Identification and Severity Assessment of COVID-19 using Lung CT Scans

ANAND THYAGACHANDRAN[1], AATHIRA BALACHANDRAN [2], AND HEMA A MURTHY [1],
[1]Computer Science and Engineering Department, Indian Institute of Technology Madras, India (e-mail: tanand@cse.iitm.ac.in, hema@cse.iitm.ac.in)
[2]Radio Diagnosis Department, Government Medical College Kollam, India (e-mail: draathirab@gmail.com)

Corresponding author: Anand Thyagachandran (e-mail: tanand@cse.iitm.ac.in), Hema A Murthy (e-mail: hema@cse.iitm.ac.in).

ABSTRACT The COVID-19 pandemic, caused by the SARS-CoV-2 virus, continues to have a significant impact on the global population. To effectively triage patients and understand the progression of the disease, a metric-based analysis of diagnostic techniques is necessary. The objective of the present study is to identify COVID-19 from chest CT scans and determine the extent of severity, defined by a severity score that indicates the volume of infection. An unsupervised preprocessing pipeline is proposed to extract relevant clinical features and utilize this information to employ a pre-trained ImageNet model to extract discriminative features. Subsequently, a shallow feed-forward neural network is trained to classify the available CT scans into three classes, namely COVID-19, Community-Acquired Pneumonia, and Normal. Through various ablation studies, we find that a domain-specific pre-processing pipeline improves classification accuracy significantly. In terms of classification accuracy, our approach, when evaluated on publicly available datasets, is seen to have an absolute improvement of 6% F1 score over the baseline model. Further, the estimated infection severity score is observed to be well correlated with radiologists' assessments. The results support the necessity of data-driven pre-processing before implementing learning algorithms.

INDEX TERMS COVID-19, CT scans, infection segmentation, semi-supervised augmentation, severity assessment.

I. INTRODUCTION

The COVID-19 pandemic, a highly contagious and primarily respiratory illness, has been of significant concern with devastating effects on public health, the world economy, and the social fabric of society. The unabated spread of infection, contrasted with other respiratory illnesses like SARS, has been ascribed to the ability of the virus to infect other people when the infected carrier is clinically asymptomatic. This poses a significant challenge in terms of early detection and containment. In severe cases, the progression of the disease often leads to respiratory problems, which can be identified by noticeable changes in chest X-rays or CT scans, such as lung fibrosis and opaqueness.

The diagnosis of COVID-19 is crucial in identifying the pathogenicity of the virus and the severity of the disease. The gold standard for this purpose in a community setting is Real-Time Polymerase Chain Reaction (RT-PCR), which uses nasal and nasopharyngeal swab samples [1]. While RT-PCR is a more reliable method for detecting infection, it

suffers from two limitations – 1. Although RT-PCR estimates the viral load, there seems to be an ambiguous correlation between viral load and the severity of the disease. This is owing to the nature of the swab taken, whereby nasopharyngeal swabs provide proof of the viral load in the upper respiratory tract, while the severe disease is usually associated with the lower respiratory tract. Also, 2. RT-PCR does not quantify the clinical features of the patient under study, i.e., it is not a test for the response of the human immune system to the pathogen. Concerning the scope of the present work with an emphasis on COVID-19 being a respiratory disease primarily, the clinical diagnosis of COVID-19 severity is analyzed through radiological techniques like chest X-ray, which can be in one angle (traditional X-ray) or along several planes to provide for tomographic reconstruction of the chest (CT Scan) [2]. The limitation of traditional X-ray imaging in being a line-of-sight integration results in poorer resolution and hence a poorer prognosis of COVID-19 compared to CT scans. In the context of image processing, CT scans

represent higher dimensional input data, with the ability of the processing algorithms being tested on the ability to recognize essential features. Although algorithms can encode high-dimensional images into a set of low-dimensional features, the overlap between features of various diseases (respiratory in the present context) results in erroneous classification as shown by [3], [4]. This can lead to a failure in managing the complications of COVID-19, such as cytokine storms, which are a major cause of fatalities in COVID-19 cases [5]. Motivated by these issues, the primary purpose of the present work is to perform better classification between a widely occurring respiratory disease – community-acquired pneumonia (CAP) and COVID-19 and also against normal CT-scan through a pre-processing routine.

Chest CT scans involve the projection of the X-ray bursts at different planes (termed as slices) to obtain high-resolution images of the chest region. The superior performance of CT scans in diagnosing respiratory illness lies in their ability to localize regions of abnormal opacity, which is usually a result of inflammation. These are often characterized in the form of ground-glass opacities (GGO), consolidation, a combination of GGOs and consolidation, halo sign (central consolidations surrounded by ground-glass opacities), reverse halo sign (central ground-glass lucent area with peripheral consolidation), and crazy paving patterns [6]–[8]. Common features of GGO include opaque foggy areas which do not obstruct the pulmonary vessels, while consolidation is marked by higher opacity, thereby rendering the visualization of the pulmonary vascular structures to be impossible. Crazy-paving patterns are linear patterns superimposed on the background of GGOs resembling irregularly shaped paving stones. The bilateral distributions of GGO with or without consolidation in the posterior or peripheral lung regions are regarded as the primary indicators for COVID-19. As the disease severity progresses, consolidations, crazy-paving patterns, and vascular enlargement [9], [10] are the hallmark features. Pleural effusion and significant mediastinal lymphadenopathy are less commonly observed findings in COVID-19 infection. Pleural effusion is accumulating excessive fluid in the pleural space surrounding each lung. In the case of pneumonia, the features are more localized with the observation of a unilateral distribution of GGOs and consolidation and associated with pleural effusion and significant mediastinal lymphadenopathy compared to COVID-19 [3]. An experienced radiologist was consulted to obtain the Hounsfield Unit (a measure of the range of radiation attenuation values – HU) corresponding to different regions in the CT scan. The clinical features observed for COVID-19 and CAP are shown in Figure. 1. The HU values of pixels in the GGO region lie between -650 HU and -50 HU, the consolidation region is between 10 HU and 45 HU, and the pleural effusion region is between 0 HU and 35 HU.

The proposed research encompasses the development of a preprocessing pipeline that enhances the features of interest from a CT scan of a COVID-19 patient, leading to a better assessment of the severity of the infection.


FIGURE 1: Clinical image findings: The imaging findings of COVID_19 and CAP in chest CT scans. Images (A)-(E) show the COVID-19 patient's chest CT scan slices, (F) and (G) show the CAP patient's chest CT scan slices. (A) Peripheral distribution of GGO in both lungs (red arrows). (B) Patchy areas of consolidations in the left lung (red arrow) and GGOs in the right lung (blue arrow). (C) Crazy paving pattern: GGOs with superimposed septal thickening (red arrows). (D) Halo sign: consolidations surrounded by GGOs (red arrow), consolidations at the lower lobes of lungs (blue arrow). (E) Subsegmental vessel enlargement near the lesion (red arrows). (F) An extensive area of consolidations with GGO in the right lung – severe disease. (G) Pleural effusion (yellow arrow) and GGOs (red arrows) in the left lung. (H) Normal chest CT scan slice.

The contributions of the paper are as follows:

- An unsupervised preprocessing module to segment the regions of infection from chest CT scans of COVID-19 patients.
- COVID-19 CT scan severity analysis (CTSS) based on the infections developed in the lungs.
- An ablation study is performed to evaluate the contribution of each stage in the preprocessing pipeline to classification accuracy.

The paper is organized as follows- Discussion on segmentation and identification of COVID-19 from chest CT scans as presented in Section II. Section IV describes the segmentation of relevant lung features in the CT scan contributing to infection. We also briefly explain the training strategy to classify CT scans into COVID-19, CAP, and Normal categories. Section V details the set of experiments and discussion on results and inferences. Section VI concludes the work.

