SEGMENTATION MODEL OF THE OPACITY REGIONS IN THE CHEST X-RAYS OF THE COVID-19 PATIENTS IN THE US RURAL AREAS AND THE APPLICATION TO THE DISEASE SEVERITY

Haiming Tang
Mercury Data Science
Houston, TX 77098
ningzhithm@gmail.com

Nanfei Sun
Department of Management Information System
University of Houston Clear Lake
Houston, TX 77058
sun@uhcl.edu

Yi Li
Fox Chase Cancer Center
Philadelphia, PA 19111
YiLiMD@gmail.com

October 20, 2020

ABSTRACT

The pandemic of Covid-19 has caused tremendous losses to lives and economy in the entire world. Up until October 2020, it has caused more than 38 million infections and 1.1 million deaths. This has created a severe burden for the health care system worldwide. The machine learning models have been applied to the radiological images of the Covid-19 positive patients for disease prediction and severity assessment. However, a segmentation model for detecting the opacity regions like haziness, ground-glass opacity and lung consolidation from the Covid-19 positive chest X-rays is still lacking. The recently published dataset of a collection of radiological images for a rural population in United States had made development of such a model a possibility due to the high quality of the radiological images and the consistency in clinical measurements. We manually annotated 221 chest X-ray images with lung fields and opacity regions and trained a segmentation model for the opacity region. The model has a good performance in regarding the overlap between predicted and manually labelled opacity regions for both the testing data set and the validation dataset from very different sources. In addition, the percentage of the opacity region over the area of the total lung fields shows a good predictive power for the patient severity. In view of the above, our model is a successful first try in developing a segmentation model for the opacity regions for the Covid-19 positive chest X-rays. However, careful manual examinations of the model predictions by experienced radiologists show mistakenly predicted opacity regions caused probably by the anatomical complexities. Thus, additional work is needed before a robust and accurate model can be developed for the ultimate goal of implementation in the clinical setting. The model, manual segmentation and other supporting materials can be found in https://github.com/haimingt/opacity_segmentation_covid_chest_X_ray.

1 Introduction

Ever since the pandemic of Covid-19 in late 2019 till October 2020, there have been 38.8 million cases infected with this virus in the entire world, causing 1.1 million deaths. In the United States alone, 21.8 thousand patients had died of the infection and its complications [1]. The pandemic of Covid-19 has caused enormous losses to the lives and economy. Accurate and fast scanning of the lung situations for assessment of the disease progression has become a global-wise ubiquitous requirement. During the peak period, more than 70,000 new cases were diagnosed almost every day in the United States, 7-10% among them were hospitalized, and 2% of the diagnosed population have to be sent to emergent departments who needed further assessment for ICU admissions. The clinical environment in the era of Covid-19 pandemic has become very challenging with high-pressure and stress.

Clinicians and researchers have been devoting tremendous amounts of work for this unprecedented challenge. The field of diagnostic radiology has been a crucial area in clinical management of this disease [2]. Clinicians have summarized the imaging characteristics of the Covid-19 infections for both the chest CT images and chest X rays[3, 4], including

ground glass opacities, consolidations, haziness and many others. Computer assisted segmentation and classification models are quickly studied[5, 6, 7] to evaluate their abilities to assist physicians and radiologists in assessing the radiological results.

Immediately after Hintington et. al. published his most substantial breakthroughs[8] of deep learning algorithms in 2006, they have been introduced to the computer vision field by many computer vision researchers[9, 10, 11, 12]. Besides the huge success of deep learning models in the computer vision field for common everyday objects, many researchers introduced deep learning into the medical imaging field[13, 14, 15], especially for the images of the diagnostic radiology field. Traditional deep learning neural networks including RNN[16], CNN[17, 18], resNet[19] and quite a few noval schemas, like FCN, Unet, SegNet, Unet-3D and Mask-RCNN[20] were widely used in segmentation studies in X-Ra[21], PET/CT[22, 23], MRI[24], and Ultrasound[25].

Researchers have developed machine-learning models for Covid-19 related imaging studies. One category of these models is the classification model of lungs infected with the Covid-19 vs. those of other diseases using either the chest CT images or chest X-rays[26, 27]. Some researchers have claimed very good results of using machine learning algorithms in differentiating Covid-19 radiological characteristics, and some even claimed better model results than junior radiologists, but just slightly worse than experienced radiologists[26]. There have also been projects on predicting the severity of the Covid-19 pneumonia[28]. For example, a research group manually curated a severity score for 96 chest X ray images[29], and used features from a model previously trained on a large data set of the non-Covid cases to develop a linear regression model to predict the severity of the opacities in the chest X-ray. Another research group came up with 6 severity scores for each chest X-ray image, separately for the upper, middle and lower regions of the left and right lung regions[30].

As far as we know about the most up to date publications, a segmentation model that can detect the location of the abnormalities for the Covid-19 positive chest X ray images is still lacking, while work has been done to successfully segment the lung regions in Covid-19 cases. However, a recently published data set has made the development of such a segmentation model a possibility.

Recently, a research group made publicly available a data set of the 115 Covid-positive patients in the rural areas of the United States[31, 32]. The data set includes the chest X-ray images of patients, chest CT images, as well as the patient's basic measurements and clinical prognosis like ICU admissions or mortality. In contrast to most of the machine-learning projects on Covid-19, which were developed from the images collected sparsely from the online resources or various publications, this data set is more consistent in image qualities and clinical data accuracy.

In this data set, the radiological data represents a population of patients who were diagnosed with covid-19 in a specific region. Most of the images are chest X-rays and only a small portion of patients ever received a CT scan. Among these Chest X-ray images, as high as 30% have clear lungs without any abnormalities. Thus, this data set avoids the problem of over-representation of the more severe cases, which could be assembled from many different areas of the world. Using such a dataset helps develop a more robust model that covers a general population.

