

This is the ACCEPTED VERSION of the following published document:

Vidal, P.L., de Moura, J., Novo, J., Ortega, M. (2022). Pulmonary-Restricted COVID-19 Informative Visual Screening Using Chest X-ray Images from Portable Devices. In: Sclaroff, S., Distante, C., Leo, M., Farinella, G.M., Tombari, F. (eds) Image Analysis and Processing – ICIAP 2022. ICIAP 2022. Lecture Notes in Computer Science, vol 13231. Springer, Cham. https://doi.org/10.1007/978-3-031-06427-2_6

Link to published version: https://doi.org/10.1007/978-3-031-06427-2_6

General rights:

©2022 This version of the conference paper has been accepted for publication, after peer review and is subject to Springer Nature's [AM terms of use](#), but is not the Version of Record and does not reflect post-acceptance improvements, or any corrections. The Version of Record is available online at: https://doi.org/10.1007/978-3-031-06427-2_6.

Pulmonary-Restricted COVID-19 Informative Visual Screening Using Chest X-ray Images from Portable Devices ^{*}

Plácido L. Vidal^{1,2}[0000–0002–6009–4737], Joaquim de Moura^{1,2}[0000–0002–2050–3786], Jorge Novo^{1,2}[0000–0002–0125–3064], and Marcos Ortega^{1,2}[0000–0002–2798–0788]

¹ Centro de investigación CITIC, Universidade da Coruña, Campus de Elviña, s/n, 15071 A Coruña, Spain

² Grupo VARPA, Instituto de Investigación Biomédica de A Coruña (INIBIC), Universidade da Coruña, Xubias de Arriba, 84, 15006 A Coruña, Spain
{placido.francisco.lizancos.vidal,joaquim.demoura,jnovo,mortega}@udc.es

Abstract. In the recent COVID-19 outbreak, chest X-rays were the main tool for diagnosing and monitoring the pathology. To prevent further spread of this disease, special circuits had to be implemented in the healthcare services. For this reason, these chest X-rays were captured with portable X-ray devices that compensate its lower quality and limitations with more deployment flexibility. However, most of the proposed computer-aided diagnosis methodologies were designed to work with traditional fixed X-ray machines and their performance is diminished when faced with these portable images. Additionally, given that the equipment needed to properly treat the disease (such as for life support and monitoring of vital signs) most of these systems learnt to identify these artifacts in the images instead of real clinically-significant variables. In this work, we present the first methodology forced to extract features exclusively from the pulmonary region of interest specially designed to work with these difficult portable images. Additionally, we generate a class activation map so the methodology also provides explainability to the results returned to the clinician. To ensure the robustness of our proposal, we tested the methodology with chest radiographs from patients diagnosed with

^{*} This research was funded by Instituto de Salud Carlos III, Government of Spain, DTS18/00136 research project; Ministerio de Ciencia e Innovación y Universidades, Government of Spain, RTI2018-095894-B-I00 research project, Ayudas para la formación de profesorado universitario (FPU), grant ref. FPU18/02271; Ministerio de Ciencia e Innovación, Government of Spain through the research project PID2019-108435RB-I00; Consellería de Cultura, Educación e Universidade, Xunta de Galicia, Grupos de Referencia Competitiva, grant ref. ED431C 2020/24 and through the postdoctoral grant contract ref. ED481B 2021/059; Axencia Galega de Innovación (GAIN), Xunta de Galicia, grant ref. IN845D 2020/38; CITIC, as Research Center accredited by Galician University System, is funded by “Consellería de Cultura, Educación e Universidade from Xunta de Galicia”, supported in an 80% through ERDF Funds, ERDF Operational Programme Galicia 2014-2020, and the remaining 20% by “Secretaría Xeral de Universidades” (Grant ED431G 2019/01).

