



OPEN

COVID-Net: a tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images

Linda Wang^{1,2,3}✉, Zhong Qiu Lin^{1,2,3} & Alexander Wong^{1,2,3}

The Coronavirus Disease 2019 (COVID-19) pandemic continues to have a devastating effect on the health and well-being of the global population. A critical step in the fight against COVID-19 is effective screening of infected patients, with one of the key screening approaches being radiology examination using chest radiography. It was found in early studies that patients present abnormalities in chest radiography images that are characteristic of those infected with COVID-19. Motivated by this and inspired by the open source efforts of the research community, in this study we introduce COVID-Net, a deep convolutional neural network design tailored for the detection of COVID-19 cases from chest X-ray (CXR) images that is open source and available to the general public. To the best of the authors' knowledge, COVID-Net is one of the first open source network designs for COVID-19 detection from CXR images at the time of initial release. We also introduce COVIDx, an open access benchmark dataset that we generated comprising of 13,975 CXR images across 13,870 patient cases, with the largest number of publicly available COVID-19 positive cases to the best of the authors' knowledge. Furthermore, we investigate how COVID-Net makes predictions using an explainability method in an attempt to not only gain deeper insights into critical factors associated with COVID cases, which can aid clinicians in improved screening, but also audit COVID-Net in a responsible and transparent manner to validate that it is making decisions based on relevant information from the CXR images. By no means a production-ready solution, the hope is that the open access COVID-Net, along with the description on constructing the open source COVIDx dataset, will be leveraged and build upon by both researchers and citizen data scientists alike to accelerate the development of highly accurate yet practical deep learning solutions for detecting COVID-19 cases and accelerate treatment of those who need it the most.

The Coronavirus Disease 2019 (COVID-19) pandemic continues to have a devastating effect on the health and well-being of the global population, caused by the infection of individuals by the severe acute respiratory syndrome coronavirus 2 (SARS-CoV-2). A critical step in the fight against COVID-19 is effective screening of infected patients, such that those infected can receive immediate treatment and care, as well as be isolated to mitigate the spread of the virus. The main screening method used for detecting COVID-19 cases is reverse transcriptase-polymerase chain reaction (RT-PCR)¹ testing, which can detect SARS-CoV-2 ribonucleic acid (RNA) from respiratory specimens (collected through a variety of means such as nasopharyngeal or oropharyngeal swabs). While RT-PCR testing is the gold standard as it is highly specific, it is a very time-consuming, laborious, and complicated manual process that is in short supply. Furthermore, the sensitivity of RT-PCR testing is highly variable and have not been reported in a clear and consistent manner to date², and initial findings in China showing relatively poor sensitivity³. Furthermore, subsequent findings showed highly variable positive rate depending on how the specimen was collected as well as decreasing positive rate with time after symptom onset^{4,5}.

An alternative screening method that has also been utilized for COVID-19 screening has been radiography examination, where chest radiography imaging (e.g., chest X-ray (CXR) or computed tomography (CT) imaging) is conducted and analyzed by radiologists to look for visual indicators associated with SARS-CoV-2 viral infection. It was found in early studies that patients present abnormalities in chest radiography images that are characteristic of those infected with COVID-19^{6–8}, with some suggesting that radiography examination could be

¹Department of Systems Design Engineering, University of Waterloo, Waterloo, Canada. ²Waterloo Artificial Intelligence Institute, Waterloo, Canada. ³DarwinAI Corp., Waterloo, Canada. ✉email: linda.wang@uwaterloo.ca

used as a primary tool for COVID-19 screening in epidemic areas⁹. For example, Huang et al.⁷ identified that the majority of the COVID-19 positive cases in their study presented bilateral radiographic abnormalities in CXR images, and Guan et al.⁸ identified COVID-19 positive cases in their study presented radiographic abnormalities such as ground-glass opacity, bilateral abnormalities, and interstitial abnormalities in CXR and CT images. While much of the recent discussion has revolved around CT imaging due to greater image detail in the acquisitions, there are several advantages to leveraging CXR imaging for COVID-19 screening amid the global COVID-19 pandemic, particularly in resource-constrained areas and heavily-affected areas:

- **Rapid triaging** CXR imaging enables rapid triaging of patients suspected of COVID-19 and can be done in parallel of viral testing (which takes time) to help relief the high volumes of patients especially in areas most affected where they have ran out of capacity (e.g., New York, Spain, and Italy), or even as standalone when viral testing isn't an option (low supplies). Furthermore, CXR imaging can be quite effective for triaging in geographic areas where patients are instructed to stay home until the onset of advanced symptoms (e.g., New York City), since abnormalities are often seen at time of presentation when patients suspected of COVID-19 arrive at clinical sites¹⁰.
- **Availability and accessibility** CXR imaging is readily available and accessible in many clinical sites and imaging centers as it is considered standard equipment in most healthcare systems. In particular, CXR imaging is much more readily available than CT imaging, especially in developing countries where CT scanners are cost prohibitive due to high equipment and maintenance costs.
- **Portability** The existence of portable CXR systems means that imaging can be performed within an isolation room, thus significantly reducing the risk of COVID-19 transmission during transport to fixed systems such as CT scanners as well as within the rooms housing the fixed imaging systems¹⁰.

As such, radiography examination can be conducted faster and have greater availability given the prevalence of chest radiology imaging systems in modern healthcare systems and the availability of portable units, making them a good complement to RT-PCR testing particularly since CXR imaging is often performed as part of standard procedure for patients with a respiratory complaint¹¹. It is also suggested that CXR imaging may have utility for situations where patients with initial negative RT-PCR testing outcomes return to the emergency department with worsening symptoms¹². Furthermore, some have suggested that as the COVID-19 pandemic progresses, there will be a greater reliance on portable CXR due to the aforementioned advantages¹³, and found portable CXR to be highly valuable for critically ill COVID-19 patients¹⁴. Hence, CXR are an integral part of some screening strategies that have been proposed¹⁵. However, one of the biggest bottlenecks faced is the need for expert radiologists to interpret the radiography images, since the visual indicators can be subtle. As such, computer-aided diagnostic systems that can aid radiologists to more rapidly and accurately interpret radiography images to detect COVID-19 cases is highly desired.

Motivated by the urgent need to develop solutions to aid in the fight against the COVID-19 pandemic, inspired by the open source and open access efforts of the research community, and intrigued in exploring the efficacy of AI systems leveraging the more readily available and accessible CXR imaging modality, this study introduces COVID-Net, a deep convolutional neural network design tailored for the detection of COVID-19 cases from CXR images that is open source and available to the general public. The main contributions of this study are:

- Introduction of a novel deep neural network architecture that is tailored for the detection of COVID-19 cases from CXR images using a human-machine collaborative design strategy. To the best of the authors' knowledge, COVID-Net is the first neural network architecture designed for COVID-19 detection to introduce a lightweight projection-expansion-projection-extension (PEPX) design, which enables enhanced representational capacity while greatly reducing computational complexity. Furthermore, to the best of the authors' knowledge, COVID-Net is one of the first open source network designs for COVID-19 detection from CXR images at the time of initial release, which encourages reproducibility.
- Introduction of COVIDx, an open access benchmark dataset that we generated comprising of 13,975 CXR images across 13,870 patient cases, created as a combination and modification of five open access data repositories containing chest radiography images (i.e.,^{16–20}), two of which we introduced^{17,18}. To the best of the authors' knowledge, COVIDx is the largest open access benchmark dataset available in terms of the number of publicly available COVID-19 positive cases.
- Introduction of an explainability approach to investigating how COVID-Net makes predictions in an attempt to not only gain deeper insights into critical factors associated with COVID cases, which can aid clinicians in improved screening, but also audit COVID-Net in a responsible and transparent manner to validate that it is making decisions based on relevant information from the CXR images.

