



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.

Design as Desired: Utilizing Visual Question Answering for Multimodal Pre-training

Tongkun Su^{1,2,*}, Jun Li^{3,*}, Xi Zhang^{1,2}, Haibo Jin⁵, Hao Chen⁵, Qiong Wang¹,
Faqin Lv⁴, Baoliang Zhao^{1,**(✉)}, and Ying Hu^{1,**(✉)}

¹ Shenzhen Institute of Advanced Technology, Chinese Academy of Sciences
{tk.su, bl.zhao, ying.hu}@siat.ac.cn

² University of Chinese Academy of Sciences

³ Technical University of Munich, Munich Center for Machine Learning

⁴ Southern Medical University

⁵ The Hong Kong University of Science and Technology

Abstract. Multimodal pre-training demonstrates its potential in the medical domain, which learns medical visual representations from paired medical reports. However, many pre-training tasks require extra annotations from clinicians, and most of them fail to explicitly guide the model to learn the desired features of different pathologies. In this paper, we utilize Visual Question Answering (VQA) for multimodal pre-training to guide the framework focusing on targeted pathological features. We leverage descriptions in medical reports to design multi-granular question-answer pairs associated with different diseases, which assist the framework in pre-training without requiring extra annotations from experts. We also propose a novel pre-training framework with a quasi-textual feature transformer, a module designed to transform visual features into a quasi-textual space closer to the textual domain via a contrastive learning strategy. This narrows the vision-language gap and facilitates modality alignment. Our framework is applied to four downstream tasks: report generation, classification, segmentation, and detection across five datasets. Extensive experiments demonstrate the superiority of our framework compared to other state-of-the-art methods. Our code is available at <https://github.com/MoramiSu/QFT>.

Keywords: Multimodal Pre-training · Visual Question Answering

1 Introduction

In recent years, multimodal pre-training has received significant attention. In clinical settings, there is a vast amount of medical images and text reports stored in the database, making them easily available for pre-training tasks. Multimodal pre-training aims to learn generalized representations through the inherent relationships between image-report pairs to benefit various downstream tasks,

* The first two authors contributed equally to this work.

** Corresponding authors.

Original Ultrasound Report: The thyroid gland appears normal in size and shape. A hypochoic nodule is observed in the left lobe at the lower pole, measuring approximately [size], displaying clear boundaries and regular morphology. CDFI shows detectable blood flow signals. Multiple nodules are present in the right lobe, with the largest one located at the mid portion, exhibiting a mixed cystic and solid echogenicity, measuring approximately [size], and displaying clear boundaries and regular morphology. CDFI shows detectable blood flow signals in the periphery. The echogenicity of the remaining gland is increased with irregularities, presenting a reticular pattern. CDFI shows no abnormal blood flow signals within the gland.	Generate ultrasound report: (The original report showing on the left side.)	
	Coarse	
	Q:	Left lobe:
	A:	In the lower pole of the left lobe, a hypochoic nodule is visible, measuring approximately [size], displaying clear boundaries and a regular shape. CDFI shows detectable blood flow signals within the nodule.
	Medium	
	Q:	Right lobe: Q: Thyroid Overview:
	A:	... A: ...
	Fine	
	Q:	How is the size and shape of the thyroid? Q: How is the echogenicity of the thyroid?
	A:	Normal. A: The echogenicity is increased with irregularities, presenting a reticular pattern.
Q:	Are there any nodules in the images? Q: Where is the nodule located in the images?	
A:	Yes. A: Left lobe lower pole and right robe.	

Fig. 1. Overview of the VQA design. Left: the original report. Right: question and answer pairs. We design three levels of VQA to enable the model to capture multi-granular features. Numerical values are replaced with a special token [size].

particularly where annotated data is limited. Generally, there are two main-stream methods in multimodal pre-training: contrastive learning [9, 29, 31] and cross-modal reconstruction [33, 37]. The first method aims to learn discriminative representations by pulling closer the representations from paired images and reports while pushing away the unpaired ones. The second method generates one modality from the other, which assists the framework in understanding the relationship between two modalities. However, these methods have some drawbacks. On the one hand, some methods [31, 33] require extra annotations from clinicians, which is often time-consuming and cost-prohibitive. On the other hand, unlike supervised learning, where labels guide the model to learn the features of interest, self-supervised methods do not explicitly direct the model to focus on the specific features associated with different diseases.

