ISSN: 2690-8662

Journal of Clinical Images

Open Access | Research Article

COVID-Net UV: An End-to-End Hybrid Deep Neural Network for Fast and Accurate COVID-19 Screening via Ultrasound Videos

Hilda Azimi[1]*; Alexander Wong[2]; Ashkan Ebadi[3]

[1]Digital Technologies, National Research Council Canada, Ottawa, ON K1A 0R6, Canada.

[2]Systems Design Engineering, University of Waterloo, Waterloo, ON N2L 3G1, Canada.

[3]Digital Technologies, National Research Council Canada, Toronto, ON M5T 3J1, Canada.

***Corresponding Author(s): Hilda Azimi**

Digital Technologies, National Research Council Canada, Ottawa, ON K1A 0R6, Canada.

Email: hilda.azimi@nrc-cnrc.gc.ca

Received: Sep 16, 2024

Accepted: Oct 07, 2024

Published Online: Oct 14, 2024

Journal: Journal of Clinical Images

Publisher: MedDocs Publishers LLC

Online edition: http://meddocsonline.org/

Abstract

Besides vaccination, as a highly effective method to mitigate the spread of COVID-19, fast and accurate screening of individuals remains crucial to ensuring public health safety, especially in high-risk or resource-limited environments. To address this need, we propose COVID-Net UV, an end-to-end hybrid spatiotemporal deep neural network architecture, to detect COVID-19 infection from lung point-of-care ultrasound videos captured by convex transducers. COVID-Net UV comprises a convolutional neural network that extracts spatial features and a recurrent neural network that learns temporal dependence. This hybrid approach enabled the network to accurately capture subtle patterns in lung dynamics associated with COVID-19. Through extensive experimentation and hyper parameter tuning, the network achieves a high average accuracy of 94.44%, with the critical advantage of producing no false-negative cases. This ensures that the model minimizes the risk of missed diagnoses, a crucial factor in controlling the spread of the virus. The goal of this work is to provide a robust tool to assist healthcare professionals in screening for COVID-19, supporting faster decision-making and improving patient outcomes through early detection. Ultimately, COVID-Net UV demonstrates the potential of combining deep learning with medical imaging to develop practical, life-saving diagnostic tools.

Introduction

The Coronavirus Disease 2019 (COVID-19) triggered a global health crisis in 2020, leading to an immense loss of life and overwhelming healthcare systems worldwide. Despite the widespread vaccination efforts that have significantly mitigated the spread of the virus, the emergence of new variants and fluctuating case numbers underscore the need for rapid and accurate screening methods to ensure public health safety [1]. Even in the post-vaccination era, fast and reliable screening remains a cornerstone in the management of COVID-19, playing a vital role in early detection, patient isolation, and prevention of outbreaks, particularly in high-risk resource-limited environments. In other words, ensuring public health safety hinges on the availability of fast, reliable, and widely accessible diagnostic tools. Traditionally, Chest X-Ray (CXR) and Computed Tomography (CT) have been among the primary imaging modalities used to diagnose COVID-19 infections [2]. These modalities provide detailed images invaluable in diagnosing respiratory conditions, including those caused by COVID-19. However, both modalities are not without limitations. CT scans, for example, while accu-

Cite this article: Azimi H, Wong A, Ebadi A. COVID-Net UV: An End-to-End Hybrid Deep Neural Network for Fast and Accurate COVID-19 Screening via Ultrasound Videos. J Clin Images. 2024; 7(2): 1161.

