# Anam-Net : Anamorphic Depth Embedding based Light-Weight CNN for Segmentation of Anomalies in COVID-19 Chest CT Images

Naveen Paluru, Aveen Dayal, Hvard Bjrke Jenssen, Tomas Sakinis, Linga Reddy Cenkeramaddi, *Senior Member, IEEE*, Jaya Prakash, and Phaneendra K. Yalavarthy, *Senior Member, IEEE*

***Abstract*—Chest computed tomography (CT) imaging has become indispensable for staging and managing of COVID-19, and current evaluation of anomalies/abnormalities associated with COVID-19 has been performed majorly by visual score. The development of automated methods for quantifying COVID-19 abnormalities in these CT images is invaluable to clinicians. The hallmark of COVID-19 in chest CT images is the presence of ground-glass opacities in the lung region, which are tedious to segment manually. We propose anamorphic depth embedding based light-weight CNN, called Anam-Net, to segment anomalies in COVID-19 Chest CT images. The proposed Anam-Net has 7.8 times fewer parameters compared to the state-of-the-art UNet (or its variants), making it light-weight capable of providing inferences in mobile or resource constraint (point-of-care) platforms. The results from chest CT images (test cases) across different experiments showed that the proposed method could provide good Dice similarity scores for abnormal as well as normal regions in the lung. We have benchmarked Anam-Net with other state-of-the-art architectures like ENet, LEDNet, UNet++, SegNet, Attention UNet and DeepLabV3+. The proposed Anam-Net was also deployed on embedded systems like Raspberry Pi 4, NVIDIA Jetson Xavier, and mobile based Android application (CovSeg) embedded with Anam-Net to demonstrate its suitability for point-of-care platforms. The generated codes, models, and the mobile application are available for enthusiastic users at https://github.com/NaveenPaluru/Segmentation-COVID-19.**

***Index Terms*—COVID-19, Coronavirus, Deep Learning, Segmentation, and Abnormalities.**

## I. Introduction

CORONAVIRUS disease 2019 (COVID-19) is an illness caused by a novel coronavirus, formerly known as 2019-nCoV. It is a constituent of a spectrum of viruses that cause respiratory diseases, like severe acute respiratory syndrome (SARS) and middle east respiratory syndrome (MERS) [1]. COVID-19 virus was first identified in Wuhan City, Hubei Province, China. The virus causes respiratory disorders with common symptoms being fever, dry cough, and shortness of breath [2]. Currently, Reverse Transcription Polymerase Chain Reaction (RT-PCR) is the standard test for diagnosing COVID-19 [3]. However, imaging-based diagnosis (Ultrasound, Chest X-ray, and Chest Computed Tomography (CT)) are also playing a crucial role in the identification and management of COVID-19 infection. Notably, a shortage of RT-PCR tests for the diagnosis of COVID-19 has led to considering Chest CT as a screening and diagnostic tool [4]. The RT-PCR tests have high false negatives, which requires serial sampling and in such cases, findings in CT were beneficial [5]. The CT based conclusions on 76 asymptomatic individuals with confirmed COVID-19 from the "Diamond Princess" cruise ship was able to identify pneumonia in 54% of these cases [6]. Also, CT is more sensitive to parenchymal lung disease, disease progression, and alternative diagnoses, including acute heart failure from COVID-19 myocardial injury [7]. Even a multinational consensus statement that was recently released confirmed that Chest CT would provide clinically actionable items in diagnosis, management, triage, and therapy of COVID-19 [3].

The common imaging features observed from a chest CT of COVID-19 patient include anomalies/abnormalities, such as Ground Glass Opacity (GGO), consolidation, and rare characteristics like pericardial effusion and pleural effusion [8]–[11], with GGO being a common feature among all chest CT images. In a retrospective study, Chung et al. [9] emphasized the identification of common features as mentioned above for better management of COVID-19. There are ongoing attempts to utilize learning-based methods to triage the patients based on chest CT. Li et al. [12] have proposed a three-dimensional (3D) deep learning model for using two-dimensional (2D) local and 3D global features to identify COVID-19 disease. Hierarchical attention and spatial pyramid networks were introduced to capture the abnormal features in the Lungs of COVID-19 patients [13]. Pan et al. [14] and Wang et al. [15] conducted a systematic study on chest CT images for understanding the changes in the lung during recovery from COVID-19. Caruso et al. [10] investigated the chest CT features of patients with COVID-19 and compared the diagnostic performance of chest CT with the gold standard RT-PCR. They reported a sensitivity of 97% while using chest

Submitted on November 15, 2020; Accepted on January 21, 2021. This work in part was supported by the WIPRO GE-CDS Collaborative Laboratory on Artificial Intelligence in Healthcare and Medical Imaging as well as Indo-Norwegian Collaboration in Autonomous Cyber-Physical Systems (INCAPS), INTPART Programme, Research Council of Norway.

N. Paluru and P. K. Yalavarthy are with the Department of Computational and Data Sciences, Indian Institute of Science, Bangalore 560 012, India (e-mail: naveenp@iisc.ac.in and yalavarthy@iisc.ac.in).

A. Dayal and L. R. Cenkeramaddi are with the Department of Information and Communication Technology, University of Agder, Grimstad, Norway (e-mail: linga.cenkeramaddi@uia.no).

H. B. Jenssen and T. Sakinis are with the Department of Radiology and Nuclear Medicine, Oslo University Hospital, Norway. They are also with the Artificial Intelligence AS, Oslo, Norway (e-mail: hbjenssen@gmail.com and sakinis.tomas@gmail.com).

J. Prakash is with the Department of Instrumentation and Applied Physics, Indian Institute of Science, Bangalore 560 012, India (e-mail: jayap@iisc.ac.in).

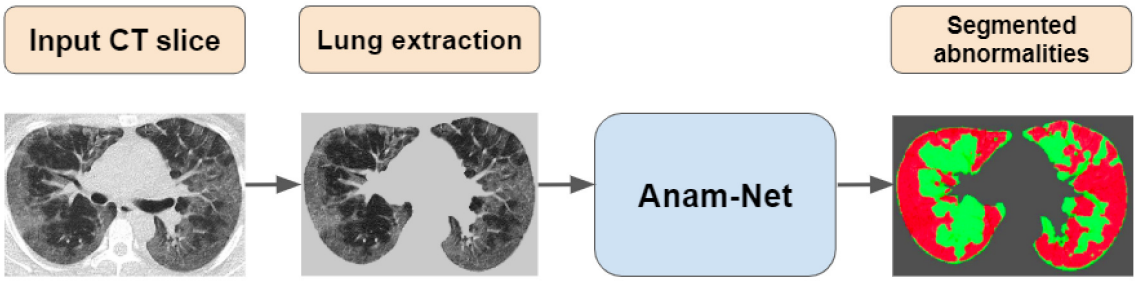


Fig. 1. Key steps of the proposed approach for automated segmentation of abnormalities in chest CT images. The details of the Anam-Net architecture are given in Fig. 2 and layer wise details are shown in Table I.

CT. The work of Pan et al. [14] defined four stages of disease progression of COVID-19 and its recovery. The remission and recovery from COVID-19 had a strong correlation with reducing the size of abnormalities in chest CT [14]. Even attempts to use lung opacity (abnormality) as a deep learning feature [16] to divide COVID-19 patients into categories of mild, moderate, severe, and critical have significantly been successful. It has also been suggested to use lung opacification as an image feature to monitor COVID-19 progression to provide better clinical management [16].

Given that chest CT imaging has an essential role in diagnosing, monitoring, and managing COVID-19, quantifying anomalies in chest CT in an automated way is the need of the hour. Current practice is semi-quantitative and performed by visual score. Manual segmentation of these abnormalities in chest CT images is a tedious Job. Automated methods, such as the one proposed here, can run on embedded as well as mobile (point-of-care) platforms with minimal hardware like Raspberry Pi, NVIDIA Jetson Xavier, or an Android device, will immensely help the clinicians to manage COVID-19 better, specifically when the patient load is very high. Needless to mention that performing annotations slice-by-slice is tedious and expensive. Further, chest CT imaging is not a standard clinical protocol for COVID-19 patients, these datasets availability in sufficiently large numbers, including annotations to develop a deep learning based methodology, will be a challenge. Given the variability in the chest CT protocols [17] and diversity in the patient population, a deep learning model that is easy to train and deployable for quantifying the abnormalities in chest CT images is of utmost importance.

Automatic Segmentation of ground-glass opacities (GGO) in lung CT images was investigated for diffuse parenchymal lung diseases (DPLD). Zhu et al. [18] have proposed an algorithm based on a Markov random field for segmenting GGO. In Ref. [19], the authors have shown that the K-NN classifier, when boosted with non-parametric density estimates, can segment the GGO nodules more accurately. Jung et al. [20] proposed an intensity-based segmentation followed by an asymmetric multi-phase deforming model for the segmentation of GGOs. All these methods were based on statistical models to segment the GGOs; however, lung infection in COVID-19 is a manifestation of GGOs and consolidation. The abnormalities present in the COVID-19 other than GGOs like consolidation can be as high as $45\%$ of the total abnormality [15]. These statistical methods may not be applicable for segmenting the COVID-19 chest CT images as the imaging features differ significantly. Recently, Fan et al. [21] proposed a pseudo label generating strategy within a deep semi-supervised network known as Inf-Net for segmenting abnormalities in COVID-19. This work only focuses on lung infection segmentation of COVID-19 patients, leading to a drop in accuracy when considering non-infected slices. Even Ouyang et al. [22] proposed a novel dual sampling attention strategy for effective mitigation of the imbalanced learning in chest CT images with three-dimensional convolutional neural networks for automatic diagnosis of COVID-19 from community-acquired pneumonia.

