

An approach based on image processing techniques to segment lung region in chest X-ray images

Luiza C. de Menezes, Augusto R.V.F. de Araújo, Aura Conci
Instituto de Computação
Universidade Federal Fluminense
Niterói, RJ 24210-346
Email: lumenezes@id.uff.br, augustorighetti@id.uff.br, aconci@ic.uff.br

Abstract—Chest X-ray (CXR) images help specialists worldwide to diagnose lung diseases, such as tuberculosis and COVID-19. A primary step in an image-based diagnostic tool is to segment the region of interest. That facilitates the disease classification problem by reducing the amount of information to be processed. However, due to the noisy nature of CXRs, identifying the lung region can be a challenging task. This paper addresses the lung segmentation problem using a less costable computational process based on image analysis and mathematical morphology techniques. The proposed method achieved a specificity of 92.92%, a Jaccard index of 77.77%, and a Dice index of 87.37% on average. All images that comprehend the dataset used and their respective ground truths are available for download at <https://github.com/mnzluiza/Lung-Segmentation>.

I. INTRODUCTION

Radiographic images are an indispensable tool regarding clinical investigation for physicians worldwide due to their speed, convenience, and simplicity. At the beginning of the pandemic of coronavirus disease 2019 (COVID-19), there was a need for a rapid response to the medical imaging community for imaging tools that could help specialists in developing and improving techniques related to chest X-rays (CXR) image-based diagnoses [1].

Cao and Zhao [2] point out that one of the most critical radiographic analysis tasks is lung segmentation. They observed that the variations in opacity in the lungs in CXRs make it a complex region to segment. Segmentation of the region of interest (ROI) is relevant because it is a primary step in an image-based diagnostic tool and directly impacts all other steps related to more complex analyses.

Rahman et al. [3] emphasize the importance of image segmentation regarding CXRs, where the authors aim to identify tuberculosis, a chronic lung disease, by training different convolutional neural networks (CNNs). Their results show that networks trained after segmentation of the ROI outperform whole X-ray image-based classification. Teixeira et al. [4], as well as Bassi and Attux [5] demonstrate the impact of lung segmentation in COVID-19 identification using CXR images. According to the results, the classification performance in models using segmented CXR images tends to achieve better results.

Therefore, this paper proposes an automatic method to identify lung regions in CXRs, as other approaches in the literature [1]–[7].

In terms of imaging segmentation, the most recent approaches [1]–[3] use machine learning networks and achieve very high similarity coefficients when comparing outputs to ground truths (GTs).

On the other hand, as pointed out by Cao and Zhao [2], the CNN method takes a long time to train due to the complexity of its network structure. Thus, this paper aims to address a second issue: is it possible to produce similar results based on less computational cost methods? Because, besides the necessary investment in equipment (not viable to many hospitals and clinics worldwide), there is a known trade-off between machine learning (ML) algorithms and their impacts on the environment. Explicitly, Garcia et al. [8], and Strubell et al. [9] showed that ML algorithms have significantly increased world energy consumption and greenhouse gas emissions. In this sense, this paper proposes to create a ROI detection process for CXR images based on computationally less expensive methods, such as Image Analysis and Mathematical Morphology techniques.

II. RELATED WORK

Current literature presents several studies that automatically detect ROI on CXR images, mainly using variations of the CNN method [1]. For example, in [3], the authors used nine different deep CNNs for ROI detection, three comparatively shallow networks (MobileNetv2, SqueezeNet and ResNet18), and six deep networks (Inceptionv3, ResNet 50, ResNet101, CheXNet, VGG19, and DenseNet201). The method consisted of training 704 CXR images and GTs in a five-fold cross-validation model.

Similarly, in [2], the authors used more than 1,000 images considering four public datasets, including the Japanese Society of Radiological Technology (JSRT), same used in our current study, for training a U-Net model. The shuffle of the training database of CXRs with patients' radiography by using a SIFT-flow registration algorithm and evaluating similarity measures of the input image with the ones in the dataset are considered in [7].

In those studies, limitations highlighted by the authors include the amount of time for training their models and, therefore, the necessity of powerful machines. For instance, in [3], the authors used a 16 GB NVIDIA GeForce GTX 1080 GPU card for training.