rate, involve significant exposure to ionizing radiation and are costly. Furthermore, CT and CXR machines are often confined to well-equipped medical centers, limiting their accessibility in rural, low-resource, or overburdened healthcare environments [3,4]. In contrast, lung Point-Of-Care Ultrasound (POCUS) has emerged as a valuable alternative imaging modality for diagnosing lung-related diseases, including COVID-19. POCUS offers several advantages over conventional imaging methods, particularly in resource-limited or emergency settings [1]. POCUS is non-invasive, portable, and cost-effective, making it ideal for use in a variety of clinical environments, including those where access to advanced imaging technologies like CT is restricted, e.g, widespread COVID-19 screening and diagnosis [5]. Its portability allows clinicians to perform bedside examinations, providing immediate diagnostic information and minimizing the need for patient transport, which can be both risky and time-consuming, especially in critically ill patients. Given these strengths, POCUS is becoming increasingly utilized for diagnosing lung-related diseases, including COVID-19 [1-6], enabling real-time decision-making in critical care situations. However, one of the main challenges in using POCUS for COVID-19 diagnosis lies in the interpretation of the ultrasound video sequences. The standard protocol for lung ultrasound examinations involves capturing multiple video sequences from various positions and angles around the chest [1-7], ensuring a comprehensive assessment of the patient's condition. While this approach increases diagnostic reliability, it also introduces complexity, as not all frames in these videos will contain relevant diagnostic information. Manual interpretation of these video sequences requires significant domain expertise and is often labor-intensive, leading to variability in diagnosis accuracy. This creates a need for automated approaches that can assist in interpreting POCUS video data, reducing the cognitive load on clinicians while ensuring consistent and accurate results. The application of Deep Learning (DL) networks to POCUS images has shown great promise in automating various tasks such as segmentation, disease classification, and detection [8]. These models can extract meaningful patterns from medical images, potentially outperforming traditional rule-based systems, especially when trained on large, well-curated datasets. These advanced models have the potential to enhance diagnostic accuracy and efficiency, particularly in settings where expert interpretation is not readily available. Existing models for POCUS-based COVID-19 diagnosis often focus on analyzing individual frames (e.g., [1-9]), neglecting the temporal context inherent in video data. Given that not all frames in an ultrasound video sequence may contain signs of suspected disease, analyzing individual frames without considering their temporal context may lead to sub-optimal diagnostic outcomes. The temporal information embedded in the video sequences is critical for understanding the progression of the disease and capturing subtle changes that might be missed in single-frame analysis. Therefore, the challenge lies in developing DL architectures that can effectively combine both spatial and temporal features present in POCUS videos for robust diagnostic performance. Motivated by this challenge, we propose COVID-Net UV, an end-to-end spatio-temporal deep neural network architecture explicitly designed to detect COVID-19-positive cases from POCUS video sequences. COVID-Net UV capitalizes on the full range of information available in POCUS videos by integrating both spatial and temporal data to enhance diagnostic accuracy. This approach not only improves the reliability of detecting COVID-19-positive cases but also offers a scalable solution that can be deployed in diverse healthcare settings, including those with limited resources. The key contributions of our work can be summarized as follows:

Automated COVID-19 screening: COVID-Net UV serves as an effective tool for the automatic detection of COVID-19-positive cases from POCUS video sequences, eliminating the need for technician intervention and additional processing steps. This can enable faster, real-time diagnostic support for front-line healthcare workers.

Spatio-temporal learning: Unlike previous approaches that analyze single frames, COVID-Net UV captures both spatial and temporal patterns from video data, leading to more comprehensive and accurate diagnoses. The model learns to detect subtle temporal dynamics, such as the progression of lung inflammation, which may be missed in frame-by-frame analysis.

Complementary to human experts: The proposed model addresses a critical gap in current diagnostic practices by reducing the reliance on time-consuming and costly training of human experts to interpret ultrasound data which typically requires extensive domain knowledge [10].

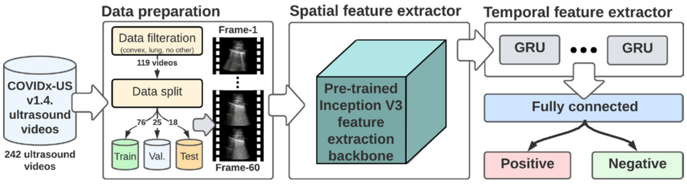

Figure 1: COVID-Net UV: A CNN-RNN architecture to classify POCUS videos into two classes of positive i.e, COVID-19 infection, and negative i.e, pneumonia or normal.

COVID-Net UV automates the diagnostic process, making it accessible to clinicians with varying levels of experience, thereby expanding its utility across a wide range of healthcare environments. By integrating cutting-edge spatiotemporal deep learning techniques with POCUS imaging, COVID-Net UV has the potential to significantly enhance the diagnostic process, making it more accessible, efficient, and accurate in screening and detecting COVID-19 infection. This work demonstrates the promise of leveraging artificial intelligence (AI) to develop practical and scalable solutions for public health crises, particularly in the context of fast-moving pandemics.

