



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Masks and Manuscripts: Advancing Medical Pre-training with End-to-End Masking and Narrative Structuring

Shreyank N Gowda¹ and David A. Clifton^{1,2}

¹ Department of Engineering Sciences, University of Oxford, OX3 7DQ Oxford, UK

² Oxford Suzhou Centre for Advanced Research, University of Oxford, Suzhou 215123, Jiangsu, China

Abstract. Contemporary medical contrastive learning faces challenges from inconsistent semantics and sample pair morphology, leading to dispersed and converging semantic shifts. The variability in text reports, due to multiple authors, complicates semantic consistency. To tackle these issues, we propose a two-step approach. Initially, text reports are converted into a standardized triplet format, laying the groundwork for our novel concept of “observations” and “verdicts”. This approach refines the Entity, Position, Exist triplet into binary questions, guiding towards a clear “verdict”. We also innovate in visual pre-training with a Meijering-based masking, focusing on features representative of medical images’ local context. By integrating this with our text conversion method, our model advances cross-modal representation in a multimodal contrastive learning framework, setting new benchmarks in medical image analysis.

1 Introduction

With the swift advancements in deep learning, computer-aided diagnosis in the medical field has made notable progress through various models. However, these models, predominantly trained on specific anatomical or disease categories, demand costly data annotation and re-training when a new disease emerges [15], limiting their practical utility. While deep learning has flourished due to large-scale labeled natural image datasets [7], the annotation of medical images remains time-consuming and expensive. A common strategy involves pre-training on large datasets like ImageNet [7] and then fine-tuning on specialized medical datasets [29], but this often falls short in terms of generalized performance due to the domain differences between natural and medical images.

Medical image analysis stands at the crossroads of technological advancements and clinical applications. The rapid growth of medical imaging modalities and the increasing volume of patient data offer unprecedented opportunities to leverage sophisticated machine learning techniques. Yet, the complexity of this data—rich in variability, granularity, and intricacies—poses unique challenges for the development of effective and reliable models. Adding to the challenge, the associated textual data, which often accompanies these images in the form of

clinical notes and radiological reports, can introduce inconsistencies due to varied semantics and writing styles across different authors [9], potentially leading to significant noise in the learning process and compromising model reliability.

Self-supervised learning has emerged as a powerful paradigm for pre-training deep neural networks without relying on human annotations, leveraging inherent structure in the data to learn useful representations in a scalable manner [23]. Drawing on these advancements, this paper introduces a novel approach that combines self-supervised learning with the nuanced requirements of medical data processing. We propose a two-step conversion process for textual data into a standardized triplet format to refine representations and an innovative visual pre-training strategy that employs Meijering-based masking for improved feature extraction [20].

Furthermore, vision-language pre-training (VLP) models like CLIP [23] have shown success in learning joint multimodal representations but struggle with the unique challenges presented by medical images and unstructured reports [30]. We address these challenges by employing a filtering-based approach for preprocessing medical images, enhancing the performance of deep learning models through classical techniques like Meijering’s vesselness filter [20,24] and adapting these for novel masking strategies to facilitate cleaner reconstructions and improved model performance.

Incorporating external knowledge into medical models has shown promise in mimicking medical decision-making processes and enhancing model training with domain-specific clinical knowledge [28]. Building on the concept of Med-KLIP [30], we extend the extraction of key data from radiology reports into triplets and encode it with domain knowledge to improve masked generative pre-training [30].

This paper aims to bridge the gap between the raw, multifaceted medical data and the refined, consistent representations needed for effective machine learning applications, pushing the boundaries of medical pre-training by tailoring advanced techniques specifically for the medical domain. Through an amalgamation of self-supervised learning, VLP models, preprocessing techniques, and the incorporation of medical knowledge, we propose a comprehensive framework for enhancing the performance and reliability of machine learning models in medical image analysis and report generation.

