Domain Adaptation Using Pseudo Labels for COVID-19 Detection

Runtian Yuan[1], Qingqiu Li[2], Junlin Hou[3], Jilan Xu[1], Yuejie Zhang[1*], Rui Feng[1*], Hao Chen[3*]

[1]School of Computer Science, Fudan University, China [2]School of Academy for Engineering and Technology, Fudan University, China [3]Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, China

{rtyuan21,qqli22,jilanxu18,yjzhang,fengrui}@fudan.edu.cn,csejlhou@ust.hk,jhc@cse.ust.hk

Abstract

Deep learning has offered advanced analytical capabilities to enhance the accuracy and efficiency of detecting COVID-19 through complex pattern recognition in medical imaging data. However, the variability across datasets from different domains poses a significant challenge to the generalization abilities of deep learning models. In this paper, we propose a novel two-stage framework for domain adaptation of COVID-19 detection. Initially, We train a model on annotated data from both domains, integrating contrastive representation learning and a modified version of CORAL loss to minimize domain discrepancies. In the subsequent stage, we employ a pseudo-labeling strategy to effectively utilize non-annotated data from the target domain, further enhancing the model's adaptability and performance. The effectiveness of our approach is demonstrated through extensive experiments, showing significant improvements in COVID-19 detection performance compared to the baseline model. On the COVID-19 domain adaptation leaderboard in the 4th COV19D Competition, our approach ranked 1st with a Macro F1 Score of 77.55%.

1. Introduction

The COVID-19 (Coronavirus Disease 2019) pandemic underscores the need for rapid and accurate diagnostic techniques, which are crucial for the early identification of COVID-19, ensuring timely intervention and better patient outcomes. Computed Tomography (CT) imaging has proven to be an essential tool in the detection of the disease, capable of revealing pulmonary manifestations such as ground-glass opacities and bilateral infiltrates characteristic of COVID-19. However, the interpretation of CT images requires significant expertise and can be time-consuming, presenting challenges in high-demand scenarios and potentially leading to delays in diagnosis and treatment.

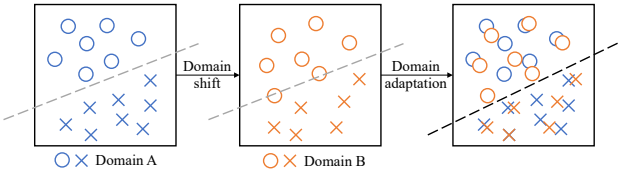

Figure 1. An illustration of domain shift and domain adaptation highlights the divergent data distributions between Domain A and Domain B. Consequently, models trained on Domain A necessitate the application of domain adaptation strategies to effectively adjust and perform optimally in Domain B.

Recent deep-learning methods have been widely explored in solving COVID-19 detection [6–8, 12, 14, 17, 22]. However, most of them require large amount of annotated data, which is challenging to acquire from various domains in real-world settings. This scarcity of data becomes more problematic when considering the phenomenon of domain shift, as illustrated in Figure 1. Domain shift occurs when the distribution of data encountered during the deployment of a machine learning or deep learning model differs from the distribution of the training data. This mismatch can lead to a decline in the model's performance because the patterns the model learned during training may not accurately represent the patterns in the new data. Domain adaptation is a set of techniques aimed at addressing this challenge by enabling a model to adapt to the new data distribution. These techniques can involve retraining the model with a mix of original and new domain data, applying transfer learning to adjust the model to the characteristics of the new data, or employing unsupervised methods to learn representations that are invariant across the two domains, thereby maintaining performance even in the face of domain shift.

Confronted with the challenge posed by the scarcity of annotated data, our approach leverages a modified version of CORAL loss and the pseudo-labeling strategy to facilitate effective domain adaptation, thereby substantially enhancing the detection of COVID-19 from CT images. By

integrating the modified CORAL loss, we aim to minimize the discrepancy in feature distributions between different domains, ensuring that the model can generalize well across varied datasets for more accurate, robust, and adaptable COVID-19 detection. Concurrently, the employment of pseudo labels for non-annotated data in the target domain serves to augment the training set, providing valuable information that enables the model to learn more comprehensive and representative features of the COVID-19 pathology.

In this paper, we present our framework for the domain adaptation challenge of the 4th COV19D Competition [18]. The proposed model aims to achieve a high level of diagnostic precision while significantly reducing the time and resources required for COVID-19 detection via CT scans. This advancement holds the promise of enabling faster, more accurate, and scalable diagnostic solutions.

2. Related Work

2.1. COVID-19 Detection

The detection and diagnosis of COVID-19 have been pivotal since the onset of the pandemic. Various methodologies have been proposed to improve detection rates and reduce diagnostic time. Several studies have focused on the application of deep learning techniques for the detection of COVID-19 from X-rays. [20] and [9] trained Convolutional Neural Networks using chest X-ray images, achieving encouraging performance for COVID-19 detection.