TABLE I
RECENT WORKS ON LUNG SEGMENTATION AND THEIR MAIN ASPECTS

Author	Year	Method	Database	Accuracy	Dice	Jaccard
[7]	2013	Fast partial Random profile similarity selection and SIFT-flow nonrigid registration	JSRT	-	0.967	0.954
[6]	2014	Edge Detection and Morphology	JSRT	-	-	0.809
[10]	2017	Convolutional neural network (CNN)	MC and CHN	0.821	-	-
[11]	2017	CNN Transfer Learning	MC and CHN	0.847	-	-
[12]	2018	Convolutional neural network (CNN) and two ensembles	JSRT, MC and CHN	0.888	-	-
[13]	2018	Transfer learning (ResNet)	MC and CHN	0.949	-	-
[14]	2019	Transfer learning (VGG 16 model)	MC and CHN	0.800	-	-
[15]	2019	3 pre-trained CNNs	MC and CHN	0.860	-	-
[16]	2019	DeTrac Class decomposition	JSRT	0.998	-	-
[3]	2020	Nine pre-trained CNN models	CHN, MC, Belarus and RSNA	0.999	-	-
[2]	2021	Decoder-encoder with variational auto-encoder in each layer	JSRT	0.978	-	0.920
[17]	2021	Image processing and Morphology	JSRT	0.917	0.873	0.777

Another challenge faced by such studies was the necessity of a high number of images (or GTs) pre-evaluated by a specialist. Notably, inherent problems in ML models applied to medicine are related to finding complete and adequate databases [18]. Among desirable characteristics of these data sets used in the analyses, we must mention: (a) balancing - the number of samples for a positive or abnormal group is expected to be almost the same as other groups; (b) costs - related to having diagnoses from different specialists to improve generalization; (c) ethical approval - the need and importance of approval by institutional and government ethics committees to allow the use of patient data.

In terms of the imaging processing technique used for solving the lung region segmentation problem, Saad et al. [6] propose a method based on Canny filter and binarization followed by the morphological operations of dilation and erosion. They used ten images from the JSRT database and manually created their correspondent GTs. For the evaluation of the method, the authors used the Jaccard index due to its easy understanding [6]. The Jaccard index is linearly related to the significance of the relationship between the output and the GT.

In our previous work [17], we considered the usage of morphological operators applied after binarization for ten images of the JSRT database and calculated different indexes for evaluation. Different from the current method, both [6] and [17] considered morphology techniques only after the binarization of the images. In the proposed approach, morphology operations are used before and after binarization to capture more details of the gray tones and facilitate the ROI identification.

Table I summarises the methods presented in previous papers and their correspondent evaluation indexes. In this table, the database abbreviations correspond to their providers:

(1) MC: Montgomery County X-ray set from National Library of Medicine, National Institutes of Health, Bethesda,

MD, USA [7]; (2) CHN: Shenzhen Chest X-ray database from the National Library of Medicine, National Institutes of Health, Bethesda, MD, USA, and Shenzhen No.3 People's Hospital, Guangdong Medical College, Shenzhen, China [19];

(3) JSRT: Japanese Society of Radiological Technology database created by the Japanese Society of Radiological Technology (JSRT) in cooperation with the Japanese Radiological Society (JRS) [20];

(4) Belarus: Belarus set [21] from the National Institute of Allergy and Infectious Diseases, Ministry of Health, Republic of Belarus;

(5) RSNA: the Radiological Society of North America, in collaboration with the US National Institutes of Health, The Society of Thoracic Radiology, and MD.ai, developed a CXR dataset publicly available in Kaggle's machine learning community [22].

Notably, the most used image sets are the MC and CHN; both of them were provided by the National Library of Medicine (NLM) in the U.S. [23].

Given the achieved accuracy of previous works on Table I, that is about 80%; our research aims to achieve similar results without the necessity of a labeled dataset and using low-cost computational methods. For this, we are considering the usage of two datasets, totalizing 24 CXR images of patients of different ages and genders for evaluation.

III. THE PROPOSED METHOD

This section describes the threads of our approach, including the datasets used, the ROI detection, and the evaluation methods. They are implemented in Python v3.7 and use open-source libraries: Scikit Image [24], Scikit-learn [25], and Numpy [26].

A. The used dataset

The datasets' choice for the current study was based on a qualitative evaluation of our team's specialist, respecting the ones that would better compare with previous works. For

example, the MC dataset [7] contains images with different sizes and acquisition exposition times, as well as tags that would need some adjustments before passing through the hither proposed method. Therefore, this dataset was put away in the current phase of this study.