2 Methodology

In Figure 1, we outline our approach, adopting MaskVLM’s [18] architecture. Our innovations primarily lie in visual and textual data pre-processing, rather than the architecture itself, which serves as a foundation for our training methodology. We employ transformer-based encoders [8,26] for concurrent image and text analysis. The image encoder extracts features as vector sets, including a class token, while the text encoder processes text with a start token. Through self-attention and cross-modality encoders featuring cross-attention, we enhance

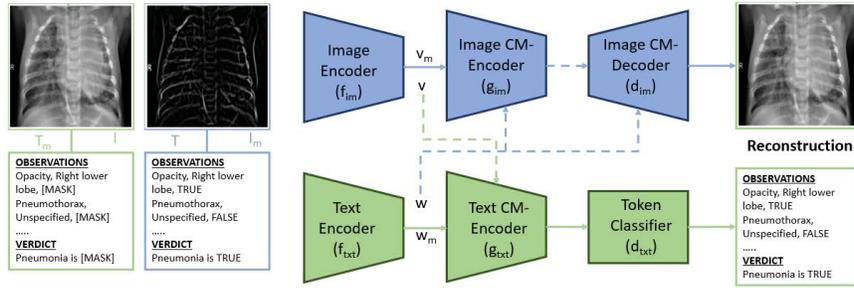


Fig. 1. An overview of the architecture for integrated modeling of masked medical visual and linguistic data. The blue and green pathways represent the flow of information for the reconstruction of images and text, respectively. The dashed lines show the intermodal contribution of exposed signals for the creation of attention.

inter-modal integration, culminating in accurate image descriptions. Further details and the rationale for our methods are elaborated in subsequent sections.

2.1 Masked Image Modeling

Masked Image Modeling (MIM) [10] revolutionizes self-supervised learning by enabling models to infer visual contexts without extensive labeled data, thereby enhancing efficiency and generalization. However, for fine-grained medical data, random masking leads to unclear reconstructions. Despite being superior to random initialization [32], improvements are needed due to the reliance on local features in medical images. Utilizing ridge filters, like the Meijering filter, converts X-ray images into intermediate forms for reconstruction, addressing the fine-grained nature of medical data [20,24]. This approach, supported by qualitative analysis in Figure 2 and further in Table 7 and Table 6, empirically outperforms random masking as evidenced by our ablation study.

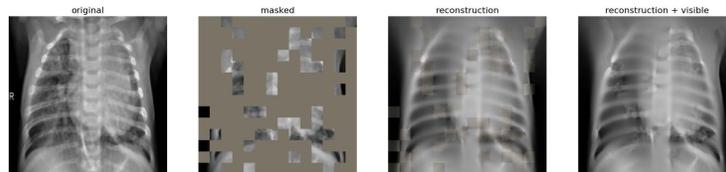


Fig. 2. Comparing masking strategies: (a) Masking random patches for reconstruction often results in blurry outputs lacking in detail. (b) Filtering the image before reconstruction preserves fine-grained details, leading to higher resolution outcomes.

Meijering Filter Overview The Meijering filter enhances tubular structures in images via the Hessian matrix’s eigenvalues [20]. For an image $I(x, y)$, the Hessian matrix \mathbf{H} and its eigenvalues λ_1, λ_2 guide the filter response R , which is zero for $\lambda_2 > 0$ and $\sqrt{\lambda_1^2 + \lambda_2^2}$ otherwise. By adjusting the scale σ , the filter adapts to features of various sizes, optimizing the response R_σ across scales to highlight relevant structures, thus aiding in medical image analysis.

2.2 Manuscript Generation

Radiology reports in publicly available datasets vary due to individual writing styles, affecting the consistency of sentence embeddings and their use in contrastive learning. We adopt triplet extraction from MedKLIP [30] for uniformity, converting extracted triplets into new reports. This approach ensures semantic consistency, facilitating the extraction of positive and negative pairs, and allows for binary answers in observations and verdicts, useful for masked pre-training. Figure 3 provides an overview of our methodology.