Another significant portion of research has been directed towards the development of AI-driven tools for the automatic detection of COVID-19 using CT scans, recognizing the potential for rapid and scalable screening. He et al. [4] proposed a Self-Trans approach, which synergistically integrates contrastive self-supervised learning to learn powerful and unbiased feature representations for reducing the risk of over-fitting, showing promising results in enhancing diagnostic accuracy and efficiency. Wang et al. [26] collected 1065 CT images of pathogen-confirmed COVID-19 cases along with those previously diagnosed with typical viral pneumonia, and modified the inception transfer-learning model to extract radiological features for timely and accurate COVID-19 diagnosis. Gupta et al. [3] developed a new lightweight, less complex deep learning model for the automated screening of COVID-19, using a repeated 10-fold holdout cross-validation scheme. These studies underscore the potential of deep learning in augmenting traditional diagnostic methods for COVID-19, leading to faster and more reliable detection processes.

2.2. Domain Adaptation

Domain adaptation is a crucial area of research in machine learning and computer vision, aimed at addressing the challenge of domain shift, where the distribution of test data differs from the training data. This discrepancy can significantly impair the performance of models when applied to real-world scenarios. A variety of strategies have been developed to mitigate these effects, focusing on enabling models to generalize across different domains.

In the realm of domain adaptation, several approaches have been developed to minimize the domain discrepancy by aligning feature distributions between the source and target domains. Sun and Saenko [23] introduced Deep CORAL, extending the CORAL [24] approach to deep learning by aligning the deep feature distributions of source and target domains directly through the minimization of the difference in their covariance matrices. Tzeng et al. [25] proposed a method called Deep Domain Confusion (DDC), which incorporates a domain confusion loss to encourage the learning of domain-invariant features by maximizing the confusion between the source and target domains. Additionally, Zhu et al. [31] designed a deep subdomain adaptation network that learns a transfer network by aligning the relevant subdomain distributions of domain-specific layer activations across different domains based on a local maximum mean discrepancy (LMMD). These methods share a common goal of reducing domain shift at the feature level, thereby improving the model's generalization capabilities across different domains.

Pseudo-labeling, as a semi-supervised learning technique, has been widely used in domain adaptation, leveraging unlabeled data to improve model generalization. A significant contribution is the work by Lee [19] on pseudo-labeling, which formalized the approach of using the model's own predictions as labels to augment the training set in semi-supervised learning. Saito et al. [21] introduced asymmetric tri-training for unsupervised domain adaptation, employing three networks where two are used to generate pseudo-labels for training the third, iteratively refining the model with less domain bias. More recently, Yu et al. [27] viewed the source data as a noisily-labeled version of the ideal target data and then proposed an semi-supervised domain adaptation model that cleans up the label noise dynamically with the help of a robust cleaner component designed from the target perspective.

In our work, we design a modified version of CORAL loss between annotated data from both domain A and B as well as non-annotated data from domain B. Additionally, we generate pseudo labels to merge with labeled data, and then finetune our model to effectively improve the classification performance of COVID-19.

3. Methodology

Figure 2 provides a schematic representation of our framework, which is structured into two stages. In the initial stage, the model is trained on annotated data from both Domain A and Domain B, employing techniques such as con-