The paper is organized as follows. First, “[Introduction](#)” discusses the strategy used to create the COVIDx dataset, the strategy leveraged to create the proposed COVID-Net, the architecture design of COVID-Net, the implementation details of COVID-Net, and the strategy leveraged to audit COVID-Net via explainability. “[Methods](#)” presents and discusses the results of experiments conducted to evaluate the efficacy of the proposed COVID-Net in both a quantitative and qualitative manner. Finally, conclusions are drawn and future directions discussed in “[Experimental results](#)”.

Related work

Motivated by the need for faster interpretation of radiography images, a number of artificial intelligence (AI) systems based on deep learning²¹ have been proposed and results have shown to be quite promising in terms of accuracy in detecting patients infected with COVID-19 via radiography imaging, with the focus primarily on CT imaging^{22–25}. However, to the best of the authors' knowledge at the time of the initial release of the proposed COVID-Net, most of the developed AI systems proposed in research literature have been closed source and unavailable to the research community to build upon for deeper understanding and extension of these systems. Furthermore, most of these systems are unavailable for public access and use. There has been significant recent efforts to push for open access and open source AI solutions for radiography-driven COVID-19 case detection^{16–18}, with an exemplary effort being the open source COVID-19 Image Data Collection, an effort by Cohen et al.¹⁶ to build a dataset consisting of COVID-19 cases (as well as severe acute respiratory syndrome (SARS) and Middle East respiratory syndrome (MERS) cases) with annotated CXR and CT images, so that the research community and citizen data scientists can leverage the dataset to explore and build AI systems for COVID-19 detection. Since the time of the initial public release of the proposed COVIDx dataset and the proposed COVID-Net, a number of studies have been conducted in the area of COVID-19 detection using CXR images^{26–38}, several of which have leveraged variants of the COVIDx dataset or COVID-Net to conduct such studies^{28–35,37}. Based on a comprehensive survey of the studies in research literature, they have focused primarily on the exploration of deep learning, in particular deep convolutional neural networks, given the significant successes and state-of-the-art achieved in a variety of computer vision tasks. It is important to note that the proposed COVID-Net and corresponding COVIDx dataset continues to evolve as new patient cases are continuously added and are made available publicly on a regular basis, and thus this study represents a snapshot of the current state of COVID-Net and COVIDx.

Methods

In this study, a human-machine collaborative design strategy is leveraged to create COVID-Net, where human-driven principled network design prototyping is combined with machine-driven design exploration to produce a network architecture tailored for the detection of COVID-19 cases from CXR images. An open access benchmark dataset called COVIDx is also created to facilitate for training and evaluating COVID-Net. Here, we will discuss in detail the process of creating the COVIDx dataset, the architecture design methodology behind the proposed COVID-Net, the resulting network architecture, the implementation details in creating COVID-Net, and the strategy for auditing COVID-Net via explainability.

COVIDx dataset. The dataset used to train and evaluate the proposed COVID-Net, which we will refer to as COVIDx, is comprised of a total of 13,975 CXR images across 13,870 patient cases. To the best of the authors' knowledge, the proposed COVIDx dataset is the largest open access benchmark dataset in terms of the number of COVID-19 positive patient cases. To generate the COVIDx dataset, we combined and modified five different publicly available data repositories: (1) COVID-19 Image Data Collection¹⁶, (2) Fig. 1 COVID-19 Chest X-ray Dataset Initiative¹⁷, established in collaboration with Fig. 1, (3) ActualMed COVID-19 Chest X-ray Dataset Initiative¹⁸, established in collaboration with ActualMed, (4) RSNA Pneumonia Detection Challenge dataset²⁰, which used publicly available CXR data from³⁹, and (5) COVID-19 radiography database¹⁹. Example CXR images from the COVIDx dataset are shown in Fig. 2 in illustrate the diversity of patient cases in the dataset. The choice of these five datasets from which to create COVIDx is guided by the fact that all five of the datasets are open source and fully accessible to the research community and the general public, and as datasets grow we will continue to grow COVIDx accordingly.

More specifically, we combined and modified the five data repositories to create the COVIDx dataset by leveraging the following types of patient cases from each of the data repositories:

- Non-COVID19 pneumonia patient cases and COVID-19 patient cases from the COVID-19 Image Data Collection¹⁶
- COVID-19 patient cases from the Fig. 1 COVID-19 Chest X-ray Dataset Initiative¹⁷,
- COVID-19 patient cases from the ActualMed COVID-19 Chest X-ray Dataset Initiative¹⁸
- Patient cases who have no pneumonia (i.e., normal) and non-COVID19 pneumonia patient cases from RSNA Pneumonia Detection Challenge dataset²⁰
- COVID-19 patient cases from COVID-19 radiography database¹⁹

The distribution of images and patient cases amongst the different infection types shown in Figs. 3 and 4, respectively. The most noticeable trend is the limited amount of COVID-19 infection cases and associated CXR images, which reflects the scarcity of COVID-19 case data available in the public domain but also highlights the need to obtain more COVID-19 data as more case data becomes available to improve the dataset. More specifically, the COVIDx dataset contains 358 CXR images from 266 COVID-19 patient cases. For CXR images with no pneumonia and non-COVID19 pneumonia, there are significantly more patient cases and corresponding CXR images. More specifically, there are a total of 8,066 patient cases who have no pneumonia (i.e., normal) and 5,538 patient cases who have non-COVID19 pneumonia. Dataset generation scripts for constructing the COVIDx dataset is available publicly for open access at <https://github.com/lindawang/COVID-Net>.

Principled network design prototyping. The first stage of the human-machine collaborative design strategy employed to create the proposed COVID-Net is a principled network design prototyping stage, where an initial network design prototype is constructed based on human-driven design principles and best practices.

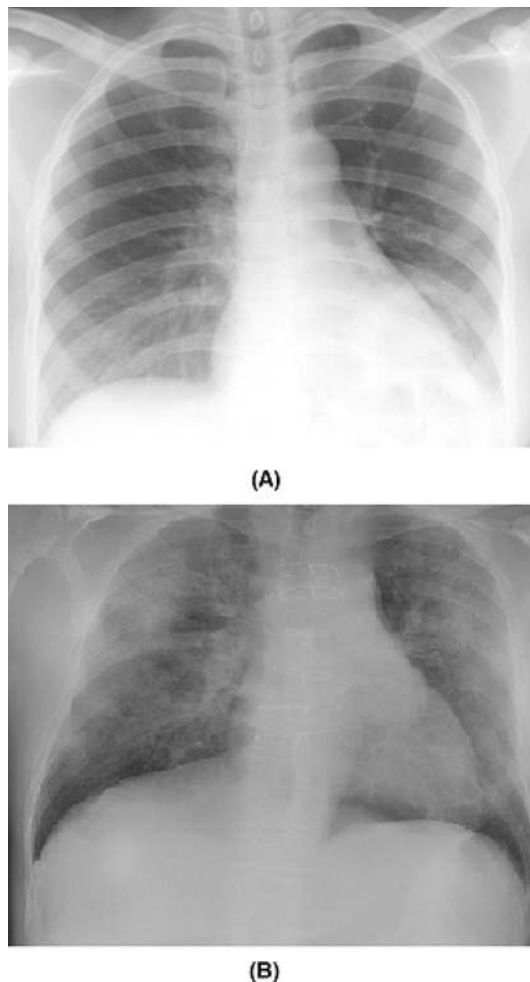


Figure 1. Example CXR images of: (A) non-COVID19 infection, and (B) COVID-19 viral infection in the COVIDx dataset.