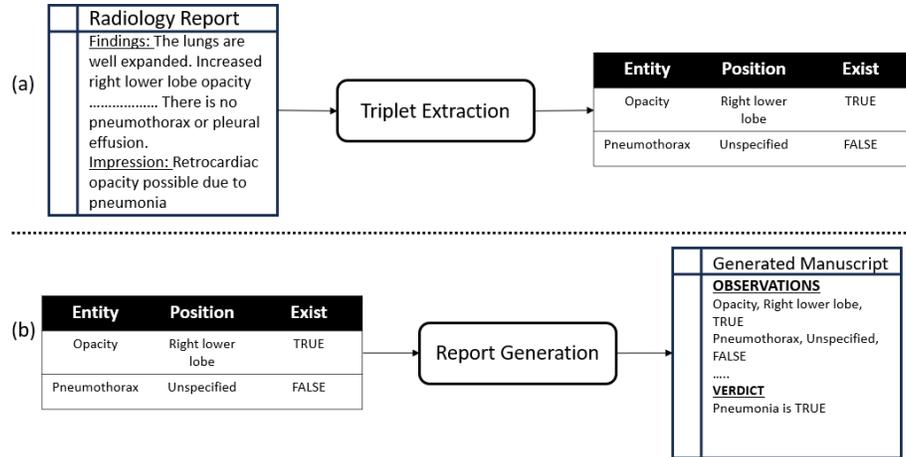


Fig. 3. Our report generation process starts with the triplet extraction method from MedKLIP, as outlined in part (a). Instead of adopting the Knowledge-enhanced Triplet Encoding, we transform these into a textual report, embedding it sentence by sentence. This approach facilitates masked pre-training by providing binary labels for each observation and the verdict.

Report Pre-processing Following MedKLIP [30], we apply RadGraph [14], a medical Named Entity Recognition (NER) tool, to streamline radiology reports into triplets. The NER identifies and classifies medical terms as "entity" (clinical findings) or "position" (body location), and assigns an "exist" status to

indicate the presence, absence, or uncertainty of clinical observations. This process, generating triplets like entity, position, exist, simplifies sentence structures in reports while retaining crucial information, leveraging the Triplet Extraction method in medical Visual Language Processing (VLP).

2.3 Conditional Reconstruction

Our method aims to reconstruct an original signal in one modality (image or text) from its masked version using unmasked information from the other modality. We follow notations used by MaskVLM [18] as we build on them. Specifically, for a medical image I and its masked form I_m , with accompanying unmasked text T , the goal is to retrieve I utilizing (I_m, T) . In reverse, with an unmasked image I and masked text T_m , the objective is to recover T with (I, T_m) . We employ image and text encoders f_{im} and f_{txt} for feature extraction. The cross-modality encoders g_{im} and g_{txt} , leveraging context from unmasked modalities, reconstruct the signals through decoders d_{im} and d_{txt} for images and texts, respectively. This reconstruction process is optimized by a multimodal joint modeling loss \mathcal{L}_{MVLM} , promoting focus on unmasked modality features for accurate reconstruction:

$$\mathcal{L}_{MVLM} = \mathbb{E}_{(I,T) \sim \mathcal{D}} \left[\mathcal{H}(y_T^M, \phi_{txt}^M(I, T|m)) + \left\| \frac{1}{\Omega(I^M)} (I^M - \phi_{im}^M(I_m, T')) \right\|_1 \right] \quad (1)$$

where ϕ_{txt} and ϕ_{im} represent the decoding processes for text and image modalities, respectively, focusing on masked regions. This process enables learning from cross-modal interactions to enhance reconstruction accuracy.