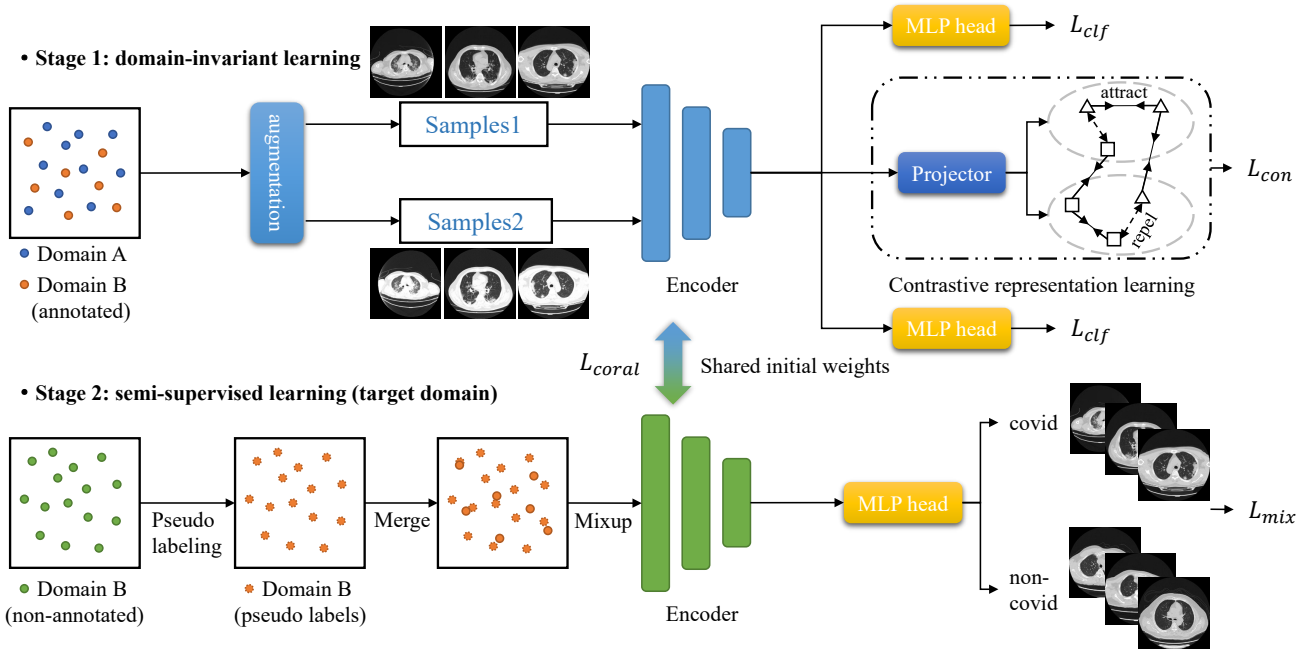

Figure 2. Overview of our two-stage framework for COVID-19 domain adaptation.

trastive representation learning and modified CORAL loss to enhance its learning efficacy. The subsequent stage involves leveraging the pre-trained model from the first stage to generate pseudo labels for the non-annotated data in Domain B. This augmented dataset, comprising both pseudo-labeled and originally annotated data, is then utilized for improved COVID-19 detection. Our methodology builds upon the foundation laid by the previous CMC network [5], which has already demonstrated its excellent performance in accurately detecting COVID-19.

Stage 1: domain-invariant learning. The initial stage of our framework is designed to develop a model that efficiently absorbs general class information from annotated data. This foundational phase is structured through the integration of five key components:

- A data augmentation module $A(\cdot)$, which transforms an input CT sample x into a randomly augmented sample $\tilde{x}$ from random crops and random changes in brightness and contrast, introducing variability and robustness into the training process.
- A base encoder $E(\cdot)$, responsible for extracting vital features from the augmented data and acting as the basis for meaningful representation extraction, maps the augmented CT sample $\tilde{x}$ to a representation vector $r = E(\tilde{x}) \in \mathbb{R}^{d_e}$ in the d_e-dimensional latent space.
- A projection network $P(\cdot)$, which further processes these representations to support contrastive learning by mapping the representation vector r to a relative low-dimension vector $z = P(r) \in \mathbb{R}^{d_p}$.
- A multi-layer perception (MLP) head $M(\cdot)$, which is followed by a Softmax operation to classify the representation vector $r \in \mathbb{R}^{d_e}$ into COVID or Non-COVID. The classification loss $\mathcal{L}_{clf}$ is implemented by the standard cross-entropy loss between predicted probabilities and ground truth labels.
- A modified CORAL loss function $\mathcal{L}_{coral}$ to align the deep feature distributions of the two domains, encompassing both annotated data from Domains A and B as well as non-annotated data from Domain B.

Contrastive representation learning. The first three elements collectively execute a contrastive representation learning approach. Given a minibatch of N CT volumes and their labels $\{(x_i, y_i)\}_{i=1,\ldots,N}$, we can generate a minibatch of $2N$ samples $\{(\tilde{x}_i, \tilde{y}_i)\}_{i=1,\ldots,2N}$ after data augmentations. We define the positives as any augmented CT samples from the same category, while those from different classes are considered as negative pairs [11]. Therefore, the contrastive loss is defined as:

$$\mathcal{L}_{con} = \frac{1}{2N} \sum_{i=1}^{2N} \mathcal{L}_{con}^{i},$$

$$\mathcal{L}_{con}^{i} = \frac{-1}{2N_{\tilde{y}_i} - 1} \sum_{j=1}^{2N} \cdot \log \frac{\exp\left(z_i^T \cdot z_j / \tau\right)}{\sum_{k=1}^{2N} \mathbb{1}_{i \neq k} \cdot \exp\left(z_i^T \cdot z_k / \tau\right)}, \quad \text{when } i \neq j \text{ and } \tilde{y}_i = \tilde{y}_j, \tag{1}$$

where $N_{\tilde{y}_i}$ is the number of samples in a minibatch that share the same label $\tilde{y}_i$, $\mathbb{1} \in \{0, 1\}$ is an indicator func-

tion, and $\tau > 0$ is a scalar temperature hyper-parameter. Contrastive representation learning prepares the model to discern and learn from the intrinsic patterns and differences within the data, setting a robust groundwork for the framework's subsequent stages.

