Deep-Adaptation: Ensembling and Test Augmentation for Covid-19 Detection and Covid-19 Domain Adaptation from 3D CT-Scans

Fares BOUGOURZI
Junia, UMR 8520, CNRS, Centrale Lille,
Univerity of Polytechnique
Hauts-de-France, 59000 Lille, France
faresbougourzi@gmail.com

Feryal WINDAL MOULAÏ
Junia, UMR 8520, CNRS, Centrale Lille,
Univerity of Polytechnique
Hauts-de-France, 59000 Lille, France
feryal.windal@junia.com

Halim Benhabiles
IMT Nord Europe, Institut Mines-Télécom,
Univ. Lille, Centre for Digital Systems,
F-59000 Lille, France
halim.benhabiles@imt-nord-europe.fr

Fadi DORNAIKA
University of the Basque Country UPV/EHU,
San Sebastian, SPAIN; IKERBASQUE, Basque
Foundation for Science, Bilbao, SPAIN
fadi.dornaika@ehu.eus

Abdelmalik Taleb-Ahmed
Université Polytechnique Hauts-de-France, Université de Lille,
CNRS, Valenciennes, 59313, Hauts-de-France, France
Abdelmalik.Taleb-Ahmed@uphf.fr

Abstract

Since the onset of the Covid-19 pandemic in late 2019, the realm of medical image analysis has seen a surge in importance, particularly with the utilization of CT-scan imaging for disease diagnosis. This paper presents findings from our participation in the 4th COV19D competition, specifically targeting the challenges of Covid-19 Detection and Covid-19 Domain Adaptation. Our methodology revolves around lung segmentation and Covid-19 infection segmentation, employing the state-of-the-art CNN-based segmentation architecture PDAtt-Unet. Unlike conventional methods, we introduce a novel approach by concatenating the input slice (grayscale) with segmented lung and infection, thereby generating three input channels reminiscent of color channels. Furthermore, we leverage three distinct 3D CNN backbones—Customized Hybrid-DeCoVNet, in addition to pretrained 3D-Resnet-18 and 3D-Resnet-50 models—to facilitate Covid-19 recognition for both challenges. To further boost performance, we explore ensemble techniques and testing augmentation. Comparison with baseline results highlights the substantial efficiency of our approach, showcasing a significant improvement in terms of F1-score (14%) on the validation data. Our approach ranked second and third in the Covid-19 Detection and Covid-19 Domain Adaptation Challenges, respectively, based on the test data results. Our approach demonstrates improvements of 9.5% and 17% compared to baseline performance in these challenges. Furthermore, our approach exhibits very promising performance compared with the approaches of other competitors, underscoring the significance of the proposed training paradigm and the utilization of ensemble and testing augmentation techniques.

1. Introduction

2. Introduction

Since the onset of the Covid-19 pandemic in late 2019, Reverse Transcription-Polymerase Chain Reaction (RT-PCR) has been widely established as the primary method for Covid-19 detection. Nevertheless, this testing modality presents several drawbacks, including limited availability of RT-PCR kits, lengthy procedures, and a notable incidence of false negative results [15]. Consequently, medical imaging techniques such as X-rays and CT-scans have gained prominence as complementary tools for Covid-19 detection [5, 10]. CT-scans not only serve in identifying Covid-19 infections but also play a crucial role in monitoring patients' conditions and predicting disease severity [4, 10].

In recent years, Deep Learning methodologies have risen to prominence in computer vision tasks, showcasing remarkable performance gains over conventional techniques [3, 6]. However, one of the primary challenges associated with Deep Learning, particularly in the realm of Convolutional Neural Networks (CNNs), lies in the necessity for extensive labeled datasets, a resource often scarce in medical domains [8, 10]. Furthermore, the majority of existing CNN architectures are tailored for processing static images, which proves inadequate in capturing the intricacies inherent in medical imaging data, especially for the volumetric scans [4]. On the other hand, domain adaptation is one of the most challenging aspects encountered in medical imaging, owing to the high variability of data from one center to another due to the variety of recording settings and scanners. Machine learning techniques used in computer-aided medical image analysis usually suffer from the domain shift problem caused by different distributions between source/reference data and target data [12].

