

Mask-Free Neuron Concept Annotation for Interpreting Neural Networks in Medical Domain

Hyeon Bae Kim[†], Yong Hyun Ahn[†], Seong Tae Kim^{*}

Kyung Hee University, Yongin, Republic of Korea
{st.kim}@khu.ac.kr

Abstract. Recent advancements in deep neural networks have shown promise in aiding disease diagnosis and medical decision-making. However, ensuring transparent decision-making processes of AI models in compliance with regulations requires a comprehensive understanding of the model’s internal workings. However, previous methods heavily rely on expensive pixel-wise annotated datasets for interpreting the model, presenting a significant drawback in medical domains. In this paper, we propose a novel medical neuron concept annotation method, named Mask-free Medical Model Interpretation (MAMMI), addresses these challenges. By using a vision-language model, our method relaxes the need for pixel-level masks for neuron concept annotation. MAMMI achieves superior performance compared to other interpretation methods, demonstrating its efficacy in providing rich representations for neurons in medical image analysis. Our experiments on a model trained on NIH chest X-rays validate the effectiveness of MAMMI, showcasing its potential for transparent clinical decision-making in the medical domain. The code is available at <https://github.com/ailab-kyunghee/MAMMI>.

Keywords: Interpretability · Explainability · Neuron-Concept annotation.

1 Introduction

Recent research has demonstrated considerable progress in employing deep neural networks (DNNs) within diverse medical domains, highlighting their potential for aiding disease diagnosis and medical decision-making. Nevertheless, when deploying AI models for real-world usage, it is important for models to provide sufficient justification for their decisions in compliance with the European Union’s General Data Protection Regulation (GDPR) law [18,21]. This regulation emphasizes the importance of transparent decision-making processes. Therefore, ensuring transparent clinical decisions of AI models requires a comprehensive understanding of the internal workings of the model [5,22].

To improve understanding of the internal workings of the model (i.e., interpretability), various methods have been proposed [2,12]. Wu *et al.* investigate

[†] Equal contribution; ^{*} Corresponding author.

the concept of neurons in the model by a subjective evaluation from radiologists. Network Dissection [2] offers a detailed explanation of the model’s internal mechanisms by automatically identifying concepts for individual neurons leveraging pixel-wise segment mask annotated datasets like Broden. Khakzar *et al.* proposed Towards Semantic Interpretation (TSI) [11] by applying Network Dissection to the medical domain leveraging pixel-wise segment mask annotated datasets. TSI is better than [25] by reducing the efforts for annotating concepts for every neuron by medical experts, which increases the annotation cost in proportion to the number of neurons. However it still heavily depends on pixel-wise annotated datasets, leading to heavy dataset collection costs [1]. It is a significant drawback, especially in medical domains where pixel-level annotation is especially expensive.

To address this issue, recent studies in computer vision (CV) domain [1,9,15] have introduced neuron concept annotation methods that leverage vision-language models, such as CLIP [16]. These methods utilize separate image datasets (probing set) and a set of text (concept set), offering a significant advantage in interpreting internal model representations without relying on pixel-wise segment mask annotated datasets. However, adaptation of the neuron concept annotation methods to the medical field necessitates careful consideration of three critical factors. Firstly, it is imperative to curate a concept set that incorporates medical-specific information, ensuring relevance to the domain. Secondly, the class-imbalanced data distribution prevalent in medical datasets might harm the performance of neuron concept annotation methods by affecting neuron representation selection. Lastly, it needs a vision-language model for the medical domain to ensure a suitable representation of medical-specific images and terms.

To address these challenges, we present a novel medical neuron concept annotation method, named **Mask-free Medical Model Interpretation (MAMMI)**, which takes into account these key aspects in this study. Comprehensive experiments have been conducted on a model trained on chest X-rays (NIH14). Experimental results demonstrate that MAMMI achieves superior performance compared to other methods, showcasing its efficacy in providing rich representations for neurons in medical image analysis. Our contributions are as follows:

- We proposed a novel medical neuron concept annotation method, MAMMI, to address the limitations of the previous neuron concept annotation methods. By leveraging vision-language models, MAMMI stands out as a mask-free method, eliminating the necessity for pixel-level datasets and thereby reducing the costs associated with collecting expensive pixel-level data.
- We explore methods for neuron-concept association in the medical field and design the method by accounting for medical-specific attributes that are often overlooked. These include factors such as class-imbalanced data distribution, which is particularly prevalent in medical datasets.
- Extensive experiments have been conducted to evaluate the effectiveness of the MAMMI. MAMMI demonstrated superior performance both qualitatively and quantitatively in a medical setting compared to other methods.