Modified CORAL loss. Deep CORAL [23] seeks to align the deep feature distributions between the source and target domains by minimizing the disparities in their covariance matrices. In our work, we have further refined this approach by computing the CORAL loss across a broader spectrum of data, encompassing both annotated data from Domains A and B as well as non-annotated data from Domain B. Given annotated training samples $D_L = \{r_i^l\}_{i=1,\dots,N_L}, r^l \in \mathbb{R}^{d_e}$ with labels $\{y_i\}_{i=1,\dots,N_L}$, and non-annotated data $D_U = \{r_i^u\}_{i=1,\dots,N_U}, r^u \in \mathbb{R}^{d_e}$, where N_L and N_U denotes the number of samples in D_L and D_U, respectively. Note that D_L consists of annotated data from both domain A and domain B. Here both r^l and r^u are the d_e-dimensional feature representations extracted by the base encoder. The covariance matrices of D_L and D_U are given by:

$$C_L = \frac{1}{N_L - 1}(D_L^\top D_L - \frac{1}{N_L}(\mathbf{1}^\top D_L)^\top(\mathbf{1}^\top D_L)), \quad (2)$$

$$C_U = \frac{1}{N_U - 1}(D_U^\top D_U - \frac{1}{N_U}(\mathbf{1}^\top D_U)^\top(\mathbf{1}^\top D_U)), \quad (3)$$

where $\mathbf{1}$ is a column vector with all elements equal to 1. Then the modified CORAL loss is defined as the distance between the covariance of the features:

$$\mathcal{L}_{coral} = \frac{1}{4d^2} \parallel C_L - C_U \parallel_F^2, \quad (4)$$

where d is the channel after the features are reshaped into two-dimensional vectors, and $\parallel \cdot \parallel_F^2$ denotes the squared matrix Frobenius norm. Suppose r_{ij}^l and r_{ij}^u indicate the j-th dimension of the i-th sample in D_L and D_U, respectively. According to [23], the gradient with respect to the features can be calculated as follows:

$$\frac{\partial \mathcal{L}_{coral}}{\partial r_{ij}^l} = \frac{1}{d^2(N_L - 1)}((r^{l\top} - \frac{1}{N_L}(\mathbf{1}^\top r^l)^\top \mathbf{1}^\top)^\top (C_L - C_U))_{ij} \quad (5)$$

$$\frac{\partial \mathcal{L}_{coral}}{\partial r_{ij}^u} = \frac{1}{d^2(N_U - 1)}((r^{u\top} - \frac{1}{N_U}(\mathbf{1}^\top r^u)^\top \mathbf{1}^\top)^\top (C_L - C_U))_{ij} \quad (6)$$

By integrating the CORAL loss in this manner, we enable the model to extract and leverage the underlying commonalities and differences more effectively, fostering a learning environment that is rich, adaptive, and ultimately more conducive to discerning the complex patterns associated with COVID-19 detection. This enhancement not only amplifies the model's ability to generalize across different domains but also enriches its interpretative depth, making its learning process significantly more meaningful and robust. Our modified CORAL loss function plays a crucial role in bridging the gap between domains, ensuring that the learned representations are domain-invariant and highly informative.

The total loss is a combination of the three loss functions with adaptive weights [10], which is defined as:

$$\mathcal{L}_{total} = \frac{1}{\sigma_1^2}\mathcal{L}_{clf} + \frac{1}{\sigma_2^2}\mathcal{L}_{con} + \mathcal{L}_{coral} + \log \sigma_1 + \log \sigma_2, \quad (7)$$

where σ_1 and σ_2 are utilized to learn the relative weights of the classification loss and contrastive loss adaptively.

Stage 2: semi-supervised learning (target domain). In the next stage, we use the trained model to tackle the non-annotated data in Domain B. Here, the model acts like a teacher, assigning "pseudo labels" to this unlabeled data, basically making educated guesses based on what it learned in the first stage. We then merge this newly labeled data with the original annotated dataset in domain B, resulting in a bigger and more diverse set of dataset.

The constructed dataset is a goldmine for improving our model's ability to detect COVID-19 in domain B. Given a minibatch of N_M CT volumes and their labels $\{(x_i, y_i)\}_{i=1,\dots,N_M}$, we then adopt the mixup [28] strategy to further boost the generalization ability of the model. Mixup can be understood as a form of data augmentation to provide a smoother estimate of uncertainty. For each CT sample x_i, the mixup sample and its label are generated as:

$$\begin{aligned} x_i^{mix} &= \lambda x_i + (1 - \lambda)x_j, \\ y_i^{mix} &= \lambda y_i + (1 - \lambda)y_j, \end{aligned} \quad (8)$$

where j is a randomly selected index in the training data, and $\lambda \in [0, 1]$. The mixup loss is defined as the cross-entropy loss of mixup samples:

$$\mathcal{L}_{mix} = \frac{1}{N_M}\sum_{i=1}^{N_M} \text{CrossEntropy}(x_i^{mix}, y_i^{mix}). \quad (9)$$

Mixup generates a newly augmented dataset that embodies the nuances and variability of the original data in a more continuous space by creating synthetic data points through the combination of pairs of images and their labels. This process not only diversifies the training data but also encourages the model to learn more generalized features rather than memorizing specific data points. We effectively broaden the model's exposure to a wide range of data interpolations. As a result, the model's predictions across different inputs become more reliable and less prone to overfitting, providing a more smooth estimation of uncertainty, making it more accurate in spotting COVID-19.