II. RELATED WORK

Prior to the classification/identification of COVID-19, CT scans primarily relied on raw images, i.e., the data obtained from a diagnostic procedure. In related fields, however, preprocessing techniques such as image enhancement and segmentation have been employed extensively for enhancing image quality through the focus on dominant features and noise removal. In our work, we amalgamate generic image preprocessing techniques to aid image classification algorithms in identifying the prevalence of COVID-19 from

TABLE 1: Different datasets used in the experiment. "*" – slice-level labels are available; "**" – slice-level labels are unavailable, but the CT scan patient-level label is available. "C" and "S" denotes the classification and segmentation datasets respectively.

Dataset	Class	COVID-19	CAP	Normal	C/S	Format
SPGC	Train	55*+116**	25*+35**	76		
	Test	28**	51**	51	C	DICOM
LDCT	Test	104**	0	56		
LDCT-PCR	Test	100*	0	0		
Mosmed	Test	854**	0	254	C-S	NIfTI
Mosmed	Test	50*	0	0		
MedSeg	Test	100* slices	0	0		NIfTI
MedSeg_1	Test	9* (638 slices)	0	0	S	
Mehta	Test	14**	0	0		DICOM

chest CT scans. Although several algorithms for preprocessing do exist, we utilize knowledge-based machine learning approaches for the said purpose before embarking on classification. Considering the limitation in obtaining slice-level labels which however is required for better classification of the CT scans, a classifier is trained using the available data, with the training data being further augmented with pseudo-labeled data to improve accuracy.

A. CT SCAN SEGMENTATION

Segmentation using sophisticated image processing techniques is required to address the issue of varying contrasts observed in CT scans procured through different sources. The proposed work extensively uses unsupervised segmentation based on domain knowledge. Unsupervised image segmentation algorithms are broadly categorized into a)threshold-based, b) region-based, c) boundary-based, d) machine learning-based, and e) deep learning-based models. [11] utilized the threshold-based segmentation algorithms such as Huang [12], Kapur [13], and Otsu [14], to binarize the CT scans and generated a region adjacency graph (RAG) [15] to demarcate COVID-19 lesions from CT scans. Threshold-based methods are very simple in approach and implementation if provided with images bearing sufficient contrast. The absence of sufficient contrast leads to a downgrade in performance while considering differences in image attenuation. Region-based methods segment an area by assessing the homogeneity of the neighboring pixels. Widely practised region-based algorithms include region-growing [16]–[18], watershed [19], graph cuts [20], [21], and fuzzy connectedness [22]. Boundary-based methods are computationally intensive but provide highly accurate segmentation when the initial iteration is in the vicinity of the actual boundary. Boundary-based methods include snakes [23], active contours [24], and level sets [25], [26]. The boundary-based and region-based methods capture variations in attenuation but fail to segment regions of infection (such as consolidations and pleural effusion) near the lung boundary owing to a similar range of HU values. [27] proposed a threshold-based approach to segment lung regions from CT scans by processing the left and right lungs separately. Further

modifications are performed using morphological operators for fine-tuning the identification of irregular boundaries of the GGOs. [28] proposed a method to extract GGOs by modeling their intensity distributions and using the Markov random field model to improve boundary identification. [29] utilized textural features from CT image intensity parameters viz. entropy, contrast, roughness, and coarseness for segmentation of infection region. These morphological features are then used to demarcate GGOs from the image but were found to be insufficient in distinguishing the consolidations from pulmonary vessels. [30] proposed a novel approach for the detection of COVID-19 features using a 3D deep convolutional neural network (CNN) called "DeCoVNet" on CT volumes. The authors employed a combination of the activation maps generated by the DeCoVNet with a 3D connected component (3DCC) algorithm to identify lesions from the CT scans. Although the model demonstrated high recall, it suffered from a high incidence of false positives. This is seen to be the result of the formulation based on a black-box approach, thereby rendering the attention of the gradient to the infection region. The attention of activation maps cannot be guaranteed to accurately focus on the region of interest. This may lead to less accurate predictions and a higher rate of false positives.

Inspired by the intensity distribution adaptive model using MAP as proposed by Zhu *et al.* [28], we propose to use a three-mixture Gaussian mixture model using adaptive thresholding instead of a single Gaussian to extract the clinical features (defined by the radiologists in Section I)) in the proposed work.

B. CT SCAN CLASSIFICATION

Previous studies have made significant progress in identifying COVID-19 using CT scan images, which can be broadly divided into two categories: 3D CT scan-based classification and 2D CT scan-based classification. In 3D CT scan-based classification, a 3D CNN is trained on volumetric CT scans, and a probability score is evaluated for each scan [30]–[35]. Among these techniques, segmentation of the lung region using image preprocessing methods is applied before performing classification. Owing to the varying dimensionality of 3D CT scans, interpolation or truncation of the slices is applied to convert them to fixed dimensionality, which might lead to information loss [36].

In 2D CT scan-based classification, a 2D CNN is trained on individual slices, generating slice-level probability scores. Further, threshold-based [36], majority voting [37], [38], weighted average methods [39], and sequence models (such as recurrent neural network (RNN) [40] and bidirectional long short term memory (BiLSTM) [41]) are used to obtain patient-level COVID-19 classification. Threshold-based and majority voting methods create higher false negatives at regions where the traces of infections are not visible, as in most CT scan slices. Considering the difficulty in obtaining annotated CT scans, transfer learning methods have been extensively employed in COVID-19 classification. Some

transfer learning works [38], [39], [41]–[43] have explored different CNN models trained on ImageNet dataset [44] for classification tasks. Transfer learning methods reduce the training requirement for every dataset and provide discriminative features for classifying COVID-19, even while being applied to raw CT scan images.

[39] employed transfer learning in generating features from raw CT images using the EfficientNet-B5 model [45]. The slice-level scores were obtained using a shallow FFNN. Further, patient-level classification was performed using a weighted average method on the slice-level scores. With higher accuracy in classifying COVID-19, this model is used as the baseline model in the present work. Inspired by this technique [39], the present work attempts to enhance the performance in classifying COVID-19 in a CT scan by using preprocessing pipeline rather than raw CT scan images. We further provide an objective understanding of the severity of a COVID-19 patient by introducing a severity analysis module in the proposed architecture.

III. DATASETS

This work uses six publicly available online datasets, summarising the details in Table 1. CT scan is a volumetric scan consisting of 'n' slices. Each slice has a dimension of 512 x 512. The classification model is based on 2D CT scan images, which use the transfer learning from ImageNet database-trained models to generate discriminative features.

The dataset SPGC [46] provides slice-wise labels for COVID-19 and CAP cases by expert radiologists and is used for building a model for the classification tasks. Since the unavailability of slice-wise labels (but volume level labels are available) for CT scans in this dataset, a semi-supervised training method is performed for the classification task, explained in detail in Section IV-C. Hereafter, this dataset is named SPGC. Three publicly available datasets (named LDCT, LDCT-PCR [47] and Mosmed [48] for further use in the paper) are used for testing the robustness of the trained model. The LDCT and Mosmed datasets are collected from different geographical locations across the globe and are available in Digital Imaging and Communications in Medicine (DICOM) and NIfTI (Neuroimaging Informatics Technology Initiative).

