MIA-COV19D: COVID-19 Detection through 3-D Chest CT Image Analysis

Dimitrios Kollias
University of Greenwich, UK
D.Kollias@greenwich.ac.uk
dkollias1@gmail.com

Anastasios Arsenos
National Technical University
Athens, Greece

Levon Soukissian
GRNET National Infrastructures
Research & Technology, Greece

Stefanos Kollias
National Technical University
Athens, Greece
stefanos@cs.ntua.gr

Abstract

Early and reliable COVID-19 diagnosis based on chest 3-D CT scans can assist medical specialists in vital circumstances. Deep learning methodologies constitute a main approach for chest CT scan analysis and disease prediction. However, large annotated databases are necessary for developing deep learning models that are able to provide COVID-19 diagnosis across various medical environments in different countries. Due to privacy issues, publicly available COVID-19 CT datasets are highly difficult to obtain, which hinders the research and development of AI-enabled diagnosis methods of COVID-19 based on CT scans.

In this paper we present the COV19-CT-DB database which is annotated for COVID-19, consisting of about 5,000 3-D CT scans, We have split the database in training, validation and test datasets. The former two datasets can be used for training and validation of machine learning models, while the latter will be used for evaluation of the developed models. We present a deep learning approach, based on a CNN-RNN network and report its performance on the COVID19-CT-DB database. Moreover, we present the results of all main techniques that were developed and used in the ICCV COV19D Competition.

1. Introduction

The Coronavirus Disease 2019 SARS-CoV-2 (COVID-19) has become a global pandemic with an exponential growth and mortality rate. The virus is harbored most commonly with little or no symptoms, but can also lead to a rapidly progressive and often fatal pneumonia [14, 12, 3].

It has become important to detect affected people as early as possible and isolate them to stop further spreading of the virus. Various methods have been proposed to diagnose COVID-19, containing a variety of medical imaging techniques, blood tests and PCR.

COVID-19 pandemic has a very severe impact on the respiratory as well as other systems of the human body. Thus, medical imaging features of chest radiography is found to be useful for rapid COVID-19 detection. The imaging features of the chest can be obtained through medical imaging modalities like CT (Computed Tomography) scans. CT images can be used for precise COVID-19 detection [1].

They provide: a) 3-D view formation of organs; CT scans provide a more detailed overview of the internal structure of lung parenchyma due to lack of overlapping tissues, b) convenient examination of disease and its location; CTs provide a window into pathophysiology that could shed light on several stages of disease detection and evolution. Radiologists report COVID-19 patterns of infection with typical features including ground glass opacities in the lung periphery, rounded opacities, enlarged intra-infiltrate vessels, and later more consolidations that are a sign of progressing critical illness.

At the time of CT scan recording, several slices are captured from each person suspected of COVID-19. The large volume of CT scan images calls for a high workload on physicians and radiologists to diagnose COVID-19. Taking this into account and also the rapid increase in number of new and suspected COVID-19 cases, it is evident that there is a need for using machine and deep learning for detecting COVID-19 in CT scans.

Such approaches require data to be trained on. Therefore, a few databases have been developed consisting of CT scans. However, new data sets with large numbers of 3-D

CT scans are needed, so that researchers can train and develop COVID-19 diagnosis systems and trustfully evaluate their performance.

The current paper presents a baseline approach for the Competition part of the Workshop "AI-enabled Medical Image Analysis Workshop and Covid-19 Diagnosis Competition (MIA-COV19D)" which occured in conjunction with the International Conference on Computer Vision (ICCV) 2021 in Montreal, Canada, October 11- 17, 2021.

The MIA-COV19D AI-enabled Medical Image Analysis (MIA) Workshop emphasizes on radiological quantitative image analysis for diagnosis of diseases. The focus is placed on Artificial Intelligence (AI), Machine and Deep Learning (ML, DL) approaches that target effective and adaptive diagnosis, as well as on approaches that enforce trustworthiness and create justifications of the decision making process.

The COV19D Competition was based on a new large database of chest CT scan series that was manually annotated for Covid-19/non-Covid-19 diagnosis. The training and validation partitions along with their annotations were provided to the participating teams to develop AI/ML/DL models for Covid-19/non-Covid-19 prediction. Performance of approaches was next evaluated on the test set. The main developed methods and their results on the COV19D test set are described and compared in this paper.

The COV19-CT-DB is a new large database with about 5,000 3-D CT scans, annotated for COVID-19 infection.

The rest of the paper is as follows. Section 2 presents former work on which the presented baseline has been based. Section 3 presents the database created and used in the Competition. The ML approach and the pre-processing steps are described in Section 4. The obtained results, are presented in Section 5. A short description of all main methods that were applied to the COV19D database, as well as their results, is provided in Section 6. Conclusions and future work are described in Section 7.