COVID-19, pathologies similar to COVID-19 (such as other types of viral pneumonias) and healthy patients in different combinations with three convolutional networks from the state of the art (for a total of 9 studied scenarios). The experimentation confirms that our proposal is able to separate COVID-19 cases, reaching a $94.7\% \pm 1.34\%$ of accuracy.

Keywords: COVID-19 · Chest X-ray · CAD system · Class activation map · Deep Learning · X-ray portable devices.

1 Introduction

Coronavirus COVID-19 disease is an affection that primarily impacts in the pulmonary region [2], predominantly resulting in viral pneumonia. This results in the appearance of pathological structures visible through chest X-rays and other imaging techniques [19], present even in asymptomatic patients [4]. Amidst the global COVID-19 pandemic of 2020, these chest radiographs were one of the main means to diagnose the aforementioned pathology. For this reason, numerous automatic computer-aided diagnosis (CAD) methodologies were developed, aimed at reducing the subjectivity of the experts as well as the needed time and resources [7, 14, 3]. These proposals were designed to work with images from either fixed X-ray devices (which allow to obtain high-definition projection of the internal tissues in different angles and depths) or CT devices (which generate a defined single slice representation of the human body).

However, during medical emergencies and cross-contamination risk, the main mean to diagnose lung diseases is by means of portable X-ray devices. These devices do not require a special platform nor room, so they can be used in dedicated circuits in hospitals and isolated areas (at the cost of lesser image quality and limited angles with worse capture conditions)[10]. However, this lesser image quality and limitations severely hinder the performance of methodologies not tested for them. For this reason, contributions specially designed to work with these portable devices were proposed, such as the work of Vidal *et al.* [20] aimed at generating a segmentation of the pulmonary region, Morís *et al.* [11, 12] generating synthetic images to palliate the lack of samples in this issue, and de Moura *et al.* [13] with the objective of performing a screening, separating normal images from COVID-19 and from pathologies with COVID-19-like patterns. In this last work (and in some works based on fixed devices [1]), they also generate a gradient flow map so they can help the experts to assess what has the deep learning model analyzed in the images.

These methodologies (as shown by the gradient flow maps generated by the networks) were taking advantage of collateral information present in the radiographs to generate their prediction. That is, for example, if foreign artifacts commonly present in afflicted patients from certain diseases (such as respirators, pacemakers, heart rate sensors, etc.) are shown in the radiographs, these networks use this information to improve the metrics [16] and guess that an intubated patient must be afflicted by some kind of pulmonary disease or at least be at higher risk of being infected by the studied disease. While it is common for neural

networks to take advantage of underlying statistical features commonly missed by human experts, this implies that the system will perform worse with different clinical protocols, stages of the pathology or even age groups as these artifacts used to their advantage may change.

For this reason, we propose the first fully automatic methodology to remove these factors by masking these extraneous artifacts and able to successfully work with images from these portable X-ray devices. This way, we obtain a robust methodology that only extracts information from the pulmonary region of interest, resulting in a system that is able to work in any stage of the progression of the patient and generate a prediction solely based on real clinical features. In addition, we take advantage of this mechanism that forces the attention of the network on the lung region to generate a representation that explains the results to the clinician. Thus, the expert will not only be able to verify that the network is obtaining the desired behavior, but may even discover factors overlooked by the expert that can help determine a better treatment and assist in their future monitoring. And, thanks to being able to work in these special scenarios with portable X-ray devices, can be seamlessly implemented in healthcare circuits specially designed to treat patients in a pandemic scenario. To the best of our knowledge, this proposal represents the only study designed to work with these portable devices able to extract information from exclusively the pulmonary regions, perform a classification of the results between COVID-19, healthy and similar pathologies to COVID-19, and generate an explainable representation of these classifications.

