Pleural Effusion Classification on Chest X-Ray Images with Contrastive Learning

Felipe André Zeiser[1][a], Ismael Garcia Santos[1], Henrique Christoph Bohn[1], Cristiano André da Costa[1][b], Gabriel de Oliveira Ramos[1][c], Rodrigo da Rosa Righi[1][d], Andreas Maier[2][e], José Rodrigo Mendes Andrade[3] and Alexandre Bacelar [3]

[1]*Software Innovation Laboratory - SOFTWARELAB, Universidade do Vale do Rio dos Sinos - Unisinos, São Leopoldo, Brazil*

[2]*Department of Computer Science, Friedrich-Alexander-Universität, Erlangen-Nürnberg, Erlangen, Germany*

[3]*Medical Physics and Radioprotection Service, Hospital de Clinicas de Porto Alegre, Porto Alegre, Brazil*

Keywords: Contrastive Learning, X-Ray, Pleural Effusion.

Abstract: Diagnosing pleural effusion is important to recognize the disease's etiology and reduce the length of hospital stay for patients after fluid content analysis. In this context, machine learning techniques have been increasingly used to help physicians identify radiological findings. In this work, we propose using contrastive learning to classify chest X-rays with and without pleural effusion. A model based on contrastive learning is trained to extract discriminative features from the images and reports to maximize the similarity between the correct image and text pairs. Preliminary results show that the proposed approach is promising, achieving an AUC of 0.900, an accuracy of 86.28%, and a sensitivity of 88.54% for classifying pleural effusion on chest X-rays. These results demonstrate that the proposed method achieves comparable or superior to state of the art results. Using contrastive learning can be a promising alternative to improve the accuracy of medical image classification models, contributing to a more accurate and effective diagnosis.

1 INTRODUCTION

Early pleural effusion detection is crucial for recognizing the etiology of the adjacent disease, choosing the ideal treatment, and reducing the worsening of the patient's health status (Hallifax et al., 2017). Pleural effusion is a fluid accumulation in the space between the parietal and visceral pleura. Various diseases and conditions, including infections, neoplasms, heart failure, and chest trauma, can cause pleural effusion. According to the literature, the late pleural effusion diagnosis may be associated with worsening the patient's health status or lead to other complications, such as respiratory failure, sepsis, and even death (Aboudara and Maldonado, 2019).

Chest radiography is an easily accessible tool to detect pleural effusion. However, this finding can be challenging in cases of small fluid volume, as the exam can detect volumes above 200 milliliters (Jany and Welte, 2019). In this way, machine learning can potentially contribute to early pleural effusion detection on chest X-rays. In addition, machine learning models can analyze large chest X-rays datasets to identify patterns and characteristics associated with the presence of chest X-ray findings (Bustos et al., 2020; Liz et al., 2023). Therefore, it is possible to develop automated classification models that can identify cases of pleural effusion, generating a double reading tool.

One of the main challenges in developing machine learning models to detect chest X-ray findings is the availability of previously labeled datasets (Cohen et al., 2020). This scenario involves the arduous process of identifying and annotating patient chest images by experienced radiologists (Bustos et al., 2020). This absence of labeled data can limit the effectiveness of machine learning models. Over the past few years, works in the literature have relied extensively

[a] https://orcid.org/0000-0002-1102-7722
[b] https://orcid.org/0000-0003-3859-6199
[c] https://orcid.org/0000-0002-6488-7654
[d] https://orcid.org/0000-0001-5080-7660
[e] https://orcid.org/0000-0002-9550-5284

Zeiser, F., Santos, I., Bohn, H., da Costa, C., Ramos, G., Righi, R., Maier, A., Andrade, J. and Bacelar, A.
Pleural Effusion Classification on Chest X-Ray Images with Contrastive Learning.
DOI: 10.5220/0012205900003584
In *Proceedings of the 19th International Conference on Web Information Systems and Technologies (WEBIST 2023)*, pages 399-405
ISBN: 978-989-758-672-9; ISSN: 2184-3252

on pre-trained models on natural image datasets, such as for the ImageNet dataset (Baltruschat et al., 2019; Zeiser et al., 2021; Guan et al., 2021). However, recent studies have demonstrated problems in capturing specific medical domain representations using pre-trained weights (Azizi et al., 2021).

