





COVIDDX: AI-based Clinical Decision Support System for Learning COVID-19 Disease Representations from Multimodal Patient Data

Veena Mayya^{1,2}, Karthik K.¹, Sowmya S. Kamath¹, Krishnananda Karadka³
and Jayakumar Jeganathan⁴

¹Healthcare Analytics and Language Engineering (HALE) Lab, Dept. of Information Technology,
National Institute of Technology Karnataka, Surathkal, Mangalore 575025, India

²Department of Information & Communication Technology, Manipal Institute of Technology,
Manipal Academy of Higher Education, Manipal, Karnataka, India

³Penzigo Technology Solutions Pvt. Ltd., NITK-Science and Technology Entrepreneurs' Park (STEP), NITK Surathkal, India

⁴Dept. of Medicine, Kasturba Medical College, Manipal Academy of Higher Education (MAHE), Karnataka, India


Keywords: Computational and Artificial Intelligence, Decision Support Systems, Automated Diagnosis, COVID-19.


Abstract: The COVID-19 pandemic has affected the world on a global scale, infecting nearly 68 million people across the world, with over 1.5 million fatalities as of December 2020. A cost-effective early-screening strategy is crucial to prevent new outbreaks and to curtail the rapid spread. Chest X-ray images have been widely used to diagnose various lung conditions such as pneumonia, emphysema, broken ribs and cancer. In this work, we explore the utility of chest X-ray images and available expert-written diagnosis reports, for training neural network models to learn disease representations for diagnosis of COVID-19. A manually curated dataset consisting of 450 chest X-rays of COVID-19 patients and 2,000 non-COVID cases, along with their diagnosis reports were collected from reputed online sources. Convolutional neural network models were trained on this multimodal dataset, for prediction of COVID-19 induced pneumonia. A comprehensive clinical decision support system powered by ensemble deep learning models (CADNN) is designed and deployed on the web*. The system also provides a relevance feedback mechanism through which it learns multimodal COVID-19 representations for supporting clinical decisions.


1 INTRODUCTION


COVID-19 has claimed the lives of more than one million people worldwide and continues to pose a severe threat to humanity (Max Roser and Hasell, 2020). Indicative clinical symptoms of COVID-19 include high fever, cough, sore throat, headache, fatigue, muscle pain, and Dyspnea or shortness of breath (SoB). Currently, the main testing procedure employed for diagnosing COVID-19 is RT-PCR (Real-time Reverse Transcription - Polymerase Chain Reaction), which primarily detects the presence of RNA in the test samples. Radiology tests like Computed Tomography (CT) and X-rays have also been used as additional diagnostic tools. Normally, CT

and X-rays show significant changes in the lung with the onset of respiratory symptoms, while some studies have reported that discernible changes occur in an infected person's scans, starting at the first onset of mild symptoms. In situations when RT-PCR kits are limited in number, medical personnel have relied on such radiography scans for confirming COVID-19 infection. This opens up a significant research scope for designing automated systems trained to process a large number of radiography scans such as CT and X-rays for testing for COVID-19 infection. Moreover, there exists a significant potential for reducing the costs associated with mass testing to a large extent and judiciously manage available RT-PCR kits. Currently, in the Indian healthcare system, RT-PCR testing costs around ₹2,000-4,000, whereas X-ray scan costs are in the range of ₹200-500 (BusinessToday, 2020). This price difference can be hugely beneficial for patients and has a high cost-benefit tradeoff.

^a  <https://orcid.org/0000-0002-8091-5053>

^b  <https://orcid.org/0000-0003-0846-2982>

^c  <https://orcid.org/0000-0002-0888-7238>

^d  <https://orcid.org/0000-0001-8385-3516>

*CADNN COVID-19 Predictor, <https://cadnn.penzigo.net>

Clinical Decision Support Systems (CDSSs) powered by learnable AI-based models, computation techniques such as image processing and big data analytics have proven to be effective in assisting health professionals in a wide array of clinical tasks such as disease prediction (Tushaar et al., 2020), anomaly analysis, patient history modeling (Gangavarapu et al., 2020) and so on. A CDSS for predicting the presence or absence of COVID-19 infection using diagnostic scans can be beneficial to healthcare professionals as well as patients. A contactless X-ray scan workflow could be achieved by using cameras for patient monitoring purposes (Scheib, 2009; Forthmann and Pfeleiderer, 2019), after a few days of onset of COVID-19-like symptoms, the patient is subjected to an X-ray. For cases where the predicted risk is high, he or she can then undergo a RT-PCR test to confirm the diagnosis. Based on the predicted risk score, the health professionals may also decide to isolate the patient and perform another X-ray in a couple more days to ensure others' safety. Such analysis can also contribute to isolating asymptomatic COVID-19 patients, who undergo chest X-ray for other reasons (e.g., pre-operation evaluation, routine medical check-up, rib fractures etc.).