Set	Type	Domain A	Domain B
Training	COVID-19	703	120
	Non-COVID-19	655	119
	Non-annotated	-	494
Validation	COVID-19	-	65
	Non-COVID-19	-	113
Testing	-	-	4055

Table 1. Data samples for training, validation and testing.

4. Dataset

We train and evaluate our framework on the COV19-CT-DB database [17]. COV19-CT-DB contains 3D chest CT scans that collected in various medical centers. The database includes 7,756 3D CT scans: 1,661 are COVID-19 samples, whilst 6,095 refer to non COVID-19 ones. There are about 2,500,000 images included in these datasets. All have been anonymized. 724,273 images refer to the COVID-19 class, whilst 1,775,727 slices belong to non COVID-19 class [1, 2, 12–16].

For the COVID-19 domain adaptation challenge, the training set contains 239 annotated 3D CT scans (120 COVID-19 cases and 119 Non-COVID-19 cases) and 494 non-annotated 3D CT scans from domain B, as well as 1358 annotated 3D CT scans (703 COVID-19 cases and 655 Non-COVID-19 cases) from domain A. The validation set consists of 178 3D CT scans (65 COVID-19 cases and 113 Non-COVID-19 cases). The testing set includes 4055 3D CT scans and the labels are not available during the challenge, as shown in Table 1.

5. Experiments

5.1. Pre-processing

Our data pre-processing procedure follows [5] and [7]. First, each sequence of 2D chest CT slices is composed into a 3D volume of shape (D, H, W), where D, H, W denotes the number of slices, height, and width, respectively. Then, each volume is resized from its original size to $(128, 256, 256)$. Finally, we transform the CT volume to the interval $[0, 1]$ for intensity normalization.

5.2. Implementation Details

We apply inflated 3D ResNeSt50 [29] as the encoder in our experiments. The value of parameter d_e is 2,048 and d_p is 128. We optimize the network using the Adam algorithm with a weight decay of 10^{-5}. The network is trained for 100 epochs. The initial learning rate is set to 0.0001 and then divided by 10 at 30% and 80% of the total number of training epochs. Our methods are implemented in PyTorch and run on four NVIDIA Tesla V100 GPUs.

5.3. Evaluation Metrics

We adopt the Macro F1 Score as the evaluation metric. The Macro F1 Score is defined as the unweighted average of the class-wise/label-wise F1 Scores. The F1 Score for the i-th class is defined as:

$$\text{F1-Score}_i = 2 \times \frac{\text{Recall}_i \times \text{Precision}_i}{\text{Recall}_i + \text{Precision}_i}. \tag{10}$$

The Macro F1 Score is the unweighted average of F1 Scores for all classes, i.e., the F1 Scores for COVID-19 and Non-COVID-19, which can be formulated as:

$$\text{Macro-F1} = \frac{1}{C} \sum_{i=1}^{C} \text{F1-Score}_i, \tag{11}$$

where C denotes the number of classes. The Macro F1 Score remains robust in the face of data imbalance, maintaining its reliability as a performance metric without being skewed by disparities in class distribution.

5.4. Results on the Validation Set

Table 2 shows the results of the baseline model and our method on the validation set of COVID-19 Domain Adaptation Challenge. The baseline model [18] employs Monte Carlo Dropout to assess uncertainty while training the CNN-RNN architecture using data from both domain A (annotated) and domain B (annotated), and then annotates the non-annotated data from domain B based on the model's predictions. We integrate the CMC architecture with modified CORAL loss and pseudo-labeling strategy, and derive an ensemble model by averaging the predictions from five individual models, achieving the detection performance with a Macro F1 Score of 92.68%. The confusion matrix on the validation set is shown in Figure 3.

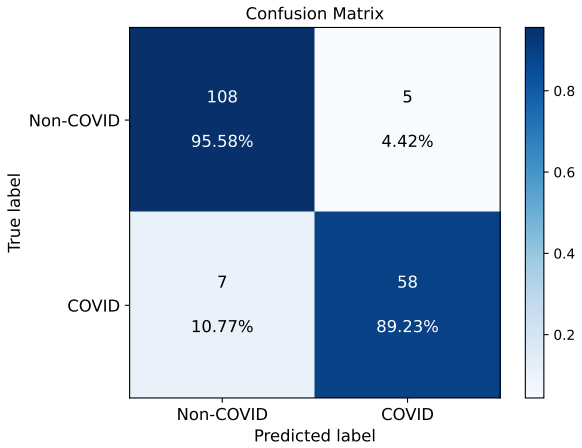

Figure 3. The confusion matrix of ensemble model's predictions.