In our paper, we present an approach to address both Covid-19 detection and domain adaptation challenges on the 4th COV19D competition. Our method revolves around lung segmentation and Covid-19 infection segmentation using the PDAtt-Unet CNN-based segmentation architecture, which concurrently segments lung regions and infections. Departing from traditional methods, we integrate the input slice with segmented lung and infection, creating three input channels akin to color channels. We utilize three 3D CNN backbones— Customized Hybrid-DeCoVNet, pretrained 3D-Resnet-18, and 3D-Resnet-50 models—to train Covid-19 recognition for both challenges. Additionally, we explore ensemble approaches and testing augmentation to enhance performance. Our main contributions are:

- We adopted a Customized Hybrid-DeCoVNet architecture for both Covid-19 Detection and Covid-19 Domain Adaptation Challenges. This architecture incorporates the concatenation of the original slice, the segmented lung, and the segmented Covid-19 infection as the three input channels.
- In addition to our proposed Customized Hybrid-DeCoVNet architecture, we leveraged two pretrained 3D-CNNs: 3D-Resnet-18 and 3D-Resnet-50.
- We explored ensemble approaches and testing augmentation techniques to enhance the robustness and performance of our method.
- Our approach demonstrated a substantial improvement in efficiency compared to baseline results, with a significant margin in F1-score (14%).
- Based on the test data results, our approach ranked second and third in the Covid-19 Detection and Covid-19 Domain Adaptation Challenges, respectively, Our approach demonstrates improvements of 9.5% and 17% compared to baseline performance in these challenges.
- We have made our codes and pretrained models publicly available in [1]

This paper is organized in following way: Section 3 describes our proposed approaches for Covid-19 Detection and Covid-19 Domain Adaptation Detection. The experiments and results are detailed in Section 4. Finally, we conclude our paper in Section 5.

3. Our Approaches

Our approach is tailored to capitalize on region of interest segmentation, specifically lung segmentation, and infection segmentation alongside input slices from CT scans. The objective is to develop a model proficient in discerning COVID-19 cases from non-COVID-19 cases. We evaluate three baseline architectures: Customized Hybrid-DeCoVNet [9], 3D-ResNet-18, and 3D-ResNet-50 [13].

3.1. Data Preprocessing

The objective of this phase is to eliminate slices that do not exhibit any lung structures and to identify lung features in the remaining slices. Following our previous approach in the 2nd COV19D competition and 3rd COV19D competition [7, 9], ResneXt-50 [29] is used to filter the CT slices that does not show lung regions, to concentrate only on the slices that may have infection.

3.2. Customized Hybrid-DeCoVNet

In this challenge, we adopted our proposed Customized Hybrid-DeCoVNet which were proposed for Covid-19 severity prediction, to perform Covid-19 recognition in this challenge. The first modification is by considering the input slice, their region of interest segmentation and the infection segmentation as input, these three images are concatenated to form 3 channels. For segmenting the lung and infection, we employed the PDAtt-Unet [8], which simultaneously segments the infection and lung regions. As illustrated in Figure 1, PDAtt-Unet comprises an Encoder with an input image pyramid, serving as attention gates for the encoder features to preserve awareness of salient parts across all encoder layers. The Decoder of PDAtt-Unet consists of two parallel decoders, each similar to the Att-Unet decoder [26]. These decoders are designed to segment the infection and lung regions (regions of interest). Both the pyramid encoder and dual decoders aim to maintain awareness of salient parts during the encoding phase and regions of interest during the decoding phase.

PDAtt-Unet is trained using three datasets Segmentation dataset nr. 2 [27], COVID-19 CT segmentation [27] and CC-CCII [25], each dataset is divided into 70%-30% as training and validation splits, then PDAtt-Unet is trained on

[1] `https://github.com/faresbougourzi/4th-COV19D-Competition`. (Last accessed on March, 17^{th} 2024).

