COVID-Net FewSE: An Open-Source Deep Siamese Convolutional Network Model for Few-Shot Detection of COVID-19 Infection from X-Ray Images

Ashkan Ebadi — National Research Council Canada
Hilda Azimi — National Research Council Canada
Pengcheng Xi — National Research Council Canada
Stéphane Tremblay — National Research Council Canada
Alexander Wong — University of Waterloo
Email: firstname.lastname@nrc-cnrc.gc.ca, alexander.wong@uwaterloo.ca

Abstract

The COVID-19 pandemic continues impacting all segments of the global population, causing many problems, from health and well-being issues to definite irretrievable damage to the society. Despite the need for a quick and accurate response for early risk stratification and diagnosis, rare and novel diseases, e.g., COVID-19, are very difficult to diagnose. Although deep learning diagnostic algorithms have shown promising results in a wide range of tasks, they require a massive amount of labelled data for training. However, due to the nature of novel diseases, availability of such huge amount of well annotated data poses a great challenge to the learning algorithms. Motivated by this, in this work, we present an open-source deep meta learning solution based on siamese convolutional networks, called COVID-Net FewSE, that is able to detect COVID-19 positive cases from a limited number of X-ray images. Trained on the COVIDx-CXR dataset, the model achieves 0.9 recall and accuracy of 0.997 in detecting COVID-19 cases from X-ray images, when only 50 training samples are available. Our experimental results confirm that the proposed model outperforms conventional machine/deep learning classifiers in COVID-19 detection when limited samples are available. The model and all the scripts are made available to the public to enable reproducibility and encourage further innovation in the field.

1 Introduction

The novel coronavirus disease (COVID-19), caused by the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2), has widely spread all over the world since the beginning of 2020 [1]. The first case of the COVID-19 disease was observed in China in December 2019 [2]. To date, there are more than 255 million confirmed coronavirus cases all around the world with ≈5.1 million deaths reported [3], with numbers increasing every day. Despite the rapid evolution and emergence, the scientific community has dynamically and actively responded to the disease spread, changing research priorities according to the needs [4]. However, the high transmissibility and mortality rate of the COVID-19 virus as well as witnessing several waves of the pandemic have made early diagnosis crucial, to instantly quarantine infected patients and start the treatment/care procedure the earliest possible [1].

The reverse transcription polymerase chain reaction (RT-PCR) is recognized as one of the main approaches in diagnosing COVID-19 [5]. The RT-PCR testing is a time-consuming process due to the complex processes involved with collecting samples, transporting, and analyzing them. Moreover, the test suffers from low sensitivity, i.e., high false negative rates [5]. Therefore, many infected patients cannot be detected in time that increases the chance of infecting other people unknowingly [6]. To overcome this inefficiency, other alternative/complementary testing solutions are required.

Earlier works on medical images have shown that specific anomalies are found in COVID-19 patients' radiography such as ground-glass opacities with rounded morphology [7]. Chest X-ray (CXR) and computed tomography (CT) scans are two common imaging modalities that are being used for COVID-19 detection, severity assessment, and monitoring response to treatment [8]. Despite the higher detection sensitivity of CT scans, CXRs are more commonly being used in clinical settings due to their many advantages such as higher availability in general or community hospitals, lower cost, lower radiation dose, and ease of operation [9], as well as the availability of portable units [10]. But, interpreting chest X-ray images to diagnose COVID-19 and distinguishing it from other non-COVID-19 infections is a challenging and time-consuming task even for radiologists, due to many similarities [11].

Deep convolutional neural networks (CNNs) have been mostly utilized so far for COVID-19 classification. Despite being powerful image processing techniques, CNNs cannot correctly capture relations between image instances in the dataset, unless they are provided with a large dataset, including all possible transformations [12]. Moreover, training such deep networks on massive datasets is often a time-consuming task. Metric-based meta learning models, e.g., siamese networks, are alternative models that are able to generalize from few examples by employing a unique structure to rank similarity between input and without necessitating extensive retraining [13]. Siamese networks can be easily scaled up to include more categories [14], an advantage that make such models even more interesting in hard-to-control situations, e.g., spread of a new disease. Motivated by these properties and to assist clinicians in the fight against the COVID-19 pandemic and as part of the COVID-Net initiative [10, 15–17], in this work, we present an open-source deep siamese network model, called COVID-Net FewSE, for COVID-19 detection from X-ray images. Although siamese networks have been studied in the literature for several applications (e.g., object tracking [18], palmprint recognition [19]), to the best of our knowledge, this is the first work that employs transfer learning and presents an open-source few-shot deep siamese network model for COVID-19 detection from X-ray images. We hope the open-source nature of the COVID-Net FewSE encourage further innovation.