The remainder of this paper is structured as follows: In the next section, we provide an overview of related work, highlighting current advancements in deep learning for medical imaging, specifically in the context of COVID-19 detection using POCUS and other modalities. Section "Data and methods" describes the architecture of COVID- Net UV in detail, and the dataset used for the analyses. Section "Results" presents the results of our experiments. Finally, in the "Discussion" section, we conclude the paper with a discussion of our findings, the potential implications for clinical practice, and directions for future work.

Related work

The application of deep learning to lung POCUS for COVID-19 diagnosis has garnered significant attention in recent years. Several studies have proposed DL models, aiming to enhance the accuracy and efficiency of automated diagnosis. While these efforts have contributed valuable insights, they also face notable challenges, including limitations related to datasets, model complexity, and generalizability. Dastider et al. [11] introduced a hybrid CNN-LSTM model to predict COVID-19 severity from Lung Ultrasound (LUS) videos. The model effectively combines spatial features from a CNN with temporal features captured by

LSTM layers, leading to improved classification accuracy. However, the study was constrained by a small dataset of only 60 LUS videos, significantly limiting the generalizability of their model. Additionally, while the model captures temporal features, its initial frame-based analysis could miss key dynamics present across the video sequence, potentially reducing its effectiveness in complex clinical scenarios that require a holistic temporal understanding of lung changes. Moreover, the incorporation of multiple components, such as autoencoders and separable convolutions, also raises concerns about overfitting due to the limited size of the dataset. In a similar vein, Barros et al. [12] proposed a hybrid model combining CNNs and LSTMs to classify LUS videos of COVID-19 patients. Their model achieved high accuracy and sensitivity, demonstrating the benefits of integrating spatial and temporal features. However, the small and specific dataset used in the study may pose a challenge to the broader applicability of the model across diverse populations and clinical settings. Furthermore, the complexity of the model, along with the computational demands of preprocessing and hyperparameter optimization, restricts its use in real-time clinical settings, which demands faster and more streamlined solutions for deployment. Erfanian Ebadi et al. [13] utilized the two-stream Inflated 3D ConvNet (I3D) architecture to detect pneumonia in LUS videos, focusing on the identification of clinical markers such as A-lines, B-lines, and consolidations. Their approach achieved high accuracy and precision. However, the requirement for optical flow extraction adds significant preprocessing complexity, potentially limiting the model's practicality in real-time settings. Additionally, variations in data capture quality and differences in imaging devices affect the model's generalizability. Handling transitional cases where multiple features coexist, also remains a challenge. Roy et al. [14] took a different approach by exploring deep learning models for classifying and localizing COVID-19 markers in LUS images, introducing novel architectures such as Spatial Transformer Networks (STN). While their models performed well in classifying COVID-19-specific lung patterns, the added architectural complexity increased the risk of overfitting. This was especially problematic given the study's limited and specific dataset, which was collected from a few hospitals in Italy. Additionally, the dataset contained noisy and subjective labels, further complicating the training process and affecting model performance. This limited geographical and institutional scope of the dataset raised concerns about the model's generalizability to broader and more diverse populations. In two recent studies, researchers further explored the use of deep learning for COVID-19 diagnosis through POCUS imaging. In one study, a framework called COVID-Net L2C-ULTRA [15] was developed to handle the heterogeneity of POCUS data captured by different probes (e.g., convex and linear). By employing an extended linear-convex ultrasound augmentation learning approach, this model improved test accuracy by 3.9% and demonstrated significant performance gains in recall and precision when trained on combined datasets, enhancing the utility of both convex and linear probe images. Another study introduced COVID- Net US-Pro [16], an explainable few-shot deep prototypical network that excels in diagnosing COVID-19 from very limited ultrasound data. Trained on five shots, the model achieved over 99% accuracy, recall, and precision, making it highly effective even with small datasets. These frameworks highlight the potential of artificial intelligence in accelerating COVID-19 diagnosis. However, both models were built on LUS images, rather than videos, limiting their ability to capture temporal dynamics. While these studies represent important strides in the development of DL-based approaches for COVID-19 detection using LUS, they highlight several challenges such as small and specialized datasets, the complexity of models, and the need for significant preprocessing, all of which can impede real-world deployment. In an aim to address these challenges and limitations, we propose COVID-Net UV, a streamlined, end-to-end spatiotemporal deep neural network tailored for robust COVID-19 detection from POCUS videos. Unlike prior models, COVID-Net UV is designed to minimize preprocessing steps and focus on efficiently capturing both spatial and temporal dependencies in POCUS video sequences. These properties make it well-suited for real-time clinical use, particularly in diverse and resource-limited settings. Furthermore, the use of a larger, more diverse dataset enhances the model's robustness, generalizability, and applicability in a variety of healthcare environments, improving its potential for widespread clinical adoption.