With the availability of this higher quality data, we aim to develop a model of segmentation for abnormal areas in chest X-ray images of the Covid-19 patients. More specifically, we want to locate the opacity regions that are characteristic of the Covid-19 infections, including haziness, ground glass opacities and consolidations. We want to answer two specific questions in this study: the first question is that is the development a model to point out the areas of the opacities for Covid-19 positive chest X-ray images possible? And other other question is that does this info of the opacities have a clinical significance in determining the severity of the patients?

2 Data and methods

2.1 Primary Data source

Recently, a collection of radiographic and CT images studies, together with patient demographics, comorbidities, prognosis and key radiology findings, as of July 2020, was published for a rural COVID-19 Positive population in the southern United States[31]. The collection was downloaded from the Cancer Imaging Archive (TCIA) Public Access[32], containing 256 studies for a total number of 105 patients. Due to the focus of this study, we filtered out CT images and also X-rays that were taken in positions other than AP or PA. In total, we have included 221 chest X-ray images from 105 patients in our study. We randomly divided the entire dataset to the training and testing data sets by a ratio of 7:3 (154 images for the training set, 67 images for the testing set), but ensured that the images of the same patient were divided to the same subset.

2.2 Validation dataset

Cohen et al. published a chest X-ray dataset together with the severity scores[29]. The images in this collection were from various sources and serve as a good outside validation dataset. Due to the time consuming process of manual curations, we randomly selected a smaller dataset of 25 images from the original 96 images as our validation dataset.

2.3 Manual curations

Manual curations of the lung regions and the opacity regions were performed for all the images in the above training, testing and validation data sets. For simplicity considerations as well as the limitation of the resolution of the X-ray images, opacity regions that include the ground glass opacities, haziness and consolidations were not differentiated. Lableme[33] was used for the curation work, with regions represented by polygons of connected dots. The curations of the lung and opacity regions were saved in the Json format. Manual curations were performed by a junior physician and were reviewed by an experienced radiologist.

2.4 Image preprocessing and augmentation

The input X-ray image is then cropped to only keep the lung regions by mapping the original image with the lung contour segmentation. The model output is an image mask that has values 1 for manually curated opacity regions and 0 for all other regions.

Although the majority of our data comes from the same data source, the images were not of the same size. For easier deployment of our model, all training and testing images were resized to 320 by 320 pixels. And all images were preprocessed through a data normalization function.

Data augmentation is a powerful technique to increase the amount of data and prevent model overfitting. Since our dataset is very small we applied a large number of different augmentations: horizontal flip, affine transforms, perspective transform, brightness/contrast/colors manipulations, image blurring and sharpening, gaussian noise and random crops. All of these transformations were performed using Albumentations, a fast augmentation library[34].

Sample images of the model input and output were included in Figure 1.

2.5 Model architecture and training

"Segmentation models" is a python library with neural networks for image segmentation based on Keras and Tensorflow[35]. It provides a high level API for segmentation model architectures and neural network backbones. For this study, we have chosen U-Net, the most popular and highly cited architecture which was developed specifically for biomedical image segmentation[36]. We have chosen the neural network backbone of ResNet18[37] for its wide usage, relatively smaller number of parameters and proven high performance. Other backbones, including MobileNet, SENet154, Densenet, EfficientNet are potential choices for further explorations. Segmentation loss was chosen to be Dice Loss[38] plus Focal Loss[39]. Dice Loss is essentially a measure of overlap between two samples, while Focal Loss is a measurement of how far off a prediction is from the truth.

2.6 Model metrics

We have chosen two popular metrics, IoU (Intersection over Union) and F1-score, to evaluate our model performance[40]. IoU is calculated using the area of overlap over the area of union between the input and prediction. F1-score is 2 * the Area of Overlap divided by the total number of pixels in both images. The model metrics were evaluated on the training dataset, the testing dataset, the training and testing dataset combined (the whole dataset), the validation dataset, the images without opacity regions in the whole dataset as well as the images with opacity regions in the whole dataset.

2.7 Correlation with the disease severity

We calculated the percentage of opacity region by simply dividing the size of both lung fields by the size of the opacity regions. We then used this metric as the opacity percentage for predicting the severity of the patients, which was extracted from the excel table of the primary data source. The prognosis of each patient was among 3 categories: recovered with no ICU admission; recovered after ICU admission or deceased after ICU admission. For this data source, all patients who were deceased had gone to the ICU. The receiver operating curve (ROC) for the ICU admission was plotted using the opacity percentage against whether a patient was admitted to the ICU. The receiver operating curve (ROC) for the mortality was plotted using the opacity percentage against whether a patient was deceased.