More specifically in this study, we leveraged residual architecture design principles⁴⁰ as they have been shown time and again to enable reliable neural network architectures that are easier to train to high performance, and enables deeper architectures to be built successfully.

In this study, we construct the initial network design prototype to make one of the following three predictions: (a) no infection (normal), (b) non-COVID19 infection (e.g., viral, bacterial, etc.), and c) COVID-19 viral infection (see Fig. 1 for example CXR images of non-COVID19 and COVID-19 infections). The rationale for choosing these three possible predictions is that it can aid clinicians to better decide not only who should be prioritized for RT-PCR testing for COVID-19 case confirmation, but also which treatment strategy to employ depending on the cause of infection, since COVID-19 and non-COVID19 infections require different treatment plans.

Machine-driven design exploration. The second stage of the human-machine collaborative design strategy employed to create the proposed COVID-Net is a machine-driven design exploration stage. More specifically, at this stage, the initial network design prototype, data, along with human specific design requirements, act as a guide to a design exploration strategy to learn and identify the optimal macroarchitecture and microarchitecture designs with which to construct the final tailor-made deep neural network architecture. Such a machine-driven design exploration stage enables much greater granularity and much greater flexibility than is possible through manual human-driven architecture design, while still ensuring that the resulting deep neural network architecture satisfies domain-specific operational requirements. This is especially important for the design of COVID-Net, where sensitivity to COVID-19 cases is significant to limit the number of missed COVID-19 cases as much as possible.

In this study, we leverage generative synthesis⁴¹ as the machine-driven design exploration strategy, which is based on an intricate interplay between a generator-inquisitor pair that work in tandem to garner insights and learn to generate deep neural network architectures that best satisfies human specified design requirements. More specifically, the following human specified design requirements were employed in this study to enable the generative synthesis process to learn and identify the optimal macroarchitecture and microarchitecture designs for the final COVID-Net network architecture: (1) COVID-19 sensitivity $\geq 80\%$, and (2) COVID-19 positive

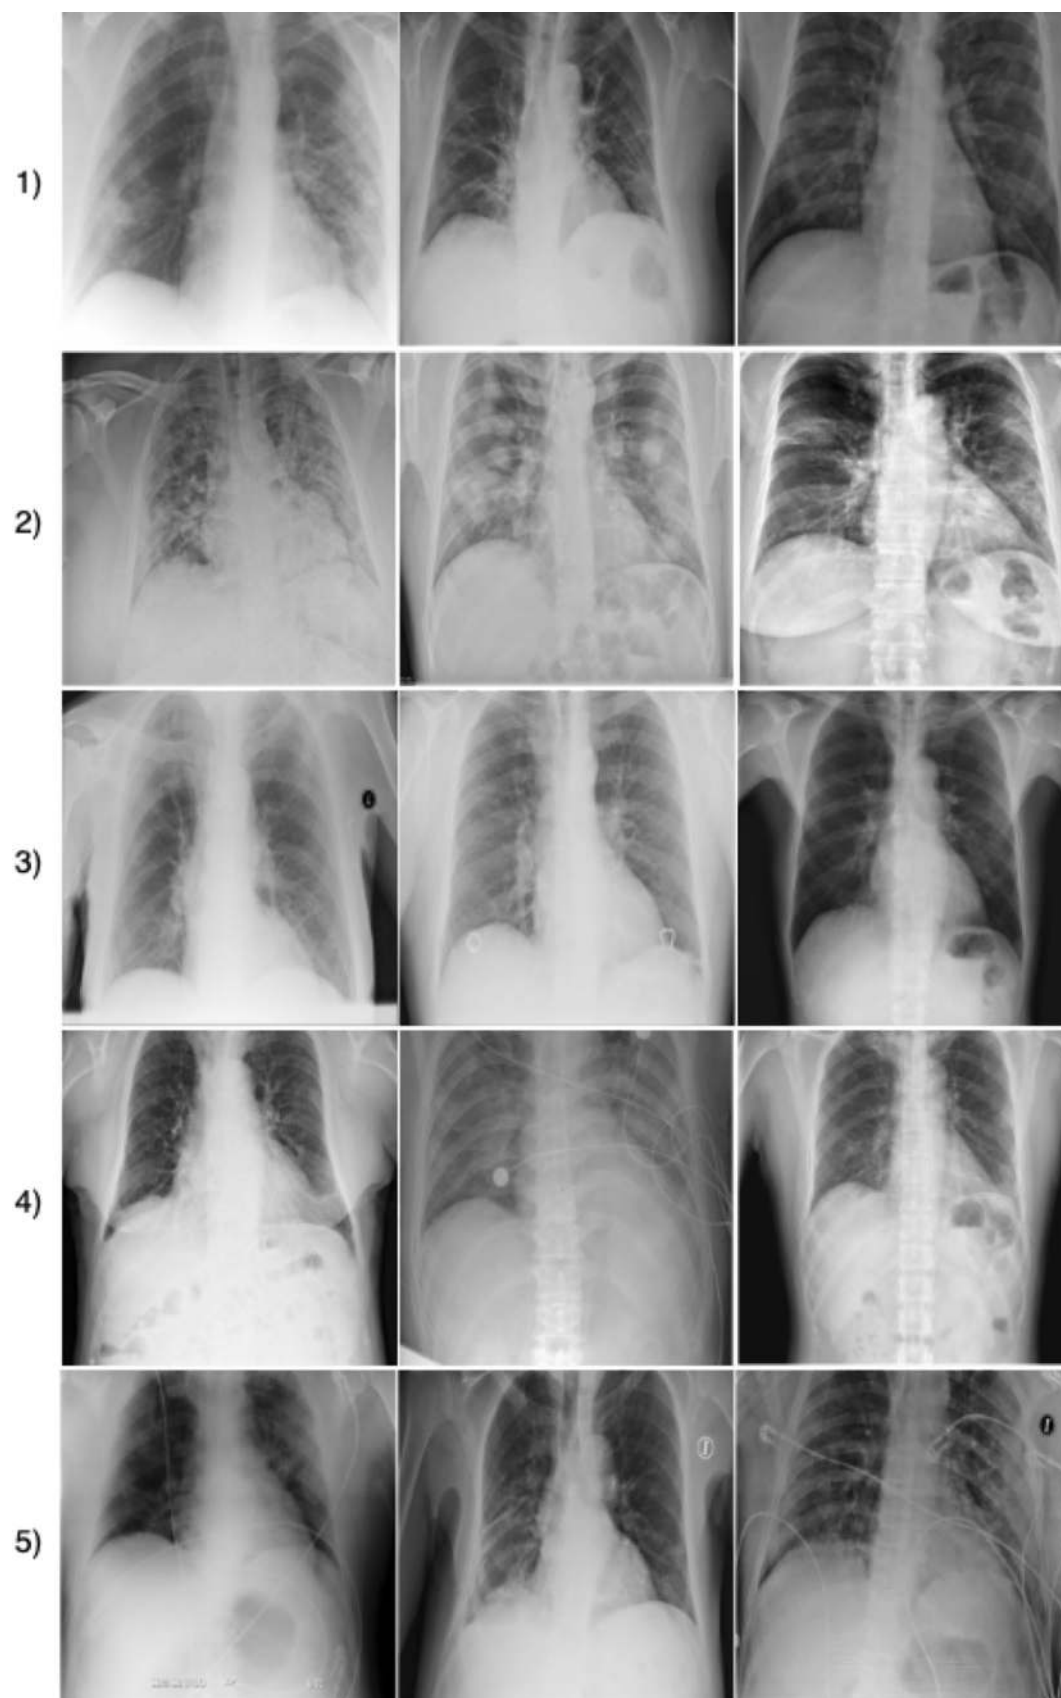


Figure 2. Example CXR images from the COVIDx dataset, which comprises of 13,975 CXR images across 13,870 patient cases from five open access data repositories: (1) COVID-19 Image Data Collection¹⁶, (2) Figure 1 COVID-19 Chest X-Ray Dataset Initiative¹⁷ (established with Fig. 1), (3) RSNA Pneumonia Detection challenge dataset²⁰, (4) ActualMed COVID-19 Chest X-Ray Dataset Initiative¹⁸ (established with ActualMed), and (5) COVID-19 radiography database¹⁹.