Recent advances address the use of self-supervision methods to reduce model dependencies on large datasets (Chen et al., 2020; Azizi et al., 2021; Zhang et al., 2022). Generally, the work methodology is based on a pre-training phase on weakly labeled datasets, and then a fine-tuning of the weights is performed on a fully labeled dataset to improve the model's capabilities. However, these methodologies only have the ability to predict explicitly annotated findings in the dataset, which still does not solve the need to perform large-scale manual annotations for less specific findings.

The availability of chest X-ray images and reports represents a rich source of information for machine learning models. However, these data cannot be used directly in supervised learning models without the production of explicit labels by radiologists. In this context, this article presents a contrastive learning model to classify chest X-ray images without explicit annotations during the training phase. The main objective is to explore the contrastive learning effectiveness in detecting pleural effusion in chest X-rays in a context with the scarcity of labeled datasets. Our method uses contrastive learning to identify key features associated with pleural effusion. In addition, this method can capture more complex characteristics by analyzing natural language present in medical reports and learning through natural supervision. Adopting contrastive learning provides a more efficient approach and reduces development costs since it eliminates the need for extensive data labeling. The results of our approach demonstrate that contrastive learning can lead to more accurate and robust classification models, even when the labeled dataset is limited or incomplete (Radford et al., 2021). Our main contributions are listed below:

- We propose a neural network model for weakly supervised pleural effusion classification on chest X-ray images. The model can identify and associate the images and reports characteristics, optimizing the similarity of the characteristics vector from two encoders based on Transformers.
- Our experiments demonstrate consistent performance improvements in pleural effusion detection. Furthermore, the results outperform methods based on supervised learning, suggesting that explicit labels are not necessary for pleural effusion detection in chest X-ray images.

The remainder of this paper is organized as follows. Section 2 presents the most significant related works to define the present study. Next, section 3 presents the methodology of the work. The section 4 details the results and discussion. Finally, Section 5 presents the conclusions of the work.

2 RELATED WORK

The detection of findings in chest X-rays plays a crucial role in the diagnosis and treatment of various diseases. However, this task heavily relies on the expertise and experience of radiologists, making it susceptible to errors and human limitations. In recent years, significant progress has been made in the field of artificial intelligence, particularly in the development of deep learning algorithms for medical image analysis. Supervised models have been extensively explored to aid in the analysis of chest X-ray images, with several studies focusing on this area (Alghamdi et al., 2021; Bustos et al., 2020; Elgendi et al., 2021; Han et al., 2021; Rajpurkar et al., 2017; Wang et al., 2017). These models have shown promising results in assisting radiologists and improving the accuracy and efficiency of chest X-ray interpretation.

Supervised learning approaches for classification based on radiological findings dependent on large data sets are frequent in the current literature. For example, (Rajpurkar et al., 2017) developed a machine learning model that achieved similar accuracy to a suite of radiologists using a labeled dataset of over 100,000 images. (Cohen et al., 2020) investigates the performance of models in different domains and the agreement between them in the X-ray images. The authors perform experiments to analyze concept similarity, using a regularization approach in a network to group tasks on different datasets, and observe the variation between tasks. One of the alternatives to increase the number of samples for training is data augmentation (Elgendi et al., 2021). However, data augmentations commonly applied to tasks in the natural imagery domain cannot be fully extended to the health domain. Therefore, labeled datasets are essential for training machine learning models for X-ray images (Goceri, 2023).

Obtaining labeled datasets can be complex and laborious, especially in resource-poor contexts (Bustos et al., 2020). In addition, the scarcity of labeled data can lead to less accurate models and impair the effectiveness of detecting radiological findings in the evaluated exams (Alghamdi et al., 2021). Some recent work has explored the use of unlabeled or partially labeled datasets to improve the detection of consolidation on chest X-ray images. For example, (Wang et al., 2017) used semi-supervised learning to improve exam consolidation detection.