Furthermore, there is a critical need for ensemble/composite models that make use of laboratory examination results to help better screening detection and diagnosis of COVID-19. In this work, we attempt to build such a model that leverages the wealth of information contained in expert reports along with the actual diagnostic scan data for learning disease representations of COVID-19. A collated dataset consisting of X-ray and the corresponding expert reports of individual patients suffering from COVID-19 and non-COVID cases were used for the experiments. Also, a complete web-based framework is deployed on the cloud, to provide a highly usable platform for expert verified screening results, with which the model is retrained based on the continuous feedback provided by the expert radiologists. Through this research, our focus is to enable a comprehensive clinical decision support framework for fast and cost-effective early screening of COVID-19, in a user-friendly and unobtrusive manner.

The rest of this paper is organized as follows: Section 2 presents a detailed discussion of the existing research works in the area of interest. In Section 3, we describe the methodology adopted for data collation and the proposed approach for diagnosing COVID-19 using patients' chest X-ray images and expert-written diagnosis reports. Section 4 elaborates the experiments performed and the observations regarding

the performance of the proposed models, followed by conclusions and directions for future work.

2 RELATED WORK

Machine learning-based medical image/text analysis and classification has seen extensive applications in healthcare for enabling medical data management and improved diagnosis. CDSSs that are built on the premise of AI-based analysis of medical data for enabling decision making has been successfully used by healthcare professional in various clinical settings. Recently, significant research interest has focused on the application of AI for diagnosing COVID-19.

Several works that make use of diagnostic scan images like chest X-rays (Ozturk et al., 2020; Narin et al., 2020; Hall et al., 2020; Ghoshal and Tucker, 2020; Hemdan et al., 2020) and computerized tomography (CT) scan images (Mishra et al., 2020; Abbas et al., 2020; Dansana et al., 2020; Mei et al., 2020) have been proposed. Most existing works make use of transfer learning using deep learning models that are pretrained on ImageNet (Deng et al., 2009) data. We observed several gaps after reviewing existing works in this area. We found that the associated metadata of patients has not been considered in most works. Also, valuable expert-written diagnosis maintained as natural language text reports after checking a patient's chest radiography images have not been explored for the task of disease prediction. Furthermore, there is ample scope for the development of a complete, easy-to-use diagnostic framework for the use of healthcare professionals. Using such tools, expert opinion & other metadata about patients can also be obtained to incorporate relevance feedback into the prediction model, building accurate CDSSs.

In this work, we incorporate multiple deep learning models for classifying X-ray images as COVID-19 positive or negative, wherein, the contributions of image features and also the latent information contained in the expert-written diagnosis text reports are modeled for the diagnosis. To alleviate the manual effort required to assess and generate diagnosis reports when a large number of diagnosed cases arrive, a content-based report generation model, to automatically generate natural language diagnosis reports is also designed, for reducing the cognitive burden of radiologists and other medical personnel involved in medical record management. The complete framework is deployed on the cloud and is made available as a web application for managing patients metadata (from the day of admission till the discharge). Functionalities like validity checks for X-ray images,

evidence-based diagnosis support through highlighting of important features learnt by the model, and automatic report generation for further processing are incorporated in the proposed framework. A feedback system is provided to verify the prediction and generated reports, which is later utilized for improving the offline training process, for fine-tuning the prediction performance of the CDSS.

3 MATERIALS & METHODS

3.1 Data Collation

Several open COVID-19 datasets are currently available (Cohen et al., 2020; Chung et al., 2020; Rahman, 2020) which are limited to only radiographical imaging data. Other pertinent information such as patient history, findings from such images etc., have not been made directly available. Thus, these datasets are not well-suited for multimodal data modeling and for multi-task learning. To address these lacunae, a multimodal patient data amenable for multimodal clinical tasks was collated from varied trusted sources.