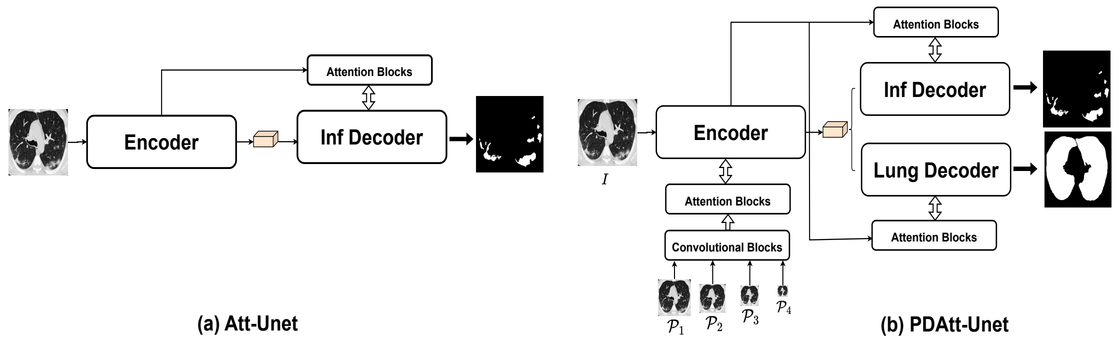

Figure 1. The comparison between Att-Unet [26] and PDAtt-Unet [8] segmentation architectures.

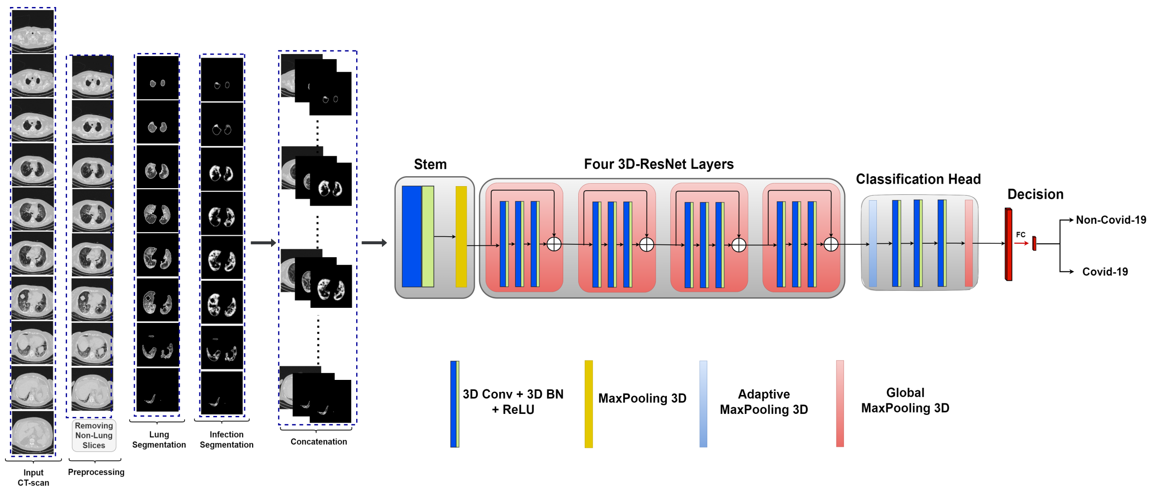

Figure 2. Customized Hybrid-DeCoVNet Approach.

the ensemble of 70% of the three datasets and evaluated on the ensemble of the their 30%.

As illustrated in Figure 2, Customized Hybrid-DeCoVNet comprises of four components. First, the three images depicting the input slice, the segmented lung and segmented infection are merged into a three-channel image. This is performed for every slice of the input CT-scans, then all of these merged 3 channels images are concatenated. For a CT-scan of N slice this will produce $S = 224 \times 224 \times 3 \times N$. Since the number of slices is different from one CT-scan to another, S is resized into a fixed size of $224 \times 224 \times 3 \times 64$. This resulting volume is fed into the Stem block, which is a 3D convolutional layer with a kernel of size (7, 7, 5) for height, width, and depth, respectively. The Stem block transforms the two input channels into 16 channels and is followed by Batch Normalization Layer (BN) and ReLU activation function. The second block of Customized Hybrid-DeCoVNet consists of four 3D-Resnet layers, which expand the channels to 64, 128, 256, and 512, respectively. The Classification Head comprises of 3D Adaptive MaxPooling, three 3D convolutional layers, and 3D Global MaxPooling. The output of the Classification Head is flattened into a single-channel deep feature map and fed into the Decision Head, which consists of one FC layer that has two outputs (Non-Covid-19 and Covid-19). Our proposed architecture is designed to enhance the performance of Covid-19 Prediction. It should be noted that Customized Hybrid-DeCoVNet does not have any pretrained weight in contrast with 3D-Resnet-18 and 3D-Resnet-50.