Three publicly available datasets named Mosmed [48], MedSeg [49], and MedSeg_1 [50] are used for testing the accuracy of infection region segmentation from the COVID-19 patients. Expert radiologists demarcated the lesion region from the CT scans. The MedSeg dataset contains 100 CT slices collected from more than 40 patients with COVID-19, ranging from minimal to severe infections. This dataset which contains random slices from different patients, is inadequate to test the lesion segmentation because the severity analysis on COVID-19 patients uses the CT scan volume. The CT scan volume may or may not contain infections in all the slices. The Mosmed and MedSeg_1 datasets contain 50 and 9 patients' CT scan volumes collected from different geographical places. The infection ranges in patients of the

Mosmed dataset are less than 25%, while infection ranges vary from minimal to severe in the MedSeg_1 dataset. The severity of COVID-19 for a patient is performed by the CT severity score (CTSS), which indicates a severity scale from 0-25. The Mehta dataset contains fourteen COVID-19 patients' CT scans collected from Mehta multispeciality hospital (native hospital). This dataset includes a detailed diagnosis report for each CT scan, such as the CT severity score (CTSS), COVID-19 Reporting and Data System (CO-RADS) score, symptoms, and infection volume developed in the lungs. The patient's details are anonymized in the CT scan metadata and the diagnosis reports. This dataset is used to identify the correlation of the CTSS predicted by the proposed model.

IV. PROPOSED SYSTEM

The proposed system is designed to accomplish three key tasks: 1) A novel preprocessing pipeline to extract relevant clinical features, 2) A Semi-supervised method for classifying CT scans, and 3) A CT severity score is generated for patients by utilizing the preprocessing pipeline. The proposed system initially preprocesses the CT scan images, which are then fed to a model trained on the ImageNet dataset for extracting high-dimensional features. These features subsequently train a shallow feed-forward neural network (FFNN) to predict slice-level scores. A weighted average method is then applied to calculate the final score for the CT scans. A severity analysis module is also implemented to determine the CT severity score (CTSS). We evaluate the system's performance on six publicly available online datasets, the details of which are summarized in Table 1.

A. PREPROCESSING PIPELINE

Radiologists often use common clinical features such as GGO, consolidation, crazy paving pattern, halo sign, reverse halo sign, and pleural effusion to distinguish between COVID-19 and CAP from healthy individuals. The abnormalities in the CT scan exhibit attenuation variations with respect to the severity of the infection. Based on the observations in Figure. 1, an image processing pipeline is proposed with primary emphasis on the differences between the three classes of interest — COVID-19, CAP, and Normal. The proposed preprocessing pipeline is shown in Figure. 2. The preprocessing pipeline is developed in an unsupervised manner by integrating traditional image processing techniques in tandem with machine learning and deep learning models.

1) Stage-I: Lung mask generation

A CT scan is a volumetric scan consisting of multiple slices with a 512 x 512 pixels resolution. In a chest CT scan, tissues, heart, stomach, blood vessels, and bones possess higher attenuation (HU) values than air and lung areas. The preprocessing pipeline is applied on the HU scale of CT scan images. CT scans are mainly available in two formats: Neuroimaging Informatics Technology Initiative (NIfTI) and Digital Imaging and Communications in Medicine (DI-
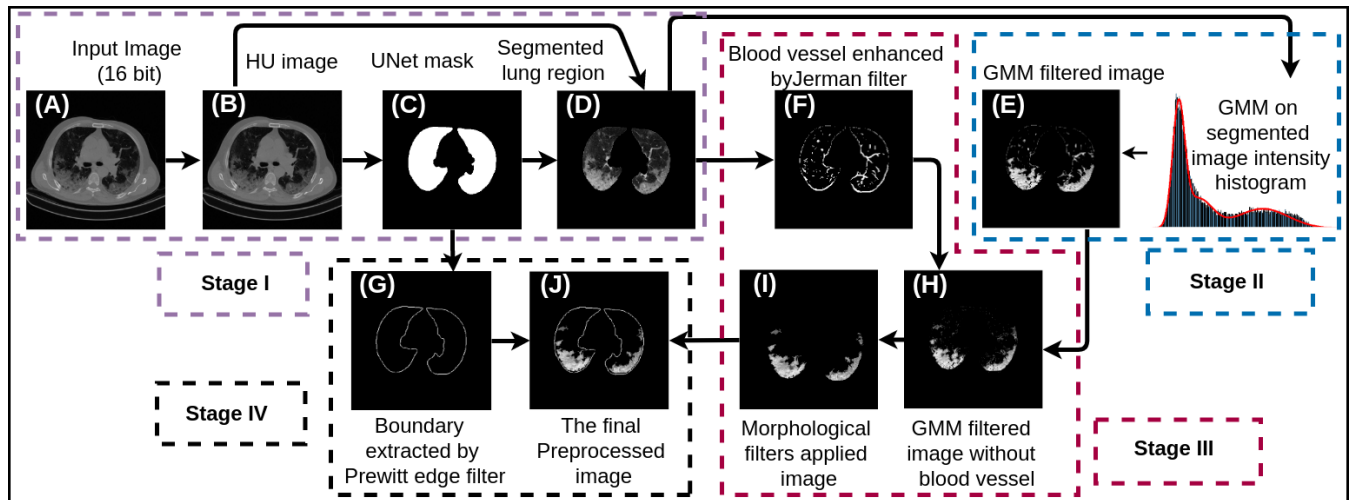
FIGURE 2: Proposed model architecture.

COM). The pixel intensity values of the CT scan slices are HU units in NIfTI format, whereas DICOM format CT scan slices are converted to HU scale images using a linear transformation. The preliminary task is to remove the unnecessary areas in the CT scan and identify the region of interest (lung region). The HU scale image is fed to the pretrained UNet model to extract the region of interest. The pretrained UNet [51] is an end-to-end, fully convolution neural network containing an encoder module that compresses input CT images using convolution and max-pooling operations into a fixed-length feature map. The decoder module (which provides a better spatial representation) upsamples these feature maps to the lung mask. The skip connections between the encoder and decoder modules enhance the semantic features for segmentation. The pre-trained UNet model segments the lung CT scan into three classes: the left lung region, the right lung region, and the background. The UNet model removes unnecessary objects such as bones, trachea, and organs and retains only the lung region. The lung region provides a clear segmented image, making identifying potential issues easier.

The initial and final slices of the CT scan typically contain structures such as bones, the trachea, the diaphragm, the heart, and the stomach, which are irrelevant to further analysis and hence not considered. A fixed threshold on the number of slices is impossible as the number of slices can vary between patients. Instead, a lung mask is used as a threshold to remove the initial and final slices. A criterion based on the degree of lung involvement is used to decide the number of slices for further analysis for every patient. This approach ensures that only the slices most relevant to COVID-19 analysis are included, improving the accuracy of the diagnosis.

2) Stage-II: GMM-based adaptive filter
The range of attenuation values (HU) of features like GGO, consolidation, and pleural effusion are discussed in Section I. The pixel intensity values of GGO, consolidation, and

pleural effusion regions can vary depending on the severity of the infection. Severe infection regions have higher pixel intensities than mild infection regions. This variability in the attenuation is captured by modeling the pixel's intensity histogram with a three-mixture GMM. Each mixture attempts to capture the GGO, consolidation, pleural effusion, and background patterns. This unsupervised clustering technique is applied to the preprocessing (Stage I) images. HU ranges for the clinical features were defined by the radiologist. Each mixture in GMM is represented by a mean μ, a standard deviation σ, and a threshold of $(\mu \pm 1.5 * \sigma)$ is considered for the pixel selection. The pixel intensities lie between the range defined above, and the proposed radiologist is considered for further analysis. While the GMM adaptive filter accurately segments the infection regions, the presence of blood vessels affects the severity analysis. To address this issue, the images are further subjected to vascular analysis.