For curating the dataset, a total of 150 confirmed COVID-19 patient cases were collected from publicly available sources^{1,2,3} and also from a local hospital with an active COVID-19 ward. Each X-ray in the collated data is associated with metadata information – demographics details like age, gender and findings in the form of plain natural language text (reports) as observed by expert radiologists. We have also collected COVID-19 X-ray images only from the open datasets mentioned earlier. In total, the dataset contained 450 chest X-ray images of COVID-19 infected patients, of which 150 images had associated metadata. In addition to this, about 2,000 normal cases were taken from the Pneumonia Detection Challenge (Radiological Society of North America, 2018). The neural models were trained on collated data and fine-tuned whenever a significant number of new cases were uploaded through the online CDSS application.

3.2 Proposed Approaches

Fig. 1 illustrates the detailed workflow of the proposed approach. The framework employs five deep neural models as discussed in Section 3.2.1. As the X-ray images were captured using different machines, there exists a large variability, mainly in pixel in-

tensities and focus on lung regions. To reduce the change in color intensities across the images, histogram matching (Gonzalez and Woods, 2008) was applied to the dataset, for which, we considered an X-ray image as a reference image (R_{img}) and then matched all the other X-ray pixel intensity histograms with R_{img} . To suppress the effect of rib shadows on the classifier’s prediction, $RIBDL$ model was used (described in Section 3.2.1). Then, the lung regions are segmented to reduce the effect of the surrounding background on the model’s prediction. This enables the classification models to learn the minute changes in lung structure, which is often missed by human experts (due to the noise created by rib shadow).

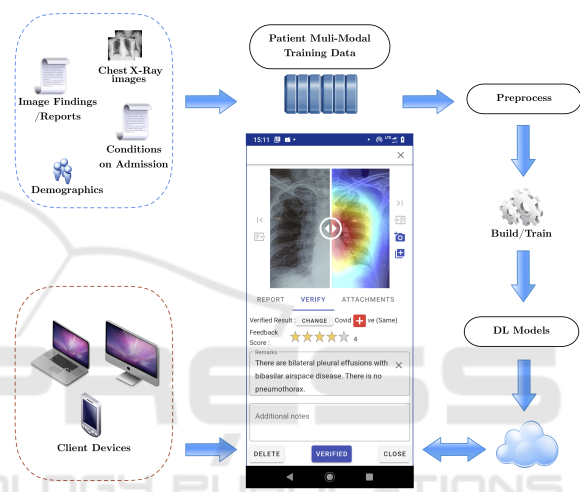


Figure 1: High-Level design of the proposed approach.

Algorithm 1 details the process of segmentation. Here, we initially segment the lung region using PIX -GAN (Section 3.2.1), then the bounding box around the lung cavity is cropped and resized to original size. The sample output after each step is depicted in Fig. 2. Cropping only the region of interest (RoI) allows the deep neural model to learn important features from within the lung regions. Crop and rotate facilities are also provided in the developed CADNN online framework so that users can select only lung region while uploading new test images. The local contrast of RoI segmented input X-ray grayscale images are further improved by applying Contrast Limited Adaptive Histogram Equalization (CLAHE) (Reza, 2004). Image augmentation is performed by applying rotation with different angles ($10^\circ - 120^\circ$) with an interval of 40° to the training images, which resulted in more than 3,000 COVID-19 and non-COVID-19 (a total of 3,474 training case samples).

¹<https://radiopaedia.org>

²<https://www.sirm.org/COVID-19-database/>

³<https://www.eurorad.org/>

Algorithm 1: Chest X-ray preprocessing pipeline.

Input: Input chest X-ray images**Output:** Preprocessed X-ray image

- 1: **for each** $img \in InputImages$ **do**
- 2: Perform histogram matching of img with reference image.
- 3: Perform rib shadow removal using *RIBDL*.
- 4: Perform lung region segmentation using *PIXGAN* by resizing img to (512, 512).
- 5: Find the contours of the generated *MaskImg*
- 6: Remove all small contours (with width < 50 and height < 50) in *MaskImg*
- 7: Dilate *MaskImg* with kernel size of (5,3) until a single contour is formed.
- 8: Draw the bounding box over the single contoured *MaskImg* ▷ *Set all other pixel intensities to zero.*
- 9: Remove all black regions from input image and resize to required image shape.
- 10: Apply CLAHE.
- 11: **end for**

