, Mahdi Saleh, Ziga Spiclin, Seong Tae Kim, and Nassir Navab. Spatio-temporal learning from longitudinal data for multiple sclerosis lesion segmentation. *arXiv preprint arXiv:2004.03675*, 2020.
- [7] Deng-Ping Fan, Tao Zhou, Ge-Peng Ji, Yi Zhou, Geng Chen, Huazhu Fu, Jianbing Shen, and Ling Shao. Inf-net: Automatic COVID-19 lung infection segmentation from CT images. *IEEE Transactions on Medical Imaging*, 39(8):2626–2637, Aug. 2020.
- [8] Abhijit Guha Roy, Sailesh Conjeti, Nassir Navab, and Christian Wachinger. Quicknat: A fully convolutional network for quick and accurate segmentation of neuroanatomy. *NeuroImage*, 186, 11 2018.
- [9] Mohamed Mohamed Hefeda. CT chest findings in patients infected with COVID-19: review of literature. *Egyptian Journal of Radiology and Nuclear Medicine*, 51(1), Nov. 2020.
- [10] Simon Jégou, Michal Drozdal, David Vazquez, Adriana Romero, and Yoshua Bengio. The one hundred layers tiramisu: Fully convolutional densenets for semantic segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition Workshops*, pages 11–19, 2017.

- [11] Seong Tae Kim, Leili Goli, Magdalini Paschali, Ashkan Khakzar, Matthias Keicher, Tobias Czempel, Egon Burian, Rickmer Braren, Nassir Navab, and Thomas Wendler. Longitudinal quantitative assessment of covid-19 infection progression from chest cts. *International Conference on Medical Image Computing and Computer Assisted Intervention*, 2021.
- [12] Titinunt Kitrungrotsakul, Qingqing Chen, Huitao Wu, Yutaro Iwamoto, Honjie Hu, Wenchao Zhu, Chao Chen, Fangyi Xu, Yong Zhou, Lanfen Lin, et al. Attention-refnet: Interactive attention refinement network for infected area segmentation of covid-19. *IEEE Journal of Biomedical and Health Informatics*, 2021.
- [13] Yan Li and Liming Xia. Coronavirus disease 2019 (COVID-19): Role of chest CT in diagnosis and management. *American Journal of Roentgenology*, 214(6):1280–1286, June 2020.
- [14] Bradley Lowekamp, David Chen, Luis Ibanez, and Daniel Blezek. The design of simpleitk. *Frontiers in neuroinformatics*, 7:45, 12 2013.
- [15] Jeffrey B Mendel, Justine T Lee, and David Rosman. Current concepts imaging in covid-19 and the challenges for low and middle income countries. *Journal of Global Radiology*, 6(1):3, 2020.
- [16] Fausto Milletari, Nassir Navab, and Seyed-Ahmad Ahmadi. V-Net: Fully Convolutional Neural Networks for Volumetric Medical Image Segmentation. *arXiv e-prints*, page arXiv:1606.04797, June 2016.
- [17] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 8024–8035. Curran Associates, Inc., 2019.
- [18] Hiba Ramadan, Chaymae Lachqar, and Hamid Tairi. A survey of recent interactive image segmentation methods. *Computational Visual Media*, 6(4):355–384, Aug. 2020.
- [19] Sashank Reddi, Satyen Kale, and Sanjiv Kumar. On the convergence of adam and beyond. In *International Conference on Learning Representations*, 2018.
- [20] Fei Shan, Yaozong Gao, Jun Wang, Weiya Shi, Nannan Shi, Miaofei Han, Zhong Xue, Dinggang Shen, and Yuxin Shi. Abnormal lung quantification in chest CT images of COVID-19 patients with deep learning and its application to severity prediction. *Medical Physics*, 48(4):1633–1645, Mar. 2021.
- [21] Heshui Shi, Xiaoyu Han, Nanchuan Jiang, Yukun Cao, Osamah Alwalid, Jin Gu, Yanqing Fan, and Chuansheng Zheng. Radiological findings from 81 patients with covid-19 pneumonia in wuhan, china: a descriptive study. *The Lancet infectious diseases*, 20(4):425–434, 2020.
- [22] Guotai Wang, Wenqi Li, Maria A. Zuluaga, Rosalind Pratt, Premal A. Patel, Michael Aertsen, Tom Doel, Anna L. David, Jan Deprest, Sebastien Ourselin, and Tom Vercauteren. Interactive medical image segmentation using deep learning with image-specific fine tuning. *IEEE Transactions on Medical Imaging*, 37(7):1562–1573, 2018.
- [23] Guotai Wang, Xinglong Liu, Chaoping Li, Zhiyong Xu, Jiugen Ruan, Haifeng Zhu, Tao Meng, Kang Li, Ning Huang, and Shaoting Zhang. A noise-robust framework for automatic segmentation of COVID-19 pneumonia lesions from CT images. *IEEE Transactions on Medical Imaging*, 39(8):2653–2663, Aug. 2020.
- [24] Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE Transactions on Image Processing*, 13(4):600–612, 2004.
- [25] Ning Xu, Brian Price, Scott Cohen, Jimei Yang, and Thomas Huang. Deep interactive object selection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, June 2016.
- [26] Huahong Zhang, Alessandra M Valcarcel, Rohit Bakshi, Renxin Chu, Francesca Bagnato, Russell T Shinohara, Kilian Hett, and Ipek Oguz. Multiple sclerosis lesion segmentation with tiramisu and 2.5 d stacked slices. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 338–346. Springer, 2019.
- [27] Bowei Zhou, Li Chen, and Zhao Wang. Interactive deep editing framework for medical image segmentation. *International Conference on Medical Image Computing and Computer Assisted Intervention*, pages 329–337, 2019.
- [28] Shuchang Zhou, Yujin Wang, Tingting Zhu, and Liming Xia. Ct features of coronavirus disease 2019 (covid-19) pneumonia in 62 patients in wuhan, china. *American Journal of Roentgenology*, 214(6):1287–1294, 2020.
- [29] Shuchang Zhou, Yujin Wang, Tingting Zhu, and Liming Xia. CT features of coronavirus disease 2019 (COVID-19) pneumonia in 62 patients in wuhan, china. *American Journal of Roentgenology*, 214(6):1287–1294, June 2020.