Therefore, our objective is to explore an effective medical multimodal pre-training task that can guide the model to focus on the desired features of different pathologies without the need for further annotations by clinicians. Our attention turns to VQA, an essential task in cross-modal generation which requires the model to understand both visual and textual knowledge [2, 10, 16]. Clinicians just need to provide the pathologies they care about in the reports, and we can design different questions according to their instructions. By answering the questions during pre-training, the framework will attempt to focus on different levels of information according to the questions. Fig. 1 illustrates the main idea of question-answer design. We design VQA tasks at three levels of granularity to guide the model in learning essential details in the image and text.

Additionally, we propose a Quasi-textual Feature Transformer (QFT) module with a contrastive learning strategy to help the framework bridge the gap between image and text. This challenge stems from the fact that discriminative pathological features and tokens often occupy only a small fraction of the medical images and reports, making it difficult for the model to learn their associations effectively. Inspired by QFormer [13], we propose a QFT module, which utilizes contrastive learning to convert visual features into a quasi-textual domain that is

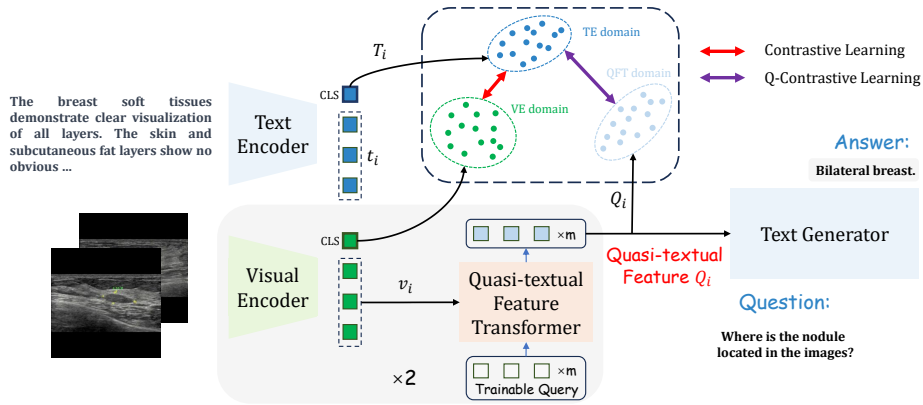


Fig. 2. Overview of our framework. The image and text features extracted by the visual and text encoders are aligned by the quasi-textual Feature transformer module with two contrastive learning tasks. Then the quasi-textual features are fed to the text generator to generate answers based on the questions during pre-training. TE domain, VE domain, QFT domain denote the latent space of global textual features T_i , global visual features V_i and quasi-textual features Q_i , respectively.

closer to the textual domain. This transformation narrows the distribution gap between two modalities and improves the model’s visual understanding capabilities. Moreover, we construct an ultrasound dataset with different organs for pre-training, comprising 10,720 ultrasound images and 5,360 reports. We transfer our model to four downstream tasks: report generation, classification, detection and segmentation. Extensive experiments demonstrate the superior performance of our framework compared to other State-Of-The-Art (SOTA) methods. Overall, our main contributions can be summarized as:

- We utilize VQA for multimodal pre-training in the medical field to assist the framework in focusing on the desired pathological features without extra expert annotations.
- We propose a QFT module with a contrastive learning strategy, which aligns the visual features into a quasi-textual domain to narrow the modality gap and facilitate modality alignment.
- Our approach demonstrates significant improvement in four downstream tasks: report generation, classification, detection and segmentation.

2 Methods

Overview. We aim to conduct multimodal pre-training through VQA, enabling the model to learn multi-granular pathological information, and thereby enhancing the performance across various downstream tasks. We employ an encoder-decoder architecture as the backbone of our model. Given a mini-batch of images

and their corresponding question-answer pairs with the batch size of B , we encode the images using the visual encoder (*e.g.*, ViT [6] or ResNet [8]) to obtain global features V_i and patch features $v_i, i \in B$. Subsequently, the text generator (*e.g.*, GPT [24]) generates answers based on the obtained image features and questions. Inspired by QFormer [13], we propose a Quasi-textual Feature Transformer (QFT) module on top of the general encoder-decoder framework to facilitate modality alignment. Our model also includes a text encoder (*e.g.*, BERT [5]), which encodes reports into global features T_i and token features t_i . Our framework is illustrated in Fig. 2.