Another promising approach is contrastive learning, which can be applied to unlabeled or partially labeled datasets (Zhang et al., 2022). This feature uses the similarity and differences between pairs of images to learn discriminative characteristics (Chen et al., 2020). Therefore, the method is beneficial when small-scale labeled data is available. For example, (Han et al., 2021) used contrastive learning in conjunction with transfer learning to improve consolidation detection on chest X-rays.

However, works in the current literature are primarily based on fully labeled datasets or require partially labeled datasets in the training process. Therefore, this work proposes to investigate the effectiveness of contrastive learning to improve the detection of pleural effusion in chest X-rays in the context of training models with weak-labeled datasets.

3 MATERIALS AND METHODS

An overview of the methodology used in this work is presented in Fig. 1. Our methodology can be divided into four main steps: pre-processing, data augmentation, training, and testing. Pre-processing consists of resizing the image, normalizing the contrast, and processing the report (Section 3.2). The data augmentation step describes the methods used to generate synthetic images (Section 3.3). Finally, in the training stage, the models and parameters used are defined (Section 3.4).

3.1 Dataset

We used three datasets to develop the model. For the training step, we retrospective collected data from patients during the COVID-19 pandemic from two hospitals in Porto Alegre/RS. The information collected comprises clinical data, X-ray images and reports of hospitalizations from 2020 to 2022. The ethics committee of each hospital approved the present study under the Certificate of Submission for Ethical Consideration (CAAE number 33540520.6.3004.5327) . In addition, this document follows the General Data Protection Law (LGPD) recommendations. In Tab. 1, the information relating to Private datasets 01 and 02 is provided.

We used the PADCHEST public dataset for the testing step. PADCHEST consists of chest X-rays with reports. The dataset includes more than 160,000 images of 67,000 patients that were interpreted and reported by radiologists at Hospital San Juan (Spain) from 2009 to 2017, encompassing six views of different incidences, additional information on image acquisition, and patient demographics. The dataset comprises 19 differential diagnoses, with 27% of reports manually annotated by physicians. In this sense, we used only manually annotated X-ray images. In Tab. 1, we present the information for each dataset.

3.2 Pre-Processing

We processed the images from the three datasets using the same methodology. Given the tenuous differences between healthy tissues and those affected by any alteration, contrast enhancement is commonly used in the literature (Diniz et al., 2018). Furthermore, contrast enhancement can help increase the performance of deep learning architectures (Pooch et al., 2020). Therefore, we applied the Contrast-Limited Adaptive Histogram Equalization (CLAHE) in the chest X-ray images. CLAHE subdivides the image into subareas using interpolation between edges. To avoid noise build-up, use a gray threshold level, redistributing pixels above that threshold in the image. The CLAHE can be defined by:

$$p = [p_{max} - p_{min}] * G(f) + p_{min} \tag{1}$$

where p is the pixel's new gray level value, the values p_{max} and p_{min} are the pixels with the lowest and highest values in the neighborhood, and $G(f)$ corresponds to the cumulative distribution function (Zuiderveld, 1994).

Human anatomy is unique to each individual. This aspect is reflected in the different anatomical sizes of patients' chests. In this way, the X-ray images have different sizes. However, to process the X-ray images in deep learning models, the images need a standard width and height. Therefore, we resized the images to $512X512$ pixels. The reduction was proportional to the width and height, with a zero padding for the axis with the smallest size. The pre-trained architecture delimited the image resolution, which will be discussed in Section 3.4. In Fig. 2, we show an example with the original and preprocessed X-ray images. We can observe (Fig. 2) that the image suffered a downsampling changing from approximately $2000X3000$ pixels to $512X512$ pixels. In addition, the area related to the lungs is highlighted about the other structures represented in the image.