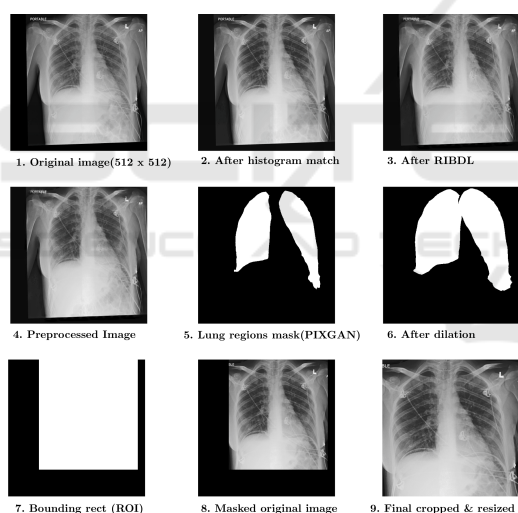



Figure 2: Chest X-ray preprocessing pipeline.

3.2.1 Deep Neural Models

The proposed framework is built on the predictive framework powered by five neural models. Transfer learning is employed to use pretrained weights for the initial layers while some of the models were trained from scratch. In the designed application, users are provided with an interface to select and upload X-ray images from those available on their smartphones/system. However, there is a possibility of them knowingly or unknowingly uploading natural photographs. To accept only valid images, we used a two-layered convolution network *ValidateDL* trained

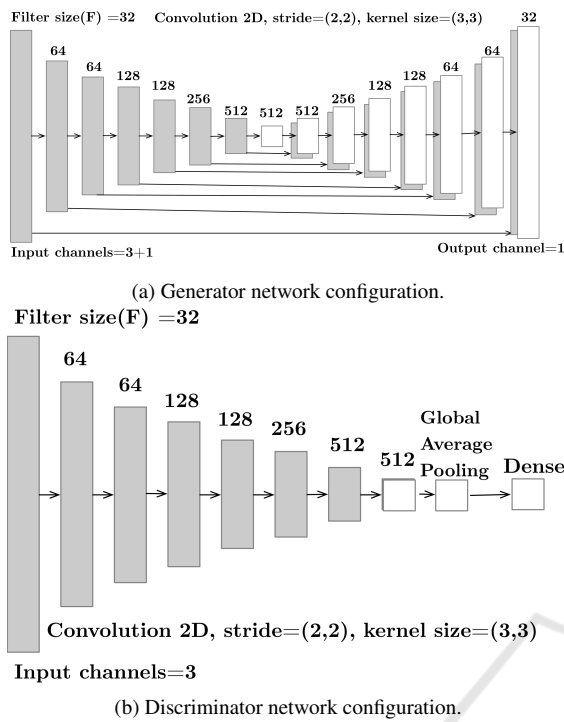
on CIFAR10 dataset (Krizhevsky, 2012) (containing 60,000 32×32 color natural images) and 5,000 X-ray images to classify between natural images and X-ray images. For every batch, randomly selected CIFAR10 images are converted to grayscale and copied to all three channels. This is done to ensure that grayscale natural images are also correctly classified by the model.

The configuration of the network is shown in Figure 3. The network is trained with stochastic gradient descent (SGD) optimizer with learning rate of 0.01. Batch size is set to 32 and trained for a maximum of 20 epochs. Early stopping was used to prevent over-fitting problems. The rib cage forms a significant part of the chest X-ray, and the rib shadow in the input images is suppressed using a pretrained autoencoder⁴ model, to ensure that the training of the classification network focuses on relevant information within the lungs region instead of the rib structure. The pre-trained model is deployed during the validation phase of the online CDSS, to generate the rib suppressed image for all collated input images.

Figure 3: *ValidateDL* network architecture.

As the collated dataset includes the X-ray images captured by different technicians using a variety of X-ray machines, substantial variation is observed in the area of focus. Some images included only lung regions, while many others covered the entire abdominal cavity. Also, several images from the collated dataset included X-ray machine labels in the form of characters/texts/numbers. To overcome variations and to restrict the classifier for effectively learning patterns from the lung region only, *PIXGAN* (Isola et al., 2018) is trained from scratch to segment the RoI. The number of convolution and deconvolution layers were increased to handle the larger input image size (512, 512) and a modified loss function (Son et al., 2017) that uses a tuning parameter (λ) was incorporated. λ is used while summing up the discriminator's binary cross entropy loss (among *predicted* and *true* labels) and generator's binary cross entropy loss (among *generated* and *true* lung masks). Thus, the generator enables the discriminator to produce outputs that are very similar to the real lung mask.