2 Related Work

Concept Activation Vectors: Kim et al. [12] proposed an interpretation method based on Concept Activation Vectors (CAVs). CAVs are defined as vectors orthogonal to a linear classifier that separate neuron activations between a positive example set and a negative non-example set for a user-defined concept. TCAV (Testing with CAVs) measures the sensitivity of the model to all concepts for a given prediction, allowing an explanation of which concepts the model’s decision is based on. Additionally, it has been demonstrated to be applicable in the medical domain by reporting results on retinal fundus images [13]. However, CAVs have limitations on detailed and diverse automatic explanations of the model’s internal workings.

Network Dissection: Bau et al. [2] proposed Network Dissection, which matches concepts by measuring the Intersection over Union (IoU) score between pixel-wise segmentation masks and individual neuron (e.g., convolution filter) feature maps. Khakzar et al. [11] proposed TSI and demonstrated the utility of detecting concepts within neurons of a thoracic disease diagnosis model, showcasing its effectiveness in the medical domain. However, the aforementioned methods require pixel-wise annotation due to their dependence on an image-concept-matched dataset like Broden [2] or NIH14 [23] with bounding boxes, which incurs expensive pixel-level annotation costs for dataset collection.

Vision-Language Model-based Methods: Vision-language models, such as CLIP [16] model, have been utilized for matching neuron concepts without the help of pixel-wise annotated datasets. Examples include FALCON [9], CLIP Dissect [15], and WWW [1]. FALCON [9] selects high-activated examples and finds similar captions in a large-scale caption dataset [19]. This approach relies on the assumption that there are captions representing neurons, which requires a diverse large-scale caption dataset, including various captions. However, building such a dataset with diverse captions, especially in the medical domain where specialized knowledge is required, can be challenging. CLIP Dissect [15] matches concepts with the concept activation matrix, calculated based on the inner product of CLIP image feature and text features. And the neuron activations of the probing set. WWW [1] matches concepts by calculating the adaptive cosine similarity in CLIP space between a fixed number of high-activated examples for each neuron and the concept set.

3 Method

There are important consideration points for designing neuron concept annotation in the medical domain. For instance, concept sets consisting of general words may not properly represent the concepts that the model learns from the medical datasets. In section 3.1, we introduce a way for constructing a concept set considering medical-specific information. Also, in the medical dataset, the number of samples between classes can be different (i.e., class imbalance), which can be another challenging point for neuron concept annotation, in particular,




Fig. 1. Overview of neuron concept annotation method of MAMMI. The representative images of neurons for concept annotation are selected by images with activation values over the neuron-wise computed adaptive threshold from the probing data. The concept set is constructed by extracting nouns from the medical report. The representative images and the text concept set are projected into the CLIP model embedding space for concept identification by the adaptive concept matching module.

for representative example selection. In section 3.2, we propose adaptive neuron representative image selection for considering class-imbalance distribution data prevalent in medical datasets. In section 3.3, we provide a description of the overall flow of MAMMI, incorporating the introduced elements.

3.1 Constructing a Medical Concept Set

Our approach involves utilizing a dataset that has medical reports related to the domain of the trained target model. Leveraging datasets that have target domain image-related reports gives great advantages in obtaining suitable concepts for interpreting the target model. We extracted nouns in medical reports of the target domain to extract meaningful concepts for model interpretation.