Finally, we pre-processed the Private datasets 01 and 02 X-ray reports. We keep only the alpha characters, removing special characters, numbers, and punctuation from the reports. For the PADCHEST dataset, it was only necessary to translate the pleural effusion

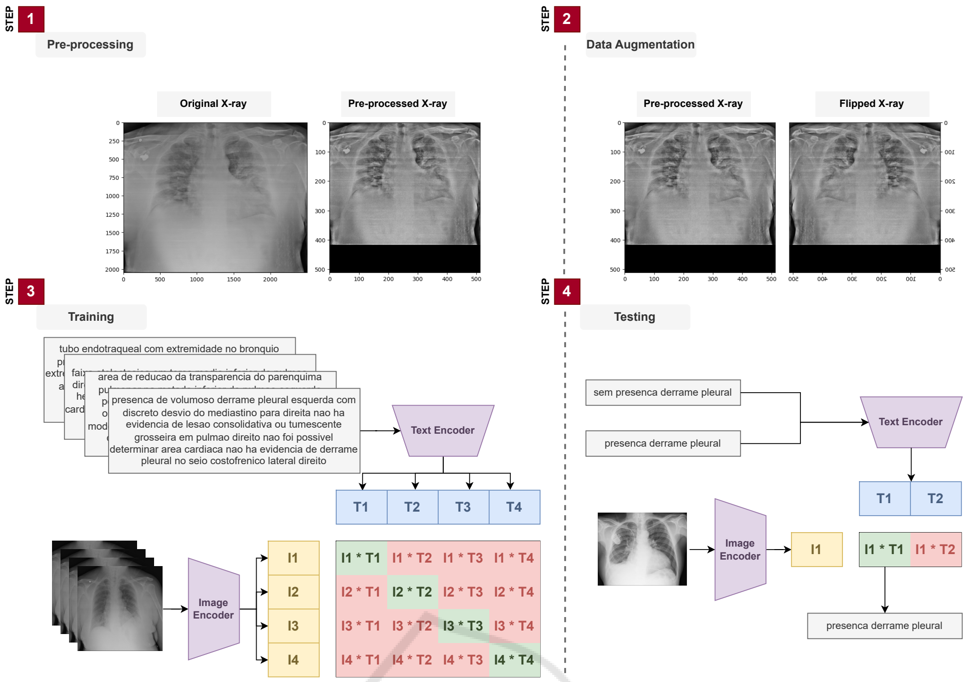

Figure 1: Summary of the proposed methodology. The text and image encoders are trained together to predict the correct pairs from a training batch. The text encoder is fed with generic text, and the image encoder with the chest X-ray for prediction.

Table 1: Chest X-ray datasets for training and testing the architecture. Total.: Total; Image: Image; Man.: Manual; and Aut.: Automatic.

Dataset	**Patients**	**Annotation**	**Label**	**Tot. Img.**	**Img. Used**	**Used in**
Private 01	697	Man.	Report	6.650	6.650	Training
Private 02	1.754	Man.	Report	7.084	7.084	Training
PADCHEST	69.882	Man./Aut.	Report and Findings	160.000	17.513	Teste
Total	**72.333**	**-**	**Report and Findings**	**-**	**173.734**	**-**

finding from Spanish into Portuguese, which was performed by automatic substitution.

3.3 Data Augmentation

We used a data augmentation methodology to increase the cases in the training set. In this context, we applied a horizontal flip to each training set image. The choice of only one data augmentation method is related to the image characteristics. X-ray exams represent anatomical structures that can be missed or distorted by other data augmentation methods such as distortion, shearing, or cropping.

3.4 Contrastive Learning Architecture

Current models used in the field of computer vision require large datasets labeled (Dosovitskiy et al., 2020). Specifically, in the context of imaging studies, obtaining large labeled datasets is a significant challenge. In this way, methods that can learn from already established information may be able to scale with greater generalization for the day-to-day use of radiologists. In this sense, recent advances in contrastive learning methods present a possibility of learning representations in images using natural language (Radford et al., 2021; Zhang et al., 2022).

In this way, our architecture shown in Fig. 1 is based on Connecting text and images (CLIP) (Radford et al., 2021). The CLIP architecture is based on the idea of zero-shot learning, seeking to find relationships between image and text pairs. This association is built using two encoders, one responsible for extracting features representing the image and another for the text. As a result, the model's objective function is that both encoders produce feature vectors for similar images and texts.