2.4 Multi-Modal Alignment

To enhance masked signal reconstruction, we incorporate cross-modal alignment through two additional objectives. The first, image-text contrastive (ITC) learning [23], projects image and text encoder outputs into a shared space, using separate fully-connected layers. This process is formulated as minimizing the contrastive loss \mathcal{L}_{ITC} , encouraging agreement between matched image-text pairs in embedding space. However, variability in radiology report styles may affect the identification of positive and negative pairs, as observed in our ablation study.

The ITC loss for the k -th image-text pair, with batch size N and temperature τ , is given by:

$$\mathcal{L}_{ITC} = -\frac{1}{N} \sum_{k=1}^N \left[\log \frac{\exp(z_{im}^k \cdot z_{txt}^k / \tau)}{\sum_{n=1}^N \exp(z_{im}^n \cdot z_{txt}^k / \tau)} + \log \frac{\exp(z_{im}^k \cdot z_{txt}^k / \tau)}{\sum_{n=1}^N \exp(z_{im}^k \cdot z_{txt}^n / \tau)} \right], \quad (2)$$

The second objective, image-text matching (ITM) classification [5], predicts the alignment of image-text pairs by element-wise multiplication of cross-modality

encoder outputs, followed by classification. The ITM loss \mathcal{L}_{ITM} ensures supervised alignment across modalities:

$$\mathcal{L}_{ITM} = \mathbb{E}_{(I,T) \sim D} [\mathcal{H}(y_{ITM}, g_{\text{cross}}^{\text{itm}}(z_{\text{im}}^{\text{cross}}, z_{\text{txt}}^{\text{cross}}))]. \quad (3)$$

The total pre-training loss \mathcal{L} combines all losses:

$$\mathcal{L} = \mathcal{L}_{MVML} + \mathcal{L}_{ITC} + \mathcal{L}_{ITM}. \quad (4)$$

Model fine-tuning is performed for each downstream task over a few epochs, alongside zero-shot evaluations following prior work [6,30].

3 Experimental Analysis

3.1 Datasets

We use datasets that have been used in recent approaches [12,6,30] namely: MIMIC-CXR v2 [16], RSNA Pneumonia Detection [25], SIIM-ACR Pneumothorax [19], NIH Chest X-Ray Dataset [27], CheXpert [13], COVIDx CXR [22] and Edema Severity [2]. We follow the evaluation protocol of MedKlip [30].

3.2 Implementation Details

The model utilizes a ViT-B pre-trained on ImageNet for image processing and ClinicalBERT for text, resizing images to 224x224 and setting a common latent dimension of 768. Training is conducted on 4 NVIDIA Tesla V100 GPUs with a batch size of 128, using AdamW optimizer with a 0.05 weight decay. Unlike fixed masking ratios, the model adapts the mask ratio based on the input, using a 15% ratio for text that increases to 30% when paired with images. The model is pre-trained for 100 epochs, considerably less than the 800 epochs in similar studies, and fine-tuned for 10 epochs with a specific learning rate strategy, including a warm-up to 3e-4 followed by a cosine scheduler for the pre-training phase, and fixed learning rates of 1e-5 during fine-tuning. Extending training to 300, 500, or 800 epochs did not improve accuracy.

3.3 Classification

Semi and Fully Supervised We conduct semi and fully supervised classification on RSNA Pneumonia, SIIM-ACR, and CheXpert datasets, varying labeled data from 1% to 100%. Results, including mean and standard deviation over 5 runs using PRIOR [6], are in Table 1. M&M surpasses SOTA by up to 2.64%.

Zero-Shot We assess zero-shot classification on RSNA Pneumonia, SIIM-ACR, and NIH Chest X-Ray datasets with state-of-the-art models, showcasing our model’s generalization across various clinical sources. Our approach achieves consistent performance enhancements, attributed to domain adaptation as discussed in MedKLIP [30]. Results in Table 2 show our M&M model outperforming existing methods by up to 3.16%, following pre-training on MIMIC-CXR and up to a 3.29% improvement on the novel COVID-19 disease in Table 3.