Data and methods

Data: To train and evaluate the COVID-Net UV, we used the COVIDx-US dataset v1.4. [17], a comprehensive collection of 242 LUS videos curated and integrated from nine distinct data sources. These videos represent four primary classes: COVID-19 infection, non-COVID-19 infection, other lung diseases/conditions, and normal control cases. We filtered out the other class, i.e., other lung diseases/conditions, due to the heterogeneity of the cases to enhance model focus and performance. Moreover, to maintain consistency in the imaging modality, we restricted our dataset to videos captured solely with convex transducers, as different probe types may yield different image characteristics, complicating the training process. We formulated the problem as a binary classification task, wherein cases of COVID-19 infection were labeled as positive, and both non-COVID-19 infections and normal control cases were grouped and labeled as negative. This approach ensured a clear distinction between COVID-19-positive cases and all other types of lung conditions or normal lung health. Following these filtering steps, the final dataset consisted of 119 videos: 60 COVID-19-positive cases and 59 negative cases (including both non-COVID-19 infections and normal controls). The dataset was split into three subsets for training, validation, and testing, as follows:

1. **Training set:** 76 videos (38 positives and 38 negatives),
2. **Validation set:** 25 videos (12 positives and 13 negatives), and
3. **Unseen test set**: 18 videos (10 positives and 8 negatives).

Model architecture: For the model architecture, we employed a hybrid deep learning model that included convolutional and recurrent layers, capable of processing both spatial and temporal dimensions of the POCUS videos, respectively (see Figure 1). Specifically, we adopted the InceptionV3 model, pre-trained on the ImageNet data set [18], as the backbone for spatial feature extraction. To process the temporal dynamics within the video sequences, we added two Gated Recurrent Unit (GRU) layers to the network. GRUs are lightweight variants of LSTM networks and offer the advantage of being computationally less expensive while retaining the ability to capture long-term dependencies [19]. This design allows the network to jointly learn both spatial features (from individual frames) and temporal patterns (from the ordered sequence of frames).

Data preprocessing: Since a video is an ordered sequence of frames, the frames can be extracted and placed on a 3D tensor. However, a key challenge in processing video data is the variability in the number of frames across different videos. This variability complicates batch processing, as it is necessary for the input data to have consistent dimensions. To address this, we employed the following preprocessing steps:

1. **Frame extraction:** Frames were first extracted from each video to form an ordered sequence. However, the number of frames varied between videos, necessitating further preprocessing.

2. **Frame truncation and padding:** We set a maximum frame count of 60, based on the distribution of frame lengths in the dataset. For videos with fewer than 60 frames, we padded the sequence with zeros to standardize the input size. For videos exceeding 60 frames, we truncated the sequence to retain only the first 60 frames. This choice balances computational efficiency with the need to capture adequate temporal information for classification. The resulting 60-frame sequences were stacked into 3D tensors to serve as input to the model. Each tensor consisted of the video's spatial information, distributed across the frame sequence.

Training strategy: We adopted several techniques to optimize the training process and avoid overfitting as follows:

1. **Learning rate scheduler:** We applied a learning rate decay strategy, reducing the learning rate by a factor of 0.5 whenever the validation loss plateaued for three consecutive epochs. This helps in fine-tuning the network and finding an optimal set of weights without overshooting.