2 Data

We used the COVIDx-CXR dataset [10], to train, evaluate, and validate our model. The dataset contains chest x-ray images divided into three categories: 1) non-COVID pneumonia infection, 2) COVID-19 positive infection, and 3) normal control cases. We applied a deep learning-based lung segmentation approach [20] on the original data to detect lung boundaries from CXR images and crop images such that they only contain the lung region. We removed segmented CXR images of low-resolution or those that the model could not correctly calculate the lung region. Table 1 shows the distribution of the segmented images per category used in our analytics pipeline.

Table 1: Distribution of the segmented images per category.

	COVID-19	Non-COVID-19	Normal
Train	279	5,451	966
Test	99	100	100

3 Methods

We defined the problem as a 3-way n-shot classification problem where n is the number of training samples per each category, and 3 is the number of categories. The goal of n-shot learning is to classify unseen data as highly accurate as possible given the limited number of training samples of size n. The model was trained on pairs of images constructed from the training samples such that images in half of the pairs belonged to the same class while in the other half images in the pair sets were taken from different classes. We named the first half as the *genuine* and the second half as the *imposite* pairs.

We tested the model on the entire test set. For this purpose, we generated genuine and imposite pairs for each single image in the test set. For example, let's consider a test image from the COVID-19 class. We generated two test pairs for this test image. In the first pair, the second image was taken randomly from the COVID-19 class in the test set. For the second pair, the second image was

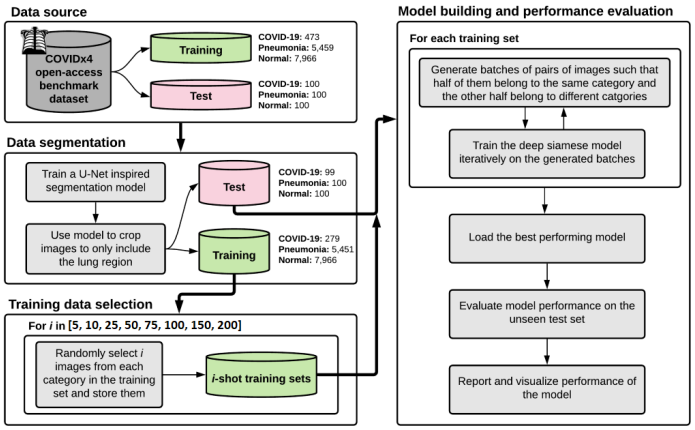

Fig. 1: High-level conceptual flow of the analysis.

taken randomly from either non-COVID pneumonia infection category or normal control images in the test set. To find the label for each test image, the pair with the lowest distance predicted by the siamese model was considered. If the class of the second image (from the pair with the lowest distance) was the same as the class of the first image, it was counted as a correct identification and incorrect otherwise. Performance is reported based on number of correct identifications. Fig. 1 shows the conceptual flow of the analysis.

We built the COVID-Net FewSE model as a deep siamese convolutional neural network that contains two identical fine-tuned pre-trained ResNet-50 models [21] as the embedding extractors, both sharing the same weights. The input is pairs of images and the model learns a distance measure between them over the training phase. We leveraged a pre-trained model as the embedding extractor since we have a limited dataset and our experiments proved it to be better performing compared with a standard convolutional neural network. Meanwhile, we tested several pre-trained models such as Resnet-50 [21], VGG-16 [22], Inception [23], InceptionResnet [23] as the base encoder and found the best performance with the ResNet-50 model. We added a customized layer to calculate the $L1$ distance between the generated embeddings. Binary cross entropy loss function is mostly used in the literature since it is a common choice for a binary classification problem. We tested both binary cross entropy and contrastive loss functions for performance evaluation and model building and found a better performance for the contrastive loss [24]. Input images were resized to 224x224 and normalized. No data augmentation was used in the pipeline. We used Adam optimizer with an initial learning rate set at $1e-4$. We reduced the learning rate by a factor of 0.1 when the performance stopped improving. Early stopping was also applied to stop the learning process if no performance improvement is observed for 10 epochs. The model was trained for 30 epochs.