The selected images from each dataset were also picked through a qualitative approach considering the ones already adequately centralized, with tags away from the lung region.

Therefore, the image set used to evaluate the proposed method consists of 24 images, with equal numbers of images selected from two databases, i.e., 12 images from each database. The used bases are:

- Japanese Society of Radiological Technology (JSRT) [20]. These images are described by pixels represented by 12-bit stored in the IMG format, with a spatial resolution of 2048 x 2048 pixels.
- Shenzhen Chest X-ray (CHN) [19]. Its images are stored in 8-bit PNG format with a resolution of approximately 3000 x 3000 pixels.

To evaluate our approach, we opt for not using GT's automatically generated [27], i.e., not produced by a human specialist, because once they are not 100% correct, they could lead to a false evaluation perspective. Hence, the radiologist of our team manually created the GTs for the 24 images. For this, he used an iPad® (6th gen) and Apple pencil® (1st gen) with the Autodesk® SketchBook app. All images and their respective GTs are available for download at <https://github.com/mnzluiza/Lung-Segmentation>.

Regarding the JSRT dataset, there is a Lung Segmentation Data Kit [28] which contains GT's provided by specialists. In this case, we evaluated the proposed method considering both the GT's created by our team's specialist and the CSIRO's one.

B. ROI detection

In a general perspective, the method's main idea is to define areas in the CXR that are not relevant for the ROI, i.e., the trachea, any part out of the patient's body, and some other neck region outliers. Once these areas are defined, the method eliminates them from the input image and applies some exposure corrections. The resulting image is binarized, and some morphology operators remove small disconnected areas and smooth the output's edge.

In this sense, the proposed ROI detection method consists of four main steps:

- (1) Input: load input image,
- (2) Identify pitfalls: define regions in which the grey tone is similar to the ones in the ROI,
- (3) Adjust and smooth tones, and binarize: adjust input image (1) exposure after elimination of the pitfalls identified in (2), and binarize it,
- (4) Apply morphology: apply morphological operations in (3) to remove noise and promote impaired areas caused by the binarization.

These steps are combining as shown in the schematic flow in Figure 1. Each of these steps consists of multiple tasks explained in the following subsections.

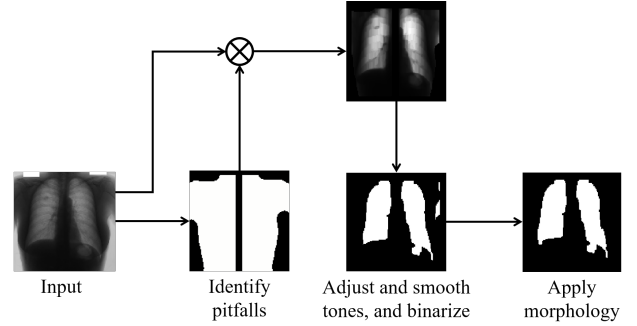


Fig. 1. ROI Detection Method Workflow

1) *Input*: Once the method reads the image, it treats it internally as an unsigned integer array for better performance. The image pixels content is normalized in a way that all its tones are between 0 and 1. The spatial resolution is re-scaled to 128 x 128 pixels independently of its original pixel number in each direction. The choice of these dimensions considers practical experience, in which we tested the code for different dimensions (from 2^6 to 2^{10}), and the one that produced the best results was the one aforementioned. Even though it may sound counterintuitive, reducing the size of the image to an optimal point produces smoother edges for the ROI, helps mitigating outliers, and improves the algorithm's performance.

2) *Identify pitfalls*: This step is composed of multiple tasks, as shown in Figure 2.