This work proposes the utility of supervised deep learning based fast and fully automated way of segmenting anomalies (primarily, GGO and consolidations) and also normal lung tissue in chest CT images of patients having COVID-19. The emphasis is also on this model being light-weight, such that it can be deployed in point-of-care platforms to have better clinical utility. We call our approach, anamorphic depth embedding based light-weight convolutional neural network, shortened as Anam-Net. The anamorphic depth of feature embeddings was obtained by anamorphic block (AD Block), as shown in Fig. 2. We also bring in a label based weighting strategy for the networks cost function for effective learning. In supervised learning, cost-sensitive networks penalize the classifiers loss by a weighting factor found from the prior information of annotated data. Several works adapt cost-sensitive training for robust segmentation [23]–[25]. In short, the main contributions of this work are **(a)** Development of novel light-weight deep learning-based robust feature learning algorithm designed for the COVID-19 anomaly segmentation in chest CT images, with the fully convolutional anamorphic depth block (AD-block) built within symmetric encoder-decoder architecture. This AD-block enabled efficient gradient flow in the network. **(b)** the adapted label weighting scheme during training makes the model highly robust during the testing phase; **(c)** the computational time required for training and testing the proposed Anam-Net to segment abnormalities is low, making it highly attractive in the clinic; **(d)** The proposed network has very few parameters thereby reducing the need for extensive annotated data for training the network, making it easy to train using site-specific data or for a specific chest CT protocol. We deployed the proposed Anam-Net on Raspberry pi 4 and NVIDIA Jetson

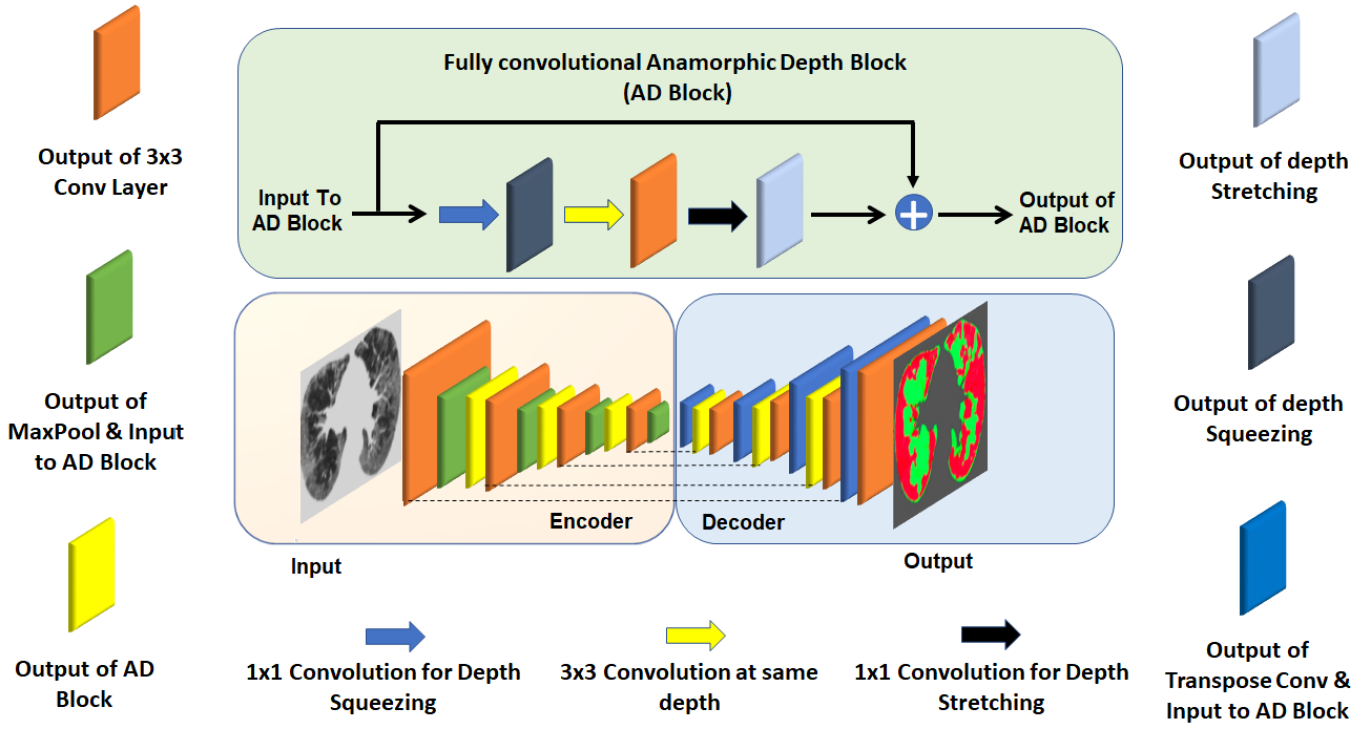


Fig. 2. Anam-Net : Network architecture utilized for segmenting abnormalities in COVID-19 chest CT images. Fully convolutional anamorphic depth blocks (AD-Block) with depth wise squeezing and stretching have been incorporated after downsampling operation in the encoder and also after upsampling operation in the decoder. Each AD Block has $1\times1$ convolution for depth wise squeezing followed by $3\times3$ convolution and finally $1\times1$ convolution for depth wise stretching. The layer-wise details of encoder and decoder are shown in Table-I and that of AD Block in Table-II. Note that these AD Blocks are independent of each other and do not share parameters.

Xavier Agx modules enabling deep learning-based embedded systems that can provide a quick initial assessment of COVID-19 lung infection. We have also developed an Android application (CovSeg) that can run on mobile devices embedded with Anam-Net for segmenting COVID-19 Anomalies. **(e)** Lastly, Anam-Net was evaluated on three datasets under different experimental conditions and benchmarked against state-of-the-art heavy and light-weight deep learning models, including UNet [25], ENet [24], UNet++ [26], SegNet [27], Attention UNet [28], LEDNet [29] as well as DeepLabV3+ [30]. Also, to the best of our knowledge, the proposed approach Anam-Net, which involves an anamorphic transformation in the depth of embeddings, has been introduced for image segmentation for the first time, resulting in a highly versatile light-weight network to perform the segmentation.

The imaging data, along with ground truth annotations, were made available as an open-source [31]. The source code, trained models, and the mobile application for the proposed segmentation scheme along with consolidated results, are available as open source at Ref. [32]. The rest of the manuscript is arranged as follows. The next section (Sec. II) provides the details of the proposed method. This section is followed by a section (Sec. III) describing the dataset utilized along with implementation details and figures of merit used for quantitatively assessing the performance of the proposed approach. The hardware deployment details are given in Sec. IV, and we provide the results in Sec. V. We present a detailed discussion of the results and the limitations of this study in Sec. VI. Finally, the conclusions were provided in Sec. VII.

## II. Methods

The key steps involved in the proposed approach for segmenting the abnormalities in COVID-19 chest CT images along with proposed Anam-Net architecture are given in Fig. 1 and Fig. 2, respectively. We discuss each of these steps in detail in the following subsections.

### A. Lung Extraction

The first step for segmenting abnormalities in chest CT images is to extract the lung region, the lung masks given in the datasets I and II (refer to Table-III) were extracted using the method described in Ref. [33], and these masks were posted for easy usage at link [31]. These lung masks were obtained using a pre-trained U-Net architecture [25] trained with batch normalization [33] consisting of 231 training samples obtained from a database of 5,300 samples. Further, the U-Net (R231) method was found to be more accurate in terms of lung segmentation compared to other trained models like Chest Imaging Platform (CIP) and Progressive Holistically Nested Networks (P-HNN) [33]. These 231 samples were obtained using random sampling, sampling from image phenotype, and manual selection of cases with different pathology's like fibrosis, trauma, and other pathology's [33]. Since the training samples used in U-Net (R231) were having a wide variety

of lung pathology and organization, U-Net (R231) method enabled accurate lung extraction on chest CT scan obtained from COVID-19 patients. Note that any process (including simple windowing) to extract the lung region is sufficient for the proposed work.

### *B. Deep Learning for Medical Image Segmentation*

Recently, techniques like neural architectural search (NAS), knowledge distillation, cross-modality adaptation, etc. were utilized for robust segmentation. Yu et al. [35] proposed a coarse to fine NAS strategy for three-dimensional biomedical segmentation by performing architectural search at the macro-level and systematic micro-level operations at each macro level topology. In a similar study, Guo et al. [36] implemented the crucial step of segmenting the organs at risk (head and neck) during radiotherapy treatment planning by adapting the NAS strategy. For the problem at hand related to COVID-19, the NAS approach may not be suited due to high computational complexity resulting in large training times. Li et al. [37] proposed a semi-supervised system to address the lack of labeled data. This approach formulates the loss function as a weighted combination of supervised component for the labeled data and a regularized detail for the unlabelled data. In contrast to this, Clough et al. [38] proposed an unsupervised strategy based on topological loss derived from prior information of the object to be segmented. The knowledge transfer from heavy models for developing the light-weight models through model pruning was explored by Zhou et al. [39]. However, this kind of knowledge transfer requires a pre-trained model. Wang et al. [40] proposed a user in the loop strategy for 2D segmentation of the placenta from fetal MRI and 3D segmentation of brain tumors from FLAIR images. In this work, the authors deployed user interactions as a hard constraint into a back propagatable conditional random field for end-to-end training. In Ref. [41], the authors proposed context encoding networks consisting of atrous convolution operations for two-dimensional medical image segmentation. Chen et al. [42] proposed a cross-modality adaption strategy between MR and CT images through a deep synergistic feature alignment module for robust medical image segmentation.

UNet architecture [25] has a deep symmetric encoder-decoder network with skip connections. Similarly, SegNet [27] also has an encoder-decoder structure embedded with a nonlinear upsampling mechanism for performing semantic segmentation. Zhou et al. [26] proposed an advanced version of UNet known as UNet++ which deploys nested dense connections [43] within the symmetric encoder-decoder framework. For enhancing the salience feature extraction, Oktay et al. [28] incorporated an attention module within the UNet framework. Chen et al. [30] introduced DeepLabV3+ formed with atrous separable convolutions for refining the segmentation results across the object boundaries. Jha et al. [44] proposed a deep residual UNet known as ResUNet++ for colonoscopic image segmentation. This architecture adds on a residual connection and attention module to the existing bottleneck of UNet. A cascade of UNet known as DoubleU-Net [45] was proposed, where the second UNet refines the segmentation results from the first UNet. However, these architectures, i.e., UNet, SegNet, UNet++, Attention UNet, DeepLabV3+, ResUNet++, and DoubleU-Net, involve many parameters resulting in heavy models (memory intensive) and have specialized compute requirements. On the other hand, an efficient neural network (ENet [24]) was developed to work with less number of parameters and was designed as a robust asymmetric encoder-decoder model for semantic segmentation in computer vision. The engineered design of ENet includes bottleneck layers similar to deep residual learning [46]. Further, it has dilated convolutions [47] for avoiding higher downsampling rates and asymmetric convolutions [48] for achieving large speedups. Note that ENet architecture consists of a large encoder and a small decoder [24]. A deep encoder was deployed to operate on images with smaller resolution and enable filtering operations, while a shallow decoder was utilized to upsample the output [24]. In similar lines, another light-weight architecture LEDNet [29] also brings in an asymmetric encoder-decoder module having ResNet [46] as backbone along with a channel-wise split and shuffle operation for fast and accurate segmentation.