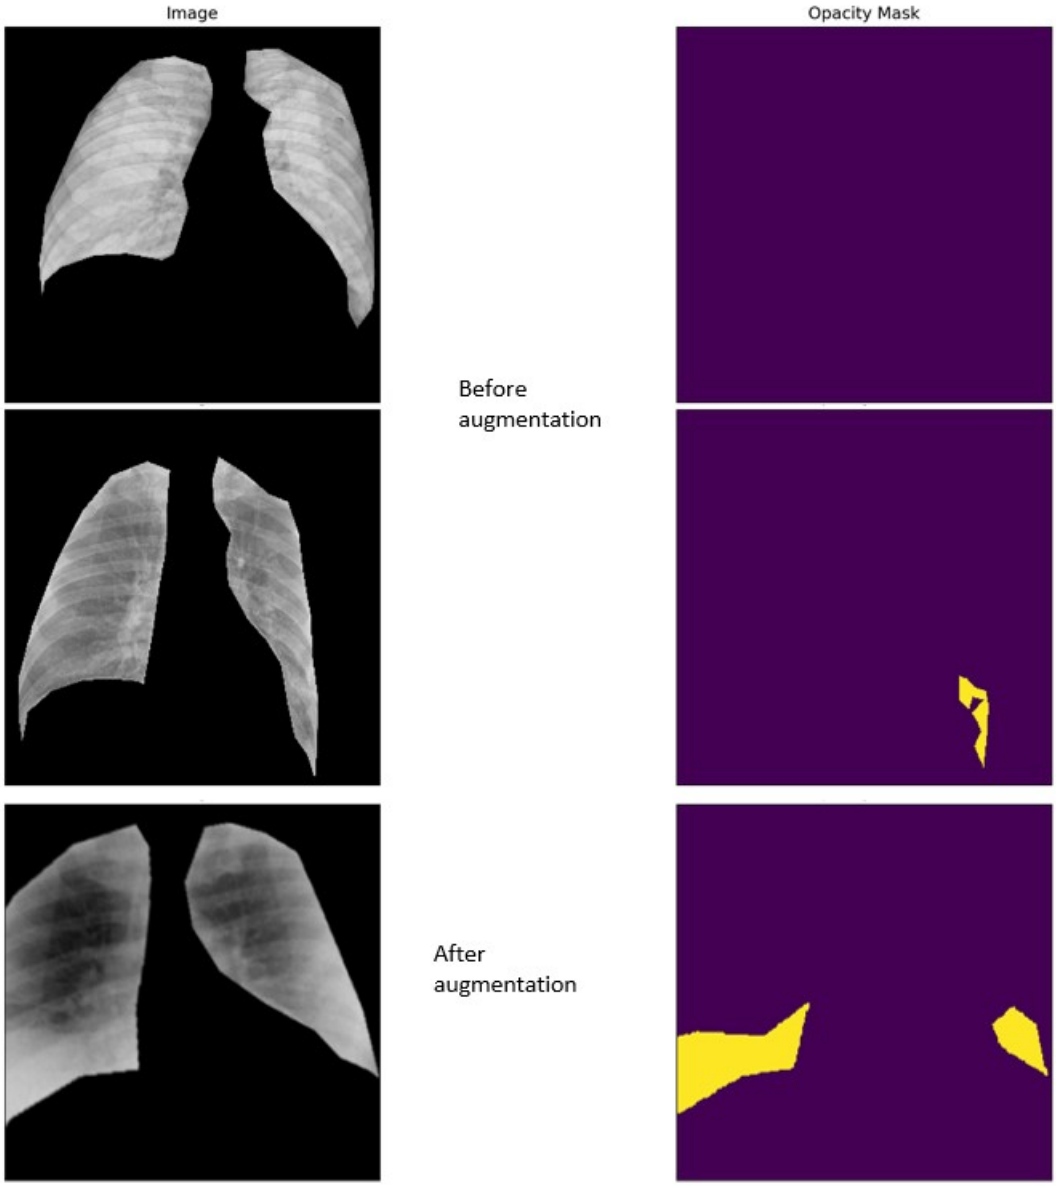

Figure 1: Sample images of the model input and output before and after augmentation.

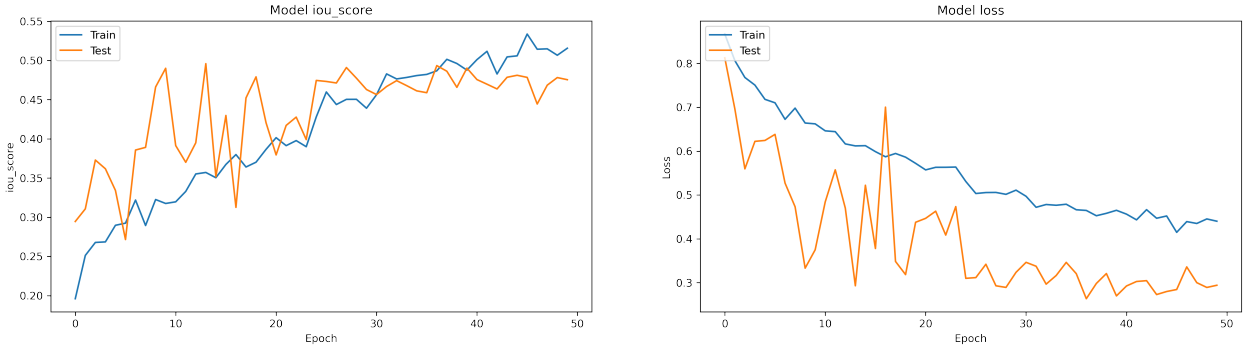

Figure 2: Model metric in the training process.

We also downloaded a recently published model [29] for predicting the severity of the chest X-ray of the Covid-19 positive patients for a comparison. ROC curves for ICU admission as well as mortality were calculated in a similar way using the output opacity score of that model.

3 Results

3.1 Metrics of the trained model

The model metrics during the training process can be found in Figure 2. We observed a steady increase of the IoU score for the training dataset and a steady decrease of the model loss with increased epochs. The IoU score and the model loss for the testing dataset showed significant fluctuations during epochs 10 to 25, with a gradual return to a more steady state during epochs 30 to 50. In general, the trend of the metrics for both the training and the testing data sets is consistent, showing a gradual improvement in the training process.

At the end of the training, the IoU score for the training dataset is 0.5157, while the IoU score for the testing dataset is 0.4755 and the IoU score the validation dataset is 0.6724. An IoU score > 0.5 is normally considered a "good" prediction[40]. But as our measurement is very close to this threshold, thus, we had doubts about the actual performance of our model.

In the meantime, the validation dataset is from an alternative source. We usually expect the same or worse performance in the validation dataset as the model trained on a dataset with a limited variability may be lacking generalization capabilities to a dataset from an alternative source. The slightly more superior performance raises strong suspicions. These issues area covered in more detail in the discussion part.

3.2 Usage of the predicted opacities for predicting the disease severity

The ROCs of the opacity percentage on the mortality and the ICU admissions can be viewed in Figure 3.a and 3.b. We see a great predictability of the opacity percentage for the severity of the patients regarding mortality and the ICU admissions. This is consistent with our expectations, as the opacity region is the most intuitive measurement for assessing the patients by the physicians and radiologists in the clinical setting. In addition, multiple papers have quantitatively measured the associations of the opacity region with the disease severity in the clinical settings[26, 41]. Thus, the results of our study re-iterate this relationship. Besides, the consistency between the performances of our model predictions and the manually curated results show a reasonably well-performing model.