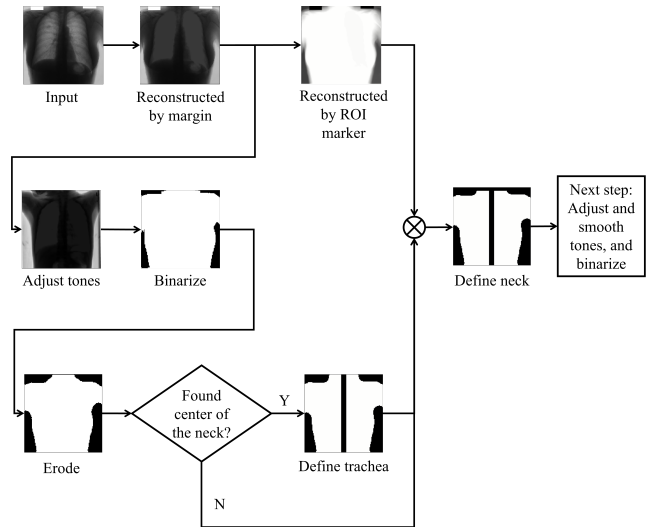


Fig. 2. Step: Identification of Pitfalls

The first tasks involve the identification of areas out of the patient's body. In this sense, a margin marker, with 1px width and tonality value of 1, morphologically reconstructs by dilation the input image. As observed in Figure 2 (in the image named Reconstructed by margin), all the area correspondent to

the patient’s body turns into darker pixels. This operation is available in the Scikit-image package [29], and it is based on the algorithm proposed by Robinson et al. [30]. The reconstruction function connects regions in the input image marked by local maxima in the defined marker image. Thus, pixels in the input image that are correspondent neighbors of the marker image and whose tonality is lesser than or greater than the marker pixels become highly connected in tonality.

Next, the grey tone image is adjusted through a sigmoid to enhance its contrast and facilitate binarization. This function transforms the input image according to the equation $1/(1 + \exp * (gain * (cutoff - I)))$, which is applied pixel by pixel [31], where *gain* was set to 5 and *cutoff* 0.7. Sequentially, the image is normalized and binarized considering a simple threshold of 50% of the image’s maximum tonality.

Once these parameters are relative to image tones that were previously normalized, their application can be generalized to various CXR images considering a standard acquisition process. Analogously, the size of the structuring elements of morphological operations can also be generalized since all images are previously re-scaled to the same resolution (128 x 128 pixels).

In the sequence, the morphological operation *remove small holes* followed by an erosion of a disk structuring element (7px diameter) is applied to reduce any slight noise caused by mismatched tones on the binarization. The *remove small holes* function is part of the Scikit-image package [32] and sets areas in the image that are smaller than a specified size to zero. In this case, areas that were smaller than 64px, i.e., half of the image’s width, were set to zero.

Considering that the patient’s silhouette is now adequately defined, it is possible to identify the trachea. That is relevant because this region may contain brighter tones that interfere with later imaging binarization. Besides, it helps to isolate the left and right lungs properly. The task of identifying the trachea consists of a sweep on the first row of the silhouette mask starting from both left and right sides until reaching a foreground pixel. With the lateral extremities of the neck defined, the trachea is expected to be centered between them. Thus, the algorithm sets a line of zeros from top to bottom of the image, removing any vestige of the trachea and vertebral column. The resulting line combined with the silhouette mask is in Figure 2 in the Define trachea step.

Nonetheless, the algorithm does not consider eliminating these regions if the center of the binarized silhouette is away from the image center more than 20% of the image’s width. This problem typically occurs when patients’ arms are upwards, or tags are located in the upper part of the image. Figure 3(b) illustrates an example in which it is not possible to identify the trachea region because there is an upward arm represented by some foreground pixels on the upright corner of the binarized image. In this case, the algorithm does not remove the trachea and vertebral column. In contrast, for the image that we are using to illustrate our method (JPCLN006), as shown in Figure 3(a), the neck silhouette is adequately defined; hence the algorithm proceeds eliminating the trachea

region.




Fig. 3. Cases in which is possible (a) and not possible (b) to identify trachea and vertebral column region based on first binarization

The following tasks of Identifying pitfalls step consist of removing other noises that appear in the neck region. Therefore, the marker shown in Figure 4 morphologically reconstructs the input image [30] already reconstructed by the margin to amplify the tone differences between the silhouette and other regions. The position of the marker considers the expected position of most silhouettes after observation of CXR images.




Fig. 4. Reconstruction by ROI Marker

After reconstruction, the image’s vertical axis’ histogram projection is calculated individually for both sides of the lungs in the fourth top portion. As shown in Figure 5, these curves begin relatively stable but going slightly down. Once they start rising, it means they reached the lung region. Therefore, any portion corresponding to the top image until this point does not belong to the ROI. For better generalization, the limit position is the average for both left and right sides.