2. **Early stopping:** To prevent overfitting and unnecessary computation, we employed an early stop- ping mechanism, halting the training process after seven consecutive epochs without improvement in validation performance. This ensures that the model does not overfit the training data by continuing to train beyond the point of optimal generalization. The network's initial learning rate was set to 0.001, and the maximum number of training epochs was set to 30. However, following the early stopping criterion, training was completed after 18 epochs.

Table 1: Performance of COVID-Net UV on the unseen test dataset.

Class	Precision	Recall	F1-score
Negative	1.0000	0.8750	0.9333
Positive	0.9091	1.0000	0.9524

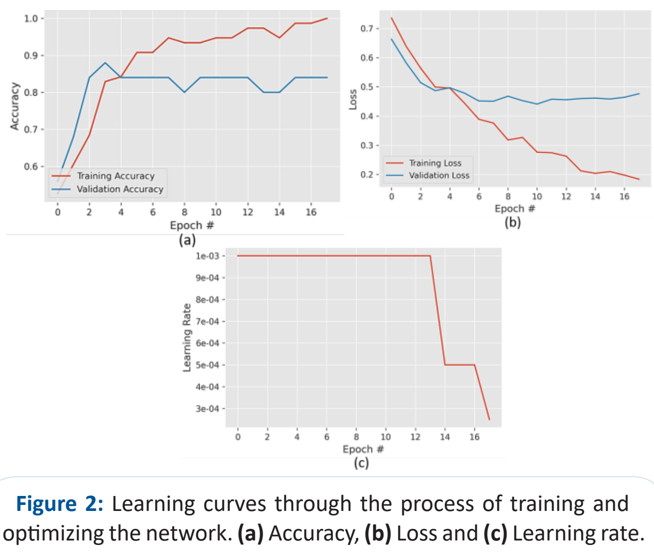

Figure 2: Learning curves through the process of training and optimizing the network. **(a)** Accuracy, **(b)** Loss and **(c)** Learning rate.

Evaluation metrics: To evaluate the performance of COVID-Net UV, we employed several standard classification metrics, including accuracy, precision, recall, and F1 score. These metrics were calculated on both the validation and test sets, providing a robust assessment of the model's generalization capabilities.

Results

The learning curves through the process of training and optimizing the network are illustrated in Figure 2. To mitigate the risk of overfitting, we carefully monitored the validation loss and applied early stopping once the validation loss stopped decreasing. Specifically, training was stopped just before the validation loss began to rise, ensuring the model captured the critical patterns without overfitting the training data (Figure 2-b). Following the learning rate scheduler strategy, during the process of training, the learning rate was decayed twice at epochs 15 and 18, which allowed for finer weight adjustments during the later stages of training. The COVID-Net UV model demonstrated robust classification performance, achieving an overall accuracy of 94.44% across both classes, see Table 1. The network's ability to correctly classify COVID-19-positive cases was particularly notable, with a sensitivity (recall) of 100%, meaning that the model produced no false negatives. This is crucial for clinical applications, as minimizing false negatives is essential in preventing missed diagnoses of COVID-19-positive patients. For the negative class (which included non-COVID-19 infections and normal cases), the sensitivity was 87.50%, reflecting the model's ability to correctly identify the absence of COVID-19 in most cases. In terms of precision, the model achieved a high score of 90.91% for the positive class, indicating that the majority of cases identified as COVID-19-positive were indeed true positives. The precision for the negative class was even higher, at 100%, signifying that the model perfectly identified all true negative cases without any false positives. This performance highlights the model's effectiveness in balancing sensitivity and precision, ensuring both accurate detection of COVID-19 cases and minimal misclassification of non-COVID cases.