2 Materials

To train, test and validate our methodology we used a specific dataset composed by 2,071 radiographs from patients diagnosed with COVID-19, 716 from patients diagnosed with non-COVID-19 pathologies (but with very similar features to it), and 797 images from healthy patients. These images were acquired with an Agfa dr100E GE and Optima Rx200 portable X-ray devices. All these images were obtained by the Radiology Service of the Complejo Hospitalario Universitario de A Coruña (CHUAC) during real live clinical practice. The images with non-COVID-19 pathologies but with very similar patterns belong to types of viral pneumonias or bacterial infections that leave a comparable trace and damage on the lungs of the afflicted. This is an specially interesting case, as the system need to discern these images from the very similar belonging to COVID-19 patients while also limited by the restrictions imposed by the portable X-ray devices. For the sake of balancing the dataset between the three available types of images, 716 images were randomly selected from each set to compose the final used dataset.

3 Methodology

Our proposal is divided into two main steps presented in Fig. 1. In the first step, we use an specially designed methodology to extract the lung region of interest

from portable chest radiographs (Section 3.1). Once we have delimited this region of interest, we mask the dataset and use it to train a convolutional neural network. Finally, using this trained network, we extract the gradient flow to generate an intuitive visualization of the relevant regions from the final classification (Section 3.2). This methodology will be followed with three different combinations of the labels that are available in the dataset to further evaluate the capabilities of the methodology and its behavior: Non-Healthy (COVID-19 + Pathologies similar to COVID-19) versus Healthy, COVID-19 versus non-COVID-19 (Pathologies similar to COVID-19 and Healthy), and every label as an independent class (COVID-19 versus Pathological versus Healthy).




Fig. 1. Workflow for each repetition during the training of the proposed methodology.

3.1 Region of interest extraction

To extract the region of interest we will use a methodology that is robust to the artifacts and structures that are commonly present in portable chest radiographs. In our case, we will use a proposal based on the work of Vidal *et al.* [20]. This

methodology performs a double-stage transfer learning: in the first step, the chosen network (in this case, an U-Net [15]) is adapted (via transfer learning) from magnetic resonance imaging for brain glioma segmentation [5] to general-purpose chest radiographs. Afterwards, the network undergoes a second transfer learning stage, where the network is further refined to be able to work also with radiographs extracted from portable devices (as well as with the pathologies considered in this work). With this specially designed segmentation methodology, we extract all the pulmonary regions of interest from the dataset. Finally, for each of these images, we apply the Hadamard product (H) to filter-out all artifacts not belonging to the lung area from the lower to the apical zone in the chest radiographs. This way, we remove any external artifacts the network could use to its advantage, artificially improving the metrics.

3.2 COVID-19 screening

In this second step, we will train three different networks for each of the case studies considered in this work. Non-Healthy versus Healthy, COVID-19 versus non-COVID-19, and every label as an independent class (COVID-19 versus Pathological versus Healthy).

To train each of the networks, we will use the masked dataset of the previous step and three architectures from the state of the art: a ResNet18[8], a Vgg19[18] and a DenseNet161[9]. Additionally, all these network architectures were pretrained with the ImageNet dataset [6] to further speed up the training and reduce the impact of the number of samples. Furthermore, to increase the effective number of available samples, a data augmentation strategy was used by randomly horizontally flipping and scaling the input lung radiographs. Finally, the dataset is randomly shuffled and divided into three sets: 60% of the images to be used for the training of the network, 20% for validation purposes of each epoch of training, and the remaining 20% to test the performance and robustness of the final models. Each model is set to be trained for a fixed number of 200 epochs, empirically determined to be enough to reach a stability point where no further improvements in the final metrics are achieved. From these 200 epochs, the model that achieved the best validation loss in a given epoch is the one chosen from the set. The models were trained using the cross-entropy as loss function, using as optimizer the Stochastic Gradient Descent (SGD) with a constant learning rate of 0.01 and a momentum of 0.9. Each network is trained independently and this process is repeated five times for each model and case study mentioned above in order to diminish the improbable selection bias and obtain more robust final metrics.