3.3. 3D-Resnet-18 and 3D-Resnet-50

In addition to our proposed Customized Hybrid-DeCoVNet, we evaluated the use of pretrained 3D CNN architectures. To this end, we used the pretrained 3D-Resnet-18 and 3D-Resnet-50 from [13]. These pretrained models were trained for action recognition from 3D-video. For 3D-Resnet-18, the pretrained weights was trained on the ensemble of Kinetics-700 and Moments in Time datasets. While, 3D-Resnet-50 was trained on the ensemble of Kinetics-700, Moments in Time, STAIR-Actions datasets. To adopt these models for

Table 1. Datasets summary of the 4th COV19D Competition. 494 Non annotated

Sub-Competition	Train		Validation	
	Covid-19	**Non-Covid-19**	**Covid-19**	**Non-Covid-19**
Covid-19 Detection Challenge	703	655	170	156
Covid-19 Domain Adaptation Challenge	120	120	56	113

Covid-19 recognition, we changed the decision layer to give 2 output which corresponds to Non-Covid-19 and Covid-19 classes.

4. Experiments and Results

4.1. The COV19-CT-DB Database

In this competition, the COVID19-CT-Database (COV19-CT-DB) [1, 2, 16–21] is used for two sub challenges, which are Covid-19 Detection Challenge and Covid-19 Domain Adaptation Challenge. In Covid-19 Detection Challenge, many CT scans have been aggregated, each one of which has been manually annotated in terms of Covid-19 and non-Covid-19 categories. The resulting dataset is split into training, validation and test partitions. The provided training and validation partitions for developing the approach are summarized in Table 1.

In the second challenge, Covid-19 Domain Adaptation Challenge, CT scans have been aggregated from various hospitals and medical centres. Each CT scan has been manually annotated with respect to Covid-19 and non-Covid-19 categories. The resulting dataset is split into training, validation and test partitions. Participants will be provided with a training set that consists of: i) the annotated data of the 1st Challenge which are aggregated from some hospitals and medical centres (case A); ii) a small number of annotated data and a larger number of non-annotated data (case B), all of which are aggregated from other hospitals and medical centres and their distribution is different from that of case A. Participants will be provided with a validation set that consists of a small number of annotated data of case B.

4.2. Experimental Setup

We utilized the Pytorch Library and four NVIDIA GPU Device GeForce TITAN RTX 16 GB for Deep Learning training and testing. The batch size of 16 CT-scan volumes was used to train the Customized Hybrid-DeCoVNet and 3D-Resnet-18 architectures for 80 epochs. While, a batch size of 8 is used to train 3D-Resnet-50 for 40 epochs. Warm up Cosine learning rate Schedule is adopted with initial learning rate of 0.0001. The following data augmentations are used for training and testing augmentation approach: random rotation with an angle between -40° to 40°, vertical and horizontal flipping with a probability of 20% for each, Multiplicative Noise, Random Brightness, Random Brightness Contrast, Random Contrast, and Random Grid Shuffle .

Table 7. Covid-19 Detection Challenge Submissions

Sub	Macro F1-score	Non-Covid-19	Covid-19
1	93.30	94.40	92.20
2	93.50	94.63	92.37
3	92.75	93.86	91.64
4	93.67	94.71	92.63
5	**94.60**	**95.53**	**93.66**

Table 8. Covid-19 Detection Challenge Final Results

Approach	F1-score	Non-Covid	Covid
MDAP [28]	**94.89**	**95.97**	**93.81**
Deep-Adaptation	94.60	95.53	93.66
ACVLAB[14]	94.39	95.52	93.26
FDVTS[23]	94.24	95.41	93.07
ViGIR Lab[11]	93.63	94.97	92.29
M2@Purdue[24]	90.14	92.06	88.22
baseline [22]	85.11	87.48	82.74

Table 10. Covid-19 Domain Adaptation Challenge Final Results

Sub	Macro F1-score	Non-Covid-19	Covid-19
1	**74.96**	**96.52**	**53.39**
2	73.67	96.10	51.25
3	64.74	92.48	37.00
4	74.33	96.23	52.44
5	63.23	91.50	34.97

