
COVIDx CXR-3: A Large-Scale, Open-Source Benchmark Dataset of Chest X-ray Images for Computer-Aided COVID-19 Diagnostics

Maya Pavlova^{1,2}, Tia Tuinstra^{1,2}, Hossein Aboutaleb^{1,3,4},
Andy Zhao^{1,2}, Hayden Gunraj^{1,2}, Alexander Wong^{1,2,4,5}

¹Vision and Image Processing Lab, University of Waterloo

²Department of Systems Design Engineering, University of Waterloo

³Cheriton School of Computer Science, University of Waterloo

⁴Waterloo AI Institute, University of Waterloo

⁵DarwinAI Corp.

{mpavlova, ttuinstra, haboutal, andy.zhao,
hayden.gunraj, a28wong}@uwaterloo.ca

Abstract

After more than two years since the beginning of the COVID-19 pandemic, the pressure of this crisis continues to devastate globally. The use of chest X-ray (CXR) imaging as a complementary screening strategy to RT-PCR testing is not only prevailing but has greatly increased due to its routine clinical use for respiratory complaints. Thus far, many visual perception models have been proposed for COVID-19 screening based on CXR imaging. Nevertheless, the accuracy and the generalization capacity of these models are very much dependent on the diversity and the size of the dataset they were trained on. Motivated by this, we introduce COVIDx CXR-3, a large-scale benchmark dataset of CXR images for supporting COVID-19 computer vision research. COVIDx CXR-3 is composed of 30,386 CXR images from a multinational cohort of 17,026 patients from at least 51 countries, making it, to the best of our knowledge, the most extensive, most diverse COVID-19 CXR dataset in open access form. Here, we provide comprehensive details on the various aspects of the proposed dataset including patient demographics, imaging views, and infection types. The hope is that COVIDx CXR-3 can assist scientists in advancing machine learning research against both the COVID-19 pandemic and related diseases.

1 Introduction

As the COVID-19 pandemic continues to progress internationally, the push for effective and robust screening methods is imperative. One method that has seen significant increase in usage in the COVID-19 clinical workflow is chest X-ray (CXR) imaging, which offers complementary screening information to transcriptase-polymerase chain reaction (RT-PCR) testing with the benefits of more readily accessible and mobile healthcare equipment [1]. In addition, CXR imaging is routinely conducted in parallel to tedious viral testing for patients with respiratory complaint [2] to reduce patient volume and provide more clinical detail. As such, multiple works have investigated deep learning-driven computer-aided diagnostics to both detect positive SARS-CoV-2 cases from CXR images [3–5] as well as provide severity assessments that can be leveraged in patient triaging [4, 6, 7]. Nevertheless, as a result of the rapid progression of the pandemic and drive for results, such studies have been limited in terms of quantity and/or diversity of patients as well as standardization between datasets for higher public accessibility. In this work, we introduce COVIDx CXR-3, a large-scale




Figure 1: Example CXR images from COVIDx CXR-3 benchmark dataset: (A) normal or no pneumonia or SARS-CoV-2 finding, (B) pneumonia infection, and (C) SARS-CoV-2 positive finding.

Table 1: Distribution of CXR images and patient cases (in parentheses) for both training and evaluation datasets in COVIDx CXR-3, split by infection type. Note the total for patients is 10 less than the sum across the training set as some patients had both SARS-CoV-2 negative and positive CXR images from different studies.

Split	Infection Type			Total
	Normal	Pneumonia	COVID-19	
Training	8,437 (8,238)	5,555 (5,612)	15,994 (2,808)	29,986 (16,648)
Test	100 (100)	100 (100)	200 (178)	400 (378)
Total	8,537 (8,338)	5655 (5,712)	16,194 (2,986)	30,386 (17,026)

benchmark dataset of CXR images that has been carefully curated with expert radiologists and clinicians for computer-aided COVID-19 research comprising of 30,386 CXR images from a multi-institutional, multinational cohort of 17,026 patients across at least 51 countries. To the best of the authors’ knowledge, COVIDx CXR-3 is the most extensive, accessible COVID-19 CXR dataset in open-source form. Additionally, we examine the dataset diversity in regards to patient findings, demographics, and imaging view, and explore areas of potential bias and dataset imbalances.