⁴<https://github.com/hmchuong/ML-BoneSuppression>


 Figure 4: *PIXGAN* for lung region segmentation.

A total of 800 X-ray images and lung region masks obtained from the Kaggle challenge⁵ was used to train the *PIXGAN*. The *PIXGAN* discriminator is fed both X-ray ($Image_{xray}$) and lung mask images ($Mask_{lung}$), which must determine whether $Mask_{lung}$ is a plausible transformation of $Image_{xray}$, as local style statistics are efficiently captured by *PIXGAN*. The generator is built on U-Net (Ronneberger et al., 2015) architecture and makes use of convolution and deconvolution layers for learning to generate realistic lungs mask from very few training X-ray images. Fig. 4 depicts the configuration of *PIXGAN* model which we used to segment the lung region for a given input X-ray image. The generator is designed using eight encoding and decoding units as shown in Fig. 4a, while the discriminator includes the complete encoding part of generator network with global average pooling and final Dense layers (shown in Fig. 4b). Once the lung region mask is generated, the original images are cropped to include only the lung region, as illustrated in Algorithm 1.

The preprocessed images are then trained using deep residual network (ResNet) (He et al., 2015). We used ResNet-18 for training the classifier to distinguish between COVID-19 and non-COVID-19 cases. A content-based technique as described in Section 3.2.2 is utilized to obtain the findings in the input

⁵<https://www.kaggle.com/nikhilpandey360/chest-xray-masks-and-labels>

X-ray image in the form of natural language textual report. The generated reports are pre-filled in the CADNN framework so that the radiologists can verify and do the changes if necessary. The collated expert X-ray reports were classified into COVID and non-COVID using the proposed explainable report prediction deep learning model $ER\mathcal{D}\mathcal{X}$. This is mainly performed to highlight the important terms in pre-filled reports of CADNN framework. $ER\mathcal{D}\mathcal{X}$ model architecture is discussed in Section 3.2.3.

3.2.2 Diagnostic Report Generation

For this task, we made use of ResNet18 for classifying COVID and non-COVID cases. The last convolution layer output of ResNet18 provides a plausible disease representation of the input X-ray image. A feature vector ($features$) is generated by summing the last convolution layer output from the trained ResNet18. A dictionary ($D_{features}$) of feature vectors and reports indexed by image names is generated for the collated X-ray images for which expert reports were available. $D_{features}$ is also updated with frontal X-ray image features and textual findings obtained from the IU dataset (Demner-Fushman et al., 2015). In total, 820 COVID and non-COVID reports along with corresponding frontal chest X-ray image features were utilized for report generation and report classification. For the given input test X-ray image, the image features ($Test_{feature}$) are extracted during classification along with the predicted label. Cosine similarity between $Test_{feature}$ and features of $D_{features}$ is computed. The report is obtained using $index(Isimax)$ of $D_{features}$ for which the maximum cosine similarity exists between $Test_{feature}$ and $D_{features}[Isimax]$. The generated reports are further processed by the X-ray report classifier.

3.2.3 X-ray Report Classification

The expert reports consisting of physician observations contain a wealth of information regarding the condition, symptoms and other details regarding the patients' status. This rich latent information can be used to model patient representations, which can then be leveraged to potentially screen COVID-19 infected patients. Each report is subjected to preprocessing using standard natural language processing (NLP) techniques. Any punctuation, digits and stop words are removed from the patient's X-ray reports. Out of vocabulary(OOV) terms are handled by including special OOV token, and the maximum allowed document length is fixed to 100.

From the preprocessed text, embeddings are generated using the Word2Vec Continuous Bag-of-Words

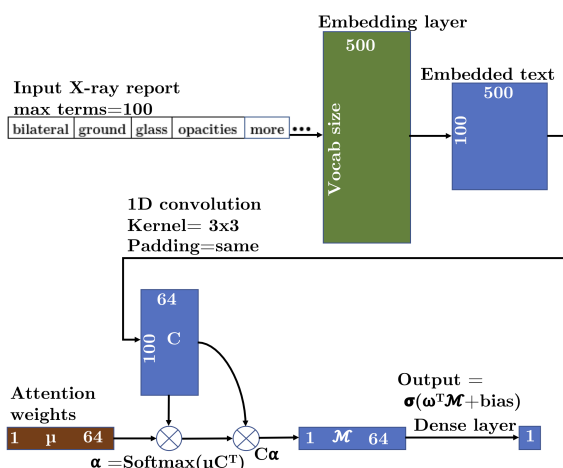


Figure 5: $R\mathcal{D}x$ X-ray report classification DL model.