Fig. 5. Vertical axis’ histogram projection for both sides of the lungs

Figure 2 (step: Define neck) shows the combination of this approach and the elimination of the other pitfalls aforementioned.

3) *Adjust and smooth tones, and binarize*: The next step is to binarize the image. Nonetheless, even after removing pitfalls, it is important to adjust some grey tone levels to enhance the contrast of the ROI due to the noisy nature of CXR images. We did numerous experiments considering different approaches for this particular step. The one that produced the best results is in Figure 6.




Fig. 6. Step: Adjusting and smoothing grey tones, and binarization

Instead of enhancing brighter tones of the ROI, we opt to enhance darker areas' values. The first task of this step consists of inverting the input image, adjusting it by a sigmoid function with a 0.6 cutoff, and morphologically closing it by a rectangle element with sides 7 and 1.

The next task is to revert the image to its original background (darker) and foreground (brighter) tone scheme and remove any values inside a 10px edge of the image.

Finally, the image's binarization task is done through local adaptive thresholding. Each pixel receives a threshold value based on the sum of its neighborhood (a block with size $2*128+1$ px) values weighted by a Gaussian window [33].

4) *Apply morphology*: After binarization, some remnants out of the ROI are still visible. Thus, a simple opening morphological operation is applied. Additionally, some images presented small portions of areas inside the ROI classified as background, but the remove small holes function of the Scikit-image package [32] correctly addressed this issue.

The last task of this step consists of applying a conditional dilation using a structuring element of 7×3 . The dilation occurs on the first-fifth portion of each side of the three-quarters superior area, as indicated in Figure 7.

The use of conditional dilation is because the cardiac region usually is already well defined by the binarization process. Therefore, dilating applied to the entire image would enhance the differences between the GT and the method results. Nonetheless, depending on the objective of using this method, it might be better to consider removing the condition to guarantee that the total ROI is in the dilation process (enhancing false-positive results).

IV. RESULTS AND DISCUSSIONS

To validate our approach and provide a better scale for comparison with previously mentioned studies, we calculated the following indexes for the 24 selected images: $Accuracy = \frac{TP+TN}{TP+TN+FP+FN}$, $Jaccard = \frac{TP}{TP+FP+FN}$, and $Dice = \frac{2TP}{2TP+FP+FN}$.




Fig. 7. Indication (in orange) of the conditional region of dilation

TP, TN, FP, and FN represent True Positive, True Negative, False Positive, and False Negative values. In Figures 8, 9 and 10 (d), TP, TN, FP, and FN are represented in white, black, grey, and yellow, respectively. The average indexes for each database are in Table II.

TABLE II
AVERAGE RESULTS FOR THE 24 IMAGES (12 FOR EACH OF THE TWO DATASETS: JSRT AND CHN) - ACCURACY, JACCARD AND DICE INDEXES

Database	Accuracy	Jaccard	Dice
JSRT	0.9355	0.8078	0.8934
CHN	0.9368	0.7463	0.8539
All	0.9361	0.7771	0.8737

The results abovementioned consider the GT's created by our team's specialist. Therefore, the results are not precisely comparable with previous works. In this sense, we also considered the GT's provided by CSIRO [28] for the JSRT database. The results in Table III compare the same indexes of Table II for both GT's regarding the same 12 images. This evaluation demonstrates that the GT's provided by our specialist is similar to those provided by CSIRO.

TABLE III
AVERAGE RESULTS FOR THE 24 IMAGES (12 FOR EACH OF THE TWO DATASETS: JSRT AND CHN) - ACCURACY, JACCARD AND DICE INDEXES

GT source	Accuracy	Jaccard	Dice
Our specialist	0.9355	0.8078	0.8934
CSIRO	0.9242	0.7680	0.8683

Figure 8 indicates that in terms of image dimensions, CHN CXRs are composed of smaller lungs than the JSRT database. Nonetheless, both databases presented similar results of around 7% of the image area in terms of misidentification. It is notable that for the CHN, most of the misidentifications happened during the classification of the false-positive regions. That is an essential consideration in medical approaches because the method prevails in considering a region as part of the ROI rather than declassifying it.

Regarding the Jaccard and Dice indexes calculated in Table II, their interpretation relies on knowing that they weigh the proportion of negative (background) and positive (foreground)




Fig. 8. Confusion Matrix Average values

values. In other words, images with small amounts of foreground pixels, such as the ones from the CHN database, will have smaller values, even though the proportion of identified pixels for the ROI is the same.