Quasi-Textual Feature Transformer. In the medical field, discriminative pathological features generally occupy only a small fraction of the medical images, which poses a challenge for modality alignment. To address this, we turn to Contrastive Learning (CL), which is widely used in text-guided image generation to facilitate alignment. For example, DALL-E [26] uses CLIP [23] to encode text, which is pre-trained with CL. The reason is that CL maps paired images and texts to the same latent space, bringing the distributions of two modalities closer and thereby narrowing the modality gap. Specifically, we propose a Quasi-textual Feature Transformer (QFT) module with a CL strategy. The backbone of QFT is a multi-layer bidirectional transformer decoder [28]. It takes m trainable queries as input. The queries interact with patch features v_i through the cross-attention mechanism. To enforce the queries to extract quasi-textual features Q_i that are closer to the textual domain, we implement a Q-Contrastive Learning (QCL) for the QFT module. Unlike CL which depends on the similarity of global features V_i and T_i , we calculate the pairwise similarity between global textual features T_i and all m tokens in the output Q_i from the QFT module, selecting the highest score in the similarity matrix as the final score, which can be formulated as $s_q(Q_i, T_i) = \max_{0 \leq l \leq m} s(Q_i[l], T_i)$, where $s(\cdot, \cdot)$ computes the cosine similarity, $Q_i[l]$ is one of the output token in Q_i . Then we use this similarity score to calculate the InfoNCE loss [18]. By minimizing this loss, the model will align the latent space of Q_i with that of T_i together. Given that medical images of the same organ share similar features [12], leading to significant visual redundancy, we introduce a bottleneck design into our model. we set the m much smaller than the number of visual tokens v_i . This bottleneck requires the QFT module to compress visual information and extract the most robust information, while also reducing the computational cost of the text generator.

Additionally, we introduce the commonly adapted vision-language CL loss to enhance the model’s visual perception. As mentioned in [7], many QFormer-based models perform poorly on visual perception tasks, especially in position recognition, which is vital in the medical domain. We assume that this phenomenon is because these models all apply a pre-training process similar to QCL. Although QCL facilitates alignment between different modalities, directly pulling Q_i to the textual domain will lead to a significant loss of fine-grained visual information. Thus we use CL as a constraint to further align the latent space of visual features V_i and that of textual features T_i together, ultimately

maintaining more visual features in Q_i . Moreover, following [14], we introduce three buffers with a buffer size of N for QCL and CL. We store the global features T_i , V_i and quasi-textual features Q_i from the most recent batches as negative samples for the current batch. Thus, each batch has $N + B$ negative samples. The QCL and CL loss can be formulated as below, where τ_q and τ_c are the temperature parameter.

$$\mathcal{L}_{qcl} = -\frac{1}{2B} \sum_i^B \left[\log \frac{\exp(s_q(Q_i, T_i)/\tau_q)}{\sum_{j=1}^{N+B} \exp(s_q(Q_i, T_j)/\tau_q)} + \log \frac{\exp(s_q(Q_i, T_i)/\tau_q)}{\sum_{j=1}^{N+B} \exp(s_q(Q_j, T_i)/\tau_q)} \right] \quad (1)$$

$$\mathcal{L}_{cl} = -\frac{1}{2B} \sum_i^B \left[\log \frac{\exp(s(V_i, T_i)/\tau_c)}{\sum_{j=1}^{N+B} \exp(s(V_i, T_j)/\tau_c)} + \log \frac{\exp(s(V_i, T_i)/\tau_c)}{\sum_{j=1}^{N+B} \exp(s(V_j, T_i)/\tau_c)} \right] \quad (2)$$

Vision Question Answering. According to the data we gathered from the hospital, each report is associated with two images. Thus, we encode two images separately with a shared visual encoder and the QFT module. We take the average of two global visual features V_i for CL, the average of two quasi-textual features Q_i for QCL, and the concatenation of Q_i for generation. For each pair of medical images, we randomly sample three question-answer pairs from coarse, medium, and fine-grained VQA tasks respectively. We then use the questions as prompts to guide the generation. Given the question p_i and ground truth answer $y_i = \{y_{i,1}, y_{i,2}, \dots, y_{i,L}\}$ composed of L tokens, the language modeling loss can be formulated as:

$$l_{clm,mlm,flm} = \frac{1}{BL} \sum_i^B \sum_j^L [-\log(p(y_{i,j}|y_{i,<j}, p_i, q_i))] \quad (3)$$