(CBoW) model (Mikolov et al., 2013). The learning rate was fixed to a default value of 0.025 (same as that of Word2Vec model), the number of iterations was set to 10 and embedding size used was 200. We used Python Gensim library (Řehůřek and Sojka, 2010) to generate the word embeddings using the preprocessed X-ray reports. As the main purpose of NLP classifier was to highlight important terms in the reports, we designed a convolutional attention explainable neural network ($R\mathcal{D}x$) to classify the X-ray reports. The model consists of a 1D convolution layer followed by an attention and a dense layer, which is depicted in Figure 5. All the reports are padded up to a maximum length of the reports in a batch, 100 being the maximum allowed length for a report. Given the test report, which is generated for the input X-ray image using the content-based approach, it is classified by the $R\mathcal{D}x$ model. The important terms contributing to the model’s prediction are highlighted using color codes. This feature is also used to enable faster verification features for experts through the functionalities provided through the CADNN interface, and to update/edit findings if necessary.

4 EXPERIMENTAL RESULTS


The proposed deep neural approaches were developed using Pytorch (Paszke et al., 2019) and Tensorflow (Abadi et al., 2016) deep learning python frameworks and trained on Ubuntu 18.04 system with NVIDIA Tesla M40 and Tesla V100-DGXS. All the test data splits are made before image augmentation. Accuracy is used as an evaluation metric for verification of classification results, which is calculated based on true positives (TP), false positives (FP), true negatives (TN) and false negatives (FN) cases predicted

by a particular neural model, and is given by Eq. (1). Here, TP is the number of cases that are correctly identified by the prediction model to be COVID-19 positives, which match with experts’ opinion, while FN are incorrectly rejected cases. TN is the number of correctly identified non-COVID-19 cases, and FP gives the total incorrectly identified COVID-19 cases.

$$Accuracy = \frac{TP + TN}{(TP + FP + FN + TN)} \quad (1)$$

The $ValidateDL$ model is evaluated on the CIFAR-10 test set and 450 test X-ray images from RSNA challenge data. The model was able to identify all X-ray images correctly and achieved 97.8% accuracy. Out of the 10,000 CIFAR-10 images, 221 images that included only cloudy sky or runway images (which appear similar to X-ray image structure in 32 x 32 dimension) were wrongly classified as X-ray images instead of natural images. For training $PIXGAN$, batch size of 32, λ value of 0.5 and Adam optimizer with 0.0002 learning rate is used. Training is performed for a maximum of 100 epochs, and the $PIXGAN$ was evaluated on 50 test images out of 800 X-ray images. The generator model that achieved the highest dice coefficient (Zijdenbos et al., 1994) (obtained at 28th epoch on validation data) was used to extract lung mask regions for the collated data.

Next, for testing the ResNet18 X-ray image classification model, 80% of input data for training, 10% for validation and rest 10% was considered. The predicted and actual class for the X-ray images are summarized in the confusion matrix shown in Fig. 6a. An accuracy of 97% was achieved using the proposed X-ray image classification model. The $R\mathcal{D}x$ model was evaluated on the 20% test split, and accuracy of 96.74% was observed. The results are summarized in the confusion matrix depicted in Fig. 6b.



(a) Chest X-ray dataset. (b) Diagnostic report dataset.
Figure 6: Confusion matrix for different datasets.

In addition to experimental benchmarking with standard datasets, a pilot study was conducted with patient data collected from a COVID ward in a private hospital. The experiments were performed on a

collection of 95 COVID and 5 non-COVID (Bacterial/Tubercular lung infection) X-ray images obtained from patients admitted for screening/treatment at the hospital. The results were very promising, as the proposed CADNN framework performed very well on this real-world data. Healthcare professionals involved in the pilot study used the CADNN framework for uploading patient data and observing the predictions of the model. In these studies, the system identified all COVID-19 cases accurately, while only one non-COVID case was wrongly predicted as COVID, achieving an overall prediction accuracy of 99%.