Discussion

In this work, we introduced COVID-Net UV, an end-to-end hybrid neural network architecture designed to classify lung POCUS videos for the diagnosis of COVID-19. The network integrates two key components: the pre-trained InceptionV3 to extract spatial features from video frames, and an RNN with GRU units to capture the temporal dependencies between video frames. The hybrid architecture, by combining spatial and temporal analysis, was tailored to extract rich information from lung ultrasound videos, enhancing the accuracy of COVID-19 diagnosis. Our results demonstrated that COVID-Net UV achieved a sensitivity of 100% for COVID-19 cases with no false negatives, significantly outperforming human experts, who achieved a sensitivity of 86.4% [20]. Furthermore, compared to models relying solely on spatial architecture (with the highest accuracy of 83.2%) [21], our approach demonstrated the benefits of incorporating temporal dynamics in clinical video-based diagnosis. The importance of fast and accurate COVID-19 diagnosis cannot be overstated, particularly in the context of pandemic preparedness and response. While various diagnostic methods, such as Reverse Transcription Polymerase Chain Reaction (RT-PCR), Chest X-Ray, and CT scans, have proven effective, they are resource-intensive and require access to specialized laboratories and equipment. In contrast, lung POCUS provides a more accessible, portable, and cost-effective alternative, particularly in resource-limited and remote settings where access to advanced imaging technology may be scarce. However, the accurate interpretation of POCUS images and videos is often contingent on the experience of the clinician, which can vary significantly. This is where COVID-Net UV can play a transformative role by providing an AI-assisted diagnostic tool that augments clinical decision-making and ensures consistent, objective, and reliable results across different healthcare settings. Due to the mentioned advantages, POCUS can be widely used in resource-limited and remote regions, and incorporating AI-based tools like COVID-Net UV can significantly enhance diagnostic capabilities in these settings. By enabling rapid and accurate identification of COVID-19 cases through POCUS, healthcare providers in under-resourced regions can make more informed decisions, allocate resources more efficiently, and prioritize treatment for patients with COVID-19, all of which are critical in mitigating the impact of the pandemic. COVID-Net UV's flexible architecture also makes it well-suited for adaptation to future pandemics and the detection of rare diseases that may emerge. By retraining the model with new data, it could be rapidly repurposed for identifying novel respiratory illnesses or other emerging pathogens, providing a crucial tool in the global response to health crises. Its ability to work with POCUS videos also means that the model can support diagnosis in challenging clinical environments where new diseases may spread, ensuring timely detection and intervention. While the results of COVID-Net UV are promising, we acknowledge the limitations posed by the relatively small size of the video dataset used in this study, comprising 119 POCUS videos. Although the model performs well on the available data, there is a need to validate its generalizability across larger and more diverse datasets. However, we believe our proposed methodology and results are the desired baseline for our future work in examining more complex models on the larger POCUS video dataset. Our future work will focus on expanding the dataset by integrating more video samples from different regions, patient demographics, and device types to further ensure the model's robustness and applicability across various clinical settings. Furthermore, we plan to explore the model's potential for multi-class classification by including additional lung diseases beyond COVID-19. This will make the model more versatile and applicable to a broader range of clinical scenarios. The architecture of COVID-Net UV is designed to be flexible and easily adaptable, allowing for the integration of new disease categories as more data becomes available. As such, the model can be retrained to classify not just COVID-19 but also other common lung pathologies.

Acknowledgment

The authors would like to acknowledge that a preliminary version of this work was presented at the 8th Annual Conference on Vision and Intelligent Systems (CVIS) in 2022. We are grateful for the valuable feedback received during the conference, which has significantly contributed to the development of this extended version. This research was supported by the Pandemic Response Challenge Program of the National Research Council Canada.

References

1. A MacLean, S Abbasi, A Ebadi, A Zhao, M Pavlova, et al. COVID-Net US: A tailored, highly efficient, self-attention deep convolutional neural network design for detection of COVID-19 patient cases from point-of-care ultrasound imaging, In Domain Adaptation and Representation Transfer, and Affordable Healthcare and AI for Resource Diverse Global Health: Third MICCAI Workshop, DART 2021, and First MICCAI Workshop, FAIR 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27 and October 1, 2021, Proceedings 3 (pp. 191-202). Springer International Publishing.

2. KR Bhatele, A Jha, D Tiwari, M Bhatele, S Sharma, MR Mithora, et al. COVID-19 detection: A systematic review of machine and deep learning-based approaches utilizing chest x-rays and CT scans, Cognitive Computation. 2024; 16: 1889-1926.

3. C Bergeron, R Fleet, FK Tounkara, I Lavall´ee-Bourget, C Turgeon-Pelchat. Lack of CT scanner in a rural emergency department increases inter-facility transfers: A pilot study, BMC Research Notes. 2017; 10: 772.