3) Stage-III: Vascular enhancement and morphological filters
After Stage II of preprocessing, the image contains essential features of COVID-19 and CAP and details of bronchi and primary pulmonary vessels, which can resemble tubular structures. The vessels share similar HU values as that of consolidations and pleural effusions. A shape-based filter is needed to retain the consolidation and pleural effusion regions and remove the blood vessels. The Jerman blood vessel enhancement filter [52] is widely used in angiographic images to enhance blood vessels. This filter can identify the local structures in the images based on the shape (elongated or circular) by evaluating the sign and magnitude of the Eigenvalues of the image's Hessian matrix. The Jerman filter examines the largest to smallest eigenvalue ratio and assigns a probability score for each pixel to be a part of elongated local structures. The output from stage-I is subjected to this filter with a threshold of 0.75 to generate a binary mask for marking and removing the blood vessels. The blood vessels underneath the GGOs are also marked and removed from
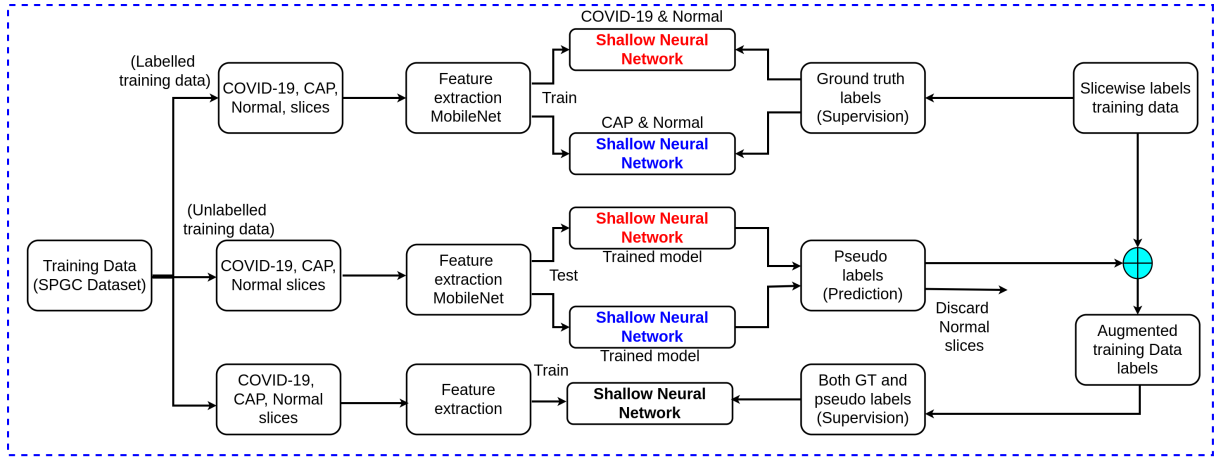
FIGURE 3: The proposed semi-supervised training method for the classifier

the image, creating holes in the GGO region. The flood fill algorithm [53] is then applied to homogenize these holes using the intensities of the pixels in the neighborhood. Following this, a morphological operation (dilation) is applied to enhance the infection boundary in the resulting image. The small white regions, which are generated due to the removal of the blood vessels, are removed using the area opening morphological method [54].

4) Stage-IV: Generate lung boundary and grayscale image
As proposed in [55], the Prewitt edge detector is employed to extract the lung boundary from the lung mask generated from Stage-I. This lung boundary aids in identifying the localization and distribution of the clinical features in the CT scan. Furthermore, the resultant HU scale image is converted to an 8-bit grayscale image using min-max normalization. This grayscale image is then resized to a three-channel image to match the input dimension of the model pre-trained on the ImageNet dataset.

B. FEATURE EXTRACTION BY IMAGENET PRE-TRAINED MODELS
Considering that the training dataset does not contain sufficient labeled 2D CT scan slices to train a deep CNN, state-of-the-art computer vision models trained on the ImageNet dataset are used to extract high-dimensional features of the preprocessed CT scan images. In this work, five deep CNN models, namely, MobileNet [56], ResNet-101 [57], ResNet-50 [57], EfficientNet-B5 [45], and EfficientNet-B1 [45] are used to extract discriminative features from the preprocessed CT scan images. The MobileNet [56] is a lightweight model that generates a 1024-dimensional feature vector for each 2D image. The novel depth-wise convolution in each layer reduces the number of parameters compared to a network with regular convolutions with the same depth. The ResNet models contain skip connections to allow for deep CNN models. These models perform better than models without skip connections. The deep CNN model with skip connections

can handle the vanishing gradient problem by following an alternate method for backpropagation. Tan *et al.* proposed a simple and effective way of scaling up CNN models to obtain better accuracy using compound coefficients. The technique allows the EfficientNet models to achieve higher accuracy with fewer parameters by focusing on the design choice of the model, thereby performing better than random scaling up of the model's width, depth, or the resolution of the feature map. Compared to other convolutional neural networks, the compound scaling method scales the input dimensions at a constant ratio to enhance classification accuracy. The ResNet and EfficientNet models generate a 2048-dimensional feature vector for each preprocessed CT scan image.

C. TRAINING
The SPGC training dataset [46] is insufficient to train a deep CNN model from scratch, so a semi-supervised approach is used to label the unannotated CT scan images in the training dataset. This method helps augment the final training dataset with more COVID-19 and CAP slices. Labeled images are preprocessed and fed to the pre-trained MobilNet to extract slice-wise feature maps. Each feature map is given to the global average pooling layer to make it a 1024-dimensional feature vector. Two shallow feed-forward neural networks (FFNN), with 1024 neurons and two neurons in the first and last layers, are trained on the 1024-dimensional extracted feature vectors. Among the two models, one FFNN is trained with an equal amount of features of COVID-19 and Normal CT scans, and the other FFNN is trained with an equal amount of features of CAP and Normal CT scans, named Mobilenet_COVID and MobileNet_CAP for future references, respectively. The SPGC training dataset contains unlabelled CT scan images; nevertheless, the entire CT scan patient-level labels are given. Unlabeled slice-wise CT scan images are preprocessed using the preprocessing pipeline developed in Section IV-A, and then fed to the pre-trained ImageNet models for feature extraction. The features corresponding to the COVID-19 slices are given to the MobileNet_COVID

model; similarly, the features of unlabelled CAP slices are given to the MobileNet_CAP model. CT scan slices with predicted labels as COVID-19 and CAP are further used to augment the final training dataset, and slices with the label as Normal are discarded. The pseudo-labeled COVID-19, CAP slices, and the labeled training dataset are used to train the final classifier. The final classifier consists of three layers: two dense layers (2048 and 1024 neurons) and a final layer (three neurons). Each neuron in the last layer gives a probability score for each class: CAP, COVID-19, and Normal. The final classifier is trained with an equal number of CT scan slices from COVID-19, CAP, and Normal categories. To avoid class imbalance, the final classifier model is trained with the same number of samples from each category. The complete pipeline of the training process is shown in Figure 3.

D. PATIENT-LEVEL CLASSIFICATION

The three-class classifier generates a probabilistic score for each CT scan slice. Since the CT scan is volumetric and has a dimension of (n * 512x512), patient-level annotation is preferable to slice-level classification. Thus, a weighted averaging method is applied to the probability score generated by the classifier. The CT scan volume (with 'n' images) is divided into three equal regions, and each part is associated with different weights (W_1, W_2, W_3), where W_1, W_2, W_3 are 0.7, 1, and 0.7 respectively. The middle region CT scan slices have a large lung region and contain more information. Hence, more weightage is assigned to the slices from this region. If the weighted sum of the probabilities of COVID-19 and CAP categories is greater than that of the Normal class, the CT scan is considered abnormal and classified as either COVID-19 or CAP based on the scores. The LDCT, Mosmed, and LDCT-PCR datasets do not contain any CAP CT scans. The CT scan slices predicted as CAP class are considered to belong to the COVID-19 category for these datasets. Consider the predicted score for the i^{th} slice is P_i. The patient label is calculated as:

$$FS = \max \left(\sum_{i=1}^{n/3} P_i * W_1 + \sum_{i=n/3}^{2n/3} P_i * W_2 + \sum_{i=2n/3}^{n} P_i * W_3 \right) \quad (1)$$

V. RESULTS AND DISCUSSIONS

The proposed pipeline is thoroughly evaluated on several aspects, including segmentation and classification accuracy. Additionally, ablation studies are conducted to evaluate the image preprocessing pipeline's significance in the proposed method's overall performance.