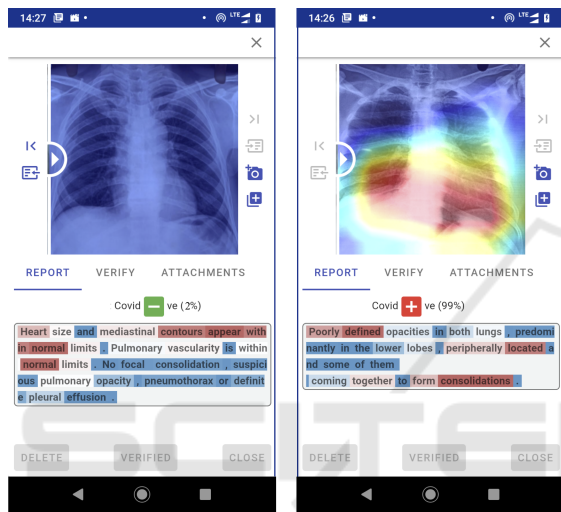


Figure 7: Proposed CDSS in action.

Qualitative Evaluation. For enabling evidence-based diagnosis, we utilised Class Activation Mapping (Grad-CAM) (Selvaraju et al., 2017) to highlight the regions of the input X-ray image that the classification model considered relevant to perform the prediction. The visualization is shown in the CADNN framework to aid the clinical decision. The regions in the image, where this gradient is predominant, are shown in Fig. 7), along with the generated report that shows the highlighted important terms. As can be observed from the attentions, the model has successfully learnt important features in the X-ray image, restricting itself mostly to within the lung region. The important radiography terms are identified (highlighted with white and red colors) by $R\mathcal{D}\mathcal{X}$.

5 CONCLUSION & FUTURE WORK

In this paper, a cost-effective, early-screening strategy for COVID-19 diagnosis based on chest X-ray

images and expert-written diagnosis reports was presented. The proposed framework has been deployed as a web-based CDSS called CADNN. The input images were subjected to extensive validation and pre-processing steps to eliminate variance and ensure effective learning by the prediction model. Pre-processed images were used to train a ResNet model for COVID-19, Pneumonia or non-COVID-19 classification and the findings obtained from the images were used to automatically generate the natural language diagnosis reports, using content-based learning approach. The proposed CADNN also includes feedback mechanisms so that the results could be verified by the experts, and feedback from experts can be utilised during offline retraining of the models. It allows users to upload additional documents like CT scan images/DICOM sequences for additional insights into the patient’s condition.

The proposed framework could be easily adapted for diagnosis of other lung related diseases and provide a comprehensive CDSS support to medical professionals. As part of future work, we aim to experiment with CT-scan images for potential improvements in performance, so that the final prediction is generated based on an ensemble of three classifiers that make use of X-ray, CT and radiography text reports. We also intend to benchmark the performance of the proposed method over existing solutions in terms of both interpretability and accuracy.

ACKNOWLEDGEMENT

The authors gratefully acknowledge the computational resources made available through the Google Cloud COVID-19 Research Grant, awarded to the third author in July 2020.

REFERENCES

Abadi, M., Barham, P., Chen, J., et al. (2016). Tensorflow: A system for large-scale machine learning. In *12th USENIX Symposium on Operating Systems Design and Implementation (OSDI 16)*, pages 265–283.

Abbas, A., Abdelsamea, M., and Gaber, M. (2020). Classification of covid-19 in chest x-ray images using detrack deep convolutional neural network. *Applied Intelligence*, pages 1–11.

BusinessToday (2020). Coronavirus testing in India.

Chung, A., Wang, L., et al. (2020). Actualmed COVID-19 Chest X-ray Dataset Initiative .

Cohen, J. P., Morrison, P., et al. (2020). COVID-19 image data collection. *arXiv 2006.11988*.