### *C. Proposed approach for Segmenting COVID-19 Anomalies*

The proposed Anam-Net brings the best of UNet (symmetric encoder-decoder architecture) and ENet (fewer number of parameters), with a shallow symmetric network along with bottleneck layers. Primarily, the presented work proposed an anamorphic depth block (AD-Block, motivated from [46]) within a mini symmetric encoder-decoder segmentation module, as shown in Fig. 2. The AD block consists of $1\times1$ convolution for depth-wise squeezing followed by $3\times3$ convolution and finally $1\times1$ convolution for depth-wise stretching (Refer to Table II). The AD blocks key idea is to squeeze (equivalent to project) the feature space dimension (depth-wise) before performing expensive $3\times3$ convolutions. Such a $1\times1$ projection-based low dimensional embeddings posses information about a relatively large input patch [49]. Following this low dimensional projection, local feature extraction by $3\times3$ convolutions was performed without reducing the depth-wise feature space dimension. Finally, the depth-wise feature space dimension was stretched to the initial stage by another $1\times1$ projection. The final output of AD block denoted by $h(x)$ is obtained by adding the feature maps $x$ at the input of AD block to the output of sequence of convolution operations $f(x)$ parameterized by $\theta_h$ performing depth squeezing, feature extraction and depth stretching. In short,

$$h(x) = f(x;\theta_h) + x \tag{1}$$

and it is easier to optimize $f(x;\theta_h)$ than to learn the underlying mapping $h(x)$ directly from $x$ [46]. To summarize, for a given spatial resolution, the encoding operation in the proposed Anam-Net is,

$$y_e = \phi_e\big(h_e(x_e);\theta_e\big) \tag{2}$$

where, $y_e$ is the output of encoding operation, $\phi_e$ is the convolution operation parameterized by $\theta_e$, $x_e$ is the input

TABLE I
ARCHITECTURE DETAILS OF PROPOSED ANAM-NET, WHERE N IS THE MINI-BATCH SIZE AND THE FLOW OF THE ARCHITECTURE IS FROM LEFT TO RIGHT IN EACH ROW (TOP TO BOTTOM FOR SUCCESSIVE STEP). THE SCHEMATIC DIAGRAM WAS PROVIDED IN FIG. 2. THE ARCHITECTURE DETAILS OF ANAMORPHIC DEPTH (AD) BLOCK ARE GIVEN IN TABLE-II.

| | | | | | | |
|---|---|---|---|---|---|---|
| **Block Size** | Input ( N, 1, 512, 512 ) | Convolution ( N, 64, 512, 512 ) | MaxPool ( N, 64, 256, 256 ) | AD Block ( N, 64, 256, 256 ) | Convolution ( N, 128, 256, 256 ) | MaxPool ( N, 128, 128, 128 ) |
| **Block Size** | AD Block ( N, 128, 128, 128 ) | Convolution ( N, 256, 128, 128 ) | MaxPool ( N, 256, 64, 64 ) | AD Block ( N, 256, 64, 64 ) | Convolution ( N, 256, 64, 64 ) | MaxPool ( N, 256, 32, 32 ) |
| **Block Size** | Trans Conv ( N, 256, 64, 64 ) | AD Block ( N, 256, 64, 64 ) | Concat ( N, 512, 64, 64 ) | Convolution ( N, 256, 64, 64 ) | Trans Conv ( N, 256, 128, 128 ) | AD Block ( N, 256, 128, 128 ) |
| **Block size** | Concat ( N, 512, 128, 128 ) | Convolution ( N, 256, 128, 128 ) | Trans Conv ( N, 128, 256, 256 ) | AD Block ( N, 128, 256, 256 ) | Concat ( N, 256, 256, 256 ) | Convolution ( N, 128, 256, 256 ) |
| **Block Size** | Trans Conv ( N, 64, 512, 512 ) | Concat ( N, 128, 512, 512 ) | Convolution ( N, 64, 512, 512 ) | Output ( N, 3, 512, 512 ) | | |

TABLE II
ARCHITECTURE DETAILS OF THE ANAMORPHIC DEPTH (AD) BLOCK WITH N BEING THE MINI-BATCH SIZE, Z IS THE DEPTH OF EMBEDDINGS (FEATURE MAPS) AND M IS THE SPATIAL EXTENT OF EMBEDDINGS.

| **Block** | Input | 1x1 Convolution | 3x3 Convolution | 1x1 Convolution | Output |
|---|---|---|---|---|---|
| **Size** | ( N, Z, M, M) | ( N, Z/4, M, M) | ( N, Z/4, M, M) | ( N, Z, M, M) | ( N, Z, M, M) |

TABLE III
DETAILS OF THE DATASETS UTILIZED IN THIS WORK

| Dataset No. | Reference | Patients | Total Slices |
|---|---|---|---|
| **I** | [31] | $> 40$ | 100 |
| **II** | [31] | 9 | 829 |
| **III** | [34] | 20 | 3410 |

TABLE IV
COMPARISON OF DEEP LEARNING MODELS UTILIZED IN THIS WORK IN TERMS OF NUMBER OF TRAINING PARAMETERS, MODEL SIZE, TRAINING TIME (FOR 100 EPOCHS IN EXPERIMENT 1) AND INFERENCE TIME.

| Model | Parameters | Model Size | Training time | Inference time |
|---|---|---|---|---|
| **UNet** | 31.07 M | 118.24 MB | 51 min | 531 ms |
| **ENet** | 343.7 K | 1.33 MB | 15 min | 248 ms |
| **UNet++** | 9.16 M | 34.95 MB | 58 min | 551 ms |
| **SegNet** | 29.44 M | 112.31 MB | 37 min | 528 ms |
| **Attention UNet** | 34.87 M | 133.05 MB | 63 min | 569 ms |
| **LEDNet** | 0.91 M | 3.8 MB | 16 min | 298 ms |
| **DeepLabV3+** | 54.70 M | 208.66 MB | 42 min | 895 ms |
| **Anam-Net (Proposed)** | 4.47 M | 17.21 MB | 27 min | 362 ms |

to be encoded and $h_e$ is the output of the AD block in the encoder. Further, $h_e$ can be written as,

$$h_e(x_e) = f_e(x_e; \theta_{h_e}) + x_e \tag{3}$$

where $f_e$ is the output of sequence of convolution operations parameterized by $\theta_{h_e}$ performing depth squeezing, feature extraction and depth stretching. In similar lines, the decoding operation at same resolution as mentioned earlier in the encoding operation of the proposed Anam-Net is,

$$y_d = \phi_d\big(\mathcal{C}[h_d(x_d), y_e]; \theta_d\big) \tag{4}$$

where $y_d$ is the output of decoding operation, $\phi_d$ is the convolution operation parameterized by $\theta_d$, $\mathcal{C}$ is the feature concatenation operation, $x_d$ is the input to be decoded, and $h_d$ is the output of the AD block in the decoder. As mentioned earlier, $h_d$ can be expressed as,

$$h_d(x_d) = f_d(x_d; \theta_{h_d}) + x_d \tag{5}$$

where $f_d$ is the output of sequence of convolution operations parameterized by $\theta_{h_d}$ performing depth squeezing, feature extraction and depth stretching. Note that, the UNet embeddings can be seen as a special case of our proposed Anam-Net embeddings when $h_e(x_e) = x_e$ and $h_d(x_d) = x_d$.

Overall, the proposed Anam-Net architecture consists of six such AD blocks (3 in the encoder and 3 in the decoder) to provide salience and robust feature learning. Each convolution layer in the proposed Anam-Net consist of convolution operation followed by batch normalization [50] and ReLU [51]. We provide the layer-wise details of the encoder-decoder module in Table I, and that of the AD Block in Table II. Given a mini-batch of size $n$, we compute the loss as,

$$\mathcal{L} = -\frac{1}{N'}\sum_{k=1}^{n}\sum_{i=1}^{rows}\sum_{j=1}^{cols}(w_{kij})\sum_{t=0}^{2} y_{kijt}\log\left(y'_{kijt}\right) \tag{6}$$

where $y$ is one hot encoded label $(3 \times 1)$, $y'$ is predicted softmax probabilities $(3 \times 1)$, $w$ is the weight associated with label $t \in \{0, 1, 2\}$ and $N'$ is $n \times rows \times cols$. The weight associated with each label in $t$ was given as,

$$w(t) = \frac{1}{p(t)} \tag{7}$$

where $p(t)$ is the fraction of samples with label $t$ in the training set. After performing end to end training, the architecture was

subjected to segment each pixel of the unseen test sample into three categories: background, abnormal-lung region, and normal-lung region. The models output is a probabilistic map (3 maps, one each for background, abnormal and normal ) having the same spatial dimension as input. Depending on the maximal probabilistic score across these three categories, each pixel was assigned a label. It is important to note that the background (non-lung region) was identified as outlined in Sec. II.A.

## III. Experiments and implementation

The summary of the datasets utilized in this study is shown in Table III. We have conducted three experiments for evaluating the performance of all discussed models. The details of the experiments performed on these datasets are as follows:

**Experiment 1:** The chest CT images that were utilized in this experiment had $929$ axial chest CT images from approximately $49$ patients with COVID-19 that were converted from openly accessible images provided by the Italian Society of Medical and Interventional Radiology [52] as well as radiopedia [53]. These images were made available in two sets. The dataset I consist of $100$ slices from $>40$ patients. From this set, we have selected $90$ slices for training. To enhance robust learning, we performed data augmentation, such as horizontal flip and vertical flip on each training sample that resulted in a training set consisting of $270$ samples. The Dataset II consists of $829$ slices from $9$ patient CT volumes. Out of these, we sampled $704$ slices for testing. The remaining $125$ slices were blank or had very minimal lung information. Note that Dataset II was exclusively utilized for testing (none of the data was utilized in training). The testing performed on all models was with test data that the network has never seen at the patient level. All these images were annotated for abnormalities and has masks for GGO, consolidations, and pleural effusion. The detailed procedure for the annotations as well as data preparation can be found at this link [54]. These annotations masks were combined together to form a single abnormal mask. The annotations along with original chest CT slices as well as lung masks as NIFTI files can be found here [31]. These annotations (original input slice and ground truth), for sample test cases, have been presented in Fig. 4 and the averaged figures of merit over all the $704$ chest CT images (test cases) considered in this experiment were given in Table-VII.

**Experiment 2:** The chest CT images that were utilized in this experiment consists of $3,410$ axial chest CT images from $20$ patients with COVID-19 (refer to Dataset III in Table-III). As mentioned earlier, we excluded the chest CT slices that did not have any visible lung region (either completely collapsed or near the end slices) from testing and training. We split this dataset at the patient level into $4$ equal folds F1, F2, F3, and F4. The fold F4 with $545$ CT images was explicitly used for testing, and we trained the deep models using a 3 fold cross-validation on the folds F1, F2, and F3. The averaged figures of merit across $3$ cross folds over all the $545$ chest CT images (test cases) considered in this experiment were shown in Table-VIII.

**Experiment 3:** In this experiment, we performed cross dataset examination, wherein, the models trained in Experiment-2 were tested on the test cases in Experiment-1. This kind of cross-examination enabled us to study the generalizability of the deep models for practical application scenarios. The averaged figures of merit over all the $704$ chest CT images (test cases) considered in this experiment were given in Table-IX.

**Implementation:** The proposed Anam-Net was trained using PyTorch [55] with a mini-batch of size $5$. The cost function was optimized using Adam optimizer [56] with an initial learning rate of $5e^{-4}$ and gradually decayed by a factor of $0.1$ once after every $33^{rd}$ epoch. All computations performed in this work including training of CNN utilized a Linux workstation having i9 9900X (CPU) with $128$ GB RAM as well as NVIDIA Quadro RTX $8000$ GPU card. For fair comparison, the state-of-the-art methods like UNet [25], ENet [24], UNet++ [26], SegNet [27], Attention UNet [28], LEDNet [29] as well as DeepLabV3+ [30] were also trained with the same training data as well as tested on the same test data in all experiments. For testing, the Anam-Net in the Fig. 1 was replaced by these trained models. The number of parameters, model size, typical training time along with the inference time were presented Table-IV for quick comparison. To quantitatively evaluate the performance of obtained segmentation results from all the models, we calculated the figures of merit, Specificity, Sensitivity, Accuracy, and the Dice Similarity Score for both abnormal and normal classes. The computed value of figures of merit will be between 0 to 1, and in all cases, the higher value (close to 1) indicates better performance of a model.