$$\mathcal{L}_{lm} = \lambda_c l_{clm} + \lambda_m l_{mlm} + \lambda_f l_{flm} \quad (4)$$

where $l_{clm,mlm,flm}$ represent language modeling loss of coarse, medium and fine-grained VQA. $\lambda_c, \lambda_m, \lambda_f$ are hyperparameters. We apply ViT-B to initialize the visual encoder, the first six layers of ERNIE [30] to initialize text encoder, and GPT2-small [25, 35, 36] to initialize the text generator. For the QFT module, we initialize it with the first three layers of GPT2-small to equip it with the same knowledge as the text generator. We weight CL with λ to balance the constraint. Our loss function is shown as below:

$$\mathcal{L} = \lambda \mathcal{L}_{cl} + \mathcal{L}_{qcl} + \mathcal{L}_{lm} \quad (5)$$

3 Experiments

Dataset. We construct a dataset for experiments with 10,720 ultrasound images and 5,360 Chinese reports for breast and thyroid. The dataset is divided into training, validation, and test sets in a ratio of 7:1:2. The training and validation sets are used for pre-training, while the test set is used for downstream report generation. For three visual recognition tasks, we evaluate the effectiveness of pre-training on three public datasets: BUSI [1], AUITD [17], and DDTI [21].

VQA Design. Fig. 1 shows the overview of our VQA design (the texts are translated from Chinese to English by ChatGPT [19]). In coarse-grained VQA, we require the model to generate complete reports. During this process, the model will focus on the report formats and writing styles. However, the medical reports are usually very long and templated, so coarse-grained VQA is not enough to highlight the details in the reports. Therefore, we further design medium and fine-grained questions. In medium-grained VQA, we require the model to generate clinical descriptions for various anatomical regions, enabling the model to effectively differentiate anatomical structures. In fine-grained VQA, we design a series of questions based on the report, targeting subtle but crucial pathological visual features. Given that nodule recognition is the most important in ultrasound, our designed questions are mainly focusing on the presence and location of the nodules. Finally, we obtain a VQA dataset with 23,572 question-answer pairs. More examples are presented in the supplementary material.

Framework Setup. Considering that report generation is a generative task, while visual recognition is a discriminative task, we separately train the framework with two settings for different downstream tasks. For the report generation, since we utilize it as coarse-grained VQA in pre-training, we directly utilize the prompt in coarse-grained to require the model to generate reports. We set $\lambda_c = 9$, $\lambda_m = 1$, and $\lambda_f = 3$ to emphasize coarse-grained VQA. We use BLEU scores [20], METEOR [3] and ROUGE-L [15] to evaluate report quality. For the other three visual recognition tasks, to guide the model focusing on the pathological features emphasized in the middle and fine-grained VQA, we set $\lambda_c = 1$, $\lambda_m = 3$, and $\lambda_f = 9$. We transfer the pre-trained visual encoder into different downstream backbones (*e.g.*, YOLOv3 [27] for detection) and only fine-tune the other parts of the backbone. We compare accuracy (ACC) for classification, average precision (AP) for detection and DICE score for segmentation. We set $\lambda = 1$ in both settings. All baselines are retrained on our ultrasound dataset. More framework setup details are illustrated in supplementary materials.

4 Results and Analysis

Visual recognition. In this section, we compare our model’s performance in three visual recognition tasks with that of other SOTA multimodal pre-training methods in Fig. 3. We observe that most of the previous SOTA methods struggle to achieve balanced performance across various downstream tasks when pre-trained on our dataset which is relatively smaller than their pre-training dataset. For instance, MRM [37] achieves excellent segmentation performance but performs poorly in classification. Our approach achieves competitive performance across all tasks. We believe that this is because VQA enhances the utilization efficiency of images. Additionally, our multi-granular VQA enables the model to learn to extract multi-granular features, thus benefiting various downstream tasks. Fig. 4 visualizes some of the examples. Compared to GLORIA [9], our method achieves more precise recognition and less misidentification.




Fig. 3. Performance on visual recognition tasks. We compare our method with GLoRIA [9], MGCA [29] and MRM [37]. We only report detection results on ResNet since we use YOLOv3 [27] as our backbone, which is based on the convolutional neural network. Despite being pre-trained on a relatively small dataset, our method demonstrates balanced and nearly the best performance across various tasks, while other methods exhibit some shortcomings. Numerical details are in the supplementary materials.