3.4.1 Training

The parameters for the encoders followed the CLIP architecture (Radford et al., 2021). The image encoder is based on the ViT-B/32 architecture pre-

Figure 2: Example of an original chest X-ray image (a) and pre-processed (b).

trained with CLIP model weights. A transformer architecture with 12 layers of size 512 and 8 attention heads was developed for the text encoder. The maximum length of the text string was 76 characters. The implementation of the deep learning model followed the hyper-parameters of the CLIP model.

The model was trained for 100 epochs using the Adam optimizer with an adaptive learning rate of 0.00001. The batch size was 32 images. The model was trained with the loss function of cosine similarity, maximizing the similarity between the correct image and text pairs.

3.4.2 Evaluation

For model evaluation, each image is processed with two synthetic texts in Portuguese: (a) presence of pleural effusion; and (b) without presence of pleural effusion. First, the respective encoders calculate the feature vector for the image and a feature vector for each synthetic text. Then, the cosine similarity is calculated with the normalization of the result performed using a softmax function. The output of the softmax function then represents the binary classification for the presence or absence of pleural effusion on the chest X-ray. We used the area under the receiver operating characteristic (AUC), precision/recall curve, accuracy, sensitivity, specificity, and F1-score as performance metrics for evaluating the model. In addition, we performed a bootstrap to generate the performance metrics confidence intervals.

4 RESULT AND DISCUSSION

This section presents the results of the proposed method and the comparison with the current literature for the detection of pleural effusion. The best weights were chosen automatically based on the validation set error. Tab. 2 presents the performance obtained for the evaluation metrics in the PADCHEST dataset and the respective comparison with the current literature.

Few studies in the current literature specifically focus on the finding of pleural effusion. Therefore, direct comparison and identifying what could be considered state o the art for this specific finding is complex. Therefore, the survey of related works was mainly based on work that detected different findings in chest X-ray images.

Despite the different objectives, the results presented in Tab. 2 show that even without labels during training, the methodology based on the CLIP (Radford et al., 2021) architecture is capable of achieving results very close or superior to supervised learning methods. At this point, it is essential to highlight the sensitivity (88.54%) for detecting pleural effusion, demonstrating that the model is capable of identifying the findings.

Regarding the ROC curve (Fig. 3), the proposed method obtained a value of 0.900. Compared with other classification methods using the same dataset (Liz et al., 2023), our approach achieved superior performance in terms of AUC-ROC. Furthermore, the method proposed by (Liz et al., 2023) used part of the PADCHEST dataset in the training process. While in our method, PADCHEST was used only for method testing. Therefore, it is possible that our method has a superior ability to generalize multi-center findings and may be more susceptible to use in clinical settings to aid in diagnosing pleural effusion.

The results of our study demonstrate that the contrastive learning method can be a practical approach to classifying pleural effusions on X-ray images. Although there is space for improvement in the classification performance metrics and extension for a radio-

Table 2: Comparison with related works. AUC: area under the receiver operating characteristic. Acc: Accuracy. Sen: Sensibility. Spe: Specificity. F1: F1-score.

Study	Dataset	AUC	Acc	Sen	Spe	F1
(Cohen et al., 2020)	Multiple	0.890	-	-	-	-
(Guan et al., 2021)	NIH ChestX-ray14	0.835	-	-	-	-
(Ibrahim et al., 2022)	NIH ChestX-ray14	-	88.86	-	-	-
(Liz et al., 2023)	PADCHEST	0.656	-	-	-	56.70
(Rajpurkar et al., 2017)	NIH ChestX-ray14	0.863	-	-	-	-
(Serte and Serener, 2021)	ChestX-ray8	0.780	75.00	100.00	67.00	-
(Zaidi et al., 2021)	NIH ChestX-ray14	0.874	98.20	-	-	-
Ours	**PADCHEST**	**0.900**	**86.28**	**88.54**	**76.91**	**68.41**

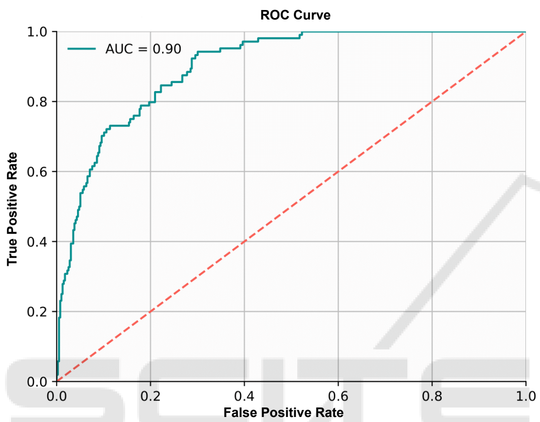

Figure 3: ROC curve for the PADCHEST dataset.

logical multi-finding classification model, the current results are promising and open possibilities for studies of our model in clinical settings.