Fig. 2 shows the high-level architecture design of the COVID-Net FewSE model. A pair of images, i.e., image 1 and 2 in the figure, is fed to the model. The model contains two identical parallel ResNet50 networks, sharing the same weights. Input images are each passed through the fine-tuned pre-trained ResNet50 model to obtain the feature embedding vectors. Then, these two embedding vectors are fed to an L_1 component-wise distance function to calculate the similarity between the two images. The similarity value is incorporated in the loss function. The embeddings are fed to the distance function, followed by 2 dense layers with 128 and 32 neurons, respectively, and dropout of 0.1. Finally, a dense layer with a sigmoid unit is used to generate the similarity score. Dropout and early stopping were used to prevent overfitting.

4 Results

We compared the performance of the siamese model in diagnosing COVID-19 positive cases with three baseline models: 1) random forests model (RF), 2) 2-layer vanilla CNN (vCNN), and 3) 1-nearest neighbors (1-NN). We were mainly interested in analyzing how does the performance of the COVID-Net FewSE model change versus the number of shots and if at those number of shots the baseline models can perform well.

Fig. 3 shows the performance of COVID-Net FewSE, RF, 1-NN, and vCNN models for various 3-way, n-shot learning settings ($n \in \{5, 10, 25, 50, 75, 100, 150, 200\}$). As seen, the COVID-Net FewSE model outperforms other baseline models in detecting COVID-19 positive cases. When there is only 50 shots available for training the model, COVID-19 recall of the COVID-Net FewSE model reaches the value of 0.9, outperforming the other models significantly. Also, as expected, performance of the vCNN model gradually increases as the number of shots available for training augments. Therefore, with having a large set of images available, performance of the CNN model would be expected to reach/pass the few-shot model.

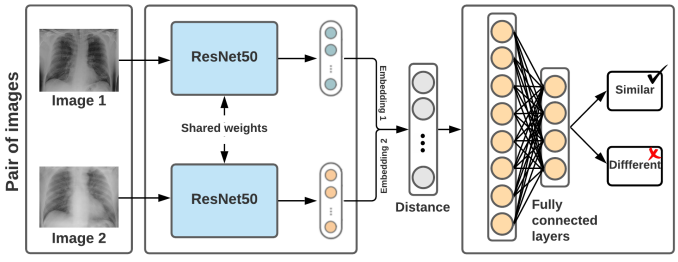

Fig. 2: The COVID-Net FewSE high-level architecture design.

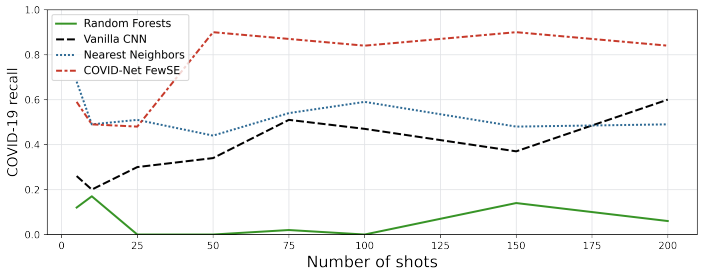

Fig. 3: Performance of the COVID-Net FewSE and three baseline models, i.e., random forests, vanilla CNN, and 1-NN, for various n-shot learning settings.

Fig. 4 shows the COVID-Net FewSE model training and validation accuracy and loss for 3-way, 50-shot learning setting. As seen, the model training and validation converge in the final epochs. During our experimentation, we also noticed that the COVID-Net FewSE model validation accuracy converge relatively fast, with few major oscillations in final epochs.

5 Conclusion and Future Work

Although deep learning diagnostic algorithms have shown promising results in a wide range of tasks, they require a massive amount of labelled data for training. However, due to the nature of novel diseases, availability of such huge amount of well annotated data poses a great challenge to the learning algorithms. As part of the COVID-Net initiative and using a highly imbalanced dataset of chest x-ray images with few COVID-19 observations, in this work, we presented a few-shot siamese convolutional neural network model, called the COVID-Net FewSE, able to detect COVID-19 positive cases with high accuracy, even if a limited number of examples is available. The model and all the scripts are made available to the public to enable reproducibility and encourage further innovation.