4. M Davidson, A Kielar, R Tonseth, K Seland, S Harvie, et al. The landscape of rural and remote radiology in Canada: Opportunities and challenges, Canadian Association of Radiologists Journal. 2024; 75: 304-312.

5. Y Amatya, J Rupp, FM Russell, J Saunders, B Bales, et al. Diagnostic use of lung ultrasound compared to chest radiograph for suspected pneumonia in a resource-limited setting, International Journal of Emergency Medicine. 2018; 11: 8.

6. A MacLean, A Ebadi, A Florea, P Xi, A Wong. An initial study into the feasibility of deep learning-based COVID-19 severity classification using point-of-care ultrasound imaging, Journal of Computational Vision and Imaging Systems. 2021; 7(1): 31-33.

7. G Volpicelli, M Elbarbary, M Blaivas, DA Lichtenstein, G Mathis, et al. International evidence-based recommendations for point-of-care lung ultrasound, Intensive Care Medicine. 2012; 38: 577-591.

8. S Liu, Y Wang, X Yang, B Lei, L Liu, et al. Deep learning in medical ultrasound analysis: A review, Engineering. 2019; 5: 261-275.

9. Ultrasound-based AI for COVID-19 detection: A comprehensive review of public and private lung ultra sound datasets and studies v1 preprints.org. 2024.

10. MY Ng, EYP Lee, J Yang, F Yang, X Li, et al. Imaging profile of the COVID-19 infection: Radiologic findings and literature review, Radiology: Cardiothoracic Imaging. 2020; 2: 200034.

11. AG Dastider, F Sadik, SA. Fattah. An integrated autoencoder based hybrid CNN-LSTM model for COVID-19 severity prediction from lung ultrasound, Computers in Biology and Medicine. 2021; 132: 104296.

12. B Barros, P Lacerda, C Albuquerque, A Conci. Pulmonary COVID-19: Learning spatiotemporal features combining CNN and LSTM networks for lung ultrasound video classification, Sensors. 2021; 21: 5486.

13. S Erfanian Ebadi, D Krishnaswamy, SES Bolouri, D Zonoobi, R Greiner, et al. Automated detection of pneumonia in lung ultrasound using deep video classification for COVID-19, Informatics in Medicine Unlocked. 2021; 25: 100687.

14. S Roy, W Menapace, S Oei, B Luijten, E Fini, et al. Deep learning for classification and localization of COVID-19 markers in point-of-care lung ultrasound, IEEE Transactions on Medical Imaging. 2020; 39: 2676-2687.

15. EZ Zeng, A Ebadi, A Florea, A Wong. Covid-net l2c-ultra: An explainable linear-convex ultrasound augmentation learning framework to improve COVID-19 assessment and monitoring, Sensors. 2024; 24: 5.

16. J Song, A Ebadi, A Florea, P Xi, S Tremblay, et al. COVID-Net US-Pro: An explainable few-shot deep prototypical network for COVID-19 screening using point-of-care ultrasound, Sensors. 2023; 23: 5.

17. A Ebadi, P Xi, A MacLean, A Florea, S Tremblay, et al. COVIDx US: An open-access benchmark dataset of ultrasound imaging data for AI-driven COVID-19 analytics, Frontiers in Bioscience-Landmark. 2022; 27(7): 198.

18. H Azimi, J Zhang, P Xi, H Asad, A Ebadi, S Tremblay, et al. Improving classification model performance on chest x-rays through lung segmentation, arXiv. 2202; 1097.

19. J Chung, C Gulcehre, K Cho, Y Bengio. Empirical evaluation of gated recurrent neural networks on sequence modeling. 2014.

20. N Islam, S Ebrahimzadeh, JP Salameh, S Kazi, N Fabiano, et al. Thoracic imaging tests for the diagnosis of COVID-19, Cochrane Database of Systematic Reviews. 2021: 3.

21. N Awasthi, A Dayal, LR Cenkeramaddi, PK Yalavarthy. Mini-covidnet: Efficient lightweight deep neural network for ultrasound based point-of-care detection of COVID-19, IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control. 2021; 68: 2023-2037.