Methods	RSNA Pneumonia			SIIM-ACR			CheXpert		
	1%	10%	100%	1%	10%	100%	1%	10%	100%
MoCo [11]	82.33	85.22	87.90	75.49	81.01	88.43	78.00	86.27	87.24
SimCLR [4]	80.18	84.60	88.07	74.97	83.21	88.72	67.41	86.74	87.97
ConVIRT [31]	83.98	85.62	87.61	84.17	85.66	91.50	85.02	87.58	88.21
GLoRIA [12]	84.12	86.83	89.13	85.05	88.51	92.11	83.61	87.40	88.34
BioViL [1]	81.95	85.37	88.62	79.89	81.62	90.48	80.77	87.56	88.41
LoVT [21]	85.51	86.53	89.27	85.47	88.50	92.16	85.13	88.05	88.27
PRIOR [6]	85.74	87.08	89.22	87.27	89.13	92.39	86.16	88.31	88.61
MedKLIP [30]	87.31	87.99	89.31	85.27	90.71	91.88	86.24	88.14	88.68
M&M (Ours)	88.11	89.44	91.91	88.81	91.15	93.88	88.45	90.02	90.88

Table 1. Comparing supervised classification results obtained by fine-tuning on RSNA Pneumonia Detection [25], SIIM [19] and CheXpert [13]. All methods are trained on different portions of the training set from 1% to 100% and evaluated using AUC-ROC. Following PRIOR [6] each reported value is the average of five runs along with the standard deviation.

Methods	RSNA Pneumonia			SIIM-ACR			NIH Chest X-Ray		
	AUC \uparrow	F1 \uparrow	ACC \uparrow	AUC \uparrow	F1 \uparrow	ACC \uparrow	AUC \uparrow	F1 \uparrow	ACC \uparrow
ConVIRT [31]	80.42	58.42	76.11	64.31	43.29	57.00	61.01	16.28	71.02
GLoRIA [12]	71.45	49.01	71.29	53.42	38.23	40.47	66.10	17.32	77.00
BioViL [1]	82.80	58.33	76.69	70.79	48.55	69.09	69.12	19.31	79.16
PRIOR [6]	85.58	62.91	77.85	86.62	70.11	84.44	74.51	23.29	84.41
MedKLIP [30]	86.94	63.42	80.02	89.24	68.33	84.28	76.76	25.25	86.19
M&M (Ours)	88.91	66.58	83.14	91.15	71.58	86.15	77.92	27.55	88.52

Table 2. Comparing recent state-of-the-art methods on zero-shot classification task. We use AUC, F1 and ACC scores for comparison. Following MedKLIP [30] for evaluation on NIH Chest X-Ray, the metrics all refer to the macro average on the 14 diseases.

Methods	AUC \uparrow	F1 \uparrow	ACC \uparrow
ConVIRT [31]	52.08	69.02	52.66
GLoRIA [12]	66.59	70.07	60.83
BioViL [1]	53.82	69.10	53.75
MedKLIP [30]	73.96	76.70	70.06
M&M (Ours)	75.15	77.89	73.35

Table 3. Performance on Covid-19 CXR.

Methods	AUC \uparrow	F1 \uparrow	ACC \uparrow
ConVIRT [31]	77.00	56.76	69.19
GLoRIA [12]	77.74	57.98	71.45
BioViL [1]	75.40	55.72	69.14
MedKLIP [30]	78.98	58.26	72.80
M&M (Ours)	80.71	60.18	73.91

Table 4. Performance comparison on Edema severity grading multi-class task.

3.4 Grading

In addition to diagnosis, determining disease severity is essential. We refine pre-trained features for a multi-category classification task using the Edema Severity dataset, which includes classes 0 to 3, each indicating a different severity stage. The average scores across all severity levels are shown in Table 4.