3.3 Comparisons with the opacity score method

The Cohen group[29] developed a method to predict the severity of the chest X-ray of the Covid-19 infections. We ran this method on our data sets and applied the predicted opacity score for plotting the ROCs using the same patient severity data as above. The results can be viewed in Figure 4. We see a surprising consistency between the performances of the opacity score of Cohen's method and the opacity percentage of our model. The consistency is a confirmation of the good performance of our segmentation model. Considering the minimal simplicity of our metric, there could be a potential further development of a model for predicting the severity of the chest X-rays. However, this is not the scope of this study.

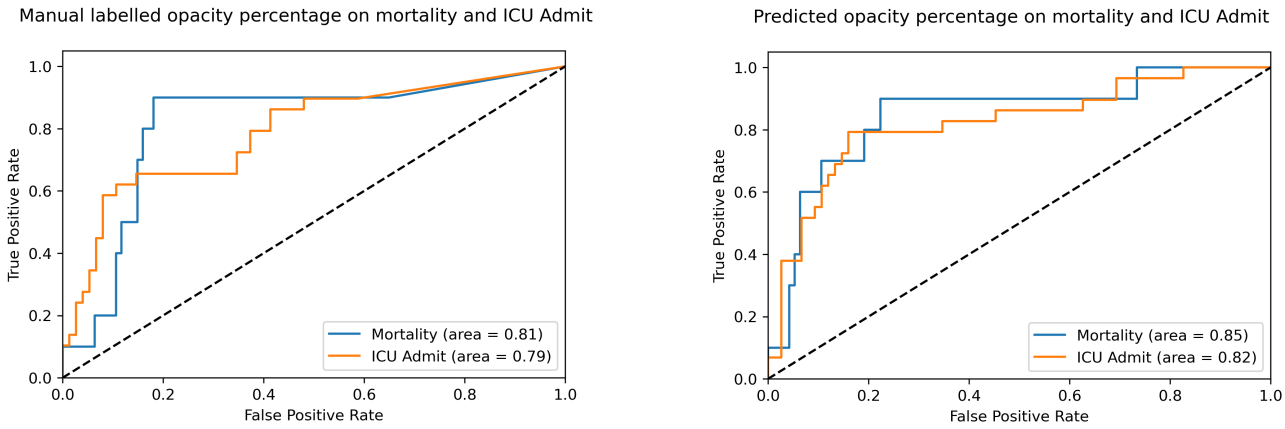

Figure 3: Performance of the opacity percentage in predicting mortality or ICU admission. 3.a, Opacity percentage calculated from manual labeling; 3.b, Opacity percentage calculated from model predictions.

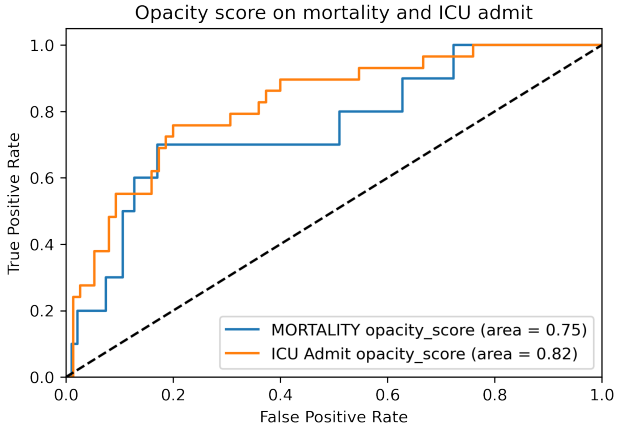

Figure 4: Performance of the opacity percentage in predicting mortality or ICU admission.

4 Discussions

4.1 The impact of the empty mask for IoU score

Due to the questions and suspicions we had for the IoU scores in the above, we further explored this issue and found a problem that significantly affected the IoU score. As defined by the formulas, the denominators are the union of the 2 regions for the IoU score. However, for our dataset, there are a significant number of images that do not have any opacity regions. These images have radiology reports of "clear lungs", and the model predicted regions are also usually none. Thus, by definition, the union of the manually labelled region and the predicted region is 0, and a meaningful IoU score can not be calculated.

To prove our hypothesis above, we extracted those images without any manually labelled opacity regions from the whole dataset combining the training and testing data sets. Then we re-applied the IoU score metrics to the new model created. It turned out that for the subset of 71 images without opacity regions, the calculated IoU score was 2.4E-10, close to 0. But for the rest 150 images with manually labelled opacity regions, the calculated IoU score increased to 0.6143 from the previous measurements of 0.5157 for the training dataset and 0.4755 for the testing set. Thus, the adjusted corrected IoU score for the entire dataset should be around (0.6143*150 + 1*71)/221 = 0.738, a much acceptable good performance for a segmentation model.

The superior performance of IoU score on the validation dataset, however, may come from the fact that most of the images in the validation dataset do contain manually labelled opacity regions. After the adjustment above, the IoU score of the validation dataset is slightly worse than that of the training and testing data sets.