2 Methods

The COVIDx CXR-3 benchmark dataset comprises carefully curated CXR images from a multi-institutional, multinational patient cohort, in collaboration with expert radiologists and clinicians. More specifically, COVIDx CXR-3 comprises patient cohorts from the following organizations and initiatives: (1) RSNA International COVID-19 Open Radiology Database (RICORD) [8], (2) Stony Brook University (COVID-19-NY-SBU) [9], (3) Valencian Region Medical Image Bank (BIMCV) [10], (4) RSNA Pneumonia Detection Challenge dataset [11], (5) COVID-19 Image Data Collection [12], (6) Fig.1 COVID-19 CXR Dataset Initiative [13], (7) ActualMed COVID-19 CXR Dataset Initiative [14], and (8) COVID-19 Radiography Database [15]. Expert feedback was used to evaluate the quality and applicability of the images from each source, leading to removal of some images from the listed sources. A test set that has no patient overlap with the training set was created from randomly sampled patients from the Radiological Society of North America (RSNA) RICORD [8] and Pneumonia Detection [11] initiatives as a result of their high image and annotation quality. Each collected patient case is documented with one of the following findings: (A) normal or no pneumonia and no SARS-CoV-2 infection, (B) non-SARS-CoV-2 pneumonia infection, and (C) positive SARS-CoV-2 (COVID-19) infection. Example images from COVIDx CXR-3 for each infection case are shown in Figure 1.

3 Results and Discussion

Table 1 shows the distribution of CXR images and patients in regard to infection type. The final COVIDx CXR-3 benchmark dataset is composed of 30,386 CXR images from 17,026 unique patients, of which 16,194 CXR images (53.3%) and 2,986 patients (17.5%) come from SARS-CoV-2 positive cases, creating a relatively balanced training and testing set for SARS-CoV-2 positive and negative detection in terms of image count. The smaller balanced test set is the result of sampling an 80%/20%

Table 2: Summary of demographic and imaging protocol variables for the COVIDx CXR-3 benchmark dataset. Age and sex statistics are expressed on a patient level, while imaging view and finding statistics are expressed on an image level.

Age		Sex	
<18	792 (4.7%)	Male	9,050 (53.2%)
[18, 59]	11,275 (66.2%)	Female	6,830 (40.1%)
(59, 74]	3,111 (18.3%)	Male	9,050 (53.2%)
(74, 90]	702 (4.1%)	Unknown	1,146 (6.7%)
>90	7 (0.04%)	Imaging view	
Unknown	1,139 (6.7%)	PA	8,553 (28.1%)
		AP	20,372 (67.0%)
		Unknown	1,461 (4.8%)




Figure 2: Distribution of demographic variables and imaging views in COVIDx CXR-3 for SARS-CoV-2 positive patient cases and CXR images respectively.

patient split from the RSNA RICORD [8] initiative for positive SARS-CoV-2 cases to ensure networks are evaluated against expertly annotated positive samples.

The distribution of demographic patient data for age and sex as well as for CXR imaging view is shown across all infection types in Table 2 and for SARS-CoV-2 positive cases only in Figure 2. The dataset does not indicate a bias for youth or senior vulnerable age groups with the majority of patients (66.2%) falling within the range of 18 to 59. Moreover, in terms of patients with documented sex, there is a relatively similar fraction of male (53.2%/37.7%) and female (40.1%/27.5%) patients across all infections as well as across SARS-CoV-2 positive cases respectively. In terms of imaging view, a majority of the CXR images are captured in anterior-posterior (AP) (67%) view as opposed to posterior-anterior (PA) (28.1%) view.

Systems trained on the COVIDx CXR-3 benchmark dataset have the opportunity to detect each individual documented infection class or group normal and pneumonia cases together to form a negative SARS-CoV-2 label for targeted SARS-CoV-2 screening. The hope is that the release of COVIDx CXR-3 in an open-source manner will help encourage researchers, clinical scientists, and citizen data scientists to accelerate advancements and innovations in the fight against the pandemic.

Finally, to serve as a baseline reference for comparison purposes, Table 3 provides benchmark test performance for three different deep neural network architecture designs on the COVIDx CXR-3 test dataset, as conducted in [4, 5]. We hope that COVIDx CXR-3 can assist scientists in advancing research against both the COVID-19 pandemic and related diseases.

Table 3: Sensitivity, positive predictive value (PPV), and accuracy of benchmark networks on the COVIDx CXR-3 test dataset [4, 5].

Architecture	Sensitivity (%)	PPV (%)	Accuracy (%)
DenseNet201 [16]	82.9	88.9	90.3
ResNet-50 [17]	88.5	92.2	90.5
COVID-Net [3]	93.5	100	94.0
COVID-Net CXR-2 [5]	95.5	97.0	96.3
MEDUSA [4]	97.5	99.0	98.3

Potential Negative Societal Impact

While the aim with this large-scale benchmark dataset release is to support researchers, clinicians, and citizen data scientists in advancing research, one negative societal impact that can potential arise from this release the misuse of the collected data. For example, the benchmark dataset may be utilized by insurance companies to construct machine learning algorithms for forecasting future medical expenses, and as such lead to increased insurance premiums by insurance companies that could be inappropriate for a given patient given imperfect predictive analytics.