A. EVALUATION METRICS

The proposed pipeline is evaluated using several widely used metrics for segmentation and classification. For segmentation, the pipeline is evaluated using metrics such as Dice score (Dice), sensitivity (Sen.), specificity (Spec.), precision (Prec.), and mean absolute error (MAE). Dice score computes the error in segmentation by computing the overlap between annotated and predicted areas. Precision denotes

the number of accurate white pixel predictions out of the overall white pixel predictions by the model. At the same time, specificity represents the number of correct predictions of black pixels from the total black pixels in the ground truth image. The MAE finds the average absolute difference between the predicted and annotated binary masks and quantifies the quality of the ROI predicted by the model. A lower MAE value indicates a better segmentation by the model. For classification, the pipeline is evaluated using metrics such as sensitivity, precision, and F1 score. Sensitivity denotes the number of CT scan volumes correctly predicted to the ground truth. It determines how well the model can discriminate patients' input CT scans with respect to the ground truth. Precision denotes the number of CT scans correctly predicted by the model's overall prediction. F1 score can be quantified as the harmonic mean of sensitivity and precision. The F1 score is similar to the Dice score in the segmentation task. The proposed pipeline also predicts the CT severity score (CTSS) and evaluates the prediction using the Pearson correlation coefficient and cosine similarity. The Pearson correlation measures the ratio between the covariance of the predicted and ground truth scores and their standard deviations, showing the trend between the proposed and ground truth CTSS. Cosine similarity calculates the angle between the proposed and ground truth CTSS vectors, measuring the prediction accuracy relative to the ground truth. The Pearson correlation coefficient and cosine similarity range from -1 to 1, with positive correlations and higher similarities indicated by values greater than zero.

B. SEGMENTATION RESULTS

The output image from Stage-III of the preprocessing pipeline (image (I) in Figure. 2) is used as the infection mask for the CT scan slices. This preprocessed image (I) is generated by adaptively filtering the clinical features using a three-mixture GMM; then, blood vessels are removed by the Jerman filter, and the extracted features are fine-tuned with the morphological filters. In a previous study, [58] evaluated five baseline models based on the different variants of the UNet [59] architecture with the MedSeg dataset (48 slices). The slices in the MedSeg dataset contain a wide range of infections, but non-infectious slices are absent. The infection segmentation results of the proposed model are evaluated with the same dataset with the same CT scan slices, and the results are summarized in Table 2. The infection segmentation of sample images from the MedSeg dataset is shown in Figure 4. The proposed model achieves better results in dice score, specificity, and MAE to the baseline, and also the models proposed by [58], namely, InfNet and SemiInfNet. Since the MedSeg dataset is a random collection of slices from different patients, they conducted infection segmentation with another dataset (MedSeg_1) comprising nine patients with 638 slices (285 non-infected and 353 infected slices). The results are summarized in Table 2. Again, the preprocessing pipeline improves performance, especially when the non-infected slices are identified more accurately

TABLE 2: The infection segmentation results with the Med-Seg Dataset (48 CT scan slices).

Model	Dice	Sen.	Spec.	MAE
UNet [59]	0.439	0.534	0.858	0.186
Attention UNet [60]	0.583	0.637	0.921	0.112
Gated UNet [61]	0.623	0.658	0.926	0.102
Dense UNet [62]	0.515	0.594	0.840	0.184
UNet++ [63]	0.581	0.672	0.902	0.120
InfNet [58]	0.682	0.692	0.943	0.082
SemiInfNet [58]	0.739	0.725	0.960	0.064
Our Method	0.673	0.678	0.9852	0.0356

TABLE 3: The infection segmentation results with the Med-Seg_1 Dataset (nine patients' real CT scan volumes).

Model	Dice	Sen.	Spec.	Prec.	MAE
UNet [59]	0.308	0.678	0.836	0.265	0.214
Attention UNet [60]	0.466	0.723	0.930	0.390	0.095
Gated UNet [61]	0.447	0.674	0.956	0.375	0.066
Dense UNet [62]	0.410	0.607	0.977	0.415	0.167
UNet++ [63]	0.444	0.877	0.929	0.369	0.106
InfNet [58]	0.579	0.870	0.974	0.500	0.047
SemiInfNet [58]	0.597	0.865	0.977	0.515	0.033
Our Method	0.726	0.775	0.996	0.711	0.0057

than the baseline, InfNet, and SemiInfNet models. Moreover, the proposed pipeline can detect even minor infections based on the statistical properties of the attenuation in the infection regions.

[64] used a nnUNet (no new UNet) based baseline model for infection segmentation. The baseline model is mainly trained with two datasets, out-of-domain, and in-domain. The out-of-domain datasets include Medical Segmentation Decathlon (MSD) lung tumor segmentation (MICCAI 2018 challenge dataset), StructSeg lung cancer volume segmentation (MICCAI 2019 challenge dataset), and NSCLC pleural effusion segmentation [65]. The in-domain datasets contain 20 COVID-19 CT scans with an infection range of 0.01% - 59% [50]. Three baseline results are provided by training the nnUNet on each out-of-domain dataset (Task 1). Two baseline results are provided by training the nnUNet with in-domain and out-of-domain datasets (Task 2). Out of the two benchmark results, one model is trained for lung and infection segmentation (Union), and the other is prepared only for infection segmentation (separate). The Mosmed dataset [48] is used as blind test data for the baseline models, and the baseline results, along with our proposed model results, are summarized in Table 4. The proposed model provides better dice score, sensitivity, and precision results. Since the infection region is less than 25% in the CT scan, it is observed that the proposed model can identify small infection regions from the CT scan slices and outperform the baseline models.

C. SEVERITY ANALYSIS

Determining the severity of infection in a COVID-19 patient is essential in determining the appropriate course of treatment. The CT severity score (CTSS) is the standard metric

TABLE 4: The infection segmentation results with the Mosmed Dataset (50 patients CT scan volume).

	Task 1			Task 2		Proposed method
	MSD	Struct seg	NSCLC	Separate	Union	Infection Mask
Dice score	0.392	0.443	0.301	0.588	0.482	0.752
Specificity	1.00	1.00	1.00	0.999	0.999	0.998
Sensitivity	0.364	0.422	0.249	0.575	0.601	0.757
Precision	0.614	0.607	0.614	0.679	0.577	0.812

used for severity analysis, and it ranges from 0 (no involvement) to 25 (maximum involvement). This score is evaluated by segmenting the right lung into three lobes and the left lung into two, respectively. The percentage of infection in each lobe is calculated based on the infection rate, as shown in Figure. 5. In the present work, we evaluate the severity analysis by projecting the volumetric CT scan into axial slices. An empirical threshold on the lung mask (C in (Figure. 2)) is used to select CT scan slices for the estimation of CTSS. The pre-trained UNet [51] categorizes each CT scan image into right, left, and background lung masks. The right and left lungs are segmented as briefed before. The infection mask (I in (Figure. 2)) is projected (element by element pixel intensity multiplication) on either lobe, and the percentage of infection in each lobe is calculated. A score of 0 to 5 is assigned to each lobe based on the percentage of infection. This process is depicted in Figure. 5. The average (across the images) lobe score is then acquired to yield the total CTSS, the sum of the averaged lobe scores.

The Mehta dataset (14 patients) and a subset from the SPGC dataset (36 patients) are used for the CTSS analysis. An expert radiologist provides the CTSS for the subset of the SPGC dataset. The CTSS predicted by the proposed method and radiologist-predicted CTSS is compared in Figure. 6. A strong correlation of 0.82 between the CTSS predicted by the proposed method and the ground truth is observed, as well as a high cosine similarity of 0.97 between the proposed and ground truth CTSS score vectors of 50 patients. These results indicate that the CTSS predicted by the proposed method follows the same trend as the ground truth provided by the radiologist.