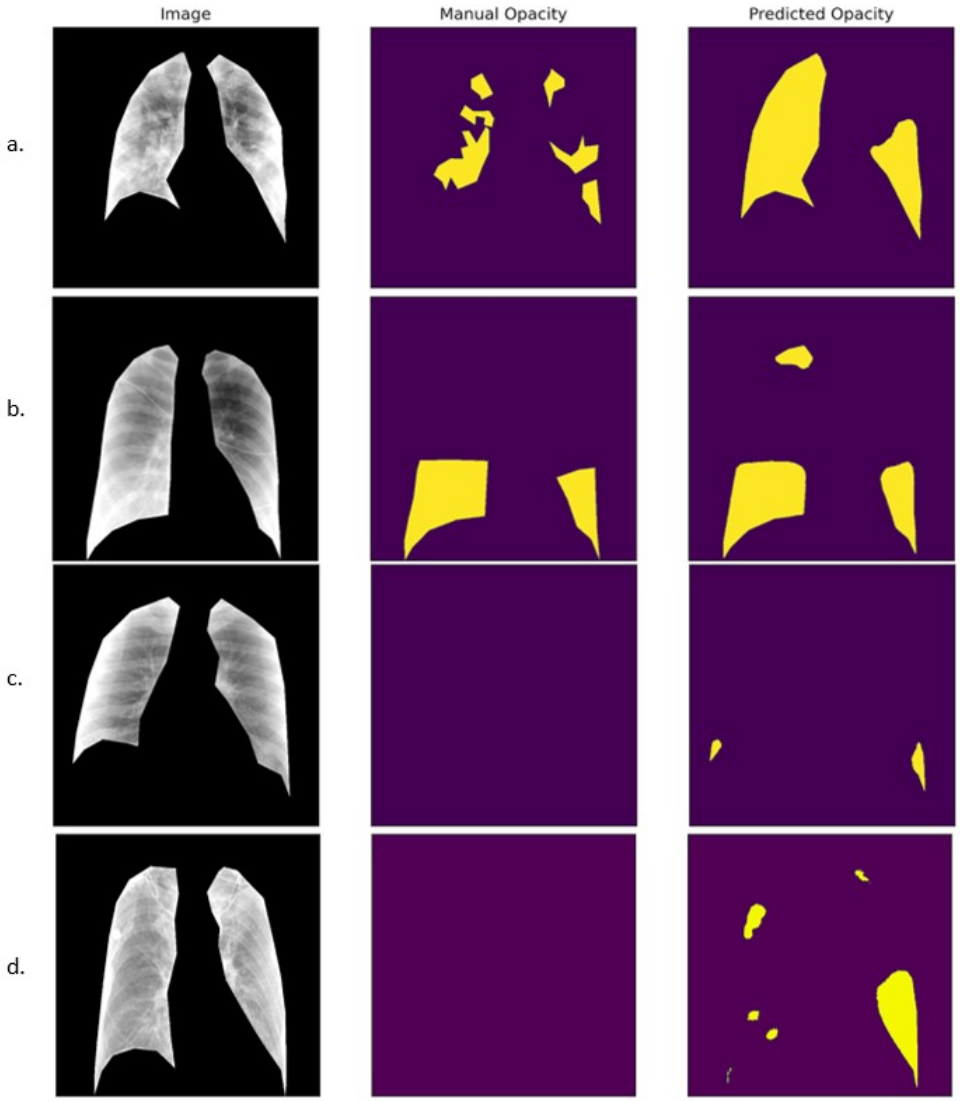

Figure 5: Representative samples that have discrepancies between manual labeling and predictions.

This discrepancy may come from the differences in these dataset: our model was trained using the images from patients in the rural US regions, while the images of the validation data set are from sources that are significantly different.

4.2 Imperfections of the models

We printed out the predicted opacity regions for each of the images in our dataset for a manual examination. The comparison images are available in the supplemental materials. The examinations revealed some discrepancies, from which we found several potential shortcomings of this segmentation model.

Figure 5 lists several representative samples that have discrepancies between manual labeling and predictions.

Firstly, the model flaw may come from the imperfections in the manual curations. Figure 5.a shows that the manual curated regions are very scattered, while the predicted regions cover larger and more consistent lung regions. This actually may represent the imperfections in the manual labeling process. This problem roots back to the ground truth for the opacity regions. This, however, is non-trivial. The boundary of the opacity region is very difficult to be clearly identified in the chest X-rays due to many issues, like image quality, patient anatomy, baseline disease etc. In the clinical setting, even experienced radiologists can identify a region of abnormality barely in a rough fashion. The clear delimitation of the abnormality boundaries are usually difficult, especially for lighter representations like ground glass opacity or haziness.

Secondly, some anatomical structures may compromise the model predictions. Figure 5.b shows that the chest X rays have opacities in both lung bases, however, the model prediction also contains a small region in the right lung apex. However, this is a wrong prediction probably caused by the rib shadow. Figure 5.c shows several tiny opacities in bilateral lung bases. However, the model may mistakenly predict the high density as an opacity. The high density is more likely due to the overlap of lung tissues instead of the ground glass opacities or consolidations caused by infections and exudates.

Thirdly, as shown in Figure 5.d, the model may mistake the markings of bronchial trees for opacities. The lungs in Figure 5.d are actually clear, but the markings of bronchial trees are prominent. They were likely to have been mistakenly predicted as opacities due to the higher densities.

5 Future improvements

As discussed above, the model predictions may be defective in a series of conditions. The segmentation of the opacity area in the chest X ray is a challenging work, even for experienced radiologists. The clear delimitation of the opacities is not as easy as delimitation of a solid mass with clear margins like tumors.

In general, the model performance is acceptable. Model metrics of IoU score and F1 score show good overlap between the manually labelled area and the predicted regions. From our examinations, correct locations can be detected more than 80% of the time, however, the exact boundaries may need further improvement.

To eliminate the influences of the anatomical structures or external instruments like ribs, scapula, wires and pacemakers, we probably need to include more chest X-ray images to more carefully label opacities. Another future direction is to develop a more robust model that takes into account all of these confounding elements. A possibility could be a multi-category segmentation model that detects the different anatomical structures separately.

6 Conclusion

In this study, we have demonstrated the development of a segmentation model for opacity regions for chest X-rays of the Covid-19 positive patients. The model performance is generally good with a precise predicted location for the opacities. Besides, the model shows a good generalization capacity when applied to a validation dataset composed of images of different sources. In addition, the percentage of the predicted opacity regions over the total lung area can predict the patient severity well regarding ICU admission and mortality. The performance of patient severity prediction is comparable or slightly better than the previously published "opacity score" method. In despite of these results, the model has a lot of imperfections in predicting the correct opacity regions. This may not only root from the lack of training data diversity, but also from the imperfections in manual labeling. Additional work is needed before a robust and accurate model can be developed for the ultimate goal of implementation in the clinical setting. In view of the above, our model is a successful first try in developing a segmentation model for the opacity regions for the Covid-19 positive chest X-rays. Our model schema and the manual segmentation data set may lay the foundation for the progress of more robust and accurate lung segmentation models in the future.