Fig. 4. Examples of visual recognition tasks. Our method achieves superior performance in both classification, detection, and segmentation compared to GLoRIA.

Report Generation. We compare our method with other SOTA report generation methods in Table 1. Our method achieves the best performance in most cases. Specifically, our method outperforms the suboptimal model (DeltaNet [32] and TF [28]) by 6.856% on the breast dataset and by 18.95% on the thyroid dataset in BLEU-4. We assume that the benefit arises from the design of medium and fine-grained VQA which enhances the precision of nodule recognition, and the QFT module contributes to generating more accurate reports. Additionally, we observe that our method does not achieve the highest METEOR score, which is more related to the correctness of token order. However, we believe that the order for describing different anatomical regions in the medical report is less crucial. Our method pays more attention to pathological features rather than the order of each region, resulting in higher BLEU and ROUGE-L scores.

Table 1. Performance of report generation. Our method shows the best in most cases compared to other SOTA methods. B1 to B4 represent BLEU-1 to BLEU-4, while MR and RL represent METEOR and ROUGE-L, respectively.

Dataset	Method	B1↑	B2↑	B3↑	B4↑	MR↑	RL↑
Breast	TriNet [34]	0.693	0.594	0.533	0.478	0.439	0.742
	R2Gen [4]	0.663	0.611	0.572	0.541	0.411	0.685
	TF [28]	0.699	0.653	0.619	0.590	0.437	0.757
	R2GenRL [22]	0.616	0.528	0.464	0.414	0.470	0.599
	DeltaNet [32]	0.716	0.661	0.628	0.598	0.517	0.758
	Ours	0.730	0.695	0.665	0.639	0.442	0.774
Thyroid	TriNet [34]	0.645	0.510	0.421	0.345	0.409	0.678
	R2Gen [4]	0.578	0.532	0.492	0.457	0.369	0.664
	TF [28]	0.709	0.642	0.585	0.538	0.425	0.701
	R2GenRL [22]	0.672	0.595	0.531	0.479	0.500	0.651
	DeltaNet [32]	0.610	0.559	0.515	0.479	0.443	0.686
	Ours	0.755	0.712	0.674	0.640	0.444	0.761

Table 2. Results of ablation studies. P and R denote precision and recall for nodule recognition in reports. The introduction of QFT improves the quality of report generation, while the addition of VQA reduces the misidentification rate of nodules.

Multi Images	QCL	CL	VQA	B1↑	B2↑	B3↑	B4↑	MR↑	RL↑	P↑ (%)	R↑ (%)
				0.652	0.619	0.591	0.568	0.396	0.717	-	-
✓				0.703	0.662	0.627	0.598	0.416	0.743	-	-
✓	✓			0.678	0.641	0.609	0.582	0.407	0.722	-	-
✓	✓	✓		0.728	0.692	0.660	0.632	0.437	0.767	85.0	94.0
✓	✓	✓	✓	0.730	0.695	0.665	0.639	0.442	0.774	87.5	100

Ablation Study. In this section, we show the ablation study of our framework on the breast report generation task in Table 2. Firstly, although only using the most relevant image is a common approach in previous research [4, 11, 34], we observe that using multiple images as input improves the quality of the generated reports (0.717 vs 0.743). As mentioned in Sec. 2, adding QCL results in a trivial solution (0.722), while the addition of CL leads to a notable increase (0.767), demonstrating that the constraint of CL can effectively reduce the loss of image information during QCL. The addition of VQA further improves the performance (0.767 vs 0.774). Besides, we conduct further analysis to evaluate the impact of VQA: we randomly sample 100 reports and calculate the precision and recall for nodule recognition. Specifically, we regard it as a binary classification task, where reports mentioning the presence of nodules are labeled as 1, and otherwise as 0. We observe a significant improvement (85.0% vs 87.5% on precision, 94.0% vs 100% on recall). This is crucial in the medical field as it can reduce misdiagnosis.

5 Conclusion

In this paper, we explore the potential of VQA in pre-training. We design different levels of VQA targeting vital pathological features according to the description in the medical report without any extra annotations from clinicians. We also propose a pre-training framework with QFT, a module used to narrow the vision-language gap with a contrastive learning strategy. We demonstrate the effectiveness of our approach in report generation and three visual recognition tasks. Experimental results indicate that VQA guides the model focusing on desired pathological features, demonstrating the potential of VQA in pre-training. This work is an initial exploration. We will further investigate how to design reasonable questions and how to efficiently utilize VQA in pre-training.