Thus, to evaluate how adequate our method is regarding the relevant area, i.e., the lung region, we propose using the specificity measure. Hence, we calculate the proportion of background pixels that are correctly classified, i.e., $\frac{TN}{TN+FP}$. By doing that, we can say that our approach classified a non ROI area correctly 92.92% on average, as shown in Table IV. The opposite of the specificity is the sensitivity measure, i.e., $\frac{TP}{TP+FN}$, and its values are also in Table IV. One relevant issue regarding this measure is that if all pixels are considered positive values, the sensitivity will always be 100%.

Another interesting evaluation is the amount of time that




Fig. 9. (a) Original, (b) Results, (c) GTs developed by our team and (d) Difference between result and GT for each image of the JSRT Database

our method demands. Cao and Zhao's [2] method presented a specificity of 0.988 for the JSRT dataset with an average running time of 44.38 ms. For the same dataset, our method's specificity is 0.934 and takes an average of 0.0006 ms to run the whole process (including the loading) for one image. Considering the amount of time to generate the ROI mask for a




Fig. 10. (a) Original, (b) Results, (c) GTs developed by our team and (d) Difference between result and GT for each image of the CHN Database

human, the radiologist of our team took on average 40 minutes to produce each of the GTs used to evaluate our method.

In our previous work [17], we also considered imaging processing techniques and morphological operations to address this issue. Nonetheless, in the current study, using morphological operators in grayscale and removing noise areas before

TABLE IV
AVERAGE RESULTS FOR THE 24 IMAGES (12 FOR EACH OF THE TWO DATASETS: JSRT AND CHN) - SPECIFICITY AND SENSITIVITY

Database	Specificity	Sensitivity
JSRT	0.9340	0.9398
CHN	0.9244	0.9878
All	0.9292	0.9638

binarizing the image improved the method's performance. In [17], we obtained a value of 0.777 for the Jaccard index for the JSRT database. The adjustments on the approach proposed here achieved 0.808.

V. CONCLUSION AND FUTURE IMPROVEMENTS

This paper presents a method based on image processing techniques for ROI segmentation in CXR images. This method serves as a pre-processing step for more complex analysis, which involves diagnosis classification problems, such as identifying tuberculosis, COVID-19, and other lung diseases. Furthermore, providing a less costable computational process without needing datasets to be trained allows the method to be more accessible and agile.

To evaluate the method, our teams' radiologist manually created the GTs for the 24 images selected, which are available for download at <https://github.com/mnzluiza/Lung-Segmentation>.

Regarding the JSRT dataset, since CSIRO's specialists provided an available GT's database [28], we also considered an evaluation with those GT's and obtained similar results.

In terms of future improvements for the presented approach, based on our experience, we can state that the lateral lung regions are visibly more challenging to delimit once they present a high overlap of adjacent tissue, including ribs edges, so this aspect could be the goal of other procedure to be designed. Also, the proposed method is highly dependent on the adequate centralization of the input image. Therefore, future works related to adjusting and rotating not centralized images and removing patients' tags (when necessary) are highly relevant for greater generalization of the ROI definitions method for lung regions.

Other future analyses include a qualitative approach from a medical group to evaluate the segmentation results. Furthermore, identifying opportunities for developing apps or software functionalities based on the proposed method. For example, to evaluate this method usability to calculate the cardiothoracic ratio in the clinical context, i.e., estimate the patient's heart size. Finally, future works related to the diagnosis of pulmonary pathologies may include more sophisticated imaging methods [34]–[36].

ACKNOWLEDGMENT

L.C.M. and A.R. thank CAPES Brazilian Foundation for supporting their graduation courses. A.C. is partially supported by MACC-INCT, CNPq Brazilian Agency (402988/2016-7

and 305416/2018-9) and FAPERJ (projects SIAD-2, e-Health Rio and Digit3D).