References

[1] Centers for Disease Control and Prevention. *COVIDView: A Weekly Surveillance Summary of U.S. COVID-19 Activity*, 2020 (accessed October 16, 202).

[2] H. Wang, Mahsen Naghavi, Craig Allen, R. M. Barber, A. Carter, D. C. Casey, F. J. Charlson, A. Z. Chen, Michael Coates, M. Coggeshall, L. Dandona, D. J. Dicker, H. E. Erskine, A. J. Ferrari, C. Fitzmaurice, Kyle Foreman, Mohammad H. Forouzanfar, M. S. Fraser, N. Fullman, E. M. Goldberg, N. Graetz, J. A. Haagsma, Simon I. Hay, C. Huynh, C. O. Johnson, N. J. Kassebaum, X. R. Kulikoff, M. J. Kutz, H. H. Kyu, H. J. Larson, J. Leung, S Lim, Marcus Lind, Rafael Lozano, N. Marquez, J. Mikesell, Ali H. Mokdad, M. D. Mooney, G. Nguyen, E. Nsoesie, David M Pigott, C. Pinho, G. A. Roth, L. Sandar, N. Silpakit, A. Sligar, R. J.D. Sorensen, J. Stanaway, C Steiner, S. Teeple, Mark A B Thomas, C. Troeger, A. VanderZanden, S. E. Vollset, V. Wanga, H. A. Whiteford, T. Wolock, L. Zoeckler, T. Achoki, A. Afshin, L. T. Alexander, G. M. Anderson, B. Bell, S. Biryukov, J. D. Blore, A. Brown, J. Brown, K. Cercy, A. Chew, A. J. Cohen, F. Daoud, E. Dossou, K. Estep, Abraham Flaxman, J. Friedman, J. Frostad, W. W. Godwin, J. Hancock, L. Kemmer, I. A. Khalil, P. Y. Liu, F. Masiye, A. Millear, M. Mirarefin, A. Misganaw, Maziar Moradi-Lakeh, K. Morgan, Maryanne Ng, Arnab Pain, J. Quame-Amaglo, P. Rao, M. B. Reitsma, K. A. Shackelford, P. J. Sur, J. A. Wagner, Theo Vos, A. D. Lopez, C. J.L. Murray, R. G. Ellenbogen, C. N. Mock, D. A. Quistberg, B. O. Anderson, C. D. Blosser, N. D. Futran, S. R. Heckbert, P. N.

Jensen, T. J. Montine, D. L. Tirschwell, D. A. Watkins, Z A Bhutta, M. I. Nisar, N. Akseer, N. K.M. Alam, L. D. Knibbs, Ratilal Lalloo, H. N. Gouda, John McGrath, P. Jeemon, R. Dandona, G. A. Kumar, Peter W. Gething, C Cooper, S. C. Darby, A. Deribew, MUHAMMAD REDZWAN S. RASHID ALI, D. A. Bennett, Vivekanand JHA, K. Rahimi, Y. Kinfu, I. D.A. Faghmous, S. M. Langan, M. McKee, G. V.S. Murthy, Neil Pearce, B. Roberts, I. R. Campos-Nonato, J. C. Campuzano, H. Gomez-Dantes, I. B. Heredia-Pi, F. Mejia-Rodriguez, J. C. Montañez Hernandez, P. Montero, M. J. Rios Blancas, E. E. Servan-Mori, S. Villalpando, L. Duan, Shanlin Liu, L Wang, P. Ye, X. Liang, S. Yu, G. A. Mensah, J. A. Salomon, A. L. Thorne-Lyman, O. N. Ajala, T. Bärnighausen, E. L. Ding, M. S. Farvid, G. R. Wagner, P. James, M. Osman, M. G. Shrime, J. R.A. Fitchett, A. K. Knudsen, C. L. Ellingsen, N. H. Krog, Maja Savic, A. D. Hailu, O. F. Norheim, K. H. Abate, T. T. Gebrehiwot, A. T. Gebremedhin, C. Abbafati, K. M. Abbas, F. Abd-Allah, S. F. Abera, Y. A. Melaku, F. H. Tesfay, G. Y. Abyu, A. F. Aregay, B. D. Betsu, A. A. Gebru, G. B. Hailu, A. Z. Yalew, H. G. Yebyo, D. M.X. Abreu, E. B. Franca, L. J. Abu-Raddad, A. L. Adelekan, R. O. Akinyemi, F. A. Ojelabi, Z. Ademi, T. Fürst, Peter Azzopardi, B. C. Cowie, K. B. Gibney, J. H. MacLachlan, A. Meretoja, K. Alam, Rohan Borschmann, S. M. Colquhoun, G. C. Patton, R. G. Weintraub, C. E.I. Szoeke, Lakshmi Vijayakumar, M. A. Bohensky, H. R. Taylor, T. Wijeratne, A. K. Adou, J. C. Adsuar, K. A. Afanvi, E. E. Agardh, Jurgen Rehm, A. Badawi, M. P. Lindsay, Svetlana Popova, A. Agarwal, A. Agrawal, P. J. Hotez, A. Ahmad, B. Norrving, A. S. Akanda, T. F. Akinyemiju, D. C. Schwebel, J. A. Singh, F. H. Al Lami, S. Alabed, Z. Al-Aly, T. R. Driscoll, A. H. Kemp, James Leigh, A. B. Mekonnen, D. Alasfoor, S. F. Aldhahri, K. A. Altirkawi, A. S. Terkawi, R. W. Aldridge, A Banerjee, T. Tillmann, M. A. Alegretti, A. V. Aleman, F. Cavalleri, V. Colistro, Z. A. Alemu, S. Alhabib, A. Alkerwi, F. Alla, P. Allebeck, J. J. Carrero, J. R. Carapetis, and GBD 2015 Mortality and Causes of Death Collaborators. Global, regional, and national life expectancy, all-cause mortality, and cause-specific mortality for 249 causes of death, 1980–2015: a systematic analysis for the global burden of disease study 2015. *Lancet*, 388(10053):1459–1544, oct 2016.