Acknowledgments. This work was supported by the Guangzhou Science and Technology Program (No. 2023B01J0022), the Key Fundamental Research Program of Shenzhen (No. JCYJ20220818101408019), NSFC General Project (No. 62072452), and the Regional Joint Fund of Guangdong (No. 2021B1515130003, 2021B1515120011).

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

References

1. Al-Dhabyani, W., Gomaa, M., Khaled, H., Fahmy, A.: Dataset of breast ultrasound images. *Data in brief* **28**, 104863 (2020)
2. Bai, L., Islam, M., Ren, H.: Cat-vil: Co-attention gated vision-language embedding for visual question localized-answering in robotic surgery. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 397–407. Springer (2023)
3. Banerjee, S., Lavie, A.: Meteor: An automatic metric for mt evaluation with improved correlation with human judgments. In: *Proceedings of the acl workshop on intrinsic and extrinsic evaluation measures for machine translation and/or summarization*. pp. 65–72 (2005)
4. Chen, Z., Song, Y., Chang, T.H., et al.: Generating radiology reports via memory-driven transformer. *Conference on Empirical Methods in Natural Language Processing* (2020)
5. Devlin, J., Chang, M.W., Lee, K., et al.: Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018)
6. Dosovitskiy, A., Beyer, L., Kolesnikov, A., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. *International Conference on Learning Representations* (2020)
7. Fu, C., Chen, P., Shen, Y., et al.: Mme: A comprehensive evaluation benchmark for multimodal large language models. *arXiv preprint arXiv:2306.13394* (2023)
8. He, K., Zhang, X., Ren, S., et al.: Deep residual learning for image recognition. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 770–778 (2016)

9. Huang, S.C., Shen, L., Lungren, M.P., et al.: Gloria: A multimodal global-local representation learning framework for label-efficient medical image recognition. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 3942–3951 (2021)
10. Huang, X., Gong, H.: A dual-attention learning network with word and sentence embedding for medical visual question answering. *IEEE Transactions on Medical Imaging* (2023)
11. Jing, B., Xie, P., Xing, E.: On the automatic generation of medical imaging reports. Annual Meeting of the Association for Computational Linguistics (2017)
12. Li, J., Li, S., Hu, Y., et al.: A self-guided framework for radiology report generation. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 588–598. Springer (2022)
13. Li, J., Li, D., Savarese, S., et al.: BLIP-2: bootstrapping language-image pre-training with frozen image encoders and large language models (2023)
14. Li, J., Selvaraju, R., Gotmare, A., et al.: Align before fuse: Vision and language representation learning with momentum distillation. *Advances in neural information processing systems* **34**, 9694–9705 (2021)
15. Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out. pp. 74–81 (2004)
16. Liu, B., Zhan, L.M., Wu, X.M.: Contrastive pre-training and representation distillation for medical visual question answering based on radiology images. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 210–220. Springer (2021)
17. Maroua, A.: Algerian ultrasound images thyroid dataset: Auitd (2022), <https://www.kaggle.com/azouzmaraoua/datasets>, accessed on February 10, 2023
18. Oord, A.v.d., Li, Y., Vinyals, O.: Representation learning with contrastive predictive coding. arXiv preprint arXiv:1807.03748 (2018)
19. OpenAI: Introducing chatgpt (2023), <https://openai.com/blog/chatgpt/>, accessed on January 10, 2023
20. Papineni, K., Roukos, S., Ward, T., et al.: Bleu: a method for automatic evaluation of machine translation (2002)
21. Pedraza, L., Vargas, C., Narváez, F., et al.: An open access thyroid ultrasound image database. In: 10th International symposium on medical information processing and analysis. vol. 9287, pp. 188–193. SPIE (2015)
22. Qin, H., Song, Y.: Reinforced cross-modal alignment for radiology report generation. In: Findings of the Association for Computational Linguistics: ACL 2022. pp. 448–458 (2022)
23. Radford, A., Kim, J.W., Hallacy, C., et al.: Learning transferable visual models from natural language supervision. In: International conference on machine learning. pp. 8748–8763. PMLR (2021)
24