## IV. Hardware Deployment

Raspberry Pi 4 Model B is the latest version among the various raspberry pi tiny dual-display computers released to date. It is a low-cost embedded system with increased connectivity, memory capacity, and processor speed compared to its predecessor Raspberry Pi 3 Model B+. The total cost of the Raspberry Pi 4 Model B embedded platform is $50. To embed the Anam-net on Raspberry Pi 4, we converted the trained Anam-Net model from PyTorch into the Tensorflow Lite version. Tensorflow Lite is a variant of TensorFlow, which helps to run a deep learning model on mobile, internet of things (IoT), and embedded devices. It acts as an accelerator to reduce the inference time of models deployed on the embedded systems. The inference time of the TensorFlow Lite version of the Anam-net model on Raspberry Pi was 23.3 seconds, whereas the inference time of the TensorFlow Lite version of the UNet model on Raspberry Pi was 43.3 seconds. We currently have the PyTorch Lite versions for Android, and PyTorch does not have any official support for Raspberry Pi. We see this high inference time in tens of seconds is due to the

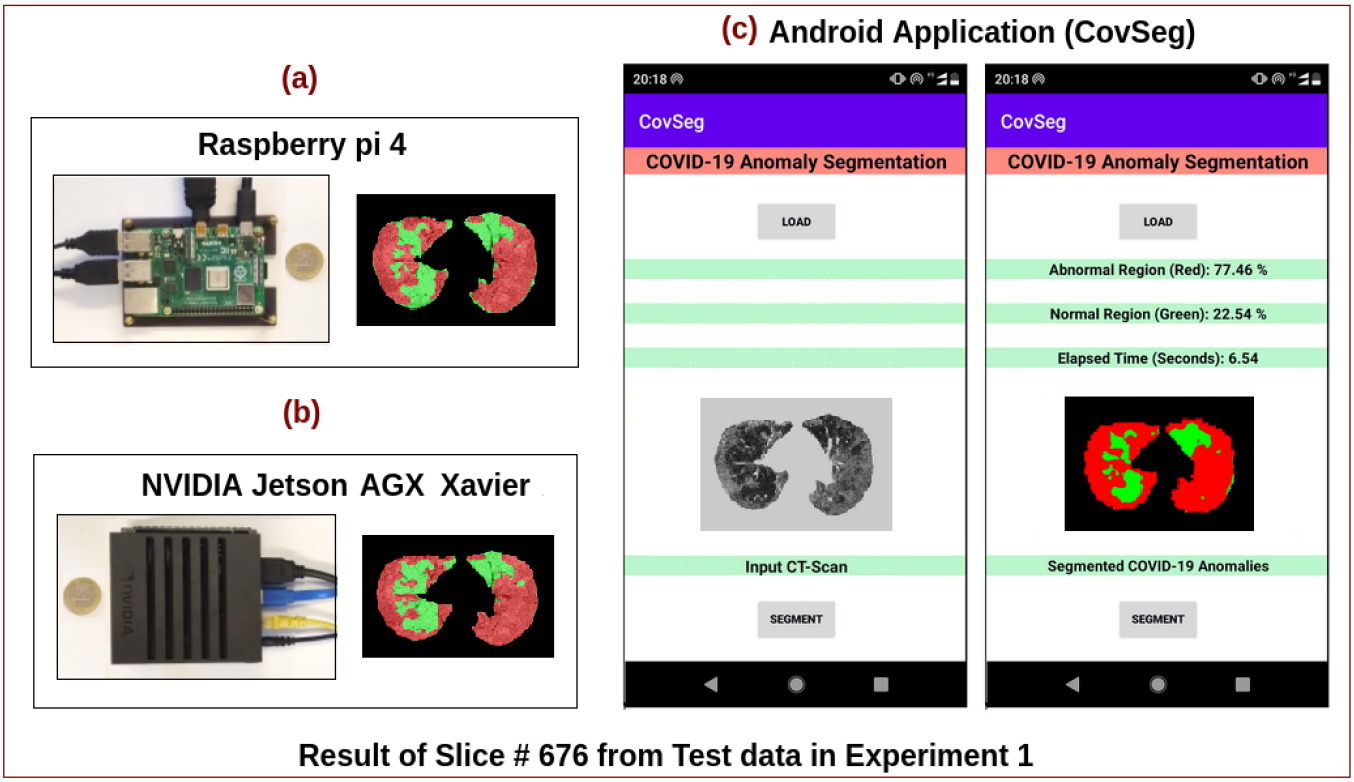


Fig. 3. The hardware devices used for deploying the proposed Anam-Net to segment the COVID-19 Anomalies in point-of-care platforms. The abnormalities in the lung region is indicated in red color and normal lung region is in green color. The Raspberry Pi 4 and the sample segmented slice #676 from the test cases (experiment 1) is shown in (a). NVIDIA Jetson Xavier along with the sample segmented slice #676 from the test cases (experiment 1) is shown in (b). One Euro coin was placed in these photographs to provide a form factor comparison. The snapshot of the developed Android application (CovSeg) showing the sample segmented slice #676 from test cases (experiment 1) was provided in (c). The inference analysis was given in Table-V.

TABLE V
COMPARISON BETWEEN THE HARDWARE DEVICES USED FOR DEPLOYING THE PROPOSED ANAM-NET IN TERMS OF COST, AVAILABLE MEMORY AND INFERENCE TIME (IN SECONDS).

| Device | Cost | Available Memory | Inference |
|---|---|---|---|
| **Raspberry Pi 4** | $50 | 4 GB RAM | Anam-Net : 23.3 s<br>UNet : 43.3 s |
| **NVIDIA Jetson** | $700 | 512-core Volta GPU<br>32 GB RAM | Anam-Net : 2.9 s<br>UNet : 5.2 s |
| **Nokia 5.1 Plus** | $95 | 3 GB RAM | Anam-Net : 6.5 s<br>UNet : 11.3 s |

model conversion from PyTorch to Tensorflow using the third-party tools. The hardware setup and the inference analysis are given Fig. 3 and Table-V respectively.

The Anam-Net model was also deployed on the NVIDIA Jetson AGX Xavier developer kit. NVIDIA Jetson AGX Xavier is the latest version among all the Jetson platforms released by NVIDIA. Jetson AGX Xavier is a deep learning model accelerator with 20 times more the performance and 10 times more energy efficiency than its predecessor Jetson TX2. Jetson AGX Xavier consists of an 8-core ARM processor CPU and 512-core Volta GPU with Tensor cores. The Anam-Net model deployed on Jetson AGX Xavier was able to perform inference within 2.9 seconds. In contrast, the UNet model, when deployed on the same, gave an inference time of 5.2 seconds. The hardware setup and the inference analysis are given Fig. 3 and Table-V respectively. The Anam-Net model and the UNet model were also trained on the Jetson AGX Xavier platform. The training time for an epoch with a batch size of 5 for the Anam-Net model was 1.49 minutes whereas for UNet it was 3.19 minutes.

We have also developed an Android Application for the mobile platforms, called it as CovSeg, for segmenting the COVID-19 Anomalies. The PyTorch trained model was converted to its lite version as given here [57]. We developed the front-end and the back-end of the CovSeg application in Android Studio [58]. The Android application snapshot on Nokia 5.1 Plus mobile phone and the inference analysis are given Fig. 3 and Table-V, respectively.

## V. RESULTS

As mentioned earlier, the segmentation results from all test cases of the three experiments were obtained and compared with the ground truth labels. The proposed method was effective in segmenting the abnormalities across all COVID-19 chest CT images. It is even useful when there is no abnormality; the proposed method provided a null result as expected. The example results, including the ground truth labels, were presented in Fig. 4 and the corresponding Dice score for the abnormal class is presented in Table-VI. The proposed Anam-Nets performance in terms of sensitivity,

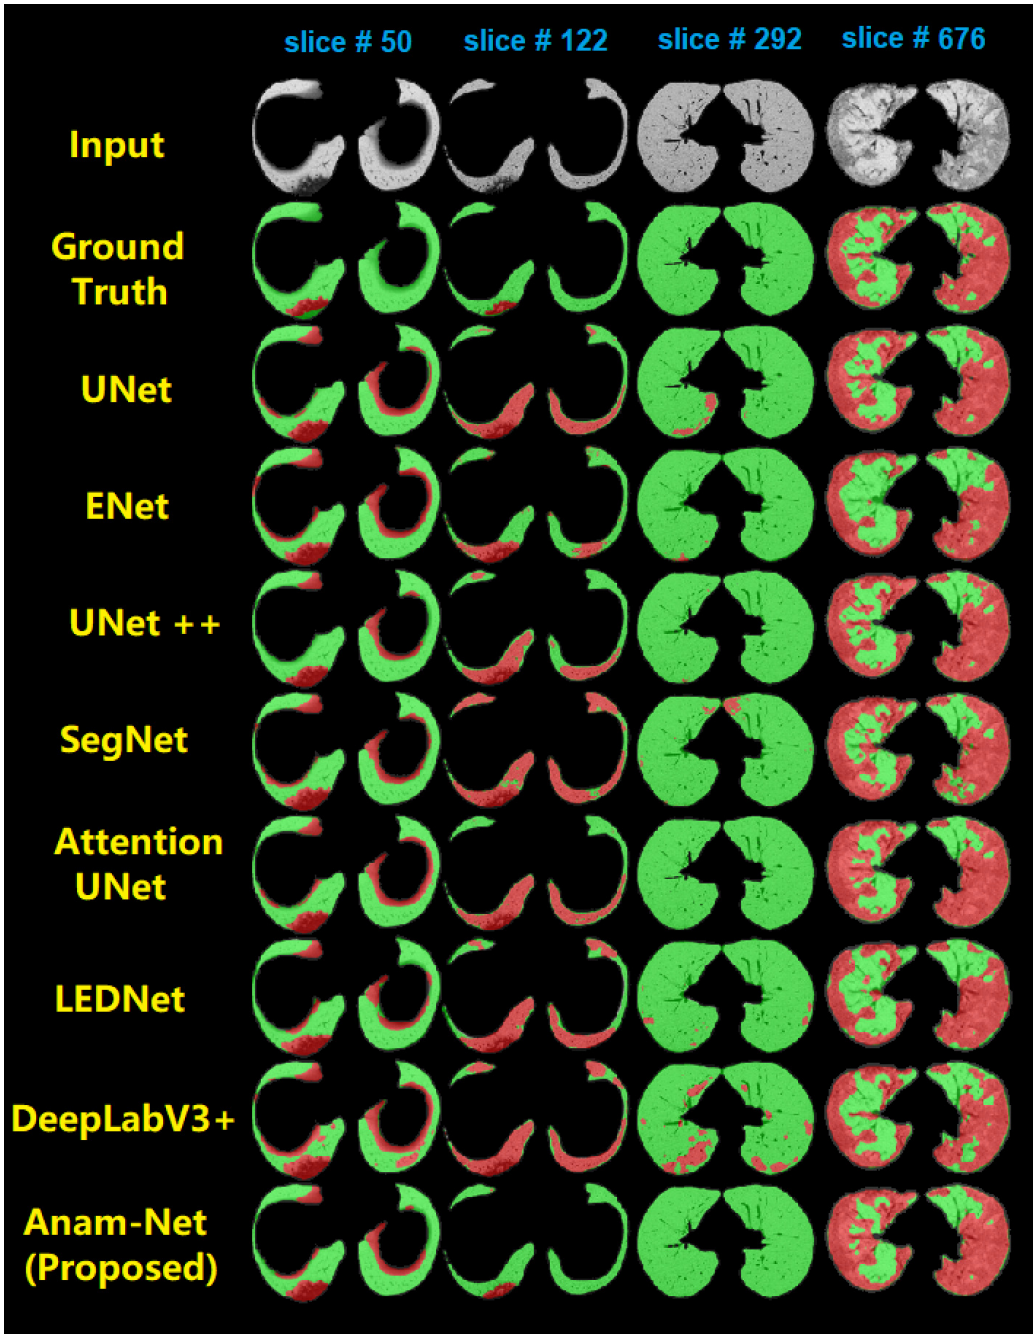


Fig. 4. Representative segmentation results of experiment 1 (key steps were given in Fig. 1). Randomly selected input slices (test cases) were shown in the top row. The annotations (ground truth) are shown in second row. The abnormalities in the lung region is indicated in red color and normal lung region is in green color. Dice Similarity scores of the abnormal lung regions for these test cases has been shown in Table VI.