A fine-tuned ResNet-50 model, pre-trained on large ImageNet [25] data, was embedded in the COVID-Net FewSE model to enhance its performance by providing feature embeddings from the input images. In addition, we used a tailored lung segmentation model [20] to localize the lung region in the the x-ray images.

Our results show that the proposed model offers an accurate solution even if the number of images for the given disease (here COVID-19 infection) is limited. In addition, the model outperforms several examined baseline models, especially when only few shots are available for training. This property is crucial for detecting new diseases/pandemics and could accelerate patients screening and treatment planning.

As future research directions, we are currently working on two main directions to improve the COVID-Net FewSE model further: 1) The model proposed in this work was trained on a portion of the

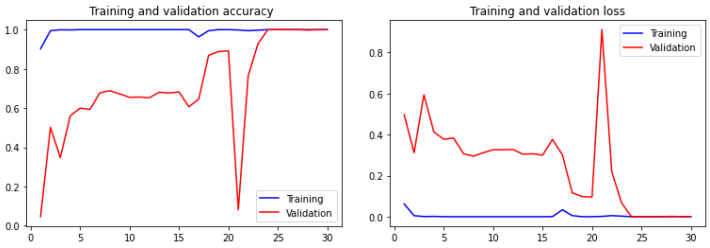

Fig. 4: Training and validation accuracy and loss for 3-way, 50-shot learning setting.

COVIDx-CXR v4 dataset to ensure the feasibility and confirm the analytics pipeline. As the next step, we will be using the latest version of the COVIDx-CXR (i.e., v8) that contains more data points. 2) We are incorporating an interpretability module to obtain a deeper understanding of the model learning process and ensure that the model's decisions are based on actual patterns. The outputs of the interpretability module will be verified by our contributing clinicians in the COVID-Net team.

Acknowledgements

We would like to thank Mr. Patrick Paul (National Research Council Canada) for helping us to access powerful computing resources. Also, we are grateful to the NRC's Pandemic Response Challenge Program for the financial support.

References

[1] D. Singh, V. Kumar, and M. Kaur, "Classification of covid-19 patients from chest ct images using multi-objective differential evolution–based convolutional neural networks." *European Journal of Clinical Microbiology & Infectious Diseases*, vol. 39, no. 7, pp. 1379–1389, 2020.

[2] E. I. Hui, D. S. amd Azhar, T. A. Madani, F. Ntoumi, R. Kock, O. Dar, G. Ippolito, T. D. Mchugh, Z. A. Memish, C. Drosten, and A. Zumla, "The continuing 2019-ncov epidemic threat of novel coronaviruses to global health—the latest 2019 novel coronavirus outbreak in wuhan, china." *International journal of infectious diseases*, vol. 91, pp. 264–266, 2020.

[3] Worldometer, "Covid-19 coronavirus pandemic." *From: https://www.worldometers.info/coronavirus/*, 2021.

[4] A. Ebadi, P. Xi, S. Tremblay, B. Spencer, R. Pall, and A. Wong, "Understanding the temporal evolution of covid-19 research through machine learning and natural language processing." *Scientometrics*, vol. 126, no. 1, pp. 725–739, 2021.

[5] T. Ai, Z. Yang, H. Hou, C. Zhan, C. Chen, W. Lv, Q. Tao, Z. Sun, and L. Xia, "Correlation of chest ct and rt-pcr testing for coronavirus disease 2019 (covid-19) in china: a report of 1014 cases." *Radiology*, vol. 296, no. 2, pp. E32–E40, 2020.

[6] Z. Islam, M. Islam, and A. Asraf, "A combined deep cnn-lstm network for the detection of novel coronavirus (covid-19) using x-ray images." *Informatics in medicine unlocked*, vol. 20, p. 100412, 2020.

[7] Y. Li and L. Xia, "Coronavirus disease 2019 (covid-19): role of chest ct in diagnosis and management." *American Journal of Roentgenology*, vol. 214, no. 6, pp. 1280–1286, 2020.