4.3. Results

4.3.1 Results of the first sub-challenge

In this part, we used the training data of Covid-19 Detection Challenge (first challenge) and the validation data to train

Table 2. Results of the first Sub-challenge without testing augmentation

Architecture	Val1		Train2		Val2	
	Accuracy	**F1-score**	**Accuracy**	**F1-score**	**Accuracy**	**F1-score**
Customized Hybrid-DeCoVNet	92.33	92.33	83.75	83.72	82.58	82.09
3D-Resnet-18	91.41	91.41	82.08	81.88	83.70	82.59
3D-Resnet-50	91.41	91.40	84.58	84.58	83.70	82.59
Ensemble	91.10	91.10	85	84.93	83.70	82.59

Table 3. Results of the second sub-challenge without testing augmentation

Architecture	Val1		Val2	
	Accuracy	**F1-score**	**Accuracy**	**F1-score**
Customized Hybrid-DeCoVNet	92.33	92.33	83.14	80.64
3D-Resnet-18	92.33	92.32	87.07	85.60
3D-Resnet-50	92.63	92.60	84.26	82.92

and save the best model on the validation data after each epoch, this two splits will be denoted as Train1 and Val1. We also used the training and validation data of the Covid-19 Domain Adaptation Challenge (the second challenge) to evaluate the performance of our approach in unseen data, these two splits will be denoted as Train2 and Val2. Table 2 summarizes the obtained results. From these results, it is noticed that the performance on the Train2 and Val2 splits decreased compared with the results on Val1, this is due to the change of data domain. However, the drop in results in not too big, this shows that our approach can achieve a good result. On the other hand, the ensembling approach achieves better performance on Train2 compared with the single architectures.

4.3.2 Results of the second sub-challenge

In the second sub-challenge, we combined the training data of Covid-19 Detection and Covid-19 Domain Adaptation challenges (Train1+Train2) in order to compare the performance of the three backbones in the scenario where the training data is augmented. The obtained results are summarized in Table 3, in which, Val1 and Val2 correspond to the validation of the first and the second challenge, respectively (correspond to the same splits used in Table 2). By comparing the results of Tables 2 and 3, it is noticed that augmenting the training data improve the performance of the three backbones, especially 3D-Resnet-18.

Table 4 depicts the results of using testing augmentation, in which, each CT-scan is augmented ten times and the CT-scan prediction corresponds to the average probabilities of the prediction of the ten augmentations. Compared with the results of Table 3, using testing augmentation further improves the performance.

4.4. Comparison with the Baseline On the Validation Data

Table 5 depicts the comparison with the baseline results from [22]. The comparison of our approach and the baseline results shows that our approach achieved better performance on both challenges. In more details, our approach achieved better performance than the baseline approach by 14.33% in terms of F1-score for Covid-19 Detection Challenge. Similarly for Covid-19 Domain Adaptation Challenge, our approach achieved better performance than the baseline by 14.52% in terms of F1-score.

4.5. Covid-19 Detection Challenge Results on the Test Data

Since it is allowed to have five submissions in the Covid-19 Detection Challenge, the details of the first four submissions are summarized in Table 6. The fifth submission is an ensemble of the four models from the previous submissions, achieved by averaging the prediction probabilities of each class.

Table 7 summarizes the results obtained on the testing data. It is noticed that the four single models have achieved similar results, with slightly better performance observed when using 3D-Resnet-50. Despite the proposed Customized Hybrid-DeCoVNet not having pretrained weights, it achieves very similar results to the pretrained ones (3D-Resnet-18 and 3D-Resnet-50). Additionally, it should be noted that the Customized Hybrid-DeCoVNet does not explore testing augmentation techniques, as they did not show effectiveness with this architecture on the validation data. The best

Table 4. Results of the second sub-challenge with testing augmentation

Architecture	Val1		Val2	
	Accuracy	**F1-score**	**Accuracy**	**F1-score**
Customized Hybrid-DeCoVNet	91.41	91.40	84.83	83.23
3D-Resnet-18	92.33	92.33	88.76	87.52
3D-Resnet-50	92.33	92.33	85.39	84.34
Ensemble	92.33	92.33	88.20	87.14