TABLE VI
DICE SIMILARITY SCORES FOR THE ABNORMAL LUNG REGION FOR TEST CASES SHOWN IN FIG. 4. THE LAST ROW REPRESENTS THE AVERAGE OF THESE RESULTS. THE BEST RESULTS ARE SHOWN IN BOLD.

| Model / Slice No. | UNet | ENet | UNet++ | SegNet | Attention UNet | LEDNet | DeepLabV3+ | Anam-Net (Proposed) |
|---|---|---|---|---|---|---|---|---|
| **50** | 0.36 | 0.30 | **0.44** | 0.35 | 0.33 | 0.37 | 0.26 | 0.38 |
| **122** | 0.16 | 0.28 | 0.19 | 0.12 | 0.16 | 0.14 | 0.13 | **0.80** |
| **292** | 0.00 | 0.00 | 0.00 | 0.00 | **1.00** | 0.00 | 0.00 | **1.00** |
| **676** | 0.90 | 0.81 | **0.90** | 0.86 | 0.89 | 0.84 | 0.86 | 0.89 |
| **Average** | 0.36 | 0.35 | 0.39 | 0.33 | 0.60 | 0.34 | 0.31 | **0.77** |

TABLE VII
AVERAGED FIGURES OF MERIT OVER ALL 704 CHEST CT IMAGES (TEST CASES) CONSIDERED IN THE EXPERIMENT 1. THE BEST RESULTS ARE SHOWN IN BOLD.

| Model | Class | Sensitivity | Specificity | Accuracy | Dice Score |
|---|---|---|---|---|---|
| **UNet** | Abnormal | 0.932 | 0.983 | 0.982 | 0.608 |
| | Normal | 0.902 | **0.997** | 0.981 | 0.943 |
| **ENet** | Abnormal | 0.857 | 0.990 | 0.988 | 0.694 |
| | Normal | 0.954 | 0.991 | 0.985 | 0.956 |
| **UNet++** | Abnormal | **0.954** | 0.986 | 0.986 | 0.674 |
| | Normal | 0.924 | **0.997** | 0.985 | 0.955 |
| **SegNet** | Abnormal | 0.936 | 0.981 | 0.980 | 0.587 |
| | Normal | 0.897 | 0.995 | 0.978 | 0.935 |
| **Attention UNet** | Abnormal | 0.951 | 0.988 | 0.987 | 0.695 |
| | Normal | 0.932 | **0.997** | 0.986 | 0.960 |
| **LEDNet** | Abnormal | 0.907 | 0.983 | 0.981 | 0.597 |
| | Normal | 0.901 | 0.991 | 0.976 | 0.929 |
| **DeepLabV3+** | Abnormal | 0.850 | 0.958 | 0.956 | 0.366 |
| | Normal | 0.762 | 0.993 | 0.954 | 0.850 |
| **Anam-Net (Proposed)** | Abnormal | 0.900 | **0.993** | **0.991** | **0.755** |
| | Normal | **0.959** | **0.997** | **0.990** | **0.972** |

TABLE VIII
AVERAGED FIGURES OF MERIT ACROSS 3 CROSS FOLDS OVER ALL 545 CHEST CT IMAGES (TEST CASES) CONSIDERED IN THE EXPERIMENT 2. THE BEST RESULTS ARE SHOWN IN BOLD.

| Model | Class | Sensitivity | Specificity | Accuracy | Dice Score |
|---|---|---|---|---|---|
| **UNet** | Abnormal | 0.907 | 0.990 | **0.988** | 0.791 |
| | Normal | 0.913 | 0.995 | 0.987 | 0.938 |
| **ENet** | Abnormal | 0.736 | 0.989 | 0.984 | 0.686 |
| | Normal | 0.914 | 0.989 | 0.912 | 0.911 |
| **UNet++** | Abnormal | 0.951 | 0.990 | **0.988** | **0.805** |
| | Normal | 0.906 | 0.995 | 0.987 | 0.937 |
| **SegNet** | Abnormal | 0.657 | 0.988 | 0.980 | 0.612 |
| | Normal | 0.888 | 0.987 | 0.977 | 0.885 |
| **Attention UNet** | Abnormal | **0.966** | 0.987 | 0.987 | 0.783 |
| | Normal | 0.884 | **0.997** | 0.986 | 0.930 |
| **LEDNet** | Abnormal | 0.782 | 0.989 | 0.984 | 0.707 |
| | Normal | 0.912 | 0.987 | 0.979 | 0.899 |
| **DeepLabV3+** | Abnormal | 0.521 | **0.992** | 0.981 | 0.563 |
| | Normal | **0.930** | 0.983 | 0.978 | 0.895 |
| **Anam-Net (Proposed)** | Abnormal | 0.918 | 0.990 | **0.988** | 0.798 |
| | Normal | 0.911 | **0.997** | **0.988** | **0.941** |

specificity, accuracy, and Dice score was superior. However, the performance of UNet++ was also on par with the proposed approach because of extensive dense connections in its design that result in hierarchical encoder-decoder modules enabling efficient feature propagation for accurate segmentation. The averaged results from Experiment-1 were presented in Table-VII, and that of Experiment-2 and Experiment-3 are in Table-VIII and Table-IX respectively. Further, higher specificity and higher accuracy for both the classes (abnormal and normal) are desirable in disease monitoring, especially in the remission of COVID-19. As it can be seen from Table-VII and Table-VIII, the proposed Anam-Net with fewer parameters was able to provide accurate segmentation results compared to already existing models. As shown in Table-IX, even in the cross dataset examination (Experiment-3), the performance of Anam-Net was reasonably good (second-best) and was comparable with the best performing method Attention UNet.

As it can be seen from results, Fig. 4, Table-VI, and Table-VII, that proposed Anam-Net provides superior performance with the utilization of anamorphic depth embeddings, which enabled the network to be light-weight. The patient-level (a sample test case from Experiment-2) segmentation results are shown in Fig. 5, wherein the average Dice score of the normal lung region was $0.95$, and that of the abnormal region was $0.68$. In most cases, while processing the initial slices of the lung we observed that the Dice score for the abnormal region was minimal($< 0.5$). However, on an average the accuracy of segmenting the anomalies was as high as $0.98$ in all the experiments (refer to Table-VII, Table-VIII and Table-IX). Overall, the next best performing network was Attention UNet,

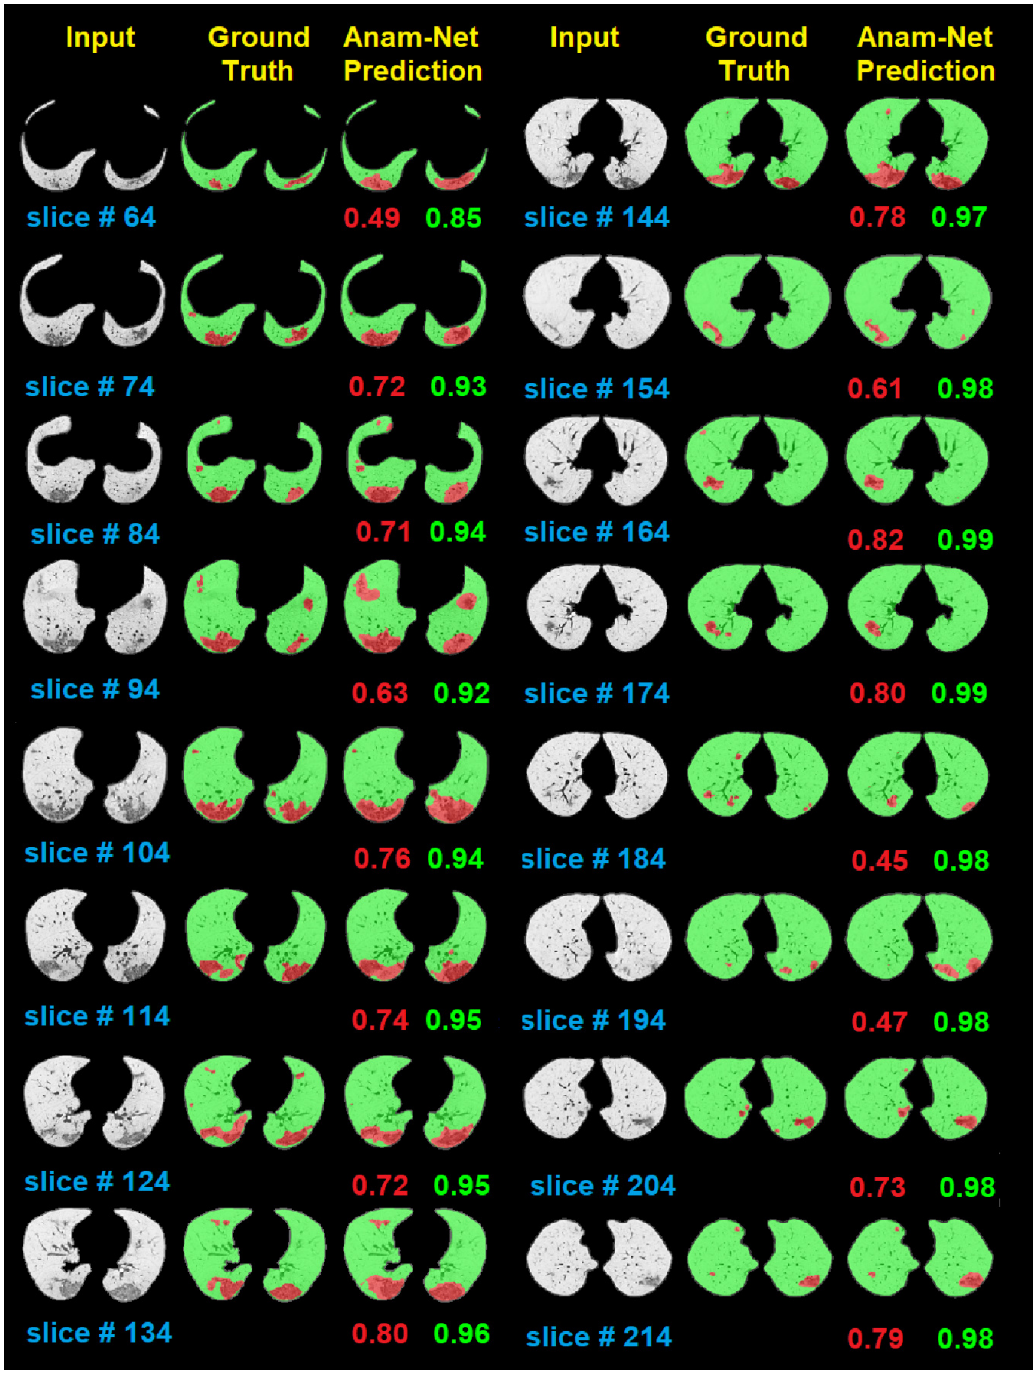


Fig. 5. Representative segmentation results on a patient volume (test cases) from experiment 2 (key steps were given in Fig. 1). The input slices (test cases) were shown in first and fourth columns. The respective annotations (ground truth) are shown in second and fifth column. The predictions of the Anam-Net are in given the third and the sixth columns. The abnormalities in the lung region are indicated in red color and normal lung region is in green color. Dice Similarity scores of the abnormal (in red color) and normal (in green color) lung regions are given below the corresponding slices.

which has at least 7.8 times more parameters (refer to Table-IV) and required 2.3 times more training time compared to Anam-Net. Also, these heavy models (sizes being in hundreds of MB) may not be well suited for the point-of-care platforms for providing quick inference. Despite the proposed model having only 4.47 million parameters (third lightest), it was able to outperform the rest of the networks in all figures of merit (Table-VII). The light-weight networks performance was much inferior compared to heavy networks and among all networks DeepLabV3+ providing the lowest Dice score (Table-VII, Table-VIII and Table-IX).