Finally, as our proposal is designed to help the clinicians with their diagnostic task, we extract the regions where the networks centered its attention. To do so, we use gradient-weighted class activation mapping in the network [17]. This strategy uses the gradient information flowing towards the last convolutional layer to determine the relevance of each neuron towards the generation of the classification, a generalization of Class Activation Mapping or CAM. This strategy is commonly used to assess the robustness of the predictions of a network, as it

allows to remove models that take advantage on circumstantial features instead of actual relevant descriptors. However, once the models were verified to be valid, the information generated by the network is also useful to the clinicians, as we generate explainable results that can, even, help to discover new biomarkers hitherto unconsidered or determine if some elements present in the image are product of the quality of the X-ray images from portable devices / actual pathological structures.

4 Results and discussion

Below we present the results that were obtained during the experimentation, divided into results during the training (Section 4.1) and the metrics of the final model and class activation maps (Section 4.2). In addition, within each section, each of the three aforementioned cases of study will be discussed separately.

4.1 Training results

The results for the training of the models studying the case Healthy versus Non-Healthy can be seen in Fig. 2. In this figure, we can see how this situation represents the easiest for the model to learn, as the accuracy during the training of all the three models quickly reached the maximum value. However, in this case, we can also see how the training loss quickly falls down while the validation loss stays on average on the same height. We see how the system is training for the initial variability, but quickly stabilises. This indicates that the issue at hand might be too easy for the methodology and a variable learning rate should be used. Because of this, we are perceiving slight traces of overfitting in the model in the final epochs (albeit we will only keep the one with the best validation loss nonetheless). This can also be seen in how the simpler architecture (Vgg19) is the one who obtains best validation accuracy and loss from the three tested models. This is normal given that healthy lungs do not present opacities, infiltrates and dense tissues present in the studied afflicted lungs.

On the other hand, in the results of the training for the study case of COVID-19 versus non-COVID-19 samples shown in Fig. 3, we can see how the results are significantly different. We are evaluating a more difficult situation than in the first scenario, and now the models present a more balanced progression between training and validation. In this case, we see how both the training loss and validation loss are proximal respective between the same model and no predominant signs of overfitting (as well as for significantly improved accuracy metrics during validation).

Finally, in Fig. 4 we can see the training and validation results for the case study where we consider each class independently in the model targets. In this case, we observe how all the models ended achieving an in-between compromise between the same models in the two previous experiments. As we will analyze in the following sections, the Healthy and the Pathological class probably present some overlapping in certain scenarios that is being exacerbated by the quality and limitations of the portable chest radiographs used in this work.

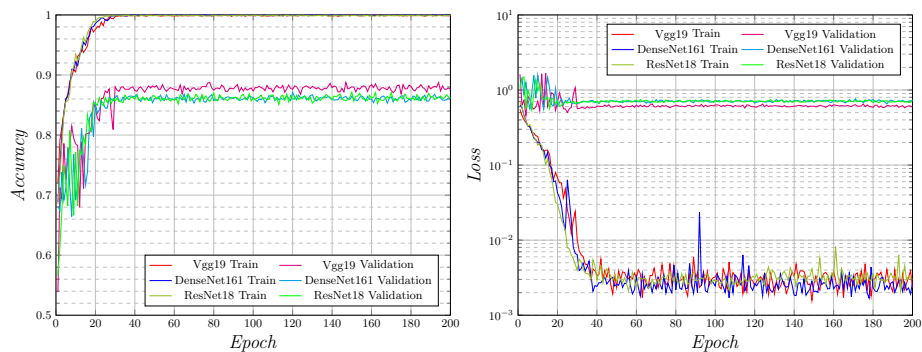


Fig. 2. Accuracy and loss during training and validation of the models for the experiment Healthy versus Non-Healthy.