D. CLASSIFICATION RESULTS

Features are extracted using pre-trained networks, as stated earlier. A shallow feed-forward neural network is trained using a semi-supervised learning method. The classification results for the ablation study, as well as the slice-level and patient-level, are summarised in Table 5, Table 6, Table 7, and Figure. 7 respectively.

1) Ablation Study

The proposed work differs significantly from previous work in its image processing pipeline. A series of ablation studies are performed to evaluate each stage's contribution in the preprocessing pipeline to classification accuracy. These studies include additional preprocessing steps in the pipeline. The different intermediate preprocessed images used in the ablation study are original image (OI), Stage II with GMM

| CT scan | U-Net | U-Net++ | Inf-Net | Semi-Inf-Net | Our Method | Ground Truth |

FIGURE 4: The region of infection extracted from the MedSeg dataset using different baseline models and the proposed model.

TABLE 5: An ablation study on the different preprocessing stages of the SPGC test dataset with four different features.

Category	Feature Extractor	OI	GMM	GMM + MO	GMM + B	GMM + MO + B	GMM+ MO + JF	GMM + MO + JF + B
	ResNet-50	0.71	0.78	**0.83**	0.73	0.81	0.68	0.77
COVID-19	ResNet-101	0.46	0.83	0.78	0.80	**0.87**	0.81	0.71
	EfficientNet-B1	0.72	**0.81**	0.77	**0.81**	0.81	0.79	0.76
	EfficientNet-B5	0.77	0.71	0.73	0.80	**0.83**	0.75	0.74
	ResNet-50	0.64	0.78	**0.89**	0.51	0.81	0.29	0.70
CAP	ResNet-101	0.65	**0.81**	**0.81**	**0.81**	0.78	0.80	0.50
	EfficientNet-B1	0.72	0.74	0.74	**0.85**	0.78	0.76	0.76
	EfficientNet-B5	**0.85**	0.57	0.73	0.79	0.78	0.68	0.73
	ResNet-50	0.64	0.90	0.86	0.85	**0.93**	0.76	0.84
Normal	ResNet-101	0.68	0.87	0.83	0.85	**0.92**	0.89	0.79
	EfficientNet-B1	0.77	0.88	0.84	0.85	**0.92**	0.87	0.82
	EfficientNet-B5	0.83	0.79	0.82	0.87	**0.93**	0.88	0.83

CTSS	
% of infection	Score
0	0
<= 5	1
5 - 25	2
25 - 50	3
50 - 75	4
>= 75	5

1) Divide the right lung into three lobes and left lung into two lobes
2) Find the percentage of infection in each lobes and assign a score based on the Table (CTSS)

FIGURE 5: A lung mask generated by UNet and an infection mask is used for estimating the severity score for a CT scan slice.

(GMM), Stage II with morphological operations (GMM + MO), Stage-II with lung region boundary (GMM + B), Stage II with morphological operations and lung boundary (GMM + MO + B), Stage-III output image (GMM + MO + JF) and

TABLE 6: The slice-level classification of the LDCT dataset with different features.

Feature	Sensitivity		Precision		F1 Score	
Extractor	COVID-19	Normal	COVID-19	Normal	COVID-19	Normal
ResNet-50	0.90	0.86	0.88	0.89	0.89	0.88
ResNet-101	0.82	0.89	0.89	0.82	0.86	0.86
EfficientNet-B5	0.88	0.88	0.89	0.87	0.88	0.87
EfficientNet-B1	0.80	0.93	0.92	0.81	0.86	0.86

Stage-IV output image (GMM + MO + JF + B). Again, we use the pre-trained ImageNet model for feature extraction and train a shallow FFNN to classify the images. Ablation studies are performed on the SPGC dataset, and the results are summarized in Table 5.

As seen in Table 5 that GMM-filtered images with morphological operations and lung region boundary (GMM + MO + B) have consistent and uniform F1 scores across all the models pre-trained on the ImageNet dataset. GMM + MO images performed well with the ResNet-50 model, but the other models with the same image produced less classification accuracy. It is observed from the ablation study

FIGURE 6: CTSS for 14 patients (Mehta) and 36 patients (a subset of SPGC datasets).


FIGURE 7: Patient-wise classification results for the baseline and proposed models with SPGC dataset.

TABLE 7: A comparative study of different features with different test datasets used in the experiments.

Feature Extractor	LDCT		RT-PCR		Mosmed	
	COVID-19	Normal	COVID-19	Normal	COVID-19	Normal
ResNet-50	0.88	0.82	0.89	0.82	0.72	0.54
ResNet-101	0.89	0.82	0.90	0.83	0.77	0.58
EfficientNet-B1	0.89	0.82	0.94	0.89	0.79	0.59
EfficientNet-B5	0.88	0.82	0.89	0.82	0.75	0.59

TABLE 8: Different severity classes and respective classification results.

Category	Prediction		Number of patients	Percentage of infection
	COVID-19	Normal		
CT-0	30	224	254	$<= 5$
CT-1	412	272	685	5 - 25
CT-2	110	15	125	25 - 50
CT-3	37	8	44	50 - 75

that blood vessel removal from the CT scan images degrades the classification performance. The blood vessel removal algorithm introduces an information loss in the CT scan images for the classification task. It is observed from the literature that blood vessel enlargement can happen in the later stages of COVID-19 infection [9], [10]. From the ablation study on the SPGC dataset, the (GMM + MO + B) image in preprocessing pipeline appears to contribute most to accuracy. The F1 score is used as the evaluation metric for the ablation study since the test dataset has a class imbalance with the different categories of CT scans.

2) Slice-level classification

An experiment is conducted with the LDCT dataset, which provides slice-level labels for 160 patients. The (GMM + MO + B) images from the preprocessing pipeline are used to train and test the models. Four different models pretrained on the ImageNet dataset extract the features. Each feature is trained with different shallow FFNN models. It is observed from Table 6 that all the features are suitable for discriminating between the classes. Since the CT scan is volumetric data, patient-level classification is preferable to slice-level identification.

3) Patient-level classification

The proposed three-class classifier model is evaluated with the SPGC test dataset and summarised the results along with the baseline model in Fig. 7. The baseline model [39] used the EfficientNet-B5 as the feature extractor from the raw CT scan images, and a shallow FFNN is trained on the extracted features. The proposed model with the EfficientNet-B5 as a feature extractor shows an average improvement of 6% F1 score for the classification task. Similarly, other features are

also performing better than the baseline model. Figure. 7 shows that the preprocessing pipeline significantly improves the classification task.

The LDCT, LDCT-PCR, and Mosmed datasets consist of two classes of CT scans: COVID-19 and Normal. The proposed model achieves better results with both LDCT and LDCT-PCR datasets. However, the Mosmed dataset provides comparatively less accuracy than the LDCT dataset. The Mosmed dataset divided the patients into four categories: CT-0, CT-1, CT-2, and CT-3, based on the volume and severity of infection in the CT scans. The CT-0 category consists of COVID-19 cases having less than 5% of infections in CT scans. Since the ground truth was not given for each CT scan, CT-0 corresponds to the Normal class, and other categories CT-1, CT-2, and CT-3 to the COVID-19 class. A category-wise experiment is conducted with ResNet-101 generated features, and the classification results are given in Table 7. Table 8 shows that patients in CT-0 are classified as Normal with less misclassification. The CT-2 and CT-3 categories have significant lung infections in the CT scans and are classified as COVID-19 with minor misclassification. Our proposed model can identify the infected slices, but the number of infected slices is very small compared to the overall slices in the CT scan volume in the case of CT-0 and CT-1. The classifier can perform well for patients with greater than 25% infections developed in the lungs.