Table 5. Results Comparison with the baseline

Architecture	sub-challenge 1		sub-challenge 2	
	Accuracy	**F1-score**	**Accuracy**	**F1-score**
Baseline [22]	-	78	-	73
Customized Hybrid-DeCoVNet	91.41	91.40	84.83	83.23
3D-Resnet-18	92.33	92.33	88.76	87.52
3D-Resnet-50	92.33	92.33	85.39	84.34
Ensemble	92.33	92.33	88.20	87.14

performance on the testing data was achieved by the ensembling approach of the four models from the first four submissions. This highlights the importance of using predictions from multiple models to achieve better performance. Table 8 summarizes the comparison with the models from other challenge teams, where our approach ranked second. Furthermore, it is observed that the narrow margin of difference, merely 0.5% between the top 4 teams, underscores the fierce competition and high level of expertise demonstrated across all participants, reflecting the intense nature of the challenge.

4.6. Covid-19 Domain Adaptation Challenge Results on the Test Data

For the Covid-19 Domain Adaptation Challenge submission, we selected three trained models and two ensembling approaches. The details of the first three submissions are summarized in Table 9. Notably, the first and second submissions correspond to the second and third submissions of the previous challenge, as outlined in Table 6. In the third submission, we utilized the second trained model (submission 2) to predict pseudo labels for the unlabelled data. Subsequently, we trained a new model using Train1, Train 2, and the unlabelled data with the pseudo labels. For the fourth and fifth submissions, we employed average probabilities ensembling and a Covid-19 domination strategy. In the latter, if any model among the first three predicts the CT-scan as Covid-19, the ensembling prediction is Covid-19.

The results of these five submissions and a comparison with other participant approaches are summarized in Tables 10 and 11, respectively. Notably, the first submission achieved better performance than the other four submissions. Additionally, using pseudo labels resulted in a reduction in performance of the 3D-Resnet-50 model compared to training the model without utilizing the unlabelled data. This decrease is likely due to inaccuracies in the predicted pseudo labels, which can mislead the model. Despite a considerable drop in performance of the third model in the third submission compared to submissions 1 and 2, ensembling did not significantly influence the overall performance from the fourth submission onwards. This underscores the importance of selecting appropriate models for ensembling to improve performance. Furthermore, the second ensembling approach proposed in the last submission was found to be unsuitable for this task. Our approach ranked third, with only four teams achieving better performance than the baseline, as shown in Table 11. This shows that domain adaptation is very challenging task in medical field, including Covid-19 Detection.

5. Conclusion

In this paper, we introduced a new approach for addressing the Covid-19 Detection and Covid-19 Domain Adaptation Challenges. Our approach primarily leveraged lung segmentation and Covid-19 infection segmentation through the utilization of state-of-the-art CNN-based segmentation architecture, namely PDAtt-Unet. This architecture enables simultaneous segmentation of lung regions and infections. Rather than feeding individual input slices to the training network, we concatenated the input slice (grayscale) with

Table 6. Covid-19 Detection Challenge Submissions

Sub	Backbone	Training Data	Test Augmentation
1	Customized Hybrid-DeCoVNet	Train1+Train2	No
2	3D-Resnet-18	Train1+Train2	Yes
3	3D-Resnet-18	Train1+Train2+Val1+Val2	Yes
4	3D-Resnet-50	Train1+Train2	Yes
5	Ensemble	-	-

Table 9. Covid-19 Domain Adaptation Challenge Submissions

Sub	Backbone	Training Data	Test Aug
1	3D-Resnet-18	Train1+Train2	Yes
2	3D-Resnet-18	Train1+Train2+Val1+Val2	Yes
3	3D-Resnet-18	Train1+Train2 +Pseudo Label	Yes
4	Ensemble 1	-	-
5	Ensemble 2	-	-

Table 11. Covid-19 Detection Challenge Submissions

Approach	F1-score	Non-Covid	Covid
FDVTS [30]	**77.55**	**96.97**	**58.14**
MDAP [28]	77.21	96.82	57.60
Deep-Adaptation	74.96	96.52	53.39
M2@Purdue [24]	65.79	91.92	39.66
baseline [22]	60.16	86.67	33.65

the segmented lung and infection, resulting in three input channels akin to color channels.