2. Related Work

In [4] a CNN plus RNN network was used, taking as input CT scan images and discriminating between COVID-19 and non-COVID-19 cases.

In [13], the authors employed a variety of 3-D ResNet models for detecting COVID-19 and distinguishing it from other common pneumonia (CP) and normal cases, using volumetric 3-D CT scans.

In [19], a weakly supervised deep learning framework was suggested using 3-D CT volumes for COVID-19 classification and lesion localization. A pre-trained UNet was utilized for segmenting the lung region of each CT scan slice; the latter was fed into a 3-D DNN that provided the classification outputs.

The presented approach is based on a CNN-RNN architecture that performs 3-D CT scan analysis. The method follows our previous work [7, 9, 8, 15] on developing deep neural architectures for predicting COVID-19, as well as neurodegenerative and other [11, 8, 6, 17, 20] diseases and medical situations, or for analyzing sequences of images [16, 18, 10].

These architectures have been applied for: a) prediction of Parkinson's, based on datasets of MRI and DaTScans, either created in collaboration with the Georgios Gennimatas Hospital (GGH) in Athens [8], or provided by the PPMI study sponsored by M. J. Fox for Parkinson's Research [11], b) prediction of COVID-19, based on CT chest scans, scan series, or x-rays, either collected from the public domain, or aggregated from various hospitals [7].

3. The COV19-CT-DB Database

The COVID19-CT-Database (COV19-CT-DB) consists of chest CT scans that are annotated for the existence of COVID-19. Data collection was conducted in the period from September 1, 2020 to March 31, 2021. Data were aggregated from many hospitals, containing anonymized human lung CT scans with signs of COVID-19 and without signs of COVID-19. Figure 1 shows some CT slices from a non-COVID-19 case and Figure 2 some CT slices from a COVID-19 case.

The COV19-CT-DB database consist of about 5000 chest CT scan series, which correspond to a high number of patients ($>$1000) and subjects ($>$2000). Annotation of each CT slice has been performed by 4 very experienced (each with over 20 years of experience) medical experts; two radiologists and two pulmonologists. Labels provided by the 4 experts showed a high degree of agreement (around 98%).

One difference of COV19-CT-DB from other existing datasets is its annotation by medical experts (labels have not been created as a result of just positive RT-PCR testing).

Each of the 3-D scans includes different number of slices, ranging from 50 to 700. The database has been split in training, validation and testing sets.

The training set contains, in total, 1560 3-D CT scans. These include 690 COVID-19 cases and 870 Non-COVID-19 cases. The validation set consists of 374 3-D CT scans. 165 of them represent COVID-19 cases and 209 of them represent Non-COVID-19 cases. Both include different numbers of CT slices per CT scan, ranging from 50 to 700.

4. The Deep Learning Approach

4.1. 3-D Analysis and COVID-19 Diagnosis

The input sequence is a 3-D signal, consisting of a series of chest CT slices, i.e., 2-D images, the number of which is varying, depending on the context of CT scanning. The context is defined in terms of various requirements, such as