REFERENCES

- [1] F. Shi, J. Wang, J. Shi, Z. Wu, Q. Wang, Z. Tang, K. He, Y. Shi, and D. Shen, "Review of artificial intelligence techniques in imaging data acquisition, segmentation and diagnosis for covid-19," *IEEE reviews in biomedical engineering*, 2020.
- [2] F. Cao and H. Zhao, "Automatic lung segmentation algorithm on chest x-ray images based on fusion variational auto-encoder and three-terminal attention mechanism," *Symmetry*, vol. 13, no. 5, p. 814, 2021.
- [3] T. Rahman, A. Khandakar, M. A. Kadir, K. R. Islam, K. F. Islam, R. Mazhar, T. Hamid, M. T. Islam, S. Kashem, Z. B. Mahbub *et al.*, "Reliable tuberculosis detection using chest x-ray with deep learning, segmentation and visualization," *IEEE Access*, vol. 8, pp. 191 586–191 601, 2020.
- [4] L. O. Teixeira, R. M. Pereira, D. Bertolini, L. S. Oliveira, L. Nanni, G. D. Cavalcanti, and Y. M. Costa, "Impact of lung segmentation on the diagnosis and explanation of covid-19 in chest x-ray images," *arXiv preprint arXiv:2009.09780*, 2020.
- [5] P. R. Bassi and R. Attux, "Covid-19 detection using chest x-rays: is lung segmentation important for generalization?" *arXiv preprint arXiv:2104.06176*, 2021.
- [6] M. N. Saad, Z. Muda, N. S. Ashaari, and H. A. Hamid, "Image segmentation for lung region in chest x-ray images using edge detection and morphology," in *2014 IEEE International Conference on Control System, Computing and Engineering (ICCSCE 2014)*. IEEE, 2014, pp. 46–51.
- [7] S. Candemir, S. Jaeger, K. Palaniappan, J. P. Musco, R. K. Singh, Z. Xue, A. Karargyris, S. Antani, G. Thoma, and C. J. McDonald, "Lung segmentation in chest radiographs using anatomical atlases with nonrigid registration," *IEEE transactions on medical imaging*, vol. 33, no. 2, pp. 577–590, 2013.
- [8] E. García-Martín, C. F. Rodrigues, G. Riley, and H. Grahn, "Estimation of energy consumption in machine learning," *Journal of Parallel and Distributed Computing*, vol. 134, pp. 75–88, 2019.
- [9] E. Strubell, A. Ganesh, and A. McCallum, "Energy and policy considerations for deep learning in nlp," *arXiv preprint arXiv:1906.02243*, 2019.
- [10] R. Hooda, S. Sofat, S. Kaur, A. Mittal, and F. Meriaudeau, "Deep-learning: A potential method for tuberculosis detection using chest radiography," in *2017 IEEE International Conference on Signal and Image Processing Applications (ICSIPA)*. IEEE, 2017, pp. 497–502.
- [11] U. Lopes and J. F. Valiati, "Pre-trained convolutional neural networks as feature extractors for tuberculosis detection," *Computers in biology and medicine*, vol. 89, pp. 138–143, 2017.
- [12] L. G. C. Evalgelista and E. B. Guedes, "Computer-aided tuberculosis detection from chest x-ray images with convolutional neural networks," in *Anais do XV Encontro Nacional de Inteligência Artificial e Computacional*. SBC, 2018, pp. 518–527.
- [13] O. Yadav, K. Passi, and C. K. Jain, "Using deep learning to classify x-ray images of potential tuberculosis patients," in *2018 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. IEEE, 2018, pp. 2368–2375.
- [14] M. Ahsan, R. Gomes, and A. Denton, "Application of a convolutional neural network using transfer learning for tuberculosis detection," in *2019 IEEE International Conference on Electro Information Technology (EIT)*. IEEE, 2019, pp. 427–433.
- [15] A. Hernández, Á. Panizo, and D. Camacho, "An ensemble algorithm based on deep learning for tuberculosis classification," in *International Conference on Intelligent Data Engineering and Automated Learning*. Springer, 2019, pp. 145–154.
- [16] A. Abbas, M. M. Abdelsamea, and M. M. Gaber, "Detrac: Transfer learning of class decomposed medical images in convolutional neural networks," *IEEE Access*, vol. 8, pp. 74 901–74 913, 2020.
- [17] A. R. V. F. Araújo, L. C. Menezes, and A. Conci, "Segmentação das áreas pulmonares em radiografias torácicas digitais," in *Anais do XXI Simpósio Brasileiro de Computação Aplicada à Saúde*. SBC, 2021.