Moreover, we employed three distinct 3D CNN backbones to train Covid-19 recognition for both challenges: Customized Hybrid-DeCoVNet, as well as pretrained 3D-Resnet-19 and 3D-Resnet-50 models. To further enhance performance, we investigated ensemble approaches and testing augmentation techniques. Comparative analysis against baseline results demonstrates the significant efficiency of our proposed approach, exhibiting a substantial margin in terms of F1-score (14%). Our approach ranked second and third in the Covid-19 Detection and Covid-19 Domain Adaptation Challenges, respectively, based on the test data results. Our approach demonstrates improvements of 9.5% and 17% compared to baseline performance in these challenges. Furthermore, our approach exhibits very promising performance compared with the approaches of other competitors, underscoring the significance of the proposed training paradigm and the utilization of ensemble and testing augmentation techniques. Thus contributing to the ongoing efforts in combatting Covid-19 pandemic and the future ones.

References

[1] Anastasios Arsenos, Dimitrios Kollias, and Stefanos Kollias. A large imaging database and novel deep neural architecture for covid-19 diagnosis. In *2022 IEEE 14th Image, Video, and Multidimensional Signal Processing Workshop (IVMSP)*, page 1–5. IEEE, 2022.

[2] Anastasios Arsenos, Andjoli Davidhi, Dimitrios Kollias, Panos Prassopoulos, and Stefanos Kollias. Data-driven covid-19 detection through medical imaging. In *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, page 1–5. IEEE, 2023.

[3] F. Bougourzi, F. Dornaika, K. Mokrani, A. Taleb-Ahmed, and Y. Ruichek. Fusing Transformed Deep and Shallow features (FTDS) for image-based facial expression recognition. *Expert Systems with Applications*, 156:113459, 2020.

[4] Fares Bougourzi, Riccardo Contino, Cosimo Distante, and Abdelmalik Taleb-Ahmed. CNR-IEMN: A Deep Learning based approach to recognise COVID-19 from CT-scan. In *ICASSP 2021-2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 8568–8572. IEEE, 2021.

[5] Fares Bougourzi, Riccardo Contino, Cosimo Distante, and Abdelmalik Taleb-Ahmed. Recognition of COVID-19 from CT Scans Using Two-Stage Deep-Learning-Based Approach: CNR-IEMN. *Sensors*, 21(17):5878, 2021.

[6] F. Bougourzi, F. Dornaika, and A. Taleb-Ahmed. Deep Learning Based Face Beauty Prediction via Dynamic Robust Losses and Ensemble Regression. *Knowledge-Based Systems*, 242: 108246, 2022.

[7] Fares Bougourzi, Cosimo Distante, Fadi Dornaika, and Abdelmalik Taleb-Ahmed. CNR-IEMN-CD and CNR-IEMN-CSD Approaches for Covid-19 Detection and Covid-19 Severity Detection from 3D CT-scans. In *Computer Vision–ECCV 2022 Workshops: Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part VII*, pages 593–604. Springer, 2023.

[8] Fares Bougourzi, Cosimo Distante, Fadi Dornaika, and Abdelmalik Taleb-Ahmed. Pdatt-unet: Pyramid dual-decoder attention unet for covid-19 infection segmentation from ct-scans. *Medical Image Analysis*, page 102797, 2023.

[9] Fares Bougourzi, Fadi Dornaika, Amir Nakib, Cosimo Distante, and Abdelmalik Taleb-Ahmed. Deep-covid-sev: an ensemble 2d and 3d cnn-based approach for covid-19 severity prediction from 3d ct-scans. In *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, pages 1–5. IEEE, 2023.

[10] Fares Bougourzi, Fadi Dornaika, Amir Nakib, and Abdelmalik Taleb-Ahmed. Emb-trattunet: a novel edge loss function and transformer-CNN architecture for multi-classes pneumonia infection segmentation in low annotation regimes. *Artificial Intelligence Review*, 57(4):90, 2024.

[11] Ramy Farag, Parth Upadhyay, and Guilhermen DeSouza. Covid-19 detection from ct scans using efficientnet and attention mechanism. *arXiv preprint arXiv:2403.11505*, 2024.

[12] Hao Guan and Mingxia Liu. Domain adaptation for medical image analysis: a survey. *IEEE Transactions on Biomedical Engineering*, 69(3):1173–1185, 2021.