3.2 Adaptive Neuron Representative Image Selection

The class imbalance is one of the prevalent attributes of medical datasets. Therefore, we propose an adaptive neuron representative image selection for each neuron representation. Let neuron activations of target model l -th layer i -th neuron regarding concept set images denoted as $A^{(l,i)}$. The adaptive threshold (τ) for neuron representative images selection is computed as follows:

$$\tau_{example}^{(l,i)} = \max(A^{(l,i)}) - (1 - \alpha) \cdot (\max(A^{(l,i)}) - \min(A^{(l,i)})), \quad (1)$$

where α denotes the hyper-parameter for neuron representation selection sensitivity. In Eq. 1, for calculating the adaptation threshold of the i -th neuron

in the l -th layer, we leverage the gap between the maximum sample activation ($\max(A^{(l,i)})$) and minimum sample activation ($\min(A^{(l,i)})$). Since each neuron has different activation patterns and ranges, leveraging the gap between maximum and minimum activation allows adaptive neurons representative image selection to consider the range of each neuron activation. By adjusting the hyperparameter α , we can select suitable examples for neuron representation while considering the data imbalance. We select l -th layer i -th neuron representative images $Imgs^{(l,i)} = \{A_k^{(l,i)} > \tau_{example}^{(l,i)}, (k = 1, 2, \dots, p)\}$ where p denotes the number of images in probing set.

3.3 Overview of MAsk-free Medical Model Interpretation: MAMMI

In this section, we provide an overview of the overall flow of MAMMI for medical neuron concept annotation. The concept identification follows the concept discovery module proposed in [1], and the overall flow is illustrated in Figure 1.

From the obtained images and concepts, the following steps are proceed for concept identification: First, we calculate CLIP text features of concepts as $T = [t_1, t_2, \dots, t_m]$ and CLIP visual features of a neuron to annotate (i.e., l -th layer i -th neuron) as $V^{(l,i)} = [v_1^{(l,i)}, v_2^{(l,i)}, \dots, v_n^{(l,i)}]$ from $Imgs^{(l,i)}$. Here, m denotes the total number of concepts, and n represents the number of selected representing images of a neuron to annotate. Then, we measure the adaptive cosine similarity of the calculated features to compute the concept score S . The concept score of j -th concept regarding l -th layer i -th neuron $s_j^{(l,i)}$ is calculated as follows:

$$s_j^{(l,i)} = \frac{1}{n} \sum_{o=1}^n \left\{ \cos(v_o^{(l,i)}, t_j) - \cos(v_o^{(l,i)}, t_{tem}) \right\}, \quad (2)$$

where t_{tem} represents the CLIP text embedding of the base template. If the calculated concept score $s_j^{(l,i)}$ is greater than the threshold $\theta_{concept}$, j -th concept is annotated as the concept for the corresponding neuron. The threshold $\theta_{concept}$ is calculated as $\theta_{concept} = \beta \times \max(s^{(l,i)})$. β denotes the ratio of selecting concepts. We used the same hyperparameter ($\beta = 0.95$) as in [1].

4 Experiments

4.1 Experimental Settings

The target model for interpretation is a DenseNet121 [6] architecture which is pre-trained with the MoCo training method and fine-tuned using the chest X-ray dataset [26]. Additionally, for qualitative results, we employed ChestX-det [14] with bounding box annotations. For quantitative evaluation, we performed it using the final layer neuron concepts and class labels as introduced in [15,1]. We calculated the similarity between the class label feature encoded by CLIP (ViT-B/16) [17] and mpnet [20] and the neuron concepts. The F1-score indicates the exactness and flexibility of neuron concepts, while the hit-rate measures the extent to which neuron concepts correctly predict class labels.

Table 1. Ablation study on the concept set selection of MAMMI.

$D_{concept}$	# of concepts	Cosine Similarity		F1-score	Hit-rate
		CLIP	mpnet		
20K [15]	20K	0.7749 ± 0.016	0.1045 ± 0.016	0.000	0.000
Wordnet nouns[1]	80K	0.7253 ± 0.022	0.1542 ± 0.014	0.000	0.000
MIMIC _{Nouns}	1361	0.8599 ± 0.024	0.5290 ± 0.084	0.2619 ± 0.113	0.2857 ± 0.121

Table 2. Ablation on the adaptive neuron representative image selection.

Example Selection	Cosine Similarity		Penultimate Unique concept
	CLIP	mpnet	
w/o Adaptive	0.8495 ± 0.026	0.5107 ± 0.085	65
Adaptive	0.8599 ± 0.024	0.5290 ± 0.084	96

Table 3. Comparison of medical CLIP models in MAMMI.