## VI. Discussion

As can be seen from Table-VII, the Dice similarity score of the proposed method (Anam-Net) was the highest compared to existing state-of-the-art architectures. Specifically, we observed that except for Attention UNet and the proposed Anam-Net, rest networks failed in identifying healthy individuals

TABLE IX
AVERAGED FIGURES OF MERIT OVER ALL 704 CHEST CT IMAGES (TEST CASES) CONSIDERED IN THE EXPERIMENT 3. THE BEST RESULTS ARE SHOWN IN BOLD.

| Model | Class | Sensitivity | Specificity | Accuracy | Dice Score |
|---|---|---|---|---|---|
| **UNet** | Abnormal | **0.960** | 0.985 | 0.985 | 0.657 |
| | Normal | 0.916 | **0.998** | 0.984 | 0.951 |
| **ENet** | Abnormal | 0.864 | 0.972 | 0.970 | 0.464 |
| | Normal | 0.845 | 0.994 | 0.968 | 0.902 |
| **UNet++** | Abnormal | 0.951 | 0.983 | 0.983 | 0.628 |
| | Normal | 0.913 | 0.997 | 0.982 | 0.947 |
| **SegNet** | Abnormal | 0.884 | 0.946 | 0.945 | 0.325 |
| | Normal | 0.699 | 0.995 | 0.945 | 0.813 |
| **Attention UNet** | Abnormal | 0.931 | **0.989** | **0.988** | **0.704** |
| | Normal | **0.941** | **0.998** | **0.988** | **0.964** |
| **LEDNet** | Abnormal | 0.872 | 0.981 | 0.979 | 0.559 |
| | Normal | 0.906 | 0.983 | 0.970 | 0.912 |
| **DeepLabV3+** | Abnormal | 0.845 | 0.906 | 0.905 | 0.209 |
| | Normal | 0.489 | 0.990 | 0.904 | 0.637 |
| **Anam-Net (Proposed)** | Abnormal | 0.937 | 0.986 | 0.986 | 0.664 |
| | Normal | 0.927 | 0.**998** | 0.985 | 0.956 |

TABLE X
ABLATION STUDY FOR ALL ANAMORPHIC DEPTH (AD) BLOCKS AND THE LABEL WEIGHTING STRATEGY DURING THE TRAINING OF ANAM-NET. THE BEST RESULTS ARE SHOWN IN BOLD. NOTE THAT THESE RESULTS ARE AVERAGED OVER 704 CHEST CT IMAGES (ALL TEST CASES) FROM EXPERIMENT 1.

| Figure of Merit | AD Blocks | Label Weighting | Abnormal | Normal |
|---|---|---|---|---|
| **Dice Similarity** | ✓ | ✓ | **0.755** | **0.972** |
| **Dice Similarity** | × | ✓ | 0.666 | 0.955 |
| **Dice Similarity** | ✓ | × | 0.558 | 0.934 |
| **Dice Similarity** | × | × | 0.295 | 0.778 |

TABLE XI
COMPARISON BETWEEN THE PROPOSED ANAM-NET AND UNET IN TERMS OF PARAMETER SENSITIVITY. THE PARAMETERS ($\theta$) OF THE TRAINED MODELS (FROM EXPERIMENT 1) WERE PERTURBED BY $\eta$ (%). NOTE THAT THESE RESULTS ARE AVERAGED OVER 704 CHEST CT IMAGES (ALL TEST CASES) FROM EXPERIMENT 1.

| Perturbation $\eta$ (%) | Anam-Net (Average Dice Score) | UNet (Average Dice Score) |
|---|---|---|
| -0.1 | 0.193 | 0.136 |
| -0.01 | 0.414 | 0.387 |
| -0.001 | 0.859 | 0.776 |
| 0.0 | 0.863 | 0.775 |
| 0.001 | 0.846 | 0.769 |
| 0.01 | 0.324 | 0.445 |
| 0.1 | 0.172 | 0.095 |

(slice 292 shown in Fig. 4). The Anam-Net successfully generated the null result; the same is evident when comparing the sensitivity values of the normal class indicated in Table-VII. In essence, the training times required for Anam-Net is lesser (roughly about 43%) compared to Attention UNet (since the number of parameters being fewer, Table-IV), and Anam-Net was able to generate more accurate results compared to others. The improved accuracy of Anam-Net compared to the UNet model can be attributed to the Anamorphic Depth (AD) block, which was confirmed by the ablation study (Refer to Table-X). Without the AD block, the performance in terms of the Dice score was comparable with UNet results (compare the second row in Table-X with UNet results in Table-VII).

Obtaining accurate labels is important while developing deep learning models for the automatic segmentation of abnormalities. The labels can be classified into two major categories, namely strong annotations (wherein the radiologist has performed a proper segmentation of the region of interest (ROI)), and weak annotations (which can simply be scribbles, sparse dots, or noisy annotated labels), for pictorial representation, please check Fig. 11 of Ref. [59]. Zheng et al. [60] have used weak annotations in the form of patient-level labels, i.e., whether the patient is COVID-positive or COVID-negative to train the network for automatically detecting COVID-19 cases. Ref. [60] has used weakly-supervised learning for COVID-19 detection, wherein a spatial global pooling layer and a temporal global pooling layer was introduced into the DeCovNet. Xu et al. [61] have developed a model that can handle multiple classes with patient-level labels. Ref. [61] included patient-level labels belonging to different classes like Influenza, COVID, and pneumonia while training the deep learning model; however, these studies did not consider annotations (i.e., segmenting the abnormality in the CT image) for training the model. Obtaining these strong annotations to train a traditional CNN is time-consuming and expensive. The proposed Anam-Net can be seen as Mini-CNN with the number of trainable parameters being 7 times less compared to UNet (or its variants) architecture. Anam-Net will also have a universal appeal to deploy for a site/protocol specific accurate segmentation network within adequate training time (Table-IV). We have also compared the proposed Anam-Net and UNet

in terms of parameter sensitivity. The parameters ($\theta$) of the trained models (from experiment 1) were perturbed by $\eta$ ($\%$) and the perturbed models were tested on 704 chest CT images (all the test cases) from Experiment 1. The averaged results of the parameter sensitivity analysis are shown in Table-XI. These results indicate that the proposed Anam-Net has the same stability (if not better) as UNet and is suitable for a critical application like the prognosis of COVID-19.

Shan et al. [62] have proposed a human-in-the-loop (HITL) strategy to improve the annotations (i.e., using image segmentation) required for training the network to quantify COVID-19 infection; this work initially used a small batch of segmentations (obtained from a radiologist) to train the VB-network. This approach utilized the trained network to generate a rough segmentation, which was then corrected by the radiologists, and the fixed segmentation improved the network performance in an iterative fashion [62]. The HITL strategy requires intervention from the radiologists for accurately training the deep learning model. Further, this strategy is computationally expensive as opposed to the proposed Anam-Net, which takes about 27 minutes for an end to end training. Note that the number of parameters involved in the UNet model is $31.07$Million as opposed to $4.47$Million in the case of Anam-Net (Table-IV), making the proposed approach computationally efficient with an added advantage of being trainable with a smaller dataset without compromising performance (Table-VII). The performance of other state-of-the-art light-weight or heavy models is sub-par compared to the proposed Anam-Net, and the observed improvement was at least 1.24 times among light-weight networks and twice compared to other massive networks.

Even radiologists performance for differentiating COVID-19 pneumonia from non-COVID-19 pneumonia could be as low as $60\%$ [63] and stand-alone chest CT images without any feature engineering might not reveal distinct patterns of COVID-19. The proposed method can quickly provide the most affected region in chest CT for a radiologist to enable faster diagnosis. Current evaluation of the COVID-19 or otherwise other pneumonia severity by the radiologist at best is semi-quantitative and typically performed by visual scoring [16]. Any automated methods will always provide unbiased estimates, which is desirable in clinical practice. Techniques such as the proposed one will pave the way for effective and better management of COVID-19 and associated morbidity. As there is a push for low-dose chest CT [17], the variation in protocols demands to retraining of CNNs; thus the proposed method will be able to meet this need and provide the versatility without compromising the accuracy of the outcome.

The number of CT scans utilized for training a deep learning model to quantify the lung infection as provided in Ref. [62] was 249, while Refs. [60] and [61] utilized 499 CT scans and 618 CT slices , respectively. The proposed Anam-Net utilized only $90$ chest CT images in experiment 1 and resulted in an average Dice score of 0.87 averaged over 704 test images (which is almost 8 times larger than the training set). Note that in this experiment, we intentionally had a smaller training set and a larger test set, as in the current pandemic situation, it is challenging to obtain accurate annotations from radiologists, which have increased clinical load. This work aimed to propose a novel network specifically designed for the task at hand and benchmark it among already existing state-of-the-art networks on large test data with a constrain of available training data being limited. In experiment 3, we conducted a cross dataset examination to further analyze the generalizability of Anam-Net and, even in this case Anam-Net gave good Dice scores for abnormal as well as normal lung regions. (Refer to Table-IX). The light-weight CNNs based on attentive hierarchical spatial pyramid modules were recently proposed to segment the abnormal regions in COVID-19 chest CT images [13]; the network had about 472.44K parameters. Even with ImageNet pre-training, this model gave only 0.84 sensitivity. Its performance was inferior to the standard UNet approach when test data was limited to only $40$ chest CT images [13]. Note that Ref. [13] utilizes the same training data as that of experiment 1 in this work. The proposed method in this work, Anam-Net, was stand-alone and did not require any pre-training; further, Anam-Net was validated with larger test data (i.e., with $704$ chest CT images in experiment 1, $545$ chest CT images in experiment 2, and $704$ chest CT images in experiment 3). The Anam-Net showed improved specificity, accuracy and Dice similarity score compared to standard UNet (Table-VII, Table-VIII and Table-IX) and other state-of-the-art deep learning models.