5 CONCLUSION

In conclusion, this study explored the use of contrastive learning to classify pleural effusion on chest X-rays. The results showed that the proposed approach could outperform reference models and achieve significantly higher accuracy. In addition, the use of contrastive learning can be helpful in scenarios with restricted information about the clinical behavior of the disease, such as outbreaks and pandemics of lesser-known or emerging diseases.

This work has a few limitations, as described next. The first is related to the technical limitation of being a model capable of finding only one finding per image. Thus, studying methods that can identify different image findings is a prospect for future work. Another limiting aspect is the data set for training. In this way, using sets with a greater variety of findings is foreseen, allowing more excellent reliability for clinical application. Finally, interpretability studies identify how the model creates the representations and how it relates to them during the prediction process. These studies can be extended to propose interpretability methods for radiologists.

The benefits of using contrastive learning to accelerate the creation of business models for real use are significant. The proposed method can help save time and valuable resources in creating machine learning models for healthcare applications, enabling faster and more accurate diagnoses, and improving the quality of the patient's hospital course. These advantages make the proposed approach a promising tool to enhance the analysis of chest X-ray images in real clinical scenarios.

ACKNOWLEDGEMENTS

The authors would like to thank the Coordination for the Improvement of Higher Education Personnel - CAPES (Financial Code 001), the National Council for Scientific and Technological Development - CNPq (Grant numbers 309537/2020-7 e 40572/2021-9), and the Research Support Foundation of the State of Rio Grande do Sul - FAPERGS (Grant number 21/2551-0000118-6) for your support in this work.

REFERENCES

Aboudara, M. and Maldonado, F. (2019). Update in the management of pleural effusions. *Medical Clinics*, 103(3):475–485.

Alghamdi, H. S., Amoudi, G., Elhag, S., Saeedi, K., and Nasser, J. (2021). Deep learning approaches for detecting covid-19 from chest x-ray images: A survey. *Ieee Access*, 9:20235–20254.

Azizi, S., Mustafa, B., Ryan, F., Beaver, Z., Freyberg, J., Deaton, J., Loh, A., Karthikesalingam, A., Kornblith, S., Chen, T., et al. (2021). Big self-supervised models advance medical image classification. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3478–3488.

Baltruschat, I. M., Nickisch, H., Grass, M., Knopp, T., and Saalbach, A. (2019). Comparison of deep learning approaches for multi-label chest x-ray classification. *Scientific reports*, 9(1):1–10.

Bustos, A., Pertusa, A., Salinas, J.-M., and de la Iglesia-Vayá, M. (2020). Padchest: A large chest x-ray image dataset with multi-label annotated reports. *Medical image analysis*, 66:101797.

Chen, T., Kornblith, S., Norouzi, M., and Hinton, G. (2020). A simple framework for contrastive learning of visual representations. In *International conference on machine learning*, pages 1597–1607. PMLR.

Cohen, J. P., Hashir, M., Brooks, R., and Bertrand, H. (2020). On the limits of cross-domain generalization in automated x-ray prediction. In *Medical Imaging with Deep Learning*, pages 136–155. PMLR.

Diniz, J. O. B., Diniz, P. H. B., Valente, T. L. A., Silva, A. C., de Paiva, A. C., and Gattass, M. (2018). Detection of mass regions in mammograms by bilateral analysis adapted to breast density using similarity indexes and convolutional neural networks. *Computer methods and programs in biomedicine*, 156:191–207.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. (2020). An image is worth 16x16 words: Transformers for image recognition at scale. *arXiv preprint arXiv:2010.11929*.