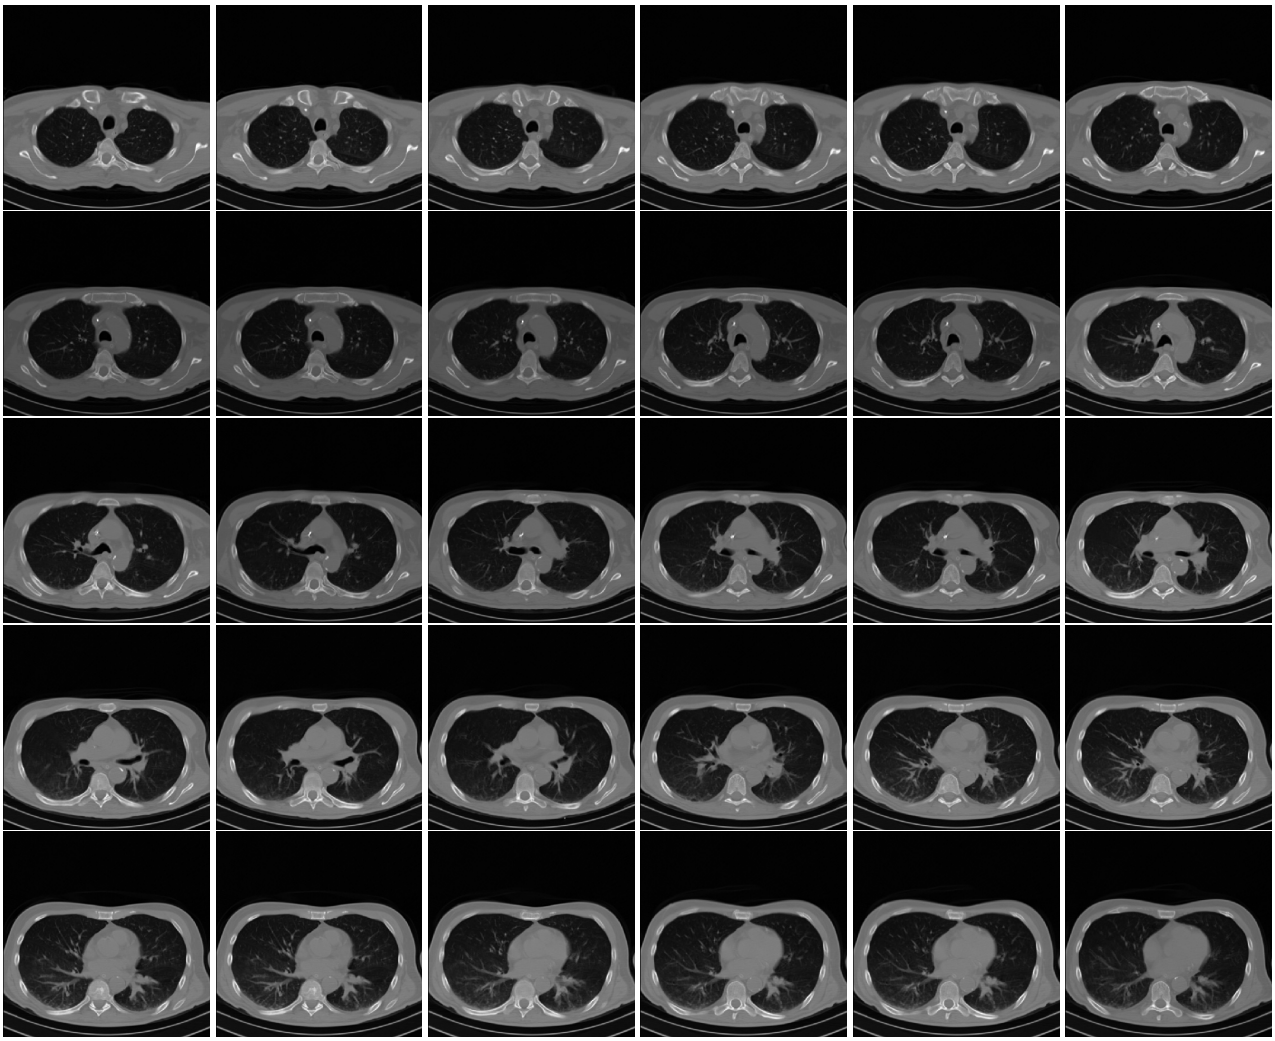

Figure 1. Slices from a non COVID-19 CT scan.

the accuracy asked by the doctor who ordered the scan, the characteristics of the CT scanner that is used, or the specific subject's features, e.g., weight and age.

The baseline approach is a CNN-RNN architecture, as shown in Figure 3. At first all input CT scans are padded to have length t (i.e., consist of t slices). The whole (unsegmented) sequence of 2-D slices of a CT-scan are fed as input to the CNN part. Thus the CNN part performs local, per 2-D slice, analysis, extracting features mainly from the lung regions. The target is to make diagnosis using the whole 3-D CT scan series, similarly to the annotations provided by the medical experts. The RNN part provides this decision, analyzing the CNN features of the whole 3-D CT scan, sequentially moving from slice 0 to slice $t-1$. The outputs of the RNN part feed the output layer -with 2 units- that uses a softmax activation function and provides the final COVID-19 diagnosis.

In this way, the CNN-RNN network outputs a probability for each CT scan slice; the CNN-RNN is followed by a voting scheme that makes the final decision; the voting scheme can be either a majority voting or an at-least one voting (i.e., if at least one slice in the scan is predicted as COVID-19, then the whole CT scan is diagnosed as COVID-19, and if all slices in the scan are predicted as non-COVID-19, then the whole CT scan is diagnosed as non-COVID-19).

4.2. Pre-Processing & Implementation Details

At first, CT images were extracted from DICOM files. Then, the voxel intensity values were clipped using a window/level of 350 Hounsfield units (HU)/-1150 HU and normalized to the range of $[0, 1]$.