The study presented here has few limitations, the first one being the dimensionality of chest CT images being restricted to two-dimensions. The operations performed here can be applied in three-dimensions, and a detailed study in this respect will be taken up as future work. The analysis of the results on the chest CT images (test cases) indicate that the proposed Anam-Net is inherently biased to the peripheral part of the lung (can be observed in Fig. 4), and most COVID-19 chest CT images have the manifestation of peripheral abnormalities [8]–[11], but these peripheral abnormalities might be absent in few cases especially in asymptomatic as well as pediatric patients, which brings down the Dice similarity score. Low Dice scores can be further improved by incorporating spatial-semantic context into the Anam-Net; this will also be explored in the future. As there is an increased variation in chest CT protocols to reduce the effective dose to the patients without compromising the diagnostic accuracy [17], the amount of fully annotated chest CT data acquired under the same protocol is still a challenge. The work presented here provides a solution to this challenge explicitly by requiring significantly less training data without compromising the accuracy of the segmentation, making them attractive and easy to deploy in the clinic. These novel methods are critical for making deep learning methods more appealing for real-time COVID-19 imaging studies.

## VII. Conclusion

This work presented an anamorphic depth embedding based light-weight CNN, called Anam-Net, to segment anomalies in COVID-19 Chest CT images. As chest CT imaging is becoming the main workhorse for staging and managing COVID-19, the methodology proposed here is the need of the hour. The results from the chest CT images (test cases) across the three

experiments showed that the proposed method could provide good Dice similarity scores for abnormal as well as normal regions. Further, the Anam-Net was benchmarked against other state-of-the-art light-weight and heavy networks, like ENet, UNet++, SegNet, Attention UNet, LEDNet, DeepLabV3+ and found to provide higher specificity, accuracy, and Dice score averaged over all the chest CT test images across three different experiments. The advantage of Anam-Net compared to other models is low computational complexity (requiring $50\%$ of the training time compared to the next best performing network and the number of parameters being 7 times fewer), making it attractive to be deployed in a clinical setting. Anam-Nets model size is in the order of tens of megabytes (to be specific, 17.2 MB), makes it easily deployable in mobile platforms to provide a quick assessment of the abnormalities in COVID-19 chest CT images. The deployment in mobile as well as embedded hardware platforms confirmed that the proposed Anam-Net is well suited for the point of care settings.

## ACKNOWLEDGEMENTS

The authors are thankful to Dr. Johannes Hofmanninger for providing the lung masks as open source, which were utilized in this work.

## REFERENCES

[1] Z. A. Memish, S. Perlman, M. D. Van Kerkhove, and A. Zumla, "Middle east respiratory syndrome," *The Lancet*, vol. 395, no. 10229, pp. 1063–1077, 2020. [Online]. Available: https://www.thelancet.com/journals/lancet/article/PIIS0140-6736(19)33221-0/fulltext

[2] N. Chen, M. Zhou, X. Dong, J. Qu, F. Gong, Y. Han, Y. Qiu, J. Wang, Y. Liu, Y. Wei, J. Xia, T. Yu, X. Zhang, and L. Zhang, "Epidemiological and clinical characteristics of 99 cases of 2019 novel coronavirus pneumonia in wuhan, china: a descriptive study," *The Lancet*, vol. 395, no. 10223, pp. 507 – 513, 2020. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0140673620302117

[3] G. D. Rubin, C. J. Ryerson, L. B. Haramati, N. Sverzellati, J. P. Kanne, S. Raoof, N. W. Schluger, A. Volpi, J.-J. Yim, I. B. K. Martin, D. J. Anderson, C. Kong, T. Altes, A. Bush, S. R. Desai, J. Goldin, J. M. Goo, M. Humbert, Y. Inoue, H.-U. Kauczor, F. Luo, P. J. Mazzone, M. Prokop, M. Remy-Jardin, L. Richeldi, C. M. Schaefer-Prokop, N. Tomiyama, A. U. Wells, and A. N. Leung, "The role of chest imaging in patient management during the covid-19 pandemic: A multinational consensus statement from the fleischner society," *Radiology*, p. 201365, 2020, pMID: 32255413. [Online]. Available: https://doi.org/10.1148/radiol.2020201365

[4] M. D. Hope, C. A. Raptis, and T. S. Henry, "Chest Computed Tomography for Detection of Coronavirus Disease 2019 (COVID-19): Don't Rush the Science," *Annals of Internal Medicine*, 04 2020. [Online]. Available: https://doi.org/10.7326/M20-1382

[5] Y. Fang, H. Zhang, J. Xie, M. Lin, L. Ying, P. Pang, and W. Ji, "Sensitivity of chest ct for covid-19: Comparison to rt-pcr," *Radiology*, p. 200432, 2020, pMID: 32073353. [Online]. Available: https://doi.org/10.1148/radiol.2020200432

[6] S. Inui, A. Fujikawa, M. Jitsu, N. Kunishima, S. Watanabe, Y. Suzuki, S. Umeda, and Y. Uwabe, "Chest ct findings in cases from the cruise ship diamond princess with coronavirus disease 2019 (covid-19)," *Radiology: Cardiothoracic Imaging*, vol. 2, no. 2, p. e200110, 2020. [Online]. Available: https://doi.org/10.1148/ryct.2020200110

[7] E. Driggin, M. V. Madhavan, B. Bikdeli, T. Chuich, J. Laracy, G. Bondi-Zoccai, T. S. Brown, C. D. Nigoghossian, D. A. Zidar, J. Haythe, D. Brodie, J. A. Beckman, A. J. Kirtane, G. W. Stone, H. M. Krumholz, and S. A. Parikh, "Cardiovascular considerations for patients, health care workers, and health systems during the coronavirus disease 2019 (covid-19) pandemic," *Journal of the American College of Cardiology*, 2020. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S0735109720346374

[8] J. P. Kanne, "Chest ct findings in 2019 novel coronavirus (2019-ncov) infections from wuhan, china: Key points for the radiologist," pp. 16–17, 2020, pMID: 32017662. [Online]. Available: https://doi.org/10.1148/radiol.2020200241

[9] M. Chung, A. Bernheim, X. Mei, N. Zhang, M. Huang, X. Zeng, J. Cui, W. Xu, Y. Yang, Z. A. Fayad *et al.*, "Ct imaging features of 2019 novel coronavirus (2019-ncov)," *Radiology*, vol. 295, no. 1, pp. 202–207, 2020. [Online]. Available: https://pubs.rsna.org/doi/10.1148/radiol.2020200230

[10] D. Caruso, M. Zerunian, M. Polici, F. Pucciarelli, T. Polidori, C. Rucci, G. Guido, B. Bracci, C. de Dominicis, and A. Laghi, "Chest ct features of covid-19 in rome, italy," *Radiology*, pp. 201–237, 2020. [Online]. Available: https://pubs.rsna.org/doi/10.1148/radiol.2020201237

[11] D. Caruso, M. Zerunian, M. Polici, F. Pucciarelli, T. Polidori, C. Rucci, G. Guido, B. Bracci, C. Dominicis, and A. Laghi, "Chest ct features of covid-19 in rome, italy," *Radiology*, vol. 0, no. 0, p. 201237, 0, pMID: 32243238. [Online]. Available: https://doi.org/10.1148/radiol.2020201237

[12] L. Li, L. Qin, Z. Xu, Y. Yin, X. Wang, B. Kong, J. Bai, Y. Lu, Z. Fang, Q. Song *et al.*, "Artificial intelligence distinguishes covid-19 from community acquired pneumonia on chest ct," *Radiology*, p. 200905, 2020. [Online]. Available: https://pubs.rsna.org/doi/10.1148/radiol.2020200905

[13] Y. Qiu, Y. Liu, and J. Xu, "Miniseg: An extremely minimum network for efficient covid-19 segmentation," *arXiv preprint arXiv:2004.09750*, 2020.

[14] F. Pan, T. Ye, P. Sun, S. Gui, B. Liang, L. Li, D. Zheng, J. Wang, R. L. Hesketh, L. Yang *et al.*, "Time course of lung changes on chest ct during recovery from 2019 novel coronavirus (covid-19) pneumonia," *Radiology*, p. 200370, 2020, pMID: 32053470. [Online]. Available: https://doi.org/10.1148/radiol.2020200370

[15] Y. Wang, C. Dong, Y. Hu, C. Li, Q. Ren, X. Zhang, H. Shi, and M. Zhou, "Temporal changes of ct findings in 90 patients with covid-19 pneumonia: a longitudinal study," *Radiology*, p. 200843, 2020. [Online]. Available: https://pubs.rsna.org/doi/full/10.1148/radiol.2020200843

[16] L. Huang, R. Han, T. Ai, P. Yu, H. Kang, Q. Tao, and L. Xia, "Serial quantitative chest ct assessment of covid-19: Deep-learning approach," *Radiology: Cardiothoracic Imaging*, vol. 2, no. 2, p. e200075, 2020. [Online]. Available: https://doi.org/10.1148/ryct.2020200075

[17] R. Iezzi, A. R. Larici, P. Franchi, R. Marano, N. Magarelli, A. Posa, B. Merlino, R. Manfredi, and C. Colosimo, "Tailoring protocols for chest ct applications: when and how?" *Diagnostic and Interventional Radiology*, vol. 23, no. 6, p. 420, 2017.

[18] Y. Zhu, Y. Tan, Y. Hua, G. Zhang, and J. Zhang, "Automatic segmentation of ground-glass opacities in lung ct images by using markov random field-based algorithms," *Journal of digital imaging*, vol. 25, no. 3, pp. 409–422, 2012. [Online]. Available: https://doi.org/10.1007/s10278-011-9435-5

[19] J. Zhou, S. Chang, D. N. Metaxas, B. Zhao, L. H. Schwartz, and M. S. Ginsberg, "Automatic detection and segmentation of ground glass opacity nodules," in *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, 2006, pp. 784–791. [Online]. Available: https://doi.org/10.1007/11866565_96

[20] J. Jung, H. Hong, and J. M. Goo, "Ground-glass nodule segmentation in chest ct images using asymmetric multi-phase deformable model and pulmonary vessel removal," *Computers in biology and medicine*, vol. 92, pp. 128–138, 2018. [Online]. Available: https://doi.org/10.1016/j.compbiomed.2017.11.013

[21] D.-P. Fan, T. Zhou, G.-P. Ji, Y. Zhou, G. Chen, H. Fu, J. Shen, and L. Shao, "Inf-net: Automatic covid-19 lung infection segmentation from ct images," *IEEE Transactions on Medical Imaging*, 2020.

[22] X. Ouyang, J. Huo, L. Xia, F. Shan, J. Liu, Z. Mo, F. Yan, Z. Ding, Q. Yang, B. Song *et al.*, "Dual-sampling attention network for diagnosis of covid-19 from community acquired pneumonia," *IEEE Transactions on Medical Imaging*, 2020.

[23] Z. Akkus, P. Kostandy, K. A. Philbrick, and B. J. Erickson, "Robust brain extraction tool for ct head images," *Neurocomputing*, 2019.

[24] A. Paszke, A. Chaurasia, S. Kim, and E. Culurciello, "Enet: A deep neural network architecture for real-time semantic segmentation," *arXiv preprint arXiv:1606.02147*, 2016.

[25] O. Ronneberger, P.Fischer, and T. Brox, "U-net: Convolutional networks for biomedical image segmentation," in *Medical Image Computing and Computer-Assisted Intervention (MICCAI)*, ser. LNCS, vol. 9351. Springer, 2015, pp. 234–241, (available on arXiv:1505.04597 [cs.CV]). [Online]. Available: http://lmb.informatik.uni-freiburg.de/Publications/2015/RFB15a

[26] Z. Zhou, M. M. R. Siddiquee, N. Tajbakhsh, and J. Liang, "Unet++: Redesigning skip connections to exploit multiscale features in image segmentation," *IEEE Transactions on Medical Imaging*, 2019.