[8] M. Heidari, S. Mirniaharikandehei, M. Heidari, A. Zargari Khuzani, G. Danala, Y. Qiu, and B. Zheng, "Improving the performance of cnn to predict the likelihood of covid-19 using chest x-ray images with preprocessing algorithms." *International journal of medical informatics*, vol. 144, p. 104284, 2020.

[9] A. Narin, C. Kaya, and Z. Pamuk, "Automatic detection of coronavirus disease (covid-19) using x-ray images and deep convolutional neural networks." *arXiv preprint arXiv:2003.10849*, 2020.

[10] L. Wang, Z. Q. Lin, and A. Wong, "Covid-net: A tailored deep convolutional neural network design for detection of covid-19 cases from chest x-ray images." *Scientific Reports*, vol. 10, no. 1, pp. 1–12, 2020.

[11] W. C. Dai, H. W. Zhang, J. Yu, H. J. Xu, H. Chen, S. P. Luo, H. Zhang, and et al., "Ct imaging and differential diagnosis of covid-19." *Canadian Association of Radiologists Journal*, vol. 71, no. 2, pp. 195–200, 2020.

[12] P. Afshar, S. Heidarian, F. Naderkhani, A. Oikonomou, K. N. Plataniotis, and A. Mohammadi, "Covid-caps: A capsule network-based framework for identification of covid-19 cases from x-ray images." *Pattern Recognition Letters*, vol. 138, pp. 638–643, 2020.

[13] G. Koch, R. Zemel, and R. Salakhutdinov, "Siamese neural networks for one-shot image recognition." in *ICML deep learning workshop, vol. 2*, 2015.

[14] N. Roberts, P. S. Purushothama, V. T. Vasudevan, S. Ravichandran, C. Zhang, W. H. Gerwick, and G. W. Garrison, "Using deep siamese neural networks to speed up natural products research." 2018.

[15] A. MacLean, S. Abbasi, A. Ebadi, A. Zhao, M. Pavlova, H. Gunraj, P. Xi, S. Kohli, and A. Wong, "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*, 2021, pp. 191–202, Springer, Cham.

[16] H. Gunraj, A. Sabri, D. Koff, and A. Wong, "Covid-net ct-2: Enhanced deep neural networks for detection of covid-19 from chest ct images through bigger, more diverse learning." *arXiv preprint arXiv:2101.07433*, 2021.

[17] A. Wong, Z. Q. Lin, L. Wang, A. G. Chung, B. Shen, A. Abbasi, M. Hoshmand-Kochi, and T. Q. Duong, "Covid-net s: Towards computer-aided severity assessment via training and validation of deep neural networks for geographic extent and opacity extent scoring of chest x-rays for sars-cov-2 lung disease severity." *Scientific Reports*, 2021.

[18] L. Bertinetto, J. Valmadre, J. F. Henriques, A. Vedaldi, and P. H. Torr, "Fully-convolutional siamese networks for object tracking." in *European conference on computer vision*, 2016, pp. 850–865, Springer, Cham.

[19] D. Zhong, Y. Yang, and X. Du, "Palmprint recognition using siamese network." in *Chinese conference on biometric recognition*, 2018, pp. 48–55, Springer, Cham.

[20] H. Azimi, J. Zhang, P. Xi, H. Asád, A. Ebadi, S. Tremblay, and A. Wong, "Improving classification model performance on chest x-rays through lung segmentation." *Submitted to BMC Medical Imaging*, 2021.

[21] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition." in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2016, pp. 770–778.

[22] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition." *arXiv preprint arXiv:1409.1556*, 2014.

[23] C. Szegedy, S. Ioffe, V. Vanhoucke, and A. A. Alemi, "Inception-v4, inception-resnet and the impact of residual connections on learning." in *Thirty-first AAAI conference on artificial intelligence*, 2017.

[24] R. Hadsell, S. Chopra, and Y. LeCun, "Dimensionality reduction by learning an invariant mapping." in *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06)*, 2006, pp. 1735–1742.

[25] J. Deng, W. Dong, R. Socher, L. J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database." in *2009 IEEE conference on computer vision and pattern recognition*, 2009, pp. 248–255.