Methods	Pseudo labels	L_{clf}	L_{con}	L_{mix}	L_{coral}	Accuracy	Macro F1 Score	F1 Non-COVID	F1 COVID
baseline [18]	-	-	-	-	-	-	73.00	-	-
model1	×	✓	×	×	×	91.01	90.43	92.79	88.06
model2	×	✓	×	×	✓	91.57	90.61	93.62	87.60
model3	×	✓	✓	✓	×	92.13	91.34	93.97	88.71
model4	✓	✓	×	✓	×	93.26	92.52	94.87	90.16
model5	✓	✓	✓	✓	×	92.70	92.10	94.27	89.92
ensemble	-	-	-	-	-	93.26	92.68	94.74	90.63

Table 2. The comparison results on the validation set of COVID-19 Domain Adaptation Challenge.

Rank	Teams	Macro F1	F1(NC)	F1(C)
1	**FDVTS (Ours)**	**77.55**	**96.97**	**58.14**
2	MDAP	77.21	96.82	57.60
3	Deep-Adaptation	74.96	96.52	53.39
4	M2@Purdue	65.79	91.92	39.66
5	baseline	60.16	86.67	33.65

Table 3. The learderboard of COVID-19 Domain Adaptation Challenge. F1(NC) and F1(C) denote the F1-Score on Non-COVID and COVID cases, respectively.

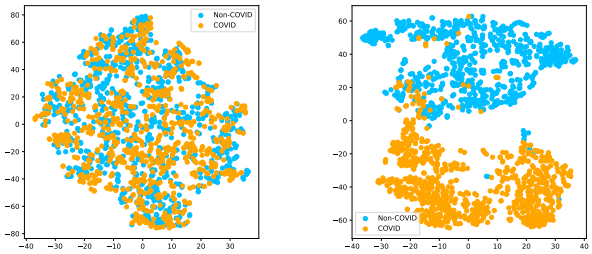

(a) Prior to Stage 1. (b) Subsequent to Stage 1.

Figure 4. t-SNE visualizations on the training set.

5.5. Results on the Domain Adaptation Challenge

In the COVID-19 domain adaptation challenge, as detailed in Table 3, our team, FDVTS, won the top position with an impressive Macro F1 Score of 77.55%. This achievement underscores our model's superior capability in handling the nuanced task of COVID-19 detection across different domains. Particularly noteworthy is our model's exceptional performance in accurately identifying Non-COVID (NC) cases with an F1 Score of 96.97%, coupled with a robust F1 Score of 58.14% in detecting COVID (C) cases, significantly outperforming the baseline's score by 24.49%. The leaderboard reflects the critical importance of domain adaptation techniques in managing class imbalances and enhancing model performance in the context of COVID-19 detection. Our team's leading position underscores the effectiveness of our approach in addressing the complexities of this task, achieving a high accuracy in Non-COVID case identification and a strong capability in COVID case detection.

5.6. Visualization Results

Figure 4 provides t-SNE visualizations that encapsulates the distribution and separation of the training dataset, encompassing data from both Domain A and Domain B. By projecting the high-dimensional feature maps generated by the encoder into a two-dimensional space, this visualization offers an intuitive and accessible representation of how the data clusters in Stage 1 and how distinct the separations are between different categories.

Figure 5 presents the qualitative results obtained from evaluating COVID-19 samples in the validation set, offering insightful visual evidence of our model's diagnostic capabilities. The images are accompanied by Class Activation Mapping (CAM) [30] that underline the areas of interest that the model focuses on when making a diagnosis, such as ground-glass opacities, consolidation, or other pulmonary abnormalities commonly associated with COVID-19. Our model can highlight specific features and patterns that are indicative of the infection.

6. Conclusion

In this paper, we present our solution for COVID-19 domain adaptation challenge in the 4th COV19D Competition. Our two-stage framework begins with training a model on annotated data, leveraging contrastive representation learning alongside a modified CORAL loss to reduce domain discrepancies by aligning feature distribution covariances. The strategy was further refined in the second stage through the use of pseudo-labeling, which capitalized on non-annotated data from the target domain to bolster the model's adaptability and overall performance. The efficacy of our framework is underscored by comprehensive experimental results, which illustrate a marked enhancement in COVID-19 detection capabilities. This superiority was further validated by our first-place ranking on the COVID-19 domain adaptation leaderboard in the 4th COV19D Competition, achieving a Macro F1 Score of 77.55%.