[27] V. Badrinarayanan, A. Kendall, and R. Cipolla, "Segnet: A deep convolutional encoder-decoder architecture for image segmentation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 39, no. 12, pp. 2481–2495, 2017.

[28] O. Oktay, J. Schlemper, L. L. Folgoc, M. Lee, M. Heinrich, K. Misawa, K. Mori, S. McDonagh, N. Y. Hammerla, B. Kainz *et al.*, "Attention u-net: Learning where to look for the pancreas," *arXiv preprint arXiv:1804.03999*, 2018.

[29] Y. Wang, Q. Zhou, J. Liu, J. Xiong, G. Gao, X. Wu, and L. J. Latecki, "Lednet: A lightweight encoder-decoder network for real-time semantic segmentation," in *2019 IEEE International Conference on Image Processing (ICIP)*. IEEE, 2019, pp. 1860–1864.

[30] L.-C. Chen, Y. Zhu, G. Papandreou, F. Schroff, and H. Adam, "Encoder-decoder with atrous separable convolution for semantic image segmentation," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 801–818.

[31] "Covid-19 ct segmentation dataset," http://medicalsegmentation.com/covid19/, accessed: 2020-05-29.

[32] "Github page," https://github.com/NaveenPaluru/Segmentation-COVID-19, accessed: 2020-06-07.

[33] J. Hofmanninger, F. Prayer, J. Pan, S. Rohrich, H. Prosch, and G. Langs, "Automatic lung segmentation in routine imaging is a data diversity problem, not a methodology problem," 2020. [Online]. Available: https://arxiv.org/abs/2001.11767

[34] M. Jun, G. Cheng, W. Yixin, A. Xingle, G. Jiantao, Y. Ziqi, Z. Minqing, L. Xin, D. Xueyuan, C. Shucheng, W. Hao, M. Sen, Y. Xiaoyu, N. Ziwei, L. Chen, T. Lu, Z. Yuntao, Z. Qiongjie, D. Guoqiang, and H. Jian, "COVID-19 CT Lung and Infection Segmentation Dataset," Apr. 2020. [Online]. Available: https://doi.org/10.5281/zenodo.3757476

[35] Q. Yu, D. Yang, H. Roth, Y. Bai, Y. Zhang, A. L. Yuille, and D. Xu, "C2fnas: Coarse-to-fine neural architecture search for 3d medical image segmentation," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4126–4135.

[36] D. Guo, D. Jin, Z. Zhu, T.-Y. Ho, A. P. Harrison, C.-H. Chao, J. Xiao, and L. Lu, "Organ at risk segmentation for head and neck cancer using stratified learning and neural architecture search," in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 2020, pp. 4223–4232.

[37] X. Li, L. Yu, H. Chen, C.-W. Fu, L. Xing, and P.-A. Heng, "Transformation-consistent self-ensembling model for semisupervised medical image segmentation," *IEEE Transactions on Neural Networks and Learning Systems*, 2020.

[38] J. Clough, N. Byrne, I. Oksuz, V. A. Zimmer, J. A. Schnabel, and A. King, "A topological loss function for deep-learning based image segmentation using persistent homology," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.

[39] Y. Zhou, G. G. Yen, and Z. Yi, "Evolutionary compression of deep neural networks for biomedical image segmentation," *IEEE transactions on neural networks and learning systems*, 2019.

[40] G. Wang, M. A. Zuluaga, W. Li, R. Pratt, P. A. Patel, M. Aertsen, T. Doel, A. L. David, J. Deprest, S. Ourselin *et al.*, "Deepigeos: a deep interactive geodesic framework for medical image segmentation," *IEEE transactions on pattern analysis and machine intelligence*, vol. 41, no. 7, pp. 1559–1572, 2018.

[41] Z. Gu, J. Cheng, H. Fu, K. Zhou, H. Hao, Y. Zhao, T. Zhang, S. Gao, and J. Liu, "Ce-net: Context encoder network for 2d medical image segmentation," *IEEE transactions on medical imaging*, vol. 38, no. 10, pp. 2281–2292, 2019.

[42] C. Chen, Q. Dou, H. Chen, J. Qin, and P. A. Heng, "Unsupervised bidirectional cross-modality adaptation via deeply synergistic image and feature alignment for medical image segmentation," *IEEE Transactions on Medical Imaging*, 2020.

[43] G. Huang, Z. Liu, L. Van Der Maaten, and K. Q. Weinberger, "Densely connected convolutional networks," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2017, pp. 4700–4708.

[44] D. Jha, P. H. Smedsrud, M. A. Riegler, D. Johansen, T. De Lange, P. Halvorsen, and H. D. Johansen, "Resunet++: An advanced architecture for medical image segmentation," in *2019 IEEE International Symposium on Multimedia (ISM)*. IEEE, 2019, pp. 225–2255.

[45] D. Jha, M. A. Riegler, D. Johansen, P. Halvorsen, and H. D. Johansen, "Doubleu-net: A deep convolutional neural network for medical image segmentation," *arXiv preprint arXiv:2006.04868*, 2020.

[46] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[47] F. Yu and V. Koltun, "Multi-scale context aggregation by dilated convolutions," *arXiv preprint arXiv:1511.07122*, 2015.

[48] J. Jin, A. Dundar, and E. Culurciello, "Flattened convolutional neural networks for feedforward acceleration.(2015)," *arXiv preprint arXiv:1412.5474*, 2015.

[49] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.

[50] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *arXiv preprint arXiv:1502.03167*, 2015.

[51] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.

[52] "Covid-19 database," https://www.sirm.org/en/category/articles/covid-19-database/, accessed: 2020-05-29.

[53] "Covid-19," https://radiopaedia.org/articles/covid-19-3, accessed: 2020-05-29.

[54] "Covid-19 radiology data collection and preparation for artificial intelligence," https://bit.ly/34t8Ih4, accessed: 2020-05-29.

[55] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan, T. Killeen, Z. Lin, N. Gimelshein, L. Antiga *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Advances in neural information processing systems*, 2019, pp. 8026–8037.

[56] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.

[57] "Pytorch mobile," https://pytorch.org/mobile/android/.

[58] "Android studio," https://developer.android.com/studio.

[59] N. Tajbakhsh, L. Jeyaseelan, Q. Li, J. N. Chiang, Z. Wu, and X. Ding, "Embracing imperfect datasets: A review of deep learning solutions for medical image segmentation," *Medical Image Analysis*, vol. 63, p. 101693, 2020. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S136184152030058X

[60] C. Zheng, X. Deng, Q. Fu, Q. Zhou, J. Feng, H. Ma, W. Liu, and X. Wang, "Deep learning-based detection for covid-19 from chest ct using weak label," *medRxiv*, 2020. [Online]. Available: https://www.medrxiv.org/content/early/2020/03/26/2020.03.12.20027185

[61] X. Xu, X. Jiang, C. Ma, P. Du, X. Li, S. Lv, L. Yu, Y. Chen, J. Su, G. Lang, Y. Li, H. Zhao, K. Xu, L. Ruan, and W. Wu, "Deep learning system to screen coronavirus disease 2019 pneumonia," 2020. [Online]. Available: https://arxiv.org/abs/2002.09334

[62] F. Shan, Y. Gao, J. Wang, W. Shi, N. Shi, M. Han, Z. Xue, D. Shen, and Y. Shi, "Lung infection quantification of covid-19 in ct images with deep learning," 2020. [Online]. Available: https://arxiv.org/abs/2003.04655

[63] H. X. Bai, B. Hsieh, Z. Xiong, K. Halsey, J. W. Choi, T. M. L. Tran, I. Pan, L.-B. Shi, D.-C. Wang, J. Mei, X.-L. Jiang, Q.-H. Zeng, T. K. Egglin, P.-F. Hu, S. Agarwal, F. Xie, S. Li, T. Healey, M. K. Atalay, and W.-H. Liao, "Performance of radiologists in differentiating covid-19 from viral pneumonia on chest ct," *Radiology*, vol. 0, no. 0, p. 200823, 0, pMID: 32155105. [Online]. Available: https://doi.org/10.1148/radiol.2020200823

**Naveen Paluru** is currently a Ph.D. student of Department of Computational and Data Sciences, Indian Institute of Science, Bangalore, India. He received a Masters by Research Degree in 2019 from the Department of Electrical Engineering, Indian Institute of Technology (IIT), Tirupati, India. His main research interests are in machine learning and deep learning for automated medical imaging analysis.

**Aveen Dayal** currently works as a visiting research student at Department of Information and communication technology, University of Agder, Grimstad, Norway. He received a bachelors degree in Computer Science and Engineering in 2020. His main research interests are in machine learning and deep learning for autonomous cyber-physical systems.

**Hvard Bjrke Jenssen** is working as a radiologist at Oslo University Hospital, Oslo, Norway. He received his MB, BCh, BAO from the Royal College of Surgeons in Ireland in 2013. His main research interests are in the segmentation of radiological imagery and the application of Artificial Intelligence in radiological practice.

**Tomas Sakinis** , MD, works as a radiologist at Oslo University Hospital - Rikshospitalet where he completed his radiological residency. He is currently a Ph.D. student in a project involving automated segmentation within neuroradiology. His main research interest is in the use of iterative and interactive AI-aided segmentation and the application of these methods in clinical practice. He is also leading the development of MedSeg - a free browser-based segmentation tool with AI capabilities.

**Linga Reddy Cenkeramaddi** (Senior Member, IEEE) received the masters degree in electrical engineering from the Indian Institute of Technology, New Delhi, India, in 2004, and the Ph.D. degree in electrical engineering from the Norwegian University of Science and Technology, Trondheim, Norway, in 2011. He worked for Texas Instruments in mixed-signal circuit design before joining the Ph.D. program at NTNU. After finishing his Ph.D., he worked in radiation imaging for an atmosphere space interaction monitor (ASIM mission to International Space Station) at the University of Bergen, Norway, from 2010 to 2012. He is currently working as an associate professor at the University of Agder, Campus Grimstad, Norway. His main scientific interests are in Cyber-Physical Systems, Autonomous Systems, and Wireless Embedded Systems.

**Jaya Prakash** received his B.Tech. in Information Technology from Amrita School of Engineering, Bengaluru, India in 2010, M.Sc. Degree in engineering from the Indian Institute of Science, Bengaluru, India in 2012, and Ph.D. Degree in medical imaging from the Indian Institute of Science, Bengaluru, India, in 2014. Prior to his current position as Assistant Professor with the Department of Instrumentation and Applied Physics, Indian Institute of Science, Bengaluru, he was a group leader of the Computational Data Analytics group at Institute for Biological and Medical Imaging in Helmholtz Zentrum Munich, Germany. His research interests are image reconstruction, inverse problems, optoacoustic imaging, biomedical instrumentation,and biomedical optics.

**Phaneendra K. Yalavarthy** (Senior Member, IEEE) received the M.Sc. Degree in engineering from the Indian Institute of Science, Bangalore, India, and Ph.D. Degree in biomedical computation from Dartmouth College, Hanover, NH, USA, in 2007. He is an Associate Professor with the Department of Computational and Data Sciences, Indian Institute of Science, Bangalore. His research interests include medical image computing, medical image analysis, and biomedical optics. He is a senior member of IEEE, SPIE, and OSA, and serves as an associate editor of IEEE Transactions on Medical Imaging.