CLIP Model	Cosine Similarity	
	CLIP	mpnet
CXR-RePaIR [3]	0.8334 ± 0.018	0.3716 ± 0.070
MedKLIP [24]	0.7704 ± 0.014	0.1780 ± 0.022
MedCLIP [24]	0.8599 ± 0.024	0.5290 ± 0.084

4.2 Evaluation of Our Method

Ablation study on the Concept Set Selection of MAMMI: To evaluate the performance of the concept set considering medical-specific information, we compared various concept sets obtained from other sources. In this experiment, we interpreted the target model fine-tuned on the NIH Chest X-ray14 (NIH14) [23]. We collect concept set from MIMIC-CXR reports [8] by extracting nouns from the datasets (i.e., MIMIC_{Nouns}). As a control group for MIMIC_{Nouns}, we compared two concept sets proposed in [15] (20k) and proposed in [4,1] (80k) from the CV domain.

In Table 1, the cosine similarity scores of the concept set from 20k and 80k sets have significantly decreased compared to the concept set of our method. Additionally, the F1 Score and hit rate, which evaluates whether annotated concepts match with labels or not, are both zero. This result indicates that by building a concept set with medical-specific information (MIMIC_{Nouns}), the model interpretability can be significantly improved.

Ablation Study on Adaptive Neuron Representative Image Selection: In Table 2, we present the results of an ablation study conducted under the same conditions for our proposed adaptive neuron representative image selection method and select a fixed number of examples (i.e., w/o adaptive neuron representative image selection). The parameter α used in our adaptive neuron representative image selection is set to 93%. We discovered an interesting finding when comparing the example distribution of adaptive neuron representative image selection with the data distribution of the probing set (NIH14). Figure 2 shows the proportion of class-wise data distribution in the probing set and the number of selected examples with and without adaptive neuron representative image selection in each neuron. Compared to the method without adaptive selection, the distribution of adaptive neuron representative image selection shows a similar distribution to the train data distribution. This result indicates that by leveraging neuron representative image selection, MAMMI can select suitable examples with less noisy representing images. As a result, using adaptive neuron representative image selection showed marginal performance improvement.




Fig. 2. The data distribution of adaptive neuron representative image selection compared to the probing set. (a) Ablation study on the value of parameter α . (b) Distribution regarding the number of selected representative images of neurons in the final layer and train set class distribution.

Table 4. Quantitative comparison with mask-based concept annotation method.

Method	$D_{concept}$	Cosine Similarity		F1-score	Hit-rate
		CLIP	mpnet		
TSI [11]	NIH14 [23]+COVID-CXR [7]	0.7872 $\pm$ 0.014	0.2554 $\pm$ 0.039	0	0
Ours	MIMIC _{Nouns}	0.8599 $\pm$ 0.024	0.5290 $\pm$ 0.084	0.2619 $\pm$ 0.113	0.2857 $\pm$ 0.121

Ablation Study on Medical CLIP Model Selection: For better annotations for the neuron concepts of the target model, we conducted an ablation study by selecting the CLIP model trained on the chest X-ray dataset. We compared three CLIP-based models, CXR-RePaIR [3], MedKLIP [24], and MedCLIP [24]. Results are listed in Table 3. Our experimental results showed that MedCLIP [24] outperforms compared to other competitive baselines.

4.3 Performance Evaluation with Other Methods

Comparison with Mask-based Method: Table 4 presents a performance comparison with the TSI [11], which utilizes mask annotation in the medical domain. TSI used a concept set that consists of COVID-19 segmentation data (COVID-CXR) [7] and class labels from NIH Chest X-ray14 (NIH14) [23]. MAMMI exhibited superior performance across all metrics, highlighting its effectiveness. Notably, TSI failed to accurately identify neuron concepts, as observed from F1-scores and Hit-rate scores. For qualitative comparison, we display two penultimate layer neurons that are important for identifying ground truth labels with their respective activation regions. Qualitative results in Figure 3, MAMMI accurately identifying neuron concepts, even at the penultimate layer, while TSI [11] failed to identify proper concepts.

Comparison with Vision-Language Model-based Methods: Table 5 provides a performance comparison by implementing Vision-Language model-based methods from the CV domain into the medical setting. WWW is not explicitly mentioned as it serves as the baseline for MAMMI. In the medical setting,

Table 5. Quantitative comparison with other concept annotation methods.