- Dansana, D., Kumar, R., et al. (2020). Early diagnosis of covid-19-affected patients based on x-ray and computed tomography images using deep learning algorithm. *Soft Computing*.
- Demner-Fushman, D., Kohli, M., et al. (2015). Preparing a collection of radiology examinations for distribution and retrieval. *JAMIA*, 23.
- Deng, J., Dong, W., et al. (2009). ImageNet: A Large-Scale Hierarchical Image Database. In *CVPR09*.
- Forthmann, P. and Pfliederer, G. (2019). Augmented display device for use in a medical imaging laboratory. US2018197337 Google Patents.
- Gangavarapu, T., Krishnan, G. S., Kamath, S., and Jegannathan, J. (2020). Farsight: Long-term disease prediction using unstructured clinical nursing notes. *IEEE Transactions on Emerging Topics in Computing*.
- Ghoshal, B. and Tucker, A. (2020). Estimating uncertainty and interpretability in deep learning for coronavirus (COVID-19) detection. *arXiv preprint arXiv:2003.10769*.
- Gonzalez, R. C. and Woods, R. E. (2008). *Digital image processing*. Prentice Hall, Upper Saddle River, N.J.
- Hall, L. O., Paul, R., Goldgof, D. B., and Goldgof, G. M. (2020). Finding covid-19 from chest x-rays using deep learning on a small dataset. *arXiv e-prints*, page arXiv:2004.02060.
- He, K., Zhang, X., et al. (2015). Deep residual learning for image recognition.
- Hemdan, E.-D., Shouman, M., and Karar, M. E. (2020). Covidx-net: A framework of deep learning classifiers to diagnose covid-19 in x-ray images. *arXiv preprint:2003.11055*.
- Isola, P., Zhu, J.-Y., Zhou, T., and Efros, A. A. (2018). Image-to-image translation with conditional adversarial networks.
- Krizhevsky, A. (2012). Convolutional deep belief networks on cifar-10.
- Max Roser, Hannah Ritchie, E. O.-O. and Hasell, J. (2020). Coronavirus pandemic (COVID-19). *Our World in Data*.
- Mei, X., Lee, H.-C., et al. (2020). Artificial intelligence-enabled rapid diagnosis of patients with covid-19. *Nature Medicine*, 26:1–5.
- Mikolov, T., Chen, K., Corrado, G., and Dean, J. (2013). Efficient Estimation of Word Representations in Vector Space. *arXiv e-prints*, page arXiv:1301.3781.
- Mishra, A., Das, S., et al. (2020). Identifying covid19 from chest ct images: A deep convolutional neural networks based approach. *Journal of Healthcare Engineering*, 2020:1–7.
- Narin, A., Kaya, C., and Pamuk, Z. (2020). Automatic detection of coronavirus disease (COVID-19) using x-ray images and deep convolutional neural networks. *arxiv preprint:2003.10849*.
- Ozturk, T., Talo, M., et al. (2020). Automated detection of COVID-19 cases using deep neural networks with x-ray images. *Computers in Biology and Medicine*.
- Paszke, A., Gross, S., Massa, F., et al. (2019). Pytorch: An imperative style, high-performance deep learning library. In *Advances in Neural Information Processing Systems* 32.
- Radiological Society of North America (2018). Rsn pneumonia detection challenge. <https://www.kaggle.com/c/rsna-pneumonia-detection-challenge>. Accessed: 2020-08-11.
- Rahman, T. (2020). COVID-19 radiography database.
- Řehůřek, R. and Sojka, P. (2010). Software Framework for Topic Modelling with Large Corpora. In *LREC 2010 Workshop on New Challenges for NLP Frameworks*. ELRA.
- Reza, A. (2004). Realization of the contrast limited adaptive histogram equalization (clahe) for real-time image enhancement. *VLSI Signal Processing*, 38:35–44.
- Ronneberger, O., Fischer, P., and Brox, T. (2015). U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 234–241. Springer.
- Scheib, S. (2009). Dosimetric end-to-end verification devices, systems, and methods. US 20150085993 Google Patents.
- Selvaraju, R., Cogswell, M., et al. (2017). Grad-cam: Visual explanations from deep networks via gradient-based localization. In *IEEE international conference on computer vision*, pages 618–626.
- Son, J., Park, S. J., and Jung, K.-H. (2017). Retinal Vessel Segmentation in Fundoscopic Images with Generative Adversarial Networks. *arXiv e-prints*, page arXiv:1706.09318.
- Tushaar, G., Jayasimha, A., Krishnan, G. S., and Kamath, S. (2020). Predicting icd-9 code groups with fuzzy similarity based supervised multi-label classification of unstructured clinical nursing notes. *Knowledge-Based Systems*, 190:105321.
- Zijdenbos, A., Dawant, B., Margolin, R., and Palmer, A. C. (1994). Morphometric analysis of white matter lesions in mr images: method and validation. *IEEE transactions on medical imaging*, 13 4:716–24.