Regarding implementation of the proposed methodology: i) we utilized ResNet50 as CNN model, stacking on top of it a global average pooling layer, a batch normalization layer and dropout (with keep probability 0.8); ii) we used a single one-directional GRU layer consisting of 128

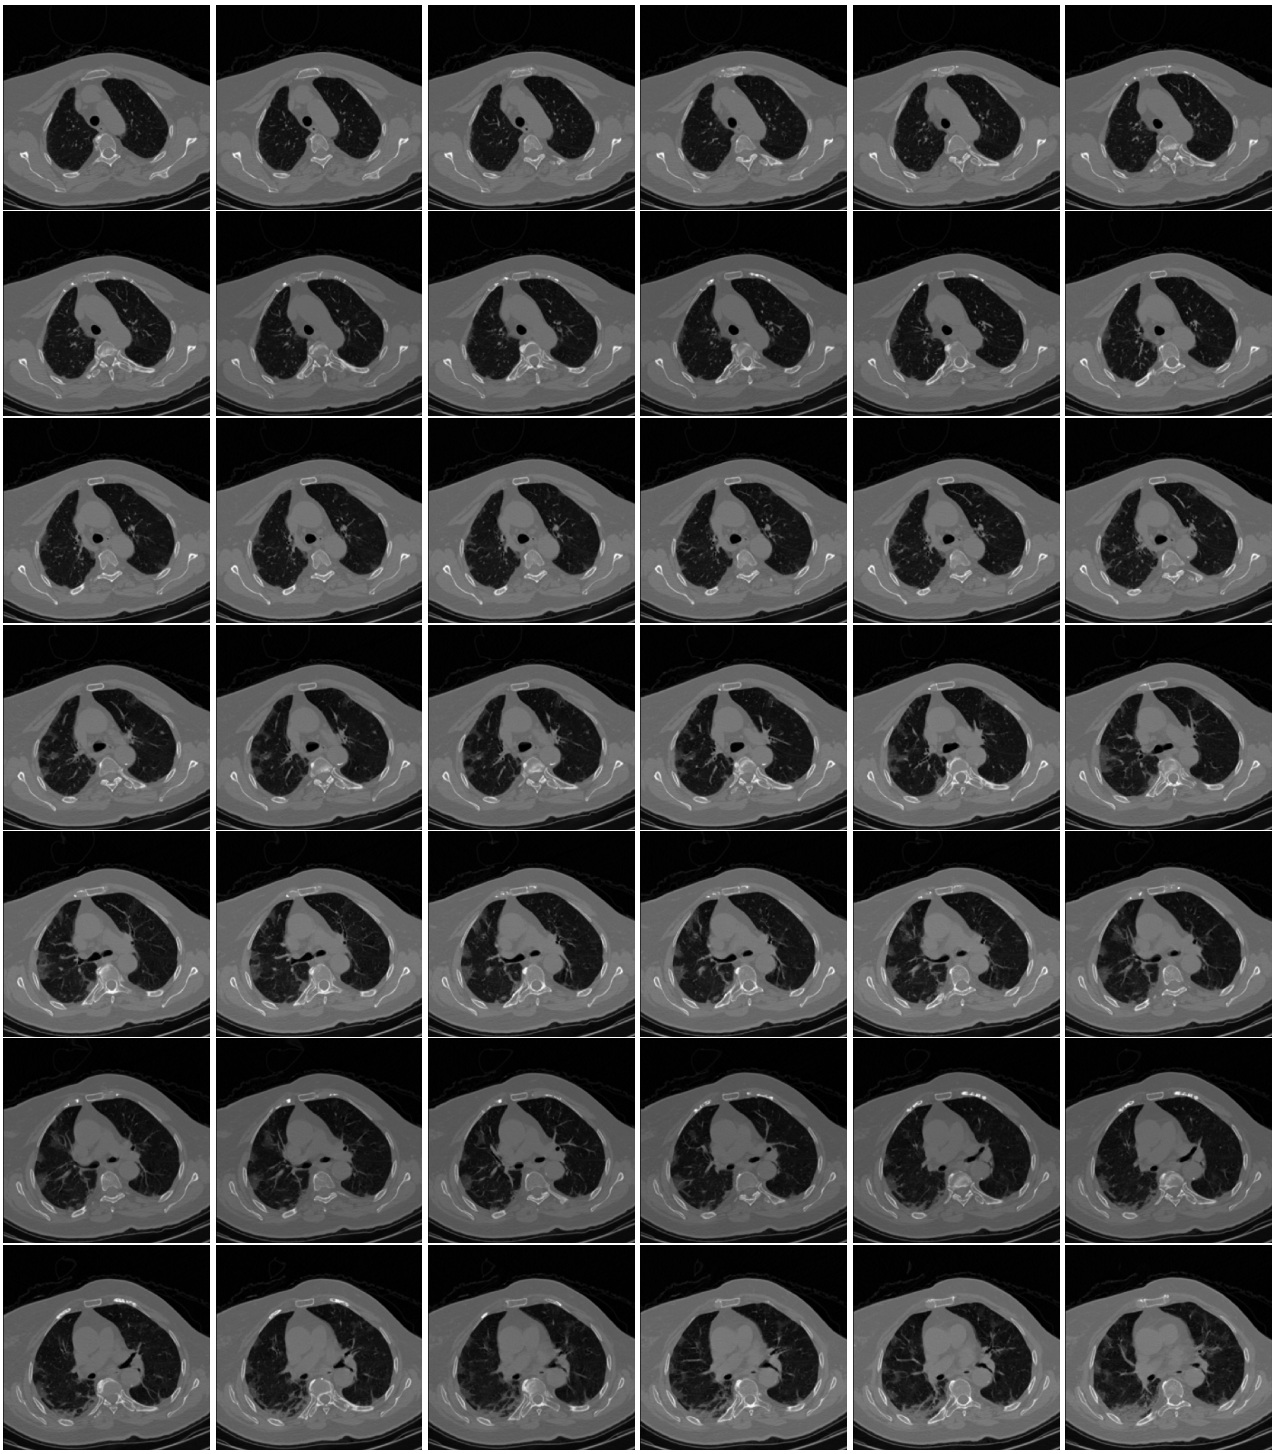

Figure 2. Slices from a COVID-19 CT scan.

units as RNN model. The model was fed with 3-D CT scans composed of the CT slices; each slice was resized from its original size of $512 \times 512 \times 3$ to $224 \times 224 \times 3$. As a voting scheme, we used the at-least one.

Batch size was equal to 5 (i.e, at each iteration our model processed 5 CT scans) and the input length 't' was 700 (the maximum number of slices found across all CT scans). Softmax cross entropy was the utilized loss function for training the model. Adam optimizer was used with learning rate 10^{-4}. Training was performed on a Tesla V100

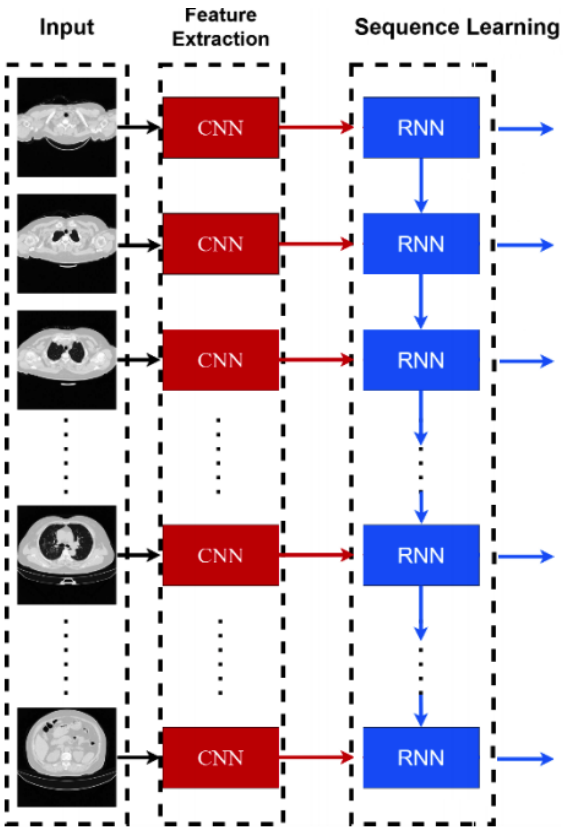

Figure 3. The CNN-RNN model

Table 1. Performance of the baseline CNN-RNN network

Method	'macro' F1 Score
Validation Set	0.70
Test Set	0.67

32GB GPU.

5. Baseline Experimental Results

This section describes a set of experiments evaluating the performance of the baseline approach.

Table 1 shows the performance of the network over the validation and the test datasets, after training with the training dataset, in terms of macro F1 score. The macro F1 score is defined as the unweighted average of the class-wise/label-wise F1-scores, i.e., the unweighted average of the COVID-19 class F1 score and of the non-COVID-19 class F1 score.

The main downside of the model is that there exists only one label for the whole CT scan and there are no labels for each CT scan slice. Thus, the presented model analyzes the whole CT scan, based on information extracted from each slice.

6. The COV19D Competition Results

Thirty five teams participated in the COV19D Competition. Eighteen teams submitted their results. Twelve teams scored higher than the baseline and made valid submissions.

The winner of the Competition was **FDVTS-COVID** from Fudan University, China. The runner-up (with a slight difference from the winning team) was **SenticLab.UAIC** from SenticLab and "Alexandru Ioan Cuza" University of Iasi, Romania. The team that ranked third was **ACVLab** from National Cheng Kung University, Taiwan. All results of the 12 teams are shown in Table 2; links are also provided to either, or both, of the related Github codes and arXiv papers.

In the following we provide a short description of the best methods that each one of the 12 teams implemented and used. The index on the right of each team represents their official rank in the classification task.