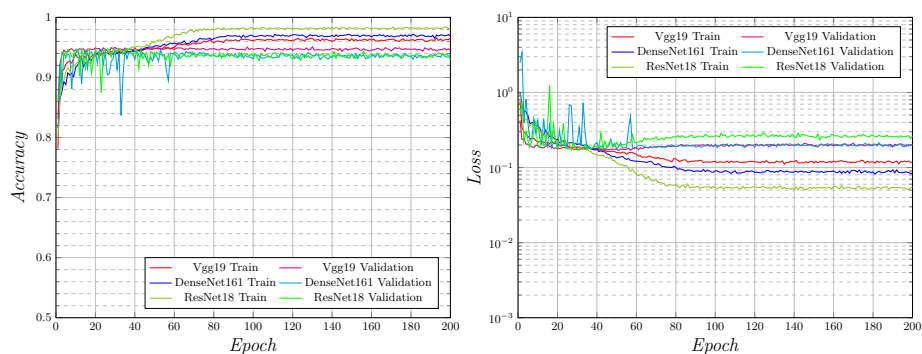


Fig. 3. Accuracy and loss during training and validation of the models for the experiment non-COVID-19 versus COVID-19.

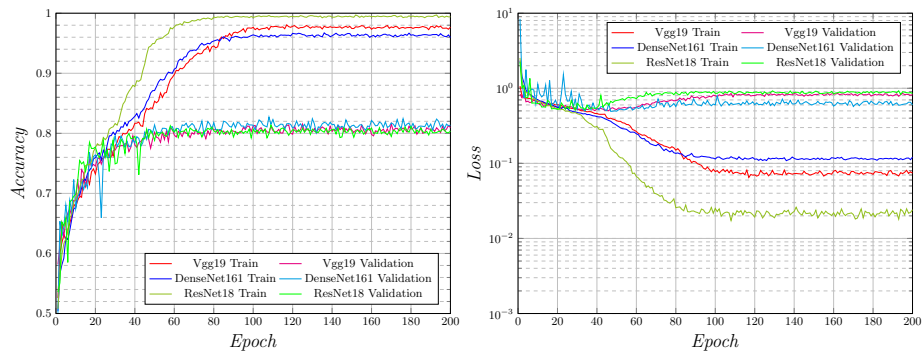


Fig. 4. Accuracy and loss during training and validation of the models for the experiment healthy versus pathological versus COVID-19.

4.2 Test results

The metrics for the first experiment (Healthy versus Non-Healthy) are shown in Table 1. In this case, we can confirm what we saw in Fig. 2, where the Vgg19 model is the one obtaining the best overall results. We can see how, in this case, the simpler model is obtaining the best results of the experiment (albeit satisfactory nonetheless in all scenarios). Additionally, we can see how the ResNet18 has particularly high standard deviation in several metrics confirming what we stated in the previous sections: the network has some traces of overfitting.

Table 1. Test results for the experiment evaluating healthy (H) versus pathological (P) and COVID-19 (C) classes.

Architecture	Accuracy	Class	Precision	Recall	F1-Score
Densenet161	87.2% \pm 2.53%	H	85.0 % \pm 2.76%	88.8% \pm 4.58%	87.0% \pm 3.33%
		P+C	89.2% \pm 2.89%	85.6% \pm 1.62%	87.4% \pm 1.85%
ResNet18	85.7% \pm 1.63%	H	85.8% \pm 3.71%	89.4% \pm 4.80%	86.2% \pm 1.94%
		P+C	86.6% \pm 4.18%	84.2% \pm 4.40%	85.2% \pm 1.47%
Vgg19	89.6% \pm 1.29%	H	87.6% \pm 1.74%	91.8% \pm 1.72%	89.8% \pm 1.47%
		P+C	91.8% \pm 1.83%	87.4% \pm 1.36%	89.6% \pm 1.29%