[3] Melina Hosseiny, Soheil Kooraki, Ali Gholamrezanezhad, Sravanthi Reddy, and Lee Myers. Radiology Perspective of Coronavirus Disease 2019 (COVID-19): Lessons From Severe Acute Respiratory Syndrome and Middle East Respiratory Syndrome. *American Journal of Roentgenology*, 214(5):1078–1082, feb 2020.

[4] Diletta Cozzi, Marco Albanesi, Edoardo Cavigli, Chiara Moroni, Alessandra Bindi, Silvia Luvarà, Silvia Lucarini, Simone Busoni, Lorenzo Nicola Mazzoni, and Vittorio Miele. Chest X-ray in new Coronavirus Disease 2019 (COVID-19) infection: findings and correlation with clinical outcome. *La Radiologia medica*, 125(8):730–737, aug 2020.

[5] Bejoy Abraham and Madhu S Nair. Computer-aided detection of covid-19 from x-ray images using multi-cnn and bayesnet classifier. *Biocybernetics and biomedical engineering*, 40(4):1436–1445, 2020.

[6] Mugahed A. Al-antari and Sungyoung Lee. Fast deep learning computer-aided diagnosis against the novel covid-19 pandemic from digital chest x-ray images, 06 2020.

[7] Arun Sharma, Sheeba Rani, and Dinesh Gupta. Artificial intelligence-based classification of chest x-ray images into covid-19 and other infectious diseases. *International Journal of Biomedical Imaging*, 10 2020.

[8] Hinton GE, Osindero S, and Teh YW. A fast learning algorithm for deep belief nets. *Neural Comput*, 18(7):1527–54, 7 2006.

[9] Li Fei-Fei, R. Fergus, and P. Perona. One-shot learning of object categories. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 28(4):594–611, 2006.

[10] Hugo Larochelle, Dumitru Erhan, Aaron Courville, James Bergstra, and Yoshua Bengio. An empirical evaluation of deep architectures on problems with many factors of variation. In *Proceedings of the 24th International Conference on Machine Learning*, ICML '07, page 473–480, New York, NY, USA, 2007. Association for Computing Machinery.

[11] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. In F. Pereira, C. J. C. Burges, L. Bottou, and K. Q. Weinberger, editors, *Advances in Neural Information Processing Systems 25*, pages 1097–1105. Curran Associates, Inc., 2012.

[12] Athanasios Voulodimos, Nikolaos Doulamis, Anastasios Doulamis, and Eftychios Protopapadakis. Deep Learning for Computer Vision: A Brief Review. *Computational Intelligence and Neuroscience*, 2018:7068349, 2018.

[13] Xiaosong Wang, Yifan Peng, Le Lu, Zhiyong Lu, Mohammadhadi Bagheri, and Ronald M. Summers. Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, Jul 2017.

[14] Ari Seff, Le Lu, Adrian Barbu, Holger Roth, Hoo-Chang Shin, and Ronald M Summers. Leveraging Mid-Level Semantic Boundary Cues for Automated Lymph Node Detection. In Nassir Navab, Joachim Hornegger, William M Wells, and Alejandro Frangi, editors, *Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015*, pages 53–61, Cham, 2015. Springer International Publishing.

[15] June-Goo Lee, Sanghoon Jun, Young-Won Cho, Hyunna Lee, Guk Bae Kim, Joon Beom Seo, and Namkug Kim. Deep Learning in Medical Imaging: General Overview. *Korean J Radiol*, 18(4):570–584, aug 2017.

[16] Yong Yu, Xiaosheng Si, Changhua Hu, and Jianxun Zhang. A Review of Recurrent Neural Networks: LSTM Cells and Network Architectures. *Neural Computation*, 31(7):1235–1270, 2019.

[17] Anamika Dhillon and Gyanendra K Verma. Convolutional neural network: a review of models, methodologies and applications to object detection. *Progress in Artificial Intelligence*, 9(2):85–112, 2020.

[18] Yu-Xing Tang, You-Bao Tang, Yifan Peng, Ke Yan, Mohammadhadi Bagheri, Bernadette A Redd, Catherine J Brandon, Zhiyong Lu, Mei Han, Jing Xiao, and Ronald M Summers. Automated abnormality classification of chest radiographs using deep convolutional neural networks. *npj Digital Medicine*, 3(1):70, 2020.

[19] Ke Zhang, Miao Sun, Tony X. Han, Xingfang Yuan, Liru Guo, and Tao Liu. Residual networks of residual networks: Multilevel residual networks. *IEEE Transactions on Circuits and Systems for Video Technology*, 28(6):1303–1314, Jun 2018.

[20] Lei Cai, Jingyang Gao, and Di Zhao. A review of the application of deep learning in medical image classification and segmentation. *Annals of translational medicine*, 8(11):713, jun 2020.

[21] Hua Ma, Pierre Ambrosini, and Theo van Walsum. Fast prospective detection of contrast inflow in x-ray angiograms with convolutional neural network and recurrent neural network. *Lecture Notes in Computer Science*, page 453–461, 2017.

[22] Justin Ker, Satya P Singh, Yeqi Bai, Jai Rao, Tchoyoson Lim, and Lipo Wang. Image Thresholding Improves 3-Dimensional Convolutional Neural Network Diagnosis of Different Acute Brain Hemorrhages on Computed Tomography Scans. *Sensors (Basel, Switzerland)*, 19(9), may 2019.

[23] Keisuke Kawauchi, Sho Furuya, Kenji Hirata, Chietsugu Katoh, Osamu Manabe, Kentaro Kobayashi, Shiro Watanabe, and Tohru Shiga. A convolutional neural network-based system to classify patients using FDG PET/CT examinations. *BMC cancer*, 20(1):227, mar 2020.