Team **FDVTS-COVID** 1 achieved the top performance in this task. They relied on a periphery-aware deep learning network with contrastive representation enhancement (CRE) mechanism. This network consisted of: a Periphery-aware Spatial Prediction network, a projection network along with a subsequent CRE module and a classification network. The Periphery-aware Spatial Prediction network's

Table 2. Competition Results: F1 Score in %; the best performing submission is in bold

Teams	Submission #	Macro F1	F1 (COVID)	F1 (NON-COVID)	Github	arXiv
FDVTS_COVID	1	84.8	74.27	95.33	link	
	2	87.47	78.56	96.38		
	3	86.42	77.07	95.77		
	4	88.37	80.04	96.70		
	5	**90.43**	**83.60**	**97.27**		
SenticLab.UAIC	1 - volumetric	**90.06**	**82.96**	**97.17**	link	link
	2 -slice	76.73	58.54	94.92		
	3 -MLP	64.35	42.46	86.25		
	4 -logistic regression	81.85	69.21	94.49		
ACVLab	1	82.6	70.8	94.4	link1 link2	link
	2	80.66	68.1	93.22		
	3	**88.74**	**80.63**	**96.84**		
	4	87.00	77.86	96.04		
	5	88.65	80.57	96.75		
DeepCam	083..MLP-0.5	87.4	78.44	96.34	link	link
	177..MLP-0.5	85.54	75.5	95.58		
	Emb-0.5	**88.22**	**79.79**	**96.64**		
	Emb-0.3	86.15	76.44	95.85		
	Bert-0.5	80.94	68.37	93.51		
TAC	thresh 10	61.8	44.85	78.75	link	link
	thresh 20	74.4	59.25	89.54		
	thresh 30	80.88	68.3	93.45		
	all	79.36	66.07	92.64		
	max	**87.77**	**78.78**	**96.75**		
LoVE	1	**84.2**	**73.65**	**94.76**	link	
Heal it	1	**78.86**	**63.65**	**94.06**	link	link
	2	74.74	57.56	91.92		
HCMUS-HGV	1	**78.13**	**63.08**	**93.18**	link	link
Blessing	1	74	56.53	91.26	link	
	2	**75.67**	**58.95**	**92.4**		
	3	73.41	56.02	90.8		
	4	71	52.66	89.28		
	5	74.5	57.25	91.75		
AvengerQ	1	**71.83**	**51.63**	**92.04**	link	link
	2	61.59	43.14	80.04		
	3	57.17	40.25	74.08		
	4	53.19	38	8.4		
	5	70.08	50.09	90.07		
Terps	1	**70.86**	**48.96**	**92.75**	link	link
xVision	1	68.65	51.37	85.94	link	link
	2	62.16	46.41	77.9		
	3	**70.5**	**53.67**	**87.33**		
	4	63.25	47.75	78.74		
	5	66.11	50	82.27		
baseline	1	**67**	**54.38**	**79.62**		link

aim was to capture the important information of the CT-Scan by predicting the boundary distance map. This map was generated to represent the location information about whether a pixel belonged to the interior of the lung region as well as the distance to the region boundary. This network was a UNet-style prediction network with an encoder-decoder architecture; ResNet was adopted as the encoder; the decoder was a mirrored version of the encoder by re-

placing the pooling layers with bilinear upsampling layers. This network was at first pre-trained; then the decoder was discarded and only the encoder was kept. Each CT image was at first augmented and was then fed into this pre-trained periphery-aware encoder, generating vector representations. A classifier was trained on top of these representations for COVID-19 classification. Meanwhile, these representations were mapped by a projection network to new representations which were further enhanced in a contrastive learning manner.

Team **SenticLab.UAIC** 2 developed a volumetric and slice-level approach; the former provided the best results. They used an inflated 3D ResNet50 model with non local operations on the second and third layers. Inflated convolutions were obtained by expanding filters and pooling kernels of 2D ConvNets into 3D, resulting in learning spatio-temporal feature extractors from 3D images while using successful ImageNet architectures. In order to overcome the problem of over-fitting they used label smoothing. In order to handle the variable length of the CT-Scans they used a sub-sampling technique, or padding, for lengths above, or under 128 respectively. During inference, parts of a single CT-Scan volume will be inputted several times in the model; then a threshold procedure is followed for eliminating some of these results; the final prediction is based on a majority voting scheme of the remaining results.

Team **ACVLab** 3 proposed two different approaches, one based on the slice-level and the other based on the 3D volume. The first model was named DWCC (Deep Wilcoxon single-rank test for COVID-19 Classification) and its main component was a vision transformer (Swin-Transformer) used for single-slice level classification followed by Wilcoxon signed-rank test. They aimed at making the predicted result more stable and explainable. The second model termed as CCAT (Convolutional CT scan-Aware Transformer) used the Within-Slice-Transformer (WST) and Between-Slice-Transformer (BST) which were based on ResNet50 for feature extraction and self-attention for context-encoded features.