3.5 Segmentation

Table 5 presents our fine-tuning experiments for segmenting three distinct diseases, where we utilize 1%, 10%, and 100% of the available data. Regardless

of the varying image distributions associated with each disease, our techniques consistently outperform current leading methods. We see significant gains in particular when data is scarce outperforming previous works by up to 2.96%.

Methods	RSNA Pneumonia			SIIM-ACR			Covid-19		
	1%	10%	100%	1%	10%	100%	1%	10%	100%
Scratch	43.47	60.47	70.68	21.33	33.23	74.47	14.81	23.67	32.28
ConVIRT [31]	57.06	64.91	72.01	54.06	61.21	73.52	19.95	27.24	37.37
GLoRIA [12]	65.55	69.07	73.28	56.73	57.78	76.94	18.89	28.09	38.69
BioVil [1]	68.24	70.38	72.49	62.67	69.98	78.49	21.13	32.39	41.62
PRIOR [6]	70.11	70.88	74.43	66.14	71.24	78.85	23.66	34.72	43.01
MedKLIP [30]	70.64	71.62	75.79	66.59	72.10	79.37	24.45	35.39	43.99
M&M (Ours)	72.28	73.11	76.68	69.55	73.47	80.28	28.25	37.32	45.04

Table 5. Evaluating Dice scores against leading methods for segmentation tasks, we analyze diseases with 1%, 10%, 100% labeled data, improving across all scenarios.

3.6 Ablation Study

We assess the effects of different image masking techniques—no masking, random patch masking [10], attention-guided masking (AttMask)[17], and salient patch selection (AutoMAE)[3]—and validate our manuscript generation approach. Our method boosts performance by as much as 8.06%, as Table 6 illustrates. For report generation, we examine four methods: original reports, extracted triplets, knowledge-enhanced triplets (KE-Triplet) via MedKLIP, and converting triplets into manuscripts, using the same masking for all (Table 7). KE-Triplet leads, surpassing even MedKLIP by up to 1.97%, highlighting our model’s efficacy on the NIH Chest X-Ray dataset.

Methods	AUC↑	F1↑	ACC↑
MaskVLM [18]	58.87	14.96	66.69
No Masking	61.48	16.33	70.54
MAE [10]	68.84	18.85	75.59
AttMask [17]	71.85	20.84	77.81
AutoMAE [3]	73.18	22.81	80.46
M&M (Ours)	77.92	27.55	88.52

Table 6. Ablation study on masking strategy, focusing on zero-shot setting to directly assess learned features.

Methods	AUC↑	F1↑	ACC↑
Original Report	69.95	20.04	77.71
Triplet	73.48	24.42	82.89
KE-Triplet	76.84	26.11	86.55
M&M (Ours)	77.92	27.55	88.52

Table 7. Ablation study on manuscript generation method selection, focusing solely on the zero-shot setting to evaluate the effectiveness of learned features.

4 Conclusion

Medical contrastive learning is challenged by inconsistent semantics and morphology in sample text pairs, causing semantic drifts. Report variability from multiple authors further complicates semantic interpretation. To address this, we propose a two-step approach. First, standardize reports into triplets. Then, convert triplets into binary “observations” that guide towards “verdicts”. For images, we use Meijering-based masking for pre-training instead of random masking to capture local context critical for medical images. Our multimodal contrastive learning framework, combining the standardized text and tailored visual representations, advances medical image analysis, achieving new state-of-the-art performances on multiple different downstream tasks. While the filter improves the performance in the discussed X-ray data, future work would be to test this on other modalities, such as MRI, which is noisier and more heterogeneous.

Disclosure of Interests. David A. Clifton was supported by the Pandemic Sciences Institute at the University of Oxford; the National Institute for Health Research (NIHR) Oxford Biomedical Research Centre (BRC); an NIHR Research Professorship; a Royal Academy of Engineering Research Chair; the Wellcome Trust funded VITAL project (grant 204904/Z/16/Z); the EPSRC (grant EP/W031744/1); and the InnoHK Hong Kong Centre for Cerebro-cardiovascular Engineering (COCHE).