[13] Kensho Hara, Hirokatsu Kataoka, and Yutaka Satoh. Can spatiotemporal 3d cnns retrace the history of 2d cnns and imagenet? In *Proceedings of the IEEE conference on Computer Vision and Pattern Recognition*, pages 6546–6555, 2018.

[14] Chih-Chung Hsu, Chia-Ming Lee, Yang Fan Chiang, Yi-Shiuan Chou, Chih-Yu Jiang, Shen-Chieh Tai, and Chi-Han Tsai. Simple 2d convolutional neural network-based approach for covid-19 detection. *arXiv preprint arXiv:2403.11230*, 2024.

[15] Ying-Hui Jin, Lin Cai, Zhen-Shun Cheng, and others. A rapid advice guideline for the diagnosis and treatment of 2019 novel coronavirus (2019-nCoV) infected pneumonia (standard version). *Military Medical Research*, 7(1):4, 2020.

[16] 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.

[17] 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*, page 251–267, 2020.

[18] Dimitrios Kollias, Anastasios Arsenos, Levon Soukissian, and Stefanos Kollias. Mia-cov19d: Covid-19 detection through 3-d chest ct image analysis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, page 537–544, 2021.

[19] Dimitrios Kollias, Anastasios Arsenos, and Stefanos Kollias. Ai-mia: Covid-19 detection & severity analysis through medical imaging. *arXiv preprint arXiv:2206.04732*, 2022.

[20] Dimitrios Kollias, Anastasios Arsenos, and Stefanos Kollias. Ai-enabled analysis of 3-d ct scans for diagnosis of covid-19 & its severity. In *2023 IEEE International Conference on Acoustics, Speech, and Signal Processing Workshops (ICASSPW)*, page 1–5. IEEE, 2023.

[21] Dimitrios Kollias, Anastasios Arsenos, and Stefanos Kollias. A deep neural architecture for harmonizing 3-d input data analysis and decision making in medical imaging. *Neurocomputing*, 542:126244, 2023.

[22] Dimitrios Kollias, Anastasios Arsenos, and Stefanos Kollias. Domain adaptation, explainability & fairness in ai for medical image analysis: Diagnosis of covid-19 based on 3-d chest ct-scans. *arXiv preprint arXiv:2403.02192*, 2024.

[23] Qingqiu Li, Runtian Yuan, Junlin Hou, Jilan Xu, Yuejie Zhang, Rui Feng, and Hao Chen. Advancing covid-19 detection in 3d ct scans. *arXiv preprint arXiv:2403.11953*, 2024.

[24] Li Lin, Yamini Sri Krubha, Zhenhuan Yang, Cheng Ren, Xin Wang, and Shu Hu. Robust covid-19 detection in ct images with clip. *arXiv preprint arXiv:2403.08947*, 2024.

[25] Xiaohong Liu, Kai Wang, Ke Wang, Ting Chen, Kang Zhang, and Guangyu Wang. Kiseg: a three-stage segmentation framework for multi-level acceleration of chest ct scans from covid-19 patients. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 25–34. Springer, 2020.

[26] Ozan Oktay, Jo Schlemper, Loic Le Folgoc, Matthew Lee, Mattias Heinrich, Kazunari Misawa, Kensaku Mori, Steven McDonagh, Nils Y Hammerla, Bernhard Kainz, et al. Attention u-net: Learning where to look for the pancreas. *arXiv preprint arXiv:1804.03999*, 2018.

[27] RADIOLOGISTS. COVID-19 CT-scans segmentation datasets, available at: `http://medicalsegmentation.com/covid19/`, 2019. Last visited: 18-08-2021.

[28] Robert Turnbull and Simon Mutch. High-confidence pseudo-labels for domain adaptation in covid-19 detection. *arXiv preprint arXiv:2403.13509*, 2024.

[29] Saining Xie, Ross Girshick, Piotr Dollár, Zhuowen Tu, and Kaiming He. Aggregated residual transformations for deep neural networks, 2017.

[30] Runtian Yuan, Qingqiu Li, Junlin Hou, Jilan Xu, Yuejie Zhang, Rui Feng, and Hao Chen. Domain adaptation using pseudo labels for covid-19 detection. *arXiv preprint arXiv:2403.11498*, 2024.