On the other hand, in the second experiment (COVID-19 versus non-COVID-19) shown in Table 2 we see how the networks were able to obtain outstanding results. Moreover, in this case, we can highlight the performance of the DenseNet161 in terms of recall for the non-COVID-19 class and precision for the COVID-19 class. Additionally, we can observe how, like in the previous experiment, the ResNet18 architecture stands out for its higher than average standard deviation. Nonetheless, the DenseNet161 in this case would be the preferable model, as it obtained the highest metrics even in the F1 Score that is resilient to the unbalances of the dataset (unlike the accuracy, for example). Additionally, its densely connected blocks that confer the network an extra layer of self-supervision are preventing the network from falling into overfitting like the ResNet18.

Finally, in Table 3 we can see the results for the last experiment, where we train the networks with all the classes independently. In this case, we can see how, in all three networks the COVID-19 class stands out from the rest, being easily distinguished from the rest. On the other hand, the precision of the healthy class and the recall of the pathological class suffer a bit when separated. In this case, we see how the pathological class and the healthy are intermixing results in some cases. This may be due to the fact that some of the patients could have been tested with an early onset stage of the diseases similar to COVID-19 that, in conjunction with the difficulties imposed by the portable X-ray devices, present very similar patterns to healthy images. This way, when these two classes are separated (such as in this experiment and the first) they return mixed results. For this reason, the desirable model in these cases is the ones generated in the

Table 2. Test results for the experiment evaluating healthy (H) and pathological (P) versus COVID-19 (C) classes.

Architecture	Accuracy	Class	Precision	Recall	F1-Score
Densenet161	94.7% \pm 1.34%	H+P	93.0% \pm 1.10%	99.4% \pm 0.8%	96.0% \pm 1.10%
		C	98.4% \pm 2.24%	86.4% \pm 2.42%	91.8% \pm 1.94%
ResNet18	94.4% \pm 1.06%	H+P	93.2% \pm 1.47%	98.6% \pm 1.74%	95.8% \pm 0.75%
		C	97.2% \pm 3.49%	85.8% \pm 3.76%	91.6% \pm 2.94%
Vgg19	94.3% \pm 1.43%	H+P	93.2% \pm 1.17%	98.6% \pm 1.36%	95.8% \pm 1.17%
		C	96.8% \pm 2.79%	86.2% \pm 1.33%	91.0% \pm 1.67%

second experiment. In this experiment the pathological and the healthy classes are combined and the confusing patterns, instead of penalizing the model, help it to achieve outstanding better results.

Table 3. Test results for the experiment evaluating healthy (H) versus pathological (P) versus COVID-19 (C) classes.

Architecture	Accuracy	Class	Precision	Recall	F1-Score
Densenet161	83.4% \pm 2.06%	H	77.2% \pm 4.07%	81.4% \pm 3.38%	79.2% \pm 2.99%
		P	80.4% \pm 3.88%	78.8% \pm 4.12%	79.4% \pm 2.73%
		C	93.8% \pm 2.64%	90.2% \pm 1.72%	91.8% \pm 0.40%
ResNet18	81.8% \pm 1.56%	H	72.4% \pm 5.39%	82.6% \pm 3.55%	77.0% \pm 2.83%
		P	81.2% \pm 1.72%	76.8% \pm 3.87%	78.8% \pm 1.60%
		C	95.0% \pm 2.68%	85.8% \pm 3.19%	90.0% \pm 1.10%
Vgg19	81.3% \pm 2.57%	H	72.8% \pm 6.85%	80.4% \pm 4.17%	75.8% \pm 3.19%
		P	82.0% \pm 3.03%	76.8% \pm 8.33%	79.0% \pm 3.74%
		C	92.4% \pm 3.77%	87.8% \pm 3.82%	89.8% \pm 2.57%