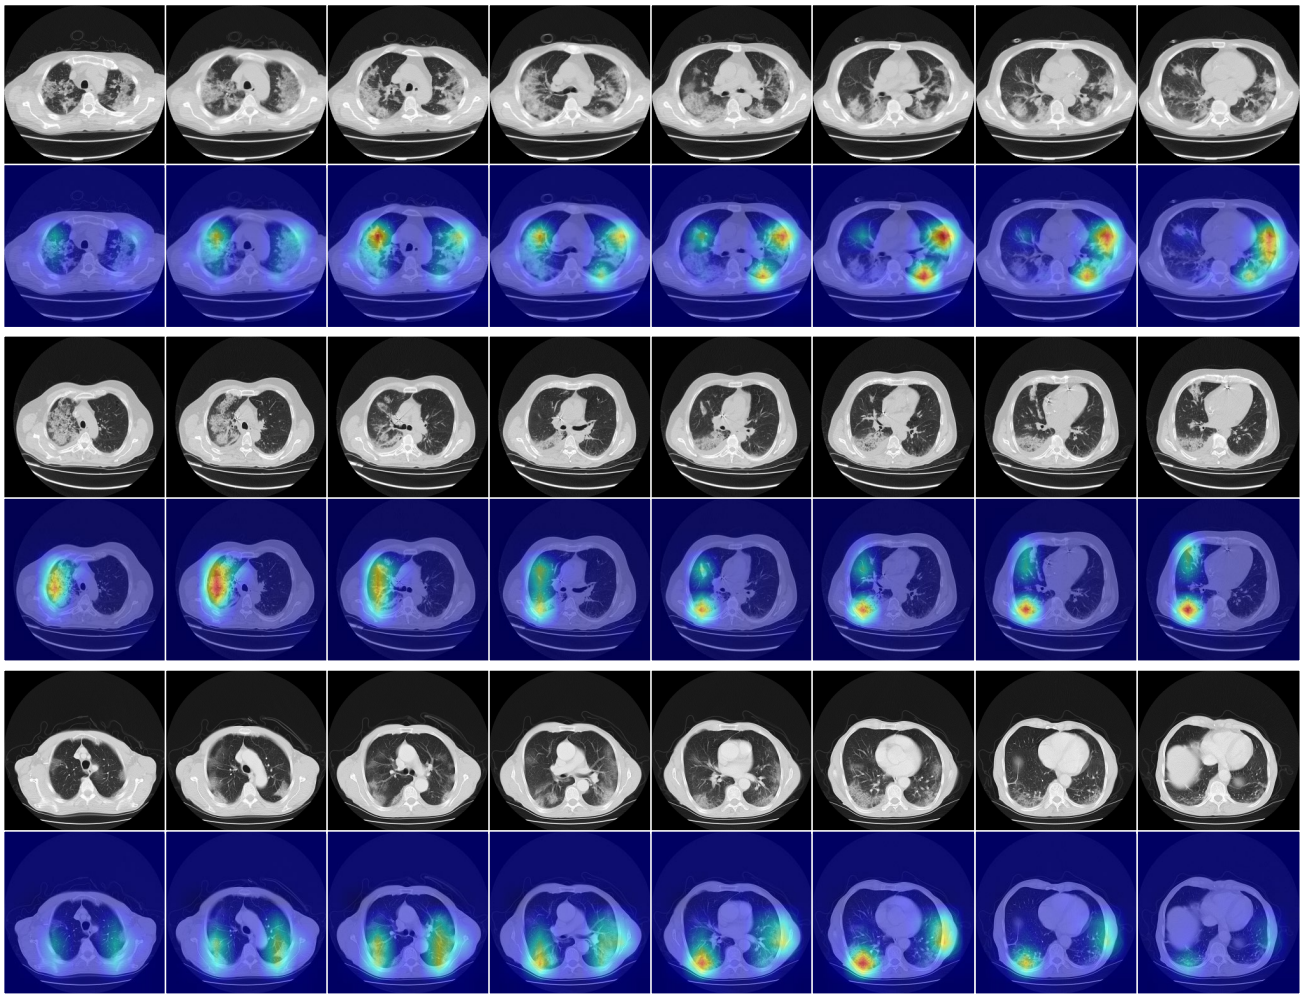

Figure 5. CAM visualization results on COVID-19 CT scans.

Acknowledgement

This work was supported by National Science and Technology Major Project (No. 2021ZD0114001), Natural Science Foundation of China (No. 62172101), the Science and Technology Commission of Shanghai Municipality (No. 21511101000; No. 23511100602), and the Hong Kong Innovation and Technology Fund (Project No. ITS/028/21FP). Yuejie Zhang, Rui Feng and Hao Chen are corresponding authors.