Method	D_concept	Cosine Similarity		F1-score	Hit-rate
		CLIP	mpnet		
FALCON _M	MIMIC report [8]	0.8054 $\pm$ 0.012	0.2420 $\pm$ 0.031	0.0635 $\pm$ 0.027	0.2857 $\pm$ 0.121
CLIP Dissect _M	MIMIC _{Nouns}	0.8205 $\pm$ 0.024	0.3700 $\pm$ 0.083	0.0714 $\pm$ 0.069	0.0714 $\pm$ 0.069
Ours	MIMIC _{Nouns}	0.8599 $\pm$ 0.024	0.5290 $\pm$ 0.084	0.2619 $\pm$ 0.113	0.2857 $\pm$ 0.121

we replace the caption dataset of FALCON [9] with MIMIC-CXR medical reports [8] for this experiment. Additionally, to ensure the identification of concepts in every neuron, the example selection threshold was modified to 0.25. CLIP Dissect [15] also adjusted into medical settings with the medical concept set (i.e., MIMIC_{Nouns}) with MedCLIP[24]. Moreover, performance was evaluated using the top-1 neuron concept, ensuring that neuron concepts were annotated in every neuron. As shown in Table 5, MAMMI achieved high scores across all metrics compared to other competitive baselines. Also, qualitative results are presented in Figure 3. As shown in the figure, MAMMI accurately identified concepts even at the penultimate layer, while other baselines [9,15] failed to identify proper concepts.




Fig. 3. Qualitative result on the penultimate layer. We display two penultimate layer neurons selected by Shapley Value [10,1] as the most important penultimate neurons for identifying the respective ground truth label in the image.

5 Conclusion

In this study, we introduce a novel approach for interpreting medical vision models via neuron concept association named MAMMI, specifically tailored for the medical domain. Recognizing the challenges associated with constructing new datasets, we devised a concept set considering medical-specific information.

Also, we employed an adaptive neuron representative image selection method to overcome the imbalanced data distribution of common attributes in medical datasets. Our extensive experiments demonstrated that MAMMI outperforms other approaches in annotating a wide variety of concepts to neurons in medical models. Furthermore, our method addresses a limitation of existing approaches, such as TSI, which requires expensive mask-annotated datasets.

Acknowledgments. This work was supported in part by the National Research Foundation of Korea(NRF) grant funded by the Korea government(MSIT) (No.RS-2024-00334321) and by the Institute of Information and Communications Technology Planning and Evaluation (IITP) Grant funded by the Korea Government (MSIT) under Grant 2022-0-00078, Grant IITP-2024-RS-2023-00258649, and Grant RS-2022-00155911.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Ahn, Y.H., Kim, H.B., Kim, S.T.: Www: A unified framework for explaining what where and why of neural networks by interpretation of neuron concepts. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 10968–10977 (2024)
2. Bau, D., Zhou, B., Khosla, A., Oliva, A., Torralba, A.: Network dissection: Quantifying interpretability of deep visual representations. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 6541–6549 (2017)
3. Endo, M., Krishnan, R., Krishna, V., Ng, A.Y., Rajpurkar, P.: Retrieval-based chest x-ray report generation using a pre-trained contrastive language-image model. In: Machine Learning for Health. pp. 209–219. PMLR (2021)
4. Fellbaum, C.: Wordnet and wordnets. encyclopedia of language and linguistics (2005)
5. Hicks, S., Riegler, M., Pogorelov, K., Anonsen, K.V., de Lange, T., Johansen, D., Jeppsson, M., Randel, K.R., Eskeland, S.L., Halvorsen, P.: Dissecting deep neural networks for better medical image classification and classification understanding. IEEE International Symposium on Computer-Based Medical Systems pp. 363–368 (2018)
6. Huang, G., Liu, Z., Van Der Maaten, L., Weinberger, K.Q.: Densely connected convolutional networks. In: Proceedings of the IEEE conference on computer vision and pattern recognition. pp. 4700–4708 (2017)
7. Inc, G.: Covid-19 chest x-ray segmentations dataset. <https://github.com/GeneralBlockchain/covid-19-chest-xray-segmentations-dataset> (2020)
8. Johnson, A.E., Pollard, T.J., Berkowitz, S.J., Greenbaum, N.R., Lungren, M.P., Deng, C.y., Mark, R.G., Horng, S.: MIMIC-CXR, a de-identified publicly available database of chest radiographs with free-text reports. Scientific data **6**(1), 317 (2019)
9. Kalibhat, N., Bhardwaj, S., Bruss, C.B., Firooz, H., Sanjabi, M., Feizi, S.: Identifying interpretable subspaces in image representations. Proceedings of the 40th International Conference on Machine Learning **202**, 15623–15638 (2023)