[24] Alexander Selvikvåg Lundervold and Arvid Lundervold. An overview of deep learning in medical imaging focusing on MRI. *Zeitschrift fur medizinische Physik*, 29(2):102–127, may 2019.

[25] Shengfeng Liu, Yi Wang, Xin Yang, Baiying Lei, Li Liu, Shawn Xiang Li, Dong Ni, and Tianfu Wang. Deep Learning in Medical Ultrasound Analysis: A Review. *Engineering*, 5(2):261–275, 2019.

[26] Kang Zhang, Xiaohong Liu, Jun Shen, Zhihuan Li, Ye Sang, Xingwang Wu, Yunfei Zha, Wenhua Liang, Chengdi Wang, Ke Wang, Linsen Ye, Ming Gao, Zhongguo Zhou, Liang Li, Jin Wang, Zehong Yang, Huimin Cai, Jie Xu, Lei Yang, Wenjia Cai, Wenqin Xu, Shaoxu Wu, Wei Zhang, Shanping Jiang, Lianghong Zheng, Xuan Zhang, Li Wang, Liu Lu, Jiaming Li, Haiping Yin, Winston Wang, Oulan Li, Charlotte Zhang, Liang Liang, Tao Wu, Ruiyun Deng, Kang Wei, Yong Zhou, Ting Chen, Johnson Yiu-Nam Lau, Manson Fok, Jianxing He, Tianxin Lin, Weimin Li, and Guangyu Wang. Clinically applicable ai system for accurate diagnosis, quantitative measurements, and prognosis of covid-19 pneumonia using computed tomography. *Cell*, 181(6):1423 – 1433.e11, 2020.

[27] Xueyan Mei, Hao-Chih Lee, Kai-yue Diao, Mingqian Huang, Bin Lin, Chenyu Liu, Zongyu Xie, Yixuan Ma, Philip M Robson, Michael Chung, Adam Bernheim, Venkatesh Mani, Claudia Calcagno, Kunwei Li, Shaolin Li, Hong Shan, Jian Lv, Tongtong Zhao, Junli Xia, Qihua Long, Sharon Steinberger, Adam Jacobi, Timothy Deyer, Marta Luksza, Fang Liu, Brent P Little, Zahi A Fayad, and Yang Yang. Artificial intelligence–enabled rapid diagnosis of patients with COVID-19. *Nature Medicine*, 26(8):1224–1228, 2020.

[28] Matthew D Li, Nishanth Thumbavanam Arun, Mishka Gidwani, Ken Chang, Francis Deng, Brent P Little, Dexter P Mendoza, Min Lang, Susanna I Lee, Aileen O'Shea, Anushri Parakh, Praveer Singh, and Jayashree Kalpathy-Cramer. Automated Assessment and Tracking of COVID-19 Pulmonary Disease Severity on Chest Radiographs using Convolutional Siamese Neural Networks. *Radiology: Artificial Intelligence*, 2(4):e200079, 2020.

[29] Joseph Paul Cohen, Lan Dao, Paul Morrison, Karsten Roth, Yoshua Bengio, Beiyi Shen, Almas Abbasi, Mahsa Hoshmand-Kochi, Marzyeh Ghassemi, Haifang Li, and Tim Q Duong. Predicting covid-19 pneumonia severity on chest x-ray with deep learning, 2020.

[30] Alberto Signoroni, Mattia Savardi, Sergio Benini, Nicola Adami, Riccardo Leonardi, Paolo Gibellini, Filippo Vaccher, Marco Ravanelli, Andrea Borghesi, Roberto Maroldi, and Davide Farina. End-to-end learning for semiquantitative rating of covid-19 severity on chest x-rays, 2020.

[31] S. Desai, A. Baghal, T. Wongsurawat, S. Al-Shukri, K. Gates, P. Farmer, M. Rutherford, G.D. Blake, T. Nolan, T. Powell, K. Sexton, W. Bennett, and F. Prior. Data from chest imaging with clinical and genomic correlates representing a rural covid-19 positive population. *The Cancer Imaging Archive (TCIA)*, 2020.

[32] Clark K, Vendt B, Smith K, Freymann J, Kirby J, Koppel P, Moore S, Phillips S, Maffitt D, Pringle M, Tarbox L, and Prior F. The cancer imaging archive (tcia): Maintaining and operating a public information repository, journal of digital imaging. *Journal of Digital Imaging*, pages 1045–1057, 2013.

[33] Kentaro Wada. labelme: Image Polygonal Annotation with Python. `https://github.com/wkentaro/labelme`, 2016.

[34] Alexander Buslaev, Vladimir I. Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A. Kalinin. Albumentations: Fast and flexible image augmentations. *Information*, 11(2), 2020.

[35] Pavel Yakubovskiy. Segmentation models. `https://github.com/qubvel/segmentation_models`, 2019.

[36] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation, 2015.

[37] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition, 2015.

[38] Carole H. Sudre, Wenqi Li, Tom Vercauteren, Sebastien Ourselin, and M. Jorge Cardoso. Generalised dice overlap as a deep learning loss function for highly unbalanced segmentations. *Lecture Notes in Computer Science*, page 240–248, 2017.

[39] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection, 2018.

[40] Md Rahman and Yang Wang. Optimizing intersection-over-union in deep neural networks for image segmentation, 12 2016.

[41] Marco Francone, Franco Iafrate, Giorgio Maria Masci, Simona Coco, Francesco Cilia, Lucia Manganaro, Valeria Panebianco, Chiara Andreoli, Maria Chiara Colaiacomo, Maria Antonella Zingaropoli, Maria Rosa Ciardi, Claudio Maria Mastroianni, Francesco Pugliese, Francesco Alessandri, Ombretta Turriziani, Paolo Ricci, and Carlo Catalano. Chest CT score in COVID-19 patients: correlation with disease severity and short-term prognosis. *European radiology*, pages 1–10, jul 2020.