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. 5

[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. 5

[3] Kapil Gupta and Varun Bajaj. Deep learning models-based ct-scan image classification for automated screening of covid-19. *Biomedical Signal Processing and Control*, 80: 104268, 2023. 2

[4] Xuehai He, Xingyi Yang, Shanghang Zhang, Jinyu Zhao, Yichen Zhang, Eric Xing, and Pengtao Xie. Sample-efficient deep learning for covid-19 diagnosis based on ct scans. *medrxiv*, pages 2020–04, 2020. 2

[5] Junlin Hou, Jilan Xu, Rui Feng, Yuejie Zhang, Fei Shan, and Weiya Shi. Cmc-cov19d: Contrastive mixup classification for covid-19 diagnosis. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 454–461, 2021. 3, 5

[6] Junlin Hou, Jilan Xu, Longquan Jiang, Shanshan Du,

Rui Feng, Yuejie Zhang, Fei Shan, and Xiangyang Xue. Periphery-aware covid-19 diagnosis with contrastive representation enhancement. *Pattern Recognition*, 118:108005, 2021. 1

[7] Junlin Hou, Jilan Xu, Nan Zhang, Yi Wang, Yuejie Zhang, Xiaobo Zhang, and Rui Feng. Cmc_v2: Towards more accurate covid-19 detection with discriminative video priors. In *European Conference on Computer Vision*, pages 485–499. Springer, 2022. 5

[8] Junlin Hou, Jilan Xu, Nan Zhang, Yuejie Zhang, Xiaobo Zhang, and Rui Feng. Boosting covid-19 severity detection with infection-aware contrastive mixup classification. In *European Conference on Computer Vision*, pages 537–551. Springer, 2022. 1

[9] Emtiaz Hussain, Mahmudul Hasan, Md Anisur Rahman, Ickjai Lee, Tasmi Tamanna, and Mohammad Zavid Parvez. Corodet: A deep learning based classification for covid-19 detection using chest x-ray images. *Chaos, Solitons & Fractals*, 142:110495, 2021. 2

[10] Alex Kendall, Yarin Gal, and Roberto Cipolla. Multi-task learning using uncertainty to weigh losses for scene geometry and semantics. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7482–7491, 2018. 4

[11] Prannay Khosla, Piotr Teterwak, Chen Wang, Aaron Sarna, Yonglong Tian, Phillip Isola, Aaron Maschinot, Ce Liu, and Dilip Krishnan. Supervised contrastive learning. *Advances in neural information processing systems*, 33:18661–18673, 2020. 3

[12] 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. 1, 5

[13] 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.

[14] 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. 1

[15] Dimitrios Kollias, Anastasios Arsenos, and Stefanos Kollias. Ai-mia: Covid-19 detection and severity analysis through medical imaging. In *European Conference on Computer Vision*, page 677–690. Springer, 2022.

[16] 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. 5

[17] 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. 1, 5

[18] 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. 2, 5, 6

[19] Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, page 896. Atlanta, 2013. 2

[20] Shervin Minaee, Rahele Kafieh, Milan Sonka, Shakib Yazdani, and Ghazaleh Jamalipour Soufi. Deep-covid: Predicting covid-19 from chest x-ray images using deep transfer learning. *Medical image analysis*, 65:101794, 2020. 2

[21] Kuniaki Saito, Yoshitaka Ushiku, and Tatsuya Harada. Asymmetric tri-training for unsupervised domain adaptation. In *International conference on machine learning*, pages 2988–2997. PMLR, 2017. 2

[22] Nandhini Subramanian, Omar Elharrouss, Somaya Al-Maadeed, and Muhammed Chowdhury. A review of deep learning-based detection methods for covid-19. *Computers in Biology and Medicine*, 143:105233, 2022. 1

[23] Baochen Sun and Kate Saenko. Deep coral: Correlation alignment for deep domain adaptation. In *European Conference on Computer Vision*, pages 443–450. Springer, 2016. 2, 4

[24] Baochen Sun, Jiashi Feng, and Kate Saenko. Return of frustratingly easy domain adaptation. In *Proceedings of the AAAI conference on artificial intelligence*, 2016. 2

[25] Eric Tzeng, Judy Hoffman, Ning Zhang, Kate Saenko, and Trevor Darrell. Deep domain confusion: Maximizing for domain invariance. *arXiv preprint arXiv:1412.3474*, 2014. 2

[26] Shuai Wang, Bo Kang, Jinlu Ma, Xianjun Zeng, Mingming Xiao, Jia Guo, Mengjiao Cai, Jingyi Yang, Yaodong Li, Xiangfei Meng, et al. A deep learning algorithm using ct images to screen for corona virus disease (covid-19). *European radiology*, 31:6096–6104, 2021. 2

[27] Yu-Chu Yu and Hsuan-Tien Lin. Semi-supervised domain adaptation with source label adaptation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 24100–24109, 2023. 2

[28] Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017. 4

[29] Hang Zhang, Chongruo Wu, Zhongyue Zhang, Yi Zhu, Haibin Lin, Zhi Zhang, Yue Sun, Tong He, Jonas Mueller, R Manmatha, et al. Resnest: Split-attention networks. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 2736–2746, 2022. 5

[30] Bolei Zhou, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba. Learning deep features for discriminative localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2921–2929, 2016. 6

[31] Yongchun Zhu, Fuzhen Zhuang, Jindong Wang, Guolin Ke, Jingwu Chen, Jiang Bian, Hui Xiong, and Qing He. Deep subdomain adaptation network for image classification. *IEEE transactions on neural networks and learning systems*, 32(4):1713–1722, 2020. 2