For this reason, the DenseNet161 model from the second experiment is the one that we will use to generate the class activation maps. In Fig. 5, we can see four chest radiographs from portable devices as an example of the class activation maps generated with our proposal. As the reader can see, the proposed methodology has achieved satisfactory results, being able to successfully discern COVID-19 patients from healthy/pathological samples. Additionally, as shown in that same figure, both for the non-COVID-19 cases and COVID-19 cases, thanks to the proposed masking strategy, the system is only focusing on lung regions. While on non-COVID-19 images the lung regions remain clean, COVID-19 detections are centered on the affected lung regions (such as the bronchus/bronchioles in the leftmost COVID-19 image). Thus, all the performed predictions are purely based on the analysis of lung structures and not other artifacts (such as the wires present in the leftmost healthy radiography).




Fig. 5. Examples of generated maps with the DenseNet161 trained to differentiate non-COVID-19 patients from COVID-19 ones.

5 Conclusions

In this work we have presented a methodology capable of working with portable X-ray images, critical for healthcare systems in emergency scenarios due to their flexibility (at the cost of inferior capture quality). Our system has been able to successfully differentiate COVID-19 from COVID-19-like pathologies and from images of healthy patients, even with the difficulties imposed by the capture quality of these portable devices. In addition, we have generated the class activation maps of the chosen model, which has demonstrated that our methodology of masking the region of interest in these portable device images ensures that the model is focused exclusively on the pulmonary region. This means that our system is not taking advantage of information foreign to the domain (such as the presence of pacemakers, respirators, cardiac pulse meters, etc.) and able to work independently of the established protocol. Moreover, these class activation maps of the network are useful to assist the clinical expert in his diagnostic task by providing an explainable representation of what is detected by the system. Thus, the expert can not only determine the validity of the results, but may also discover factors that went unnoticed during the overloaded clinical practice of crowded emergency services. As a result, the system will provide robust measurements regardless of clinical procedure, patient status or stage of the pathology development. As future work, we plan to further improve the experiments to reduce the impact of overfitting in some of the models with a dynamic loss paradigm (to improve the resolution of the gradient descent) and early stopping strategies, as well as the implementation of a data augmentation strategy by synthesis of new samples with generative adversarial networks to complement the variability of the dataset. Finally, as mentioned, some of the samples presented are probably from early onset stages of different pathologies with similar symptoms to COVID-19. A study with clinicians to reverse-engineer the models and discover these early markers in images could be also of interest.