VI. CONCLUSION

The proposed work facilitates the identification of COVID-19 infection and the prediction of the CTSS for assessing the severity of infections developed in the lungs. The task is challenging since the slice level has limited labeled data. Deep learning models use raw images for classification but require much data. On the other hand, pre-trained ImageNet models enable classification with small amounts of data but are not trained solely for the purpose of COVID-19 detection. Hence, an image preprocessing module is introduced to enhance the lung's infection regions, resulting in better classification. The domain expertise of radiologists helps to identify the lung infection regions during the preprocessing stages. The preprocessing module is unsupervised, which helps to extract the clinical features from the CT scans per the radiologist's suggestions. These preprocessed images are used to train the COVID-19 classifier in a semi-supervised way. The semi-supervised approach is adopted along with transfer learning to overcome the limitation of small amounts of training data. The importance of the preprocessing module is established as an absolute improvement of 6% is observed in the F1 score compared to the baseline model. The CTSS achieves a correlation of 82% with the manual assessment of radiologists. Thus, the work forms the framework with a possibility of further extension to a real-time COVID-19 support tool for clinicians.

REFERENCES

[1] M. Dramé et al., "Should rt-pcr be considered a gold standard in the diagnosis of covid-19?," Journal of medical virology, 2020.

[2] E. Benmalek et al., "Comparing ct scan and chest x-ray imaging for covid-19 diagnosis," Biomedical Engineering Advances, vol. 1, p. 100003, 2021.

[3] A. Nambu et al., "Imaging of community-acquired pneumonia: roles of imaging examinations, imaging diagnosis of specific pathogens and discrimination from noninfectious diseases," World Journal of Radiology, vol. 6, no. 10, p. 779, 2014.

[4] C. Hani et al., "Covid-19 pneumonia: a review of typical ct findings and differential diagnosis," Diagnostic and interventional imaging, vol. 101, no. 5, pp. 263–268, 2020.

[5] M. I. Mustafa et al., "Cytokine storm in covid-19 patients, its impact on organs and potential treatment by qty code-designed detergent-free chemokine receptors," Mediators of Inflammation, vol. 2020, 2020.

[6] T. C. Kwee and R. M. Kwee, "Chest ct in covid-19: what the radiologist needs to know," Radiographics, vol. 40, no. 7, pp. 1848–1865, 2020.

[7] M. Chung et al., "Ct imaging features of 2019 novel coronavirus (2019-ncov)," Radiology, vol. 295, no. 1, pp. 202–207, 2020.

[8] L. Lin et al., "Ct manifestations of coronavirus disease (covid-19) pneumonia and influenza virus pneumonia: A comparative study," American Journal of Roentgenology, vol. 216, no. 1, pp. 71–79, 2021.

[9] A. Esposito et al., "Chest ct–derived pulmonary artery enlargement at the admission predicts overall survival in covid-19 patients: insight from 1461 consecutive patients in italy," European radiology, vol. 31, no. 6, pp. 4031–4041, 2021.

[10] R. Aoki et al., "Pulmonary vascular enlargement and lesion extent on computed tomography are correlated with covid-19 disease severity," Japanese journal of radiology, vol. 39, no. 5, pp. 451–458, 2021.

[11] A. Khan et al., "A novel threshold-based segmentation method for quantification of covid-19 lung abnormalities," Signal, Image and Video Processing, pp. 1–8, 2022.

[12] L.-K. Huang and M.-J. J. Wang, "Image thresholding by minimizing the measures of fuzziness," Pattern recognition, vol. 28, no. 1, pp. 41–51, 1995.

[13] J. N. Kapur et al., "A new method for gray-level picture thresholding using the entropy of the histogram," Computer vision, graphics, and image processing, vol. 29, no. 3, pp. 273–285, 1985.

[14] N. Otsu, "A threshold selection method from gray-level histograms," IEEE transactions on systems, man, and cybernetics, vol. 9, no. 1, pp. 62–66, 1979.

[15] A. Trémeau and P. Colantoni, "Regions adjacency graph applied to colour image segmentation," IEEE Transactions on image processing, vol. 9, no. 4, pp. 735–744, 2000.

[16] M. S. Brown et al., "Method for segmenting chest ct image data using an anatomical model: preliminary results," IEEE transactions on medical imaging, vol. 16, no. 6, pp. 828–839, 1997.

[17] R. Adams and L. Bischof, "Seeded region growing," IEEE Transactions on pattern analysis and machine intelligence, vol. 16, no. 6, pp. 641–647, 1994.

[18] J. Dehmeshki et al., "Segmentation of pulmonary nodules in thoracic ct scans: A region growing approach," IEEE Transactions on Medical Imaging, vol. 27, no. 4, pp. 467–480, 2008.

[19] A. P. Mangan and R. T. Whitaker, "Partitioning 3d surface meshes using watershed segmentation," IEEE Transactions on Visualization and Computer Graphics, vol. 5, no. 4, pp. 308–321, 1999.

[20] Y. Boykov and M.-P. Jolly, "Interactive organ segmentation using graph cuts," in International conference on medical image computing and computer-assisted intervention, pp. 276–286, Springer, 2000.

[21] G. Li et al., "Automatic liver segmentation based on shape constraints and deformable graph cut in ct images," IEEE Transactions on Image Processing, vol. 24, no. 12, pp. 5315–5329, 2015.

[22] J. K. Udupa and S. Samarasekera, "Fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation," Graphical models and image processing, vol. 58, no. 3, pp. 246–261, 1996.

[23] C. Xu and J. L. Prince, "Snakes, shapes, and gradient vector flow," IEEE Transactions on image processing, vol. 7, no. 3, pp. 359–369, 1998.

[24] S. C. Zhu and A. Yuille, "Region competition: Unifying snakes, region growing, and bayes/mdl for multiband image segmentation," IEEE transactions on pattern analysis and machine intelligence, vol. 18, no. 9, pp. 884–900, 1996.

[25] L. A. Vese and T. F. Chan, "A multiphase level set framework for image segmentation using the mumford and shah model," International journal of computer vision, vol. 50, no. 3, pp. 271–293, 2002.

[26] J. S. Schildkraut et al., "Level-set segmentation of pulmonary nodules in radiographs using a ct prior," in Medical Imaging 2009: Image Processing, vol. 7259, pp. 1103–1116, SPIE, 2009.

[27] S. Hu et al., "Automatic lung segmentation for accurate quantitation of volumetric x-ray ct images," IEEE transactions on medical imaging, vol. 20, no. 6, pp. 490–498, 2001.

[28] Y. Zhu et al., "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.

[29] B. Ye et al., "Severity assessment of covid-19 based on feature extraction and v-descriptors," IEEE Transactions on Industrial Informatics, 2021.

[30] X. Wang et al., "A weakly-supervised framework for covid-19 classification and lesion localization from chest ct," IEEE Transactions on Medical Imaging, vol. 39, no. 8, pp. 2615–2625, 2020.

[31] X. Wu et al., "Covid-al: The diagnosis of covid-19 with deep active learning," Medical Image Analysis, vol. 68, p. 101913, 2021.

[32] S. Xue and C. Abhayaratne, "Covid-19 diagnostic using 3d deep transfer learning for classification of volumetric computerised tomography chest scans," in ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 8573–8577, 2021.

[33] J. Wang et al., "Prior-attention residual learning for more discriminative covid-19 screening in ct images," IEEE Transactions on Medical Imaging, vol. 39, no. 8, pp. 2572–2583, 2020.

[34] T. Javaheri et al., "Covidctnet: an open-source deep learning approach to diagnose covid-19 using small cohort of ct images," NPJ digital medicine, vol. 4, no. 1, pp. 1–10, 2021.

[35] Z. Han et al., "Accurate screening of covid-19 using attention-based deep 3d multiple instance learning," IEEE Transactions on Medical Imaging, vol. 39, no. 8, pp. 2584–2594, 2020.