References

1. Boecking, B., Usuyama, N., Bannur, S., Castro, D.C., Schwaighofer, A., Hyland, S., Wetscherek, M., Naumann, T., Nori, A., Alvarez-Valle, J., et al.: Making the most of text semantics to improve biomedical vision–language processing. In: *Eur. Conf. Comput. Vis.* pp. 1–21. Springer (2022) [7](#), [8](#)
2. Chauhan, G., Liao, R., Wells, W., Andreas, J., Wang, X., Berkowitz, S., Horng, S., Szolovits, P., Golland, P.: Joint modeling of chest radiographs and radiology reports for pulmonary edema assessment. In: *International Conference on Medical Image Computing and Computer Assisted Intervention*. pp. 529–539. Springer (2020) [6](#)
3. Chen, H., Zhang, W., Wang, Y., Yang, X.: Improving masked autoencoders by learning where to mask. *arXiv:2303.06583* (2023) [8](#)
4. Chen, T., Kornblith, S., Norouzi, M., Hinton, G.: A simple framework for contrastive learning of visual representations. In: *Int. Conf. Mach. Learn.* pp. 1597–1607. PMLR (2020) [7](#)
5. Chen, Y.C., Li, L., Yu, L., El Kholy, A., Ahmed, F., Gan, Z., Cheng, Y., Liu, J.: Uniter: Universal image-text representation learning. In: *Eur. Conf. Comput. Vis.* pp. 104–120. Springer (2020) [5](#)
6. Cheng, P., Lin, L., Lyu, J., Huang, Y., Luo, W., Tang, X.: Prior: Prototype representation joint learning from medical images and reports. In: *Int. Conf. Comput. Vis.* pp. 21361–21371 (2023) [6](#), [7](#), [8](#)
7. Deng, J., Dong, W., Socher, R., Li, L.J., Li, K., Fei-Fei, L.: Imagenet: A large-scale hierarchical image database. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 248–255. Ieee (2009) [1](#)

8. Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Deghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition scale. arXiv:2010.11929 (2020) **2**
9. Flanders, A.E., Lakhani, P.: Radiology reporting and communications: a look forward. *Neuroimaging Clinics* **22**(3), 477–496 (2012) **2**
10. He, K., Chen, X., Xie, S., Li, Y., Dollár, P., Girshick, R.: Masked autoencoders are scalable vision learners. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 16000–16009 (2022) **3, 8**
11. He, K., Fan, H., Wu, Y., Xie, S., Girshick, R.: Momentum contrast for unsupervised visual representation learning. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 9729–9738 (2020) **7**
12. Huang, S.C., Shen, L., Lungren, M.P., Yeung, S.: Gloria: A multimodal global-local representation learning framework for medical image recognition. In: *Int. Conf. Comput. Vis.* pp. 3942–3951 (2021) **6, 7, 8**
13. Irvin, J., Rajpurkar, P., Ko, M., Yu, Y., Ciurea-Ilcus, S., Chute, C., Marklund, H., Haghighi, B., Ball, R., Shpanskaya, K., et al.: Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In: *AAAI*. vol. 33, pp. 590–597 (2019) **6, 7**
14. Jain, S., Agrawal, A., Saporta, A., Truong, S.Q., Duong, D.N., Bui, T., Chambon, P., Zhang, Y., Lungren, M.P., Ng, A.Y., et al.: Radgraph: Extracting clinical entities and relations from radiology reports. arXiv:2106.14463 (2021) **4**
15. Ji, W., Chung, A.C.: Unsupervised domain adaptation for medical image segmentation using transformer with meta attention. *IEEE Trans. Med. Imag.* (2023) **1**
16. 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* (2019) **6**
17. Kakogeorgiou, I., Gidaris, S., Psomas, B., Avrithis, Y., Bursuc, A., Karantzas, K., Komodakis, N.: What to hide from your students: Attention-guided masked image modeling. In: *Eur. Conf. Comput. Vis.* pp. 300–318. Springer (2022) **8**
18. Kwon, G., Cai, Z., Ravichandran, A., Bas, E., Bhotika, R., Soatto, S.: Masked vision and language modeling for multi-modal representation learning. In: *Int. Conf. Learn. Represent.* (2022) **2, 5, 8**
19. for imaging informatics in medicine, S.: Siim-acr pneumothorax segmentation (2019), <https://www.kaggle.com/c/siim-acr-pneumothorax-segmentation> **6, 7**
20. Meijering, E., Jacob, M., Sarria, J.C., Steiner, P., Hirling, H., Unser, M.: Neurite tracing in fluorescence microscopy images using ridge filtering and graph searching: principles and validation. In: *ISBI*. pp. 1219–1222. IEEE (2004) **2, 3, 4**
21. Müller, P., Kaissis, G., Zou, C., Rueckert, D.: Joint learning of localized representations from medical images and reports. In: *Eur. Conf. Comput. Vis.* pp. 685–701. Springer (2022) **7**
22. Pavlova, M., Terhjan, N., Chung, A.G., Zhao, A., Surana, S., Aboutaleb, H., Gunraj, H., Sabri, A., Alaref, A., Wong, A.: Covid-net cxr-2: An enhanced deep convolutional neural network design for detection of covid-19 cases from chest x-ray images. *Frontiers in Medicine* **9**, 861680 (2022) **6**
23. 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: *Int. Conf. Mach. Learn.* pp. 8748–8763. PMLR (2021) **2, 5**
24. Sato, Y., Nakajima, S., Shiraga, N., Atsumi, H., Yoshida, S., Koller, T., Gerig, G., Kikinis, R.: 3-d multi-scale line filter for segmentation and visualization of curvilinear structures in medical images. *Medical image analysis* (1998) **2, 3**

25. Shih, G., Wu, C.C., Halabi, S.S., Kohli, M.D., Prevedello, L.M., Cook, T.S., Sharma, A., Amorosa, J.K., Arteaga, V., Galperin-Aizenberg, M., et al.: Augmenting the national institutes of health chest radiograph dataset with expert annotations of possible pneumonia. *Radiology: Artificial Intelligence* **1**(1) (2019) [6](#), [7](#)
26. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I.: Attention is all you need. *Adv. Neural Inform. Process. Syst.* **30** (2017) [2](#)
27. 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: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 2097–2106 (2017) [6](#)
28. Wang, X., Peng, Y., Lu, L., Lu, Z., Bagheri, M., Summers, R.M.: Tienet:text-image embedding network for common thorax disease classification and reporting in chest x-rays. In: *IEEE Conf. Comput. Vis. Pattern Recog.* pp. 9049–9058 (2018) [2](#)
29. Wen, Y., Chen, L., Deng, Y., Zhou, C.: Rethinking pre-training on medical imaging. *Journal of Visual Communication and Image Representation* **78**, 103145 (2021) [1](#)
30. Wu, C., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: Medklip: Medical knowledge enhanced language-image pre-training. *Int. Conf. Comput. Vis.* (2023) [2](#), [4](#), [6](#), [7](#), [8](#)
31. Zhang, Y., Jiang, H., Miura, Y., Manning, C.D., Langlotz, C.P.: Contrastive learning of medical visual representations from paired images and text. In: *Machine Learning for Healthcare Conference.* pp. 2–25. PMLR (2022) [7](#), [8](#)
32. Zhou, L., Liu, H., Bae, J., He, J., Samaras, D., Prasanna, P.: Self pre-training with masked autoencoders for medical image classification and segmentation. In: *2023 IEEE 20th International Symposium on Biomedical Imaging (ISBI).* pp. 1–6. IEEE (2023) [3](#)