References

1. Explainable deep learning for pulmonary disease and coronavirus covid-19 detection from x-rays. *Comp. Methods and Programs in Biomed.* **196**, 105608 (2020). <https://doi.org/https://doi.org/10.1016/j.cmpb.2020.105608>
2. Aguiar, D., Lobrinus, J.A., Schibler, M., Fracasso, T., Lardi, C.: Inside the lungs of COVID-19 disease. *International Journal of Legal Medicine* **134**(4), 1271–1274 (May 2020). <https://doi.org/10.1007/s00414-020-02318-9>
3. Alom, M.Z., Rahman, M.M.S., Nasrin, M.S., Taha, T.M., Asari, V.K.: Covid mtnet: Covid-19 detection with multi-task deep learning approaches (2020)
4. Bandirali, M., Sconfienza, L.M., Serra, R., Brembilla, R., Albano, D., Pregliasco, F.E., Messina, C.: Chest radiograph findings in asymptomatic and minimally symptomatic quarantined patients. *Radiology* **295**(3), E7–E7 (Jun 2020). <https://doi.org/10.1148/radiol.2020201102>
5. Buda, M., Saha, A., Mazurowski, M.A.: Association of genomic subtypes of lower-grade gliomas with shape features automatically extracted by a deep learning algorithm. *Computers in Biology and Medicine* **109**, 218–225 (Jun 2019). <https://doi.org/10.1016/j.compbiomed.2019.05.002>
6. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: 2009 IEEE Conf. on comp. vision and pattern recognition. pp. 248–255. IEEE (2009)
7. Fan, D.P., Zhou, T., Ji, G.P., Zhou, Y., Chen, G., Fu, H., Shen, J., Shao, L.: Inf-net: Automatic COVID-19 lung infection segmentation from CT images. *IEEE Transactions on Medical Imaging* **39**(8), 2626–2637 (Aug 2020). <https://doi.org/10.1109/tmi.2020.2996645>
8. He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. *CoRR abs/1512.03385* (2015), <http://arxiv.org/abs/1512.03385>
9. Huang, G., Liu, Z., Weinberger, K.Q.: Densely connected convolutional networks. *CoRR abs/1608.06993* (2016), <http://arxiv.org/abs/1608.06993>
10. Jacobi, A., Chung, M., Bernheim, A., Eber, C.: Portable chest x-ray in coronavirus disease-19 (COVID-19): A pictorial review. *Clinical Imaging* **64**, 35–42 (Aug 2020). <https://doi.org/10.1016/j.clinimag.2020.04.001>
11. Morís, D.I., de Moura, J., Novo, J., Ortega, M.: Cycle generative adversarial network approaches to produce novel portable chest x-rays images for covid-19 diagnosis. In: ICASSP 2021 - 2021 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP). IEEE (Jun 2021). <https://doi.org/10.1109/icassp39728.2021.9414031>, <https://doi.org/10.1109/icassp39728.2021.9414031>
12. Morís, D.I., de Moura Ramos, J.J., Buján, J.N., Hortas, M.O.: Data augmentation approaches using cycle-consistent adversarial networks for improving COVID-19 screening in portable chest x-ray images. *Expert Systems with Applications* **185**, 115681 (Dec 2021). <https://doi.org/10.1016/j.eswa.2021.115681>, <https://doi.org/10.1016/j.eswa.2021.115681>
13. de Moura, J., García, L.R., Vidal, P.F.L., Cruz, M., López, L.A., Lopez, E.C., Novo, J., Ortega, M.: Deep convolutional approaches for the analysis of covid-19 using chest x-ray images from portable devices. *IEEE Access* **8**, 195594–195607 (2020). <https://doi.org/10.1109/ACCESS.2020.3033762>
14. de Moura, J., Novo, J., Ortega, M.: Fully automatic deep convolutional approaches for the analysis of covid-19 using chest x-ray images. *medRxiv* (May 2020). <https://doi.org/10.1101/2020.05.01.20087254>

15. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. CoRR **abs/1505.04597** (2015), <http://arxiv.org/abs/1505.04597>
16. Sadre, R., Sundaram, B., Majumdar, S., Ushizima, D.: Validating deep learning inference during chest x-ray classification for COVID-19 screening. *Scientific Reports* **11**(1) (Aug 2021). <https://doi.org/10.1038/s41598-021-95561-y>, <https://doi.org/10.1038/s41598-021-95561-y>
17. Selvaraju, R.R., Cogswell, M., Das, A., Vedantam, R., Parikh, D., Batra, D.: Grad-cam: Visual explanations from deep networks via gradient-based localization. In: 2017 IEEE International Conference on Computer Vision (ICCV). pp. 618–626 (2017). <https://doi.org/10.1109/ICCV.2017.74>
18. Simonyan, K., Zisserman, A.: Very deep convolutional networks for large-scale image recognition. CoRR **abs/1409.1556** (2015), <http://arxiv.org/abs/1409.1556>
19. Velavan, T.P., Meyer, C.G.: The COVID-19 epidemic. *Tropical Medicine & International Health* **25**(3), 278–280 (Feb 2020). <https://doi.org/10.1111/tmi.13383>
20. Vidal, P.L., de Moura, J., Novo, J., Ortega, M.: Multi-stage transfer learning for lung segmentation using portable x-ray devices for patients with COVID-19. *Expert Systems with Applications* **173**, 114677 (Jul 2021). <https://doi.org/10.1016/j.eswa.2021.114677>