[36] T. Anwar, "Covid19 diagnosis using automl from 3d ct scans," in Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 503–507, 2021.

[37] M. N. Teli, "Telinet: Classifying ct scan images for covid-19 diagnosis," in Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 496–502, 2021.

[38] P. Afshar et al., "Covid-ct-md, covid-19 computed tomography scan dataset applicable in machine learning and deep learning," Scientific Data, vol. 8, no. 1, pp. 1–8, 2021.

[39] S. Chaudhary et al., "Detecting covid-19 and community acquired pneumonia using chest ct scan images with deep learning," in ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 8583–8587, IEEE, 2021.

[40] D. Kollias et al., "Mia-cov19d: Covid-19 detection through 3-d chest ct image analysis," in Proceedings of the IEEE/CVF International Conference on Computer Vision, pp. 537–544, 2021.

[41] P. Garg et al., "Multi-scale residual network for covid-19 diagnosis using ct-scans," in ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), pp. 8558–8562, IEEE, 2021.

[42] T. D. Pham, "A comprehensive study on classification of covid-19 on computed tomography with pretrained convolutional neural networks," Scientific reports, vol. 10, no. 1, pp. 1–8, 2020.

[43] M. Singh et al., "Transfer learning–based ensemble support vector machine model for automated covid-19 detection using lung computerized tomography scan data," Medical & biological engineering & computing, vol. 59, no. 4, pp. 825–839, 2021.

[44] J. Deng et al., "Imagenet: A large-scale hierarchical image database," in 2009 IEEE conference on computer vision and pattern recognition, pp. 248–255, Ieee, 2009.

[45] M. Tan and Q. Le, "Efficientnet: Rethinking model scaling for convolutional neural networks," in International conference on machine learning, pp. 6105–6114, PMLR, 2019.

[46] S. Heidarian et al., "SPGC-COVID Dataset." {https://figshare.com/articles/dataset/SPGC-COVID_Dataset/166-32397/1},DOI={10.6084/m9.figshare.16632397.v1}, 2021.

[47] P. Afshar et al., "COVID-19 Low-Dose and Ultra-Low-Dose CT Scans." https://dx.doi.org/10.21227/sed8-6r15, 2021. IEEE Dataport.

[48] S. P. Morozov et al., "Mosmeddata: data set of 1110 chest ct scans performed during the covid-19 epidemic," Digital Diagnostics, vol. 1, no. 1, pp. 49–59, 2020.

[49] S. T. MedSeg, Jenssen HB, "MedSeg Covid Dataset." https://doi.org/10.6084/m9.figshare.13521488.v2, 2021. figshare Dataset, January 2021.

[50] M. Jun et al., "COVID-19 CT Lung and Infection Segmentation Dataset." {https://doi.org/10.5281/zenodo.3757476},DOI={10.5281/zenodo.3757476}, 2020. April, 2020.

[51] J. Hofmanninger et al., "Automatic lung segmentation in routine imaging is primarily a data diversity problem, not a methodology problem," European Radiology Experimental, vol. 4, no. 1, pp. 1–13, 2020.

[52] T. Jerman et al., "Enhancement of vascular structures in 3d and 2d angiographic images," IEEE transactions on medical imaging, vol. 35, no. 9, pp. 2107–2118, 2016.

[53] P. Soille, Morphological image analysis: principles and applications. Springer Science & Business Media, 2013.

[54] L. Vincent, "Morphological area openings and closings for grey-scale images," in Shape in Picture, pp. 197–208, Springer, 1994.

[55] J. M. Prewitt et al., "Object enhancement and extraction," Picture processing and Psychopictorics, vol. 10, no. 1, pp. 15–19, 1970.

[56] A. G. Howard et al., "Mobilenets: Efficient convolutional neural networks for mobile vision applications," arXiv preprint arXiv:1704.04861, 2017.

[57] K. He et al., "Deep residual learning for image recognition," in Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 770–778, 2016.

[58] D.-P. Fan et al., "Inf-net: Automatic covid-19 lung infection segmentation from ct images," IEEE Transactions on Medical Imaging, vol. 39, no. 8, pp. 2626–2637, 2020.

[59] O. Ronneberger et al., "U-net: Convolutional networks for biomedical image segmentation," in International Conference on Medical image computing and computer-assisted intervention, pp. 234–241, Springer, 2015.

[60] O. Oktay et al., "Attention u-net: Learning where to look for the pancreas," arXiv preprint arXiv:1804.03999, 2018.

[61] J. Schlemper et al., "Attention gated networks: Learning to leverage salient regions in medical images," Medical image analysis, vol. 53, pp. 197–207, 2019.

[62] X. Li et al., "H-denseunet: hybrid densely connected unet for liver and tumor segmentation from ct volumes," IEEE transactions on medical imaging, vol. 37, no. 12, pp. 2663–2674, 2018.

[63] Z. Zhou et al., "Unet++: A nested u-net architecture for medical image segmentation," in Deep learning in medical image analysis and multimodal learning for clinical decision support, pp. 3–11, Springer, 2018.

[64] J. Ma et al., "Toward data-efficient learning: A benchmark for covid-19 ct lung and infection segmentation," Medical physics, vol. 48, no. 3, pp. 1197–1210, 2021.

[65] K. J. Kiser et al., "Plethora: Pleural effusion and thoracic cavity segmentations in diseased lungs for benchmarking chest ct processing pipelines," Medical physics, vol. 47, no. 11, pp. 5941–5952, 2020.

ANAND THYAGACHANDRAN (Student Member, IEEE) received the bachelor's degree in computer science and engineering from Cochin University for Science and Technology, Kerala, India, in 2011. He received a master's degree in information technology from the University of Hyderabad, Telangana, India, in 2016. From 2016 to 2018, he worked as an Assistant Professor with the Computer Science Department, Vignan's Foundation for Science, Technology & Research, Andhra Pradesh, India. He is pursuing a Ph.D. with the Computer Science Department, IIT Madras, Chennai, India. Since 2019, he has been serving as a Teaching Assistant with IIT Madras. His research interests include medical image processing, computational brain research, and other areas of machine learning and signal processing.

AATHIRA BALACHANDRAN received the bachelor of medicine, bachelor of surgery degree from Government Medical College Trivandrum, Kerala, India, in 2011. She received a master's in Radiodiagnosis from the Government Thirumala Devaswom Medical College, Alappuzha, Kerala, India 2016. From 2017 to 2018, she worked as a consultant radiologist with the Radio Diagnosis Department, Government Medical College Kollam, Kerala, India. From 2018 to 2020, she worked as a consultant radiologist with Muthoot Hospitals Kozhencherry, and Metro Scans & Laboratory in Kottarakara, Kollam, Kerala, India. She is currently an Assistant Professor at the Radio Diagnosis Department, Government Medical College Kollam. Her research interests include diagnostic values of ultrasound in resource-limited settings in the area of head and neck malignancies, gynecological malignancies, and antenatal ultrasound scans.

HEMA A MURTHY (Senior Member, IEEE) received the bachelor's degree from Osmania University, Hyderabad, India, in 1980, the master's degree from McMaster University, Hamilton, ON, Canada, in 1986, and the Ph.D. degree from IIT Madras, Chennai, India, in 1992. From 1980 to 1983, she worked as a Scientific Officer (SO/SC) with the Speech and Digital Systems Group, TIFR, Mumbai. She is currently a Professor at the Department of Computer Science and Engineering, IIT Madras. She has over 39 journal publications, two book publications, and over 220 articles. Her research interests include speech processing, computer networks, music information retrieval, computational brain research, and other areas of machine learning and signal processing. Dr. Murthy was a recipient of the Manthan Award, in 2012, and the IBM Faculty Award, in 2006. Her awards include being elected as a Fellow of the Indian National Academy of Engineering, in November 2017, and to the board of the International Speech Communication Association for the duration of 2017–2021.

• • •