- [18] X. Wan, J. Liu, W. K. Cheung, and T. Tong, "Learning to improve medical decision making from imbalanced data without a priori cost," *BMC medical informatics and decision making*, vol. 14, no. 1, pp. 1–9, 2014.
- [19] S. Jaeger, A. Karargyris, S. Candemir, L. Folio, J. Siegelman, F. Callaghan, Z. Xue, K. Palaniappan, R. K. Singh, S. Antani *et al.*, "Automatic tuberculosis screening using chest radiographs," *IEEE transactions on medical imaging*, vol. 33, no. 2, pp. 233–245, 2013.
- [20] J. Shiraishi, S. Katsuragawa, J. Ikezoe, T. Matsumoto, T. Kobayashi, K.-i. Komatsu, M. Matsui, H. Fujita, Y. Kodera, and K. Doi, "Development of a digital image database for chest radiographs with and without a lung nodule: receiver operating characteristic analysis of radiologists' detection of pulmonary nodules," *American Journal of Roentgenology*, vol. 174, no. 1, pp. 71–74, 2000.
- [21] N. I. of Allergy and I. Diseases, "Belarus tuberculosis portal," 2021. [Online]. Available: <https://tbportals.niaid.nih.gov/>
- [22] N. Pandey, "Rsna pneumonia detection challenge," 2018. [Online]. Available: <https://www.kaggle.com/c/rsnapneumonia-detection-challenge/data>
- [23] S. Jaeger, S. Candemir, S. Antani, Y.-X. J. Wang, P.-X. Lu, and G. Thoma, "Two public chest x-ray datasets for computer-aided screening of pulmonary diseases," *Quantitative imaging in medicine and surgery*, vol. 4, no. 6, p. 475, 2014.
- [24] S. Van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, and T. Yu, "scikit-image: image processing in python," *PeerJ*, vol. 2, p. e453, 2014.
- [25] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay, "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [26] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant, "Array programming with NumPy," *Nature*, vol. 585, no. 7825, pp. 357–362, Sep. 2020. [Online]. Available: <https://doi.org/10.1038/s41586-020-2649-2>
- [27] N. Pandey, "Chest xray masks and labels," 2019. [Online]. Available: <https://www.kaggle.com/nikhilpandey360/chest-xray-masks-and-labels>
- [28] F. Rusak, D. Wang, and Y. Arzhaeva, "Lung segmentation data kit. v1," in *Data Collection*. CSIRO, 2018. [Online]. Available: <https://doi.org/10.25919/5c49548be0551>
- [29] Scikit-image, "Module: Morphology - reconstruction," 2021. [Online]. Available: <https://scikit-image.org/docs/dev/api/skimage.morphology.html#skimage.morphology.reconstruction>
- [30] K. Robinson and P. F. Whelan, "Efficient morphological reconstruction: a downhill filter," *Pattern Recognition Letters*, vol. 25, no. 15, pp. 1759–1767, 2004.
- [31] Scikit-image, "Module: Exposure - adjust sigmoid," 2021. [Online]. Available: https://scikit-image.org/docs/dev/api/skimage.exposure.html#skimage.exposure.adjust_sigmoid
- [32] —, "Module: Morphology - remove small holes," 2021. [Online]. Available: https://scikit-image.org/docs/0.13.x/api/skimage.morphology.html#skimage.morphology.remove_small_holes
- [33] —, "Module: Filters - threshold local," 2021. [Online]. Available: https://scikit-image.org/docs/stable/api/skimage.filters.html#skimage.filters.threshold_local
- [34] B. Barros, P. Lacerda, C. Albuquerque, and A. Conci, "Pulmonary covid-19: Learning spatiotemporal features combining cnn and lstm networks for lung ultrasound video classification," *Sensors*, vol. 21, no. 16, pp. 5486–5523, 2021.
- [35] P. Lacerda, B. Barros, C. Albuquerque, and A. Conci, "Hyperparameter optimization for covid-19 pneumonia diagnosis based on chest ct," *Sensors*, vol. 21, no. 6, pp. 2174–2190, 2021.
- [36] P. Lacerda, J. Gonzalez, N. Rocha, F. Seixas, C. Albuquerque, E. Clua, and A. Conci, "A parallel method for anatomical structure segmentation based on 3d seeded region growing," in *2020 International Joint Conference on Neural Networks (IJCNN)*, 2020, pp. 1–6.