10. Khakzar, A., Baselizadeh, S., Khanduja, S., Rupprecht, C., Kim, S.T., Navab, N.: Neural response interpretation through the lens of critical pathways. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 13528–13538 (2021)
11. Khakzar, A., Musatian, S., Buchberger, J., Valeriano Quiroz, I., Pinger, N., Baselizadeh, S., Kim, S.T., Navab, N.: Towards semantic interpretation of thoracic disease and covid-19 diagnosis models. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part III 24. pp. 499–508. Springer (2021)
12. Kim, B., Wattenberg, M., Gilmer, J., Cai, C., Wexler, J., Viegas, F., et al.: Interpretability beyond feature attribution: Quantitative testing with concept activation vectors (tcav). In: International conference on machine learning. pp. 2668–2677. PMLR (2018)
13. Krause, J., Gulshan, V., Rahimy, E., Karth, P., Widner, K., Corrado, G.S., Peng, L., Webster, D.R.: Grader variability and the importance of reference standards for evaluating machine learning models for diabetic retinopathy. *Ophthalmology* **125**(8), 1264–1272 (2018)
14. Liu, J., Lian, J., Yu, Y.: Chestx-det10: chest x-ray dataset on detection of thoracic abnormalities. arXiv preprint arXiv:2006.10550 (2020)
15. Oikarinen, T., Weng, T.W.: CLIP-dissect: Automatic description of neuron representations in deep vision networks. International Conference on Learning Representations (2023)
16. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
17. Radford, A., Kim, J.W., Hallacy, C., Ramesh, A., Goh, G., Agarwal, S., Sastry, G., Askell, A., Mishkin, P., Clark, J., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
18. Salahuddin, Z., Woodruff, H.C., Chatterjee, A., Lambin, P.: Transparency of deep neural networks for medical image analysis: A review of interpretability methods. *Computers in biology and medicine* **140**, 105111 (2022)
19. Schuhmann, C., Kaczmarczyk, R., Komatsuzaki, A., Katta, A., Vencu, R., Beaumont, R., Jitsev, J., Coombes, T., Mullis, C.: Laion-400m: Open dataset of clip-filtered 400 million image-text pairs. In: NeurIPS Workshop Datacentric AI. No. FZJ-2022-00923, Jülich Supercomputing Center (2021)
20. Song, K., Tan, X., Qin, T., Lu, J., Liu, T.Y.: MpNet: Masked and permuted pre-training for language understanding. *Advances in Neural Information Processing Systems* **33**, 16857–16867 (2020)
21. Temme, M.: Algorithms and transparency in view of the new general data protection regulation. *Eur. Data Prot. L. Rev.* **3**, 473 (2017)
22. Tjoa, E., Guan, C.: A survey on explainable artificial intelligence (xai): Toward medical xai. *IEEE transactions on neural networks and learning systems* **32**(11), 4793–4813 (2020)
23. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: 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. pp. 2097–2106 (2017)

24. Wang, Z., Wu, Z., Agarwal, D., Sun, J.: Medclip: Contrastive learning from unpaired medical images and text. In: Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing. pp. 3876–3887 (2022)
25. Wu, J., Peck, D., Hsieh, S., Dialani, V., Lehman, C.D., Zhou, B., Syrgkanis, V., Mackey, L., Patterson, G.: Expert identification of visual primitives used by cnns during mammogram classification. In: SPIE Medical Imaging. vol. 10575, pp. 633–641 (2018)
26. Xiao, J., Bai, Y., Yuille, A., Zhou, Z.: Delving into masked autoencoders for multi-label thorax disease classification. In: Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision. pp. 3588–3600 (2023)