Acknowledgements

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

References

- [1] Adam Jacobi, Michael Chung, Adam Bernheim, and Corey Eber. Portable chest X-ray in coronavirus disease-19 (COVID-19): A pictorial review. *Clin. Imaging*, 2020.
- [2] A. Nair, J. Rodrigues, S. Hare, A. Edey, A. Devaraj, J. Jacob, A. Johnstone, R. McStay, Erika Denton, and G. Robinson. A British Society of Thoracic Imaging statement: considerations in designing local imaging diagnostic algorithms for the COVID-19 pandemic. *Clin. Radiol.*, 2020.
- [3] Linda Wang, Zhong Qiu Lin, and Alexander Wong. COVID-Net: A tailored deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images. *Scientific Reports*, 2020.
- [4] Hossein Aboutaleb, Maya Pavlova, Hayden Gunraj, Mohammad Javad Shafiee, Ali Sabri, Amer Alaref, and Alexander Wong. MEDUSA: Multi-scale encoder-decoder self-attention deep neural network architecture for medical image analysis. *Frontiers in Medicine*, 8:821120, 2022.
- [5] Maya Pavlova, Naomi Terhlan, Audrey G. Chung, Andy Zhao, Siddharth Surana, Hossein Aboutaleb, Hayden Gunraj, Ali Sabri, Amer Alaref, and Alexander Wong. COVID-Net CXR-2: An enhanced deep convolutional neural network design for detection of COVID-19 cases from chest X-ray images. *Frontiers in Medicine*, 9:861680, 2022.
- [6] Alexander Wong, Zhong Qiu Lin, Linda Wang, Audrey G. Chung, Beiyi Shen, Almas Abbasi, Mahsa Hoshmand-Kochi, and Timothy Q. Duong. COVIDNet-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, 2020.
- [7] Hossein Aboutaleb, Maya Pavlova, Mohammad Javad Shafiee, Ali Sabri, Amer Alaref, and Alexander Wong. COVID-Net CXR-S: Deep convolutional neural network for severity assessment of COVID-19 cases from chest X-ray images, 2021.
- [8] E. Tsai, S. Simpson, et al. The RSNA international COVID-19 open annotated radiology database (RICORD). *Radiology*, page 203957, 2021.

- [9] Joel Saltz, Mary Saltz, Prateek Prasanna, Richard Moffitt, Janos Hajagos, Erich Bremer, Joseph Balsamo, and Tahsin Kurc. Stony Brook University COVID-19 positive cases [data set]. *The Cancer Imaging Archive*, 2021. doi: 10.7937/TCIA.BBAG-2923.
- [10] Maria de la Iglesia Vayá, Jose Manuel Saborit-Torres, Joaquim Angel Montell Serrano, Elena Oliver-Garcia, Antonio Pertusa, Aurelia Bustos, Miguel Cazorla, Joaquin Galant, Xavier Barber, Domingo Orozco-Beltrán, Francisco García-García, Marisa Caparrós, Germán González, and Jose María Salinas. BIMCV COVID-19+: a large annotated dataset of RX and CT images from COVID-19 patients, 2021. URL <https://dx.doi.org/10.21227/w3aw-rv39>.
- [11] Radiological Society of North America. RSNA pneumonia detection challenge. <https://www.kaggle.com/c/rsna-pneumonia-detection-challenge/data>, 2019.
- [12] Joseph Paul Cohen, Paul Morrison, and Lan Dao. COVID-19 image data collection. *arXiv 2003.11597*, 2020. URL <https://github.com/ieee8023/covid-chestxray-dataset>.
- [13] Audrey Chung. Figure 1 COVID-19 chest X-ray data initiative. <https://github.com/agchung/Figure1-COVID-chestxray-dataset>, 2020.
- [14] Audrey Chung. Actualmed COVID-19 chest X-ray data initiative. <https://github.com/agchung/Actualmed-COVID-chestxray-dataset>, 2020.
- [15] Radiological Society of North America. COVID-19 radiography database. <https://www.kaggle.com/tawsifurrahman/covid19-radiography-database>, 2019.
- [16] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q. Weinberger. Densely connected convolutional networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2261–2269, 2017. doi: 10.1109/CVPR.2017.243.
- [17] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *Computer Vision - ECCV 2016*, pages 630–645, Cham, 2016. Springer International Publishing. ISBN 978-3-319-46493-0.