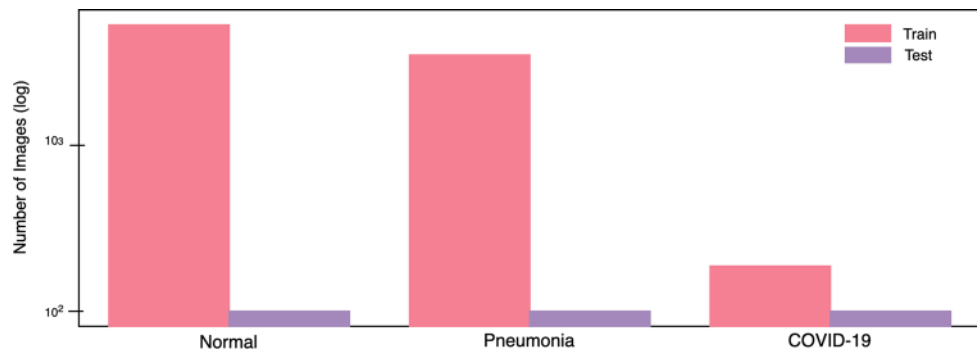


Figure 3. CXR images distribution for each infection type of the COVIDx dataset (normal means no infection). (Left bar) number of training images, (right bar) number of test images.

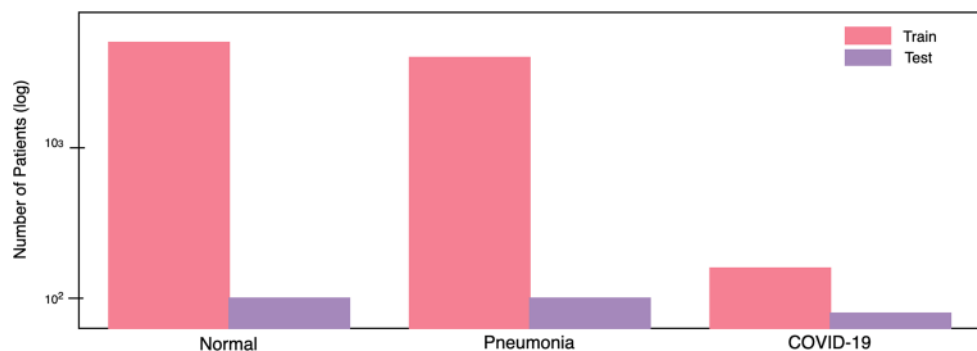


Figure 4. Number of patient cases for each infection type of the COVIDx dataset (normal means no infection). (Left bar) number of patient cases for training, (right bar) number of patient cases for testing.

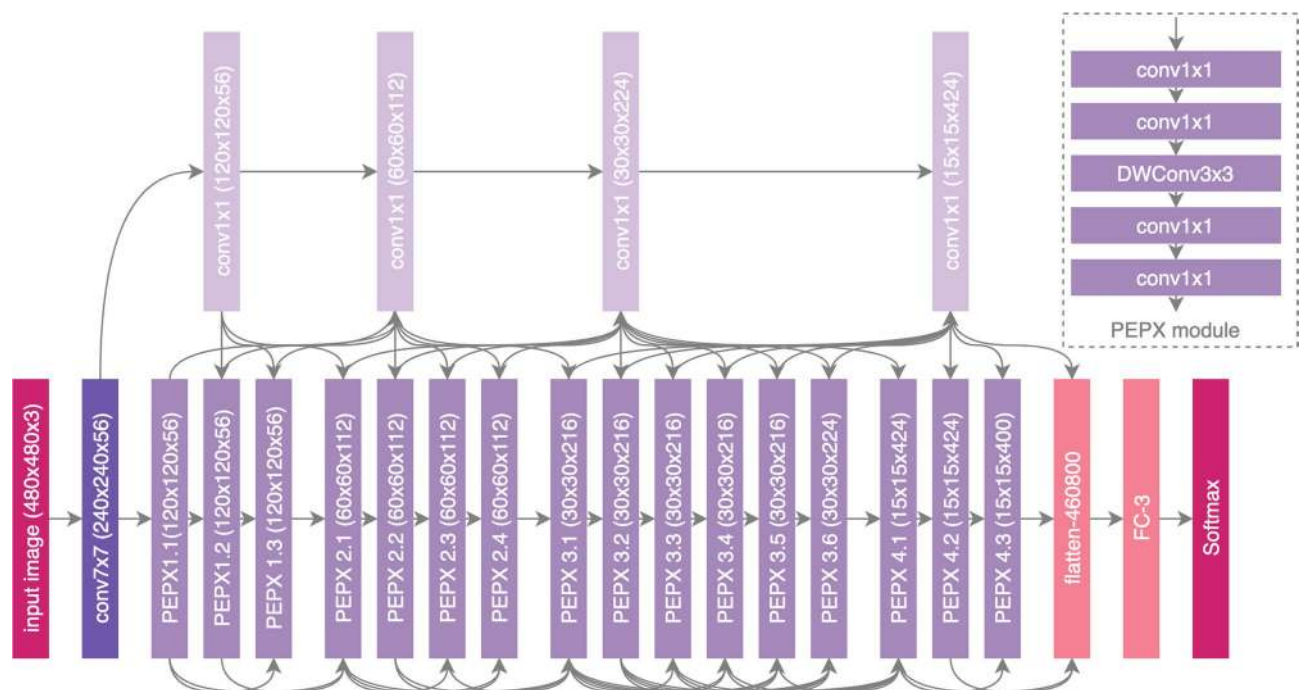


Figure 5. COVID-Net architecture. High architectural diversity and selective long-range connectivity can be observed as it is tailored for COVID-19 case detection from CXR images. The heavy use of a projection-expansion-projection design pattern in the COVID-Net architecture can also be observed, which provides enhanced representational capacity while maintaining computational efficiency.

predictive value (PPV) $\geq 80\%$. By employing the aforementioned human specified design requirements, the machine-driven design exploration can be conducted in a way that ensures that the resulting COVID-Net is able to detect COVID-19 positive cases while limiting the number of false positive COVID-19 detections to avoid overwhelming clinical sites with unnecessary burden.

COVID-Net network architecture. The proposed COVID-Net network architecture is shown in Fig. 5, and available publicly for open access at <https://github.com/lindawangg/COVID-Net>. A number of interesting observations can be made with regards to the COVID-Net network architecture design that was created via a human-machine collaborative design strategy.