Team **DeepCam** 4 used a 3D CNN with BERT to classify the CT-scan volumes. In particular, they segmented the initial CT-scans through morphological transformation and a pre-trained UNET network. Following that, they used a re-sampling method to select a set of fixed number of slice images for training and validation. This selection depended on the use of a threshold for the percentage of lung masks in the whole image. The fixed size volume was then fed to the classification model which consisted of a 3D CNN with BERT. This part of the network extracted an embedding feature vector for each CT-scan and then this vector was passed to an auxiliary classification layer which included an MLP classifier.

Team **TAC** 5 used an AutoML automated pipeline, aiming at fewer resources and time to develop their approach. Their pipeline used different 2D CNN pre-trained models such as VGG, ResNet, DenseNet. 2D CNNs were trained on slice level instead of 3D volume level. Evaluation was made on slice level. For 3D volume level, predictions were made on 2D slices, and then most occurred predictions were taken and assigned as 3D image labels.

Team **LoVe** 6 expanded a vision transformer as a robust feature learner of the 3D CT-scans so as to diagnose the COVID-19. Their network consisted of two main stages. At first, lung segmentation was applied using pre-trained UNET followed by the classification model in which the features were extracted from each slice using Swin transformer and then aggregated into a 3D volume level feature by a max pooling layer.

Team **Heal it** 7 presented a hybrid deep learning model named CTNet. This network consisted of three components, the input module with a data resampling strategy, the CNN feature extractor module with SE attention module and the information aggregation module with the transformer and fully connected (FC) layer.

Team **HCMUS-HGV** 8 proposed a method based on ResNet and DenseNet networks. More specifically, they combined these two networks in an ensemble that classified each slice of the CT-scan and then the most frequent result was assigned as the volume's prediction.

Team **Blessing** 9 used 3D convolutional neural networks based on large-scale pre-trained parameters for extracting and classifying the rich spatial information of the 3-dimensional CT images. In addition, they utilized a resampling technique, in order to unify the size and number of channels of all CT images.

Team **AvengerQ** 10 modified the RegNet (2D) neural network into a 3D convolution neural network. For inference, they tested a variety of thresholds in order to achieve the best performance in the classification of covid and non-covid CT-scans.

Team **Terps** 11 used a shallow Convolutional neural network named TeliNet. It consisted of 2D convolutions, batch normalization, LeakyReLU activations and max pooling layers. In this approach the classification was done at slice level.

Team **xVision** 12 used two deep learning methods, vision transformer (ViT) based on attention models and DenseNet built upon a convolutional neural network. Comparing these two different architectures they concluded that the vision transformer outperformed the CNN based approach.

7. Conclusions and Future Work

In this paper we have introduced a new large database of chest 3-D CT scans, obtained in various contexts and consisting of different numbers of CT slices. We have also developed a deep neural network, based on a CNN-RNN ar-

chitecture and used it, as baseline, for COVID-19 diagnosis on this database.

The paper presented the results of a large number of methodologies applied to the database of chest 3-D CT scans , in the context of the ICCV 2021 COV19D Competition; these methods provided performances that were superior to that of the baseline method.

The database and the models presented in the paper will form the basis for expansion towards more transparent modelling of COVID-19 diagnosis. In particular, the baseline method is extended with attention modules so as to provide higher prediction performance; it has been also enriched with transparent visualization capabilities, inspired by the techniques in [7, 9]. Future work includes blending of the deep learning methodology with knowledge based encoding, extending former work in [2, 5].

References

[1] Roohallah Alizadehsani, Zahra Alizadeh Sani, Mohaddeseh Behjati, Zahra Roshanzamir, Sadiq Hussain, Niloofar Abedini, Fereshteh Hasanzadeh, Abbas Khosravi, Afshin Shoeibi, Mohamad Roshanzamir, et al. Risk factors prediction, clinical outcomes, and mortality in covid-19 patients. *Journal of medical virology*, 93(4):2307–2320, 2021.

[2] Birte Glimm, Yevgeny Kazakov, Ilianna Kollia, and Giorgos B Stamou. Using the tbox to optimise sparql queries.

[3] Sebastian Hoehl, Holger Rabenau, Annemarie Berger, Marhild Kortenbusch, Jindrich Cinatl, Denisa Bojkova, Pia Behrens, Boris Böddinghaus, Udo Götsch, Frank Naujoks, et al. Evidence of sars-cov-2 infection in returning travelers from wuhan, china. *New England Journal of Medicine*, 382(13):1278–1280, 2020.