Elgendi, M., Nasir, M. U., Tang, Q., Smith, D., Grenier, J.-P., Batte, C., Spieler, B., Leslie, W. D., Menon, C., Fletcher, R. R., et al. (2021). The effectiveness of image augmentation in deep learning networks for detecting covid-19: A geometric transformation perspective. *Frontiers in Medicine*, 8:629134.

Goceri, E. (2023). Medical image data augmentation: techniques, comparisons and interpretations. *Artificial Intelligence Review*, pages 1–45.

Guan, Q., Huang, Y., Luo, Y., Liu, P., Xu, M., and Yang, Y. (2021). Discriminative feature learning for thorax disease classification in chest x-ray images. *IEEE Transactions on Image Processing*, 30:2476–2487.

Hallifax, R., Talwar, A., Wrightson, J., Edey, A., and Gleeson, F. (2017). State-of-the-art: Radiological investigation of pleural disease. *Respiratory medicine*, 124:88–99.

Han, Y., Chen, C., Tewfik, A., Ding, Y., and Peng, Y. (2021). Pneumonia detection on chest x-ray using radiomic features and contrastive learning. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 247–251. IEEE.

Ibrahim, R. F., Yhiea, N., Mohammed, A. M., and Mohamed, A. M. (2022). Pleural effusion detection using machine learning and deep learning based on computer vision. In *Proceedings of the 8th International Conference on Advanced Intelligent Systems and Informatics 2022*, pages 199–210. Springer.

Jany, B. and Welte, T. (2019). Pleural effusion in adults—etiology, diagnosis, and treatment. *Deutsches Ärzteblatt International*, 116(21):377.

Liz, H., Huertas-Tato, J., Sánchez-Montañés, M., Del Ser, J., and Camacho, D. (2023). Deep learning for understanding multilabel imbalanced chest x-ray datasets. *Future Generation Computer Systems*.

Pooch, E. H. P., Alva, T. A. P., and Becker, C. D. L. (2020). A deep learning approach for pulmonary lesion identification in chest radiographs. In *Brazilian Conference on Intelligent Systems*, pages 197–211. Springer.

Radford, A., Kim, J. W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al. (2021). Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PMLR.

Rajpurkar, P., Irvin, J., Zhu, K., Yang, B., Mehta, H., Duan, T., Ding, D., Bagul, A., Langlotz, C., Shpanskaya, K., et al. (2017). Chexnet: Radiologist-level pneumonia detection on chest x-rays with deep learning. *arXiv preprint arXiv:1711.05225*.

Serte, S. and Serener, A. (2021). Classification of covid-19 and pleural effusion on chest radiographs using cnn fusion. In *2021 International Conference on INnovations in Intelligent SysTems and Applications (INISTA)*, pages 1–6. IEEE.

Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., and Summers, R. M. (2017). Chestx-ray8: Hospital-scale chest x-ray database and benchmarks on weakly-supervised classification and localization of common thorax diseases. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2097–2106.

Zaidi, S. Z. Y., Akram, M. U., Jameel, A., and Alghamdi, N. S. (2021). Lung segmentation-based pulmonary disease classification using deep neural networks. *IEEE Access*, 9:125202–125214.

Zeiser, F. A., Costa, C. A. d., Ramos, G. d. O., Bohn, H., Santos, I., and Righi, R. d. R. (2021). Evaluation of convolutional neural networks for covid-19 classification on chest x-rays. In *Intelligent Systems: 10th Brazilian Conference, BRACIS 2021, Virtual Event, November 29–December 3, 2021, Proceedings, Part II 10*, pages 121–132. Springer.

Zhang, Y., Jiang, H., Miura, Y., Manning, C. D., and Langlotz, C. P. (2022). Contrastive learning of medical visual representations from paired images and text. In *Machine Learning for Healthcare Conference*, pages 2–25. PMLR.

Zuiderveld, K. (1994). Graphics gems iv. In Heckbert, P. S., editor, *Graphics Gems*, chapter Contrast Limited Adaptive Histogram Equalization, pages 474–485. Academic Press Professional, Inc., San Diego, CA, USA.