Lightweight design pattern. It can be observed that the COVID-Net network architecture makes heavy use of a lightweight residual projection-expansion-projection-extension (PEPX) design pattern, which consists of:

- **First-stage projection** 1×1 convolutions for projecting input features to a lower dimension,
- **Expansion:** 1×1 convolutions for expanding features to a higher dimension that is different than that of the input features,
- **Depth-wise representation** efficient 3×3 depth-wise convolutions for learning spatial characteristics to minimize computational complexity while preserving representational capacity,
- **Second-stage projection** 1×1 convolutions for projecting features back to a lower dimension, and
- **Extension:** 1×1 convolutions that finally extend channel dimensionality to a higher dimension to produce the final features.

Such a customized lightweight design pattern, which is discovered by the machine-driven design exploration strategy and not previously introduced in literature to the best of the authors' knowledge, enables enhanced representational capacity while maintaining reduced computational complexity.

Selective long-range connectivity. It can also be observed that the COVID-Net architecture possesses selective long-range connectivity at various areas of the network architecture. The use of long-range connectivity has been found previously to enable improved representational capacity as well as make it easier to train, with an exemplary network architecture leveraging long-range connectivity being densely-connected deep neural network architectures⁴². However, a disadvantage of leveraging a large number of long-range connections as done in densely-connected deep neural network architectures is a noticeable increase in computational complexity as well as memory overhead. To alleviate this, it can be observed that the COVID-Net network architecture only leverages long range connections in a sparing manner where necessary, as exhibited by the existence of four densely connected 1×1 convolution layers that act as central hubs of long-range connectivity for earlier layers to be connected to much later layers in the network architecture. As such, the COVID-Net network architecture is able to achieve high representational capacity and improve ease of training while still maintaining computational and memory efficiency. This also highlights the advantages of a machine-driven design strategy for identifying the optimal connectivity within a deep neural network architecture, tailored in this case specifically for the task of COVID-19 detection based on CXR images.

Architectural diversity. Finally, it can be observed that there is considerable architectural diversity in the COVID-Net architecture. In particular, the COVID-Net network architecture is comprised of a heterogeneous mix of convolution layers with a diversity of kernel sizes (ranging from 7×7 to 1×1), and grouping configurations (ranging from ungrouped to depth-wise). The considerable architectural diversity exhibited by the COVID-Net architecture further reinforces the fact that the machine-driven design exploration strategy has tailored the network architecture at a very fine level of granularity for COVID-19 case detection from CXR images to achieve strong representational capacity for a specific task.

Implementation details. The proposed COVID-Net was pretrained on the ImageNet⁴³ dataset and then trained on the COVIDx dataset using the Adam optimizer using a learning rate policy where the learning rate decreases when learning stagnates for a period of time (i.e., 'patience'). The following hyperparameters were used for training: learning rate = $2e-4$, number of epochs = 22, batch size = 64, factor = 0.7, patience = 5. Finally, we introduce a batch re-balancing strategy to promote better distribution of each infection type at a batch level. The initial COVID-Net prototype was built and evaluated using the Keras deep learning library with a TensorFlow backend. The proposed COVID-Net architecture was built using generative synthesis⁴¹, as described in "Machine-driven design exploration".

COVID-Net auditing via explainability. Due to the mission-critical nature of clinical applications such as COVID-19 detection that can affect the health and well-being of patients, it is important to design deep neural network architectures such as COVID-Net with responsibility and transparency in mind. Therefore, in this study, we perform an explainability-driven audit on COVID-Net to validate that it is making detection decisions based on relevant information rather than improper information (e.g., erroneous visual indicators outside of the body, embedded markup symbols, imaging artifacts, etc.). More specifically, we audit COVID-Net via an qualitative analysis to study the critical factors leveraged by COVID-Net in making detection decisions. Here, we leveraged GSInquire⁴⁴, an explainability method that is a critical aspect of the generative synthesis strategy⁴¹

Architecture	Params (M)	MACs (G)	Acc. (%)
VGG-19	20.37	89.63	83.0
ResNet-50	24.97	17.75	90.6
COVID-Net	11.75	7.50	93.3

Table 1. Performance of tested deep neural network architectures on COVIDx test dataset. Best results highlighted in bold.

Architecture	Normal	Non-COVID19	COVID-19
Sensitivity (%)			
VGG-19	98.0	90.0	58.7
ResNet-50	97.0	92.0	83.0
COVID-Net	95.0	94.0	91.0

Table 2. Sensitivity for each infection type. Best results highlighted in bold.

leveraged in the machine-driven exploration strategy used to create the proposed COVID-Net network architecture. A brief summary of the GSInquire is provided as follows.

GSInquire revolves around the notion of an inquisitor \mathcal{I} within a generator-inquisitor pair $\{\mathcal{G}, \mathcal{I}\}$, with \mathcal{G} denoting a generator, that work in tandem to obtain improved insights about deep neural networks as well as learn to generate networks. The insights gained by \mathcal{I} can not only be used to improve \mathcal{G} to generate better networks, but also be subsequently transformed into an interpretation of decisions made by a network. More specifically, a deep neural network is defined as a graph $N = \{V, E\}$, comprising a set V of vertices $v \in V$ and a set E of edges $e \in E$ that form the network. A generator function is defined as $\mathcal{G}(s; \theta_{\mathcal{G}})$ parameterized by $\theta_{\mathcal{G}}$ that, given a seed $s \in S$, generates a deep neural network $N_s = \{V_s, E_s\}$ (i.e., $N_s = \mathcal{G}(s)$), where S is the set of possible seeds. Finally, an inquisitor function is defined as $\mathcal{I}(\mathcal{G}; \theta_{\mathcal{I}})$ parameterized by $\theta_{\mathcal{I}}$ that, given a generator \mathcal{G} , produces a set of parameter changes $\Delta\theta_{\mathcal{G}}$ (i.e., $\Delta\theta_{\mathcal{G}} = \mathcal{I}(\mathcal{G})$).

In the scenario where the underlying goal is to obtain an interpretation z of a decision made by a reference network N_{ref} (in this case, COVID-Net) for an input signal x (in this case, a CXR image), both $\theta_{\mathcal{G}}$ and $\theta_{\mathcal{I}}$ are initialized based on $\{V_{ref}, E_{ref}\}$, a universal performance function \mathcal{U} (e.g.,⁴⁵), and an indicator function $1_r(\cdot)$ s.t. $\{V_s, E_s\} = \{V_{ref}, E_{ref}\}$ to ensure interpretation consistency for N_{ref} . Given the generated $N_s = \mathcal{G}(s)$, the inquisitor \mathcal{I} probes $\{V_s, E_s\}$, where $V_s \subseteq V_s$ and $E_s \subseteq E_s$, with the targeted stimulus signal as x and the corresponding set $Y_{\mathcal{G}(s)}$ of reactionary response signals $y \in Y_{\mathcal{G}(s)}$ are observed. The parameters $\theta_{\mathcal{I}}$ are updated based on $Y_{\mathcal{G}(s)}$, $\mathcal{U}(\mathcal{G}(s))$, and $1_r(\mathcal{G}(s))$, leading to the inquisitor \mathcal{I} learning from the insights that are derived from $Y_{\mathcal{G}(s)}$. Following the update of $\theta_{\mathcal{I}}$, set of parameters $\Delta\theta_{\mathcal{G}} = \mathcal{I}(\mathcal{G})$ is generated which can not only be leveraged to update $\theta_{\mathcal{G}}$ to improve \mathcal{G} , but can also be transformed and projected into same subspace as x via a transformation $\mathcal{T}(\Delta\theta_{\mathcal{G}(s)})$ to produce an interpretation $z(x; N_{ref})$. In this study, the produced interpretation indicates the critical factors leveraged by COVID-Net in making a detection decision based on a CXR image, and can be visualized spatially relative to the CXR image for greater insights into whether COVID-Net is making the right decisions for the right reasons and validate its performance.

Ethics approval. The study has received ethics clearance from the University of Waterloo (42235).

Experimental results

To evaluate the efficacy of the proposed COVID-Net, we perform both quantitative and qualitative analysis to get a better understanding of its detection performance and decision-making behaviour.

Data pre-processing details. The COVIDx dataset was used to train all tested deep neural network architectures. As a pre-processing step, the chest CXR images were cropped (top 8% of the image) prior to training in order to mitigate commonly-found embedded textual information in the CXR images. Furthermore, to train the tested deep neural network architectures, data augmentation was leveraged with the following augmentation types: translation ($\pm 10\%$ in x and y directions), rotation ($\pm 10^\circ$), horizontal flip, zoom ($\pm 15\%$), and intensity shift ($\pm 10\%$). Data pre-processing scripts are available publicly for open access at <https://github.com/lindawang/COVID-Net>.

Quantitative analysis. To investigate the proposed COVID-Net in a quantitative manner, we computed the test accuracy, as well as sensitivity and positive predictive value (PPV) for each infection type, on the aforementioned COVIDx dataset. The test accuracy, along with the architectural complexity (in terms of number of parameters) and computational complexity (in terms of number of multiply-accumulation (MAC) operations) are shown in Table 1. It can be observed that COVID-Net achieves good accuracy by achieving 93.3% test accuracy, thus highlighting the efficacy of leveraging a human-machine collaborative design strategy for creating

Architecture	Normal	Non-COVID19	COVID-19
Positive predictive value (%)			
VGG-19	83.1	75.0	98.4
ResNet-50	88.2	86.8	98.8
COVID-Net	90.5	91.3	98.9

Table 3. Positive predictive value (PPV) for each infection type. Best results highlighted in bold.

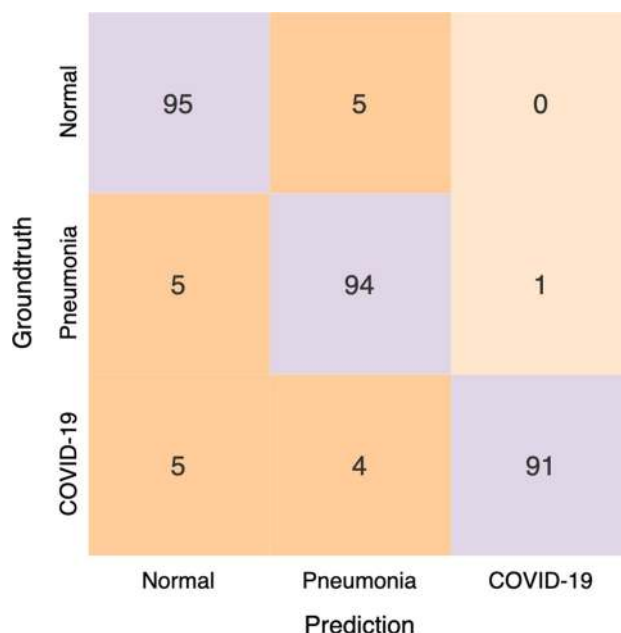


Figure 6. Confusion matrix for COVID-Net on the COVIDx test dataset.

highly-customized deep neural network architectures in an accelerated manner, tailored around task, data, and operational requirements. This is especially important for scenarios such as disease detection, where new cases and new data are collected continuously and the ability to rapidly generate new deep neural network architectures tailored to the ever-evolving knowledge base over time is highly desired.

Next, we take a deeper exploration into the current limitations of the proposed COVID-Net by studying the sensitivity and PPV for each infection type, which is shown in Tables 2 and 3, respectively, and the confusion matrix in Fig. 6. A number of interesting observations can be made about how COVID-Net performs under the different scenarios. First, it can be observed that COVID-Net can achieve good sensitivity for COVID-19 cases (91.0% sensitivity), which is important since we want to limit the number of missed COVID-19 cases as much as possible. Second, it can be observed that COVID-Net achieves high PPV for COVID-19 cases (98.9% PPV), which indicates very few false positive COVID-19 detections (for example, as seen in Fig. 6, one patient with non-COVID19 infection was misidentified as having COVID-19 viral infections). This high PPV is important given that too many false positives would increase the burden for the healthcare system due to the need for additional PCR testing and additional care. Therefore, based on these results, it can be seen that while COVID-Net performs well as a whole in detecting COVID-19 cases from CXR images, there are several areas of improvement that can benefit from collecting additional data, as well as improving the underlying training methodology to generalize better across such scenarios.

Architecture comparisons. We now take a deep exploration into the impact of architectural design choices made by generative synthesis during the machine-driven design exploration process on the resulting COVID-Net network architecture being able to achieve in terms of balance between computational efficiency and performance. It is important to note that, based on a thorough survey of all machine learning methods for COVID-19 detection using chest x-rays in research literature^{26–37}, the machine learning methods leveraged in these studies have been deep learning approaches. As such, we have focused this study to quantitatively compare between different deep neural network architectures to study the impact on computational efficiency and performance.

In order to perform this analysis, we evaluated the performances of the following deep neural network architectures for comparative purposes:

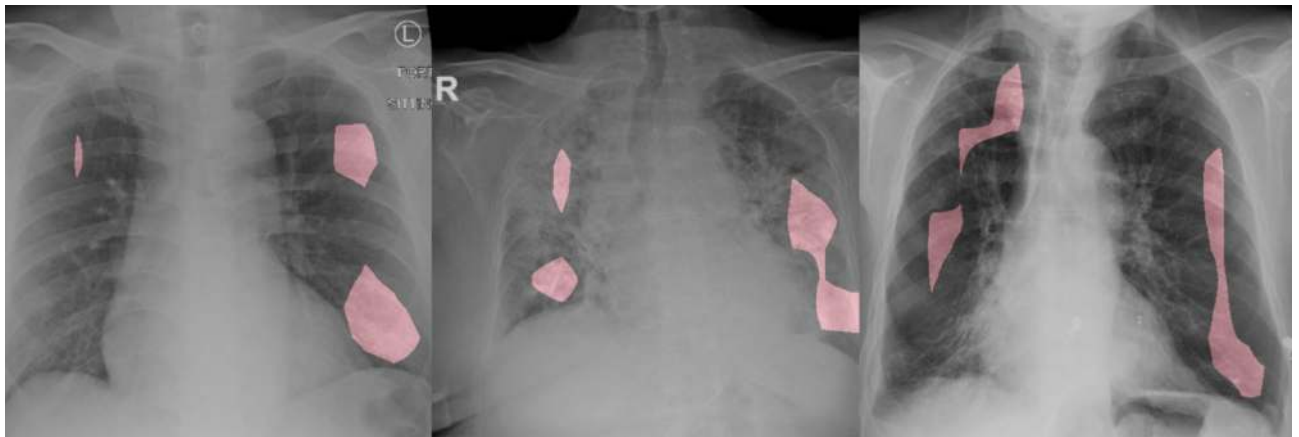


Figure 7. Example CXR images of COVID-19 cases from several different patients and their associated critical factors (highlighted in red) as identified by GSInquire⁴⁴.

- **VGG-19**⁴⁶ A deep neural network architecture that does not leverage residual design principles, lightweight design patterns, and have very low architectural diversity.
- **ResNet-50**⁴⁰ A deep neural network architecture that leverages residual design principles and lightweight design patterns (e.g., bottleneck design patterns) and have moderate architectural diversity, but does not leverage lightweight PEPX design patterns and selective long-range connectivity. This architecture was leveraged in a study³⁷ on an early variant of COVIDx to good effect. As such, a quantitative performance evaluation of this architecture enables a more direct comparison with³⁷ and highlight the benefits of the uniqueness of the COVID-Net architecture.

The choice of these two deep neural network architectures is based on the fact that they do not possess the key defining traits of COVID-Net, which are lightweight PEPX design patterns, selective long-range connectivity, and high architectural diversity, thus allowing for a better understanding of the benefits of the unique traits of COVID-Net.

It can be observed in Table 1 that COVID-Net had noticeably lower architectural complexity and computational complexity than the VGG-19 and ResNet-50 architectures. For example, the proposed COVID-Net requires $\sim 12\times$ fewer MAC operations than the VGG-19 architecture and $\sim 2.37\times$ fewer MAC operations than the ResNet-50 architecture, respectively. This illustrates the benefits of the lightweight PEPX design patterns within the COVID-Net architecture compared to not using lightweight design patterns (e.g., VGG-19) and using other types of lightweight design patterns (e.g., bottleneck patterns as used in ResNet-50). It can also be observed in Tables 1, 2, and 3 that COVID-Net achieved noticeably higher test accuracy and COVID-19 sensitivity than the VGG-19 and ResNet-50 network architectures. For example, COVID-19 sensitivity of COVID-Net is $> 32\%$ higher than VGG-19 and 8% higher than ResNet-50. These results illustrate the benefits of the selective long range connectivity and high architectural diversity found in COVID-Net, which enables strong representational capacity that is tailored for the task as well as making it easier to train. Therefore, these results demonstrate the benefits of the different design choices made during the machine-driven design exploration stage of the human-machine collaborative design strategy employed to create COVID-Net.

Qualitative analysis. As mentioned earlier, we performed an audit on the proposed COVID-Net to gain better insights into how COVID-Net makes decisions, and validate whether it is making detection decisions based on relevant information rather than erroneous information that bias decisions based on irrelevant visual indicators. The critical factors identified by GSInquire⁴⁴ in several example CXR images of COVID-19 cases are shown in Fig. 7. It can be observed that, based on the interpretation produced by GSInquire, the proposed COVID-Net primarily leverages areas in the lungs in the CXR images as the main critical factors in determining whether a CXR image is of a patient with a SARS-CoV-2 viral infection, as shown in red in Fig. 7. As such, we were able to validate that COVID-Net was not relying on improper information to make decisions (e.g., erroneous visual indicators outside the body, embedded markup symbols, imaging artifacts, etc.), which could lead to scenarios where the right decisions are made for the wrong reasons. Such ‘right decision, wrong reason’ scenarios are very difficult to track and identify without the use of such an explainability-driven auditing strategy, and thus highlight the value of explainability in improving the reliability of deep neural networks for clinical applications.

We further explored the relationship between the critical factors identified by GSInquire and the experimental results, particularly in relationship to clinical visual indicators leveraged for radiography examination of chest CXR images for COVID-19. In the cases in the experimental results where the proposed COVID-Net correctly detects COVID-19 infections (see Fig. 6), the critical factors that are identified by GSInquire often corresponds to clinical visual factors such as ground-glass opacities, bilateral abnormalities, and interstitial abnormalities, which have been found to be useful for COVID-19 radiography examinations^{6–8}. What this illustrates is that the proposed COVID-Net can learn to leverage visual indicators in its decision-making process that relate to those leveraged by clinicians in their decision-making process. Therefore, these observations are insightful for

establishing a deeper understanding into the overlaps between the decision-making process of deep neural networks for clinical applications and the decision-making process of clinicians during radiography examination.

In addition to performance validation for more responsible and transparent design, the ability to interpret and gain insights into how the proposed COVID-Net detects COVID-19 infections is also important for a number of other reasons:

- **Transparency** By understanding the critical factors being leveraged in COVID-19 case detection, the predictions made by the proposed COVID-Net become more transparent and trustworthy for clinicians to leverage during their screening process to aid them in making faster yet accurate assessments.
- **New insight discovery** The critical factors leveraged by the proposed COVID-Net could potentially help clinicians discover new insights into the key visual indicators associated with SARS-CoV-2 viral infection, which they can then leverage to improve screening accuracy.

Conclusion

In this study, we introduced COVID-Net, a deep convolutional neural network design for the detection of COVID-19 cases from CXR images that is open source and available to the general public. We also introduce COVIDx, an open access benchmark dataset that is comprised of 13,975 CXR images across 13,870 patient cases from five open access data repositories. Moreover, we investigated how COVID-Net makes predictions using an explainability method in an attempt to gain deeper insights into critical factors associated with COVID cases, which can aid clinicians in improved screening as well as improve trust and transparency when leveraging COVID-Net for accelerated computer-aided screening.

By no means a production-ready solution, the hope is that the promising results achieved by COVID-Net on the COVIDx test dataset, along with the fact that it is available in open source format alongside the description on constructing the open source dataset, will lead it to be leveraged and build upon by both researchers and citizen data scientists alike to accelerate the development of highly accurate yet practical deep learning solutions for detecting COVID-19 cases from CXR images and accelerate treatment of those who need it the most. Future directions include continuing to improve sensitivity and PPV to COVID-19 infections as new data is collected, as well as extend the proposed COVID-Net to risk stratification for survival analysis, predicting risk status of patients, and predicting hospitalization duration which would be useful for triaging, patient population management, and individualized care planning.

Data availability

The model, data and scripts are all available at <https://github.com/lindawangg/COVID-Net>.

Received: 20 May 2020; Accepted: 26 October 2020

Published online: 11 November 2020

References

1. Wang, W. *et al.* Detection of SARS-CoV-2 in different types of clinical specimens. *JAMA* **323**(18), 1843–1844 (2020).
2. West, C. P., Montori, V. M., & Sampathkumar, P. Covid-19 testing: the threat of false-negative results. In *Mayo Clinic Proceedings* (2020).
3. Fang, Y. *et al.* Sensitivity of chest CT for covid-19: comparison to RT-PCR. *Radiology*. <https://doi.org/10.1148/radiol.20200432> (2020).
4. Yang, Y. *et al.* Evaluating the accuracy of different respiratory specimens in the laboratory diagnosis and monitoring the viral shedding of 2019-ncov infections. *medRxiv:2020.02.11.20021493* (2020).
5. Wikramaratna, P., Paton, R. S., Ghafari, M., & Lourenco, J. Estimating false-negative detection rate of sars-cov-2 by rt-pcr. *medRxiv:2020.04.05.20053355* (2020).
6. Ng, M.-Y. *et al.* Imaging profile of the COVID-19 infection: radiologic findings and literature review. *Radiol. Cardiothorac. Imaging* **2**(1), e200034 (2020).
7. Huang, C. *et al.* Clinical features of patients infected with 2019 Novel Coronavirus in Wuhan China. *The Lancet* **395**, 497–506 (2020).
8. Guan, W. J., Hu Y., & Ni Z. Y. Clinical characteristics of Coronavirus disease 2019 in China. *N. Engl. J. Med.* **382**(18), 1708–1720 (2020).
9. Ai, T. *et al.* Correlation of chest CT and RT-PCR testing in coronavirus disease 2019 (COVID-19) in china: a report of 1014 cases. *Radiology* (2020). <https://doi.org/10.1148/radiol.20200642>.
10. Rubin, G. D. *et al.* The role of chest imaging in patient management during the COVID-19 pandemic: a multinational consensus statement from the fleischner society. *Radiology* (2020). <https://doi.org/10.1016/j.chest.2020.04.003>.
11. Nair, A. *et al.* A British Society of thoracic imaging statement: considerations in designing local imaging diagnostic algorithms for the COVID-19 pandemic. *Clin. Radiol.* **75**(5), 329–334 (2020).
12. Dennie, C. *et al.* The Canadian Society of Thoracic Radiology (CSTR) and Canadian Association of Radiologists (CAR) consensus statement regarding chest imaging in suspected and confirmed COVID-19. *Can. Assoc. Radiol. J.* **71**(4), 470–481 (2020).
13. Jacobi, A., Chung, M., Bernheim, A. & Eber, C. Portable chest x-ray in coronavirus disease-19 (covid-19): a pictorial review. *Clin. Imaging* <https://doi.org/10.1016/j.clinimag.2020.04.001> (2020).
14. Wu, G., & Li, X. Mobile x-rays are highly valuable for critically ill covid patients. *Eur. Radiol.* <https://doi.org/10.1007/s00330-020-06918-2> (2020).
15. Mao, B. *et al.* Assessing risk factors for SARS-CoV-2 infection in patients presenting with symptoms in Shanghai, China: a multicentre observational cohort study. *Lancet Dig. Health.* [https://doi.org/10.1016/S2589-7500\(20\)30109-6](https://doi.org/10.1016/S2589-7500(20)30109-6) (2020).
16. Cohen, J. P., Morrison, P., & Dao, L. COVID-19 image data collection. *arXiv:2003.11597* (2020).
17. Chung, A. Figure 1 COVID-19 chest x-ray data initiative. <https://github.com/agchung/Figure1-COVID-chestxray-dataset> (2020).
18. Chung, A. Actualmed COVID-19 chest x-ray data initiative. <https://github.com/agchung/Actualmed-COVID-chestxray-dataset> (2020).
19. Radiological Society of North America. COVID-19 radiography database. <https://www.kaggle.com/tawsifurrahman/covid19-radiography-database> (2019).

20. Radiological Society of North America. RSNA pneumonia detection challenge. <https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/data> (2019).
21. LeCun, Y., Bengio, Y., & Hinton, G. Deep learning. *Nature* **521**, 436–444 (2015).
22. Gozes, O. *et al.* Rapid AI development cycle for the coronavirus (COVID-19) pandemic: initial results for automated detection and patient monitoring using deep learning ct image analysis. [arXiv:2003.05037](https://arxiv.org/abs/2003.05037) (2020).
23. Xu, X. *et al.* Deep learning system to screen coronavirus disease 2019 pneumonia. [arXiv:2002.09334](https://arxiv.org/abs/2002.09334) (2020).
24. Li, L. *et al.* Artificial intelligence distinguishes COVID-19 from community acquired pneumonia on chest CT. *Radiology*. <https://doi.org/10.1148/radiol.20200905> (2020).
25. Shi, F. *et al.* Large-scale screening of covid-19 from community acquired pneumonia using infection size-aware classification. [arXiv:2003.09860](https://arxiv.org/abs/2003.09860) (2020).
26. Li, X., Li, C., & Zhu, D. Covid-mobilexpert: On-device covid-19 screening using snapshots of chest x-ray. [arXiv:2004.03042](https://arxiv.org/abs/2004.03042) (2020).
27. Minaee, S., Kafieh, R., Sonka, M., Yazdani, S., & Soufi, G. J. Deep-covid: predicting covid-19 from chest x-ray images using deep transfer learning. [arXiv:2004.09363](https://arxiv.org/abs/2004.09363) (2020).
28. Afshar, P. *et al.* Covid-caps: a capsule network-based framework for identification of covid-19 cases from x-ray images. [arXiv:2004.02696](https://arxiv.org/abs/2004.02696) (2020).
29. Luz, E. *et al.* Towards an effective and efficient deep learning model for covid-19 patterns detection in x-ray images. [arXiv:2004.05717](https://arxiv.org/abs/2004.05717) (2020).
30. Khobahi, S., Agarwal, C., & Soltanian, M. Coronet: a deep network architecture for semi-supervised task-based identification of covid-19 from chest x-ray images. *medRxiv*. <https://doi.org/10.1101/2020.04.14.20065722> (2020).
31. Ucar, F., & Korkmaz, D. COVIDDiagnosis-net: deep bayes-squeezenet based diagnosis of the coronavirus disease 2019 (COVID-19) from x-ray images. *Med. Hypotheses*. <https://doi.org/10.1016/j.mehy.2020.109761> (2020).
32. Tartaglione, E., Barbano, C. A., Berzovini, C., Calandri, M., & Grangetto, M. Unveiling covid-19 from chest x-ray with deep learning: a hurdles race with small data. [arXiv:2004.05405](https://arxiv.org/abs/2004.05405) (2020).
33. Yeh, C.-F. *et al.* A cascaded learning strategy for robust covid-19 pneumonia chest x-ray screening. [arXiv:2004.12786](https://arxiv.org/abs/2004.12786) (2020).
34. Zhang, Y. *et al.* Covid-da: deep domain adaptation from typical pneumonia to covid-19. [arXiv:2005.01577](https://arxiv.org/abs/2005.01577) (2020).
35. Karim, M. R. *et al.* Deepcovidexplainer: explainable covid-19 predictions based on chest x-ray images (2020).
36. Apostolopoulos, I. D., & Mpesiana, T. A. Covid-19: automatic detection from x-ray images utilizing transfer learning with convolutional neural networks. *Phys. Eng. Sci. Med.* **43**, 635–640. <https://doi.org/10.1007/s13246-020-00865-4> (2020).
37. Farooq, M., & Hafeez, A. Covid-resnet: a deep learning framework for screening of covid19 from radiographs (2020).
38. Majeed, T., Rashid, R., Ali, D., & Asaad, A. Covid-19 detection using cnn transfer learning from x-ray images. *medRxiv* (2020).
39. Wang, X. *et al.* Chestx-ray8: hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 3462–3471 (2017).
40. He, K., Zhang, X., Ren, S., & Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition* 770–778 (2016).
41. Wong, A., Shafiee, M. J., Chwyl, B., & Li, F. Ferminets: learning generative machines to generate efficient neural networks via generative synthesis. *arXiv preprint* [arXiv:1809.05989](https://arxiv.org/abs/1809.05989) (2018).
42. Huang, G., Liu, Z., & Weinberger, K. Q. Densely connected convolutional networks. *arXiv preprint* [arXiv:1608.06993](https://arxiv.org/abs/1608.06993) (2016).
43. Deng, J. *et al.* Imagenet: a large-scale hierarchical image database. In *2009 IEEE Conference on Computer Vision and Pattern Recognition* 248–255 (IEEE, 2009).
44. Lin, Z. Q. *et al.* Do explanations reflect decisions? A machine-centric strategy to quantify the performance of explainability algorithms. *arXiv preprint* [arXiv:1910.07387](https://arxiv.org/abs/1910.07387) (2019).
45. Wong, A. Netscore: towards universal metrics for large-scale performance analysis of deep neural networks for practical usage. *CoRR*. *arXiv preprint* [arXiv:abs/1806.05512](https://arxiv.org/abs/1806.05512) (2018).
46. Simonyan, K., & Zisserman, A. Very deep convolutional networks for large-scale image recognition (2014).

Acknowledgements

We would like to thank Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Research Chairs program, DarwinAI Corp., and Audrey Chung.

Author contributions

L.W. and A.W. conceived the experiment, L.W., Z.L. and A.W. conducted the experiment, L.W. and A.W. analysed the results. All authors reviewed the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Correspondence and requests for materials should be addressed to L.W.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2020