[4] Adil Khadidos, Alaa O Khadidos, Srihari Kannan, Yuvaraj Natarajan, Sachi Nandan Mohanty, and Georgios Tsaramirsis. Analysis of covid-19 infections on a ct image using deepsense model. *Frontiers in Public Health*, 8, 2020.

[5] Ilianna Kollia, Birte Glimm, and Ian Horrocks. Answering queries over owl ontologies with sparql.

[6] Ilianna Kollia, Andreas-Georgios Stafylopatis, and Stefanos Kollias. Predicting parkinson's disease using latent information extracted from deep neural networks. In *2019 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2019.

[7] Dimitrios Kollias, N Bouas, Y Vlaxos, V Brillakis, M Seferis, Ilianna Kollia, Levon Sukissian, James Wingate, and S Kollias. Deep transparent prediction through latent representation analysis. *arXiv preprint arXiv:2009.07044*, 2020.

[8] Dimitrios Kollias, Athanasios Tagaris, Andreas Stafylopatis, Stefanos Kollias, and Georgios Tagaris. Deep neural architectures for prediction in healthcare. *Complex & Intelligent Systems*, 4(2):119–131, 2018.

[9] Dimitris Kollias, Y Vlaxos, M Seferis, Ilianna Kollia, Levon Sukissian, James Wingate, and Stefanos D Kollias. Transparent adaptation in deep medical image diagnosis. In *TAILOR*, pages 251–267, 2020.

[10] Dimitrios Kollias and Stefanos Zafeiriou. Training deep neural networks with different datasets in-the-wild: The emotion recognition paradigm. In *2018 International Joint Conference on Neural Networks (IJCNN)*, pages 1–8. IEEE, 2018.

[11] Stefanos Kollias, Luc Bidaut, James Wingate, Ilianna Kollia, et al. A unified deep learning approach for prediction of parkinson's disease. *IET Image Processing*, 2020.

[12] Chih-Cheng Lai, Yen Hung Liu, Cheng-Yi Wang, Ya-Hui Wang, Shun-Chung Hsueh, Muh-Yen Yen, Wen-Chien Ko, and Po-Ren Hsueh. Asymptomatic carrier state, acute respiratory disease, and pneumonia due to severe acute respiratory syndrome coronavirus 2 (sars-cov-2): Facts and myths. *Journal of Microbiology, Immunology and Infection*, 53(3):404–412, 2020.

[13] Yan Li and Liming Xia. Coronavirus disease 2019 (covid-19): role of chest ct in diagnosis and management. *American Journal of Roentgenology*, 214(6):1280–1286, 2020.

[14] Catrin Sohrabi, Zaid Alsafi, Niamh O'Neill, Mehdi Khan, Ahmed Kerwan, Ahmed Al-Jabir, Christos Iosifidis, and Riaz Agha. World health organization declares global emergency: A review of the 2019 novel coronavirus (covid-19). *International journal of surgery*, 76:71–76, 2020.

[15] Athanasios Tagaris, Dimitrios Kollias, Andreas Stafylopatis, Georgios Tagaris, and Stefanos Kollias. Machine learning for neurodegenerative disorder diagnosis—survey of practices and launch of benchmark dataset. *International Journal on Artificial Intelligence Tools*, 27(03):1850011, 2018.

[16] Nicolas Tsapatsoulis and Stefanos Kollias. Face detection in color images and video sequences. In *2000 10th Mediterranean Electrotechnical Conference. Information Technology and Electrotechnology for the Mediterranean Countries. Proceedings. MeleCon 2000 (Cat. No. 00CH37099)*, volume 2, pages 498–502. IEEE, 2000.

[17] Paraskevi Tzouveli, Andreas Schmidt, Michael Schneider, Antonis Symvonis, and Stefanos Kollias. Adaptive reading assistance for the inclusion of students with dyslexia: The agent-dysl approach. In *2008 Eighth IEEE International Conference on Advanced Learning Technologies*, pages 167–171. IEEE, 2008.

[18] Manolis Wallace, Nicolas Tsapatsoulis, and Stefanos Kollias. Intelligent initialization of resource allocating rbf networks. *Neural Networks*, 18(2):117–122, 2005.

[19] Xinggang Wang, Xianbo Deng, Qing Fu, Qiang Zhou, Jiapei Feng, Hui Ma, Wenyu Liu, and Chuansheng Zheng. A weakly-supervised framework for covid-19 classification and lesion localization from chest ct. *IEEE transactions on medical imaging*, 39(8):2615–2625, 2020.

[20] Miao Yu, Dimitrios Kollias, James Wingate, Niro Siriwardena, and Stefanos Kollias. Machine learning for predictive modelling of ambulance calls. *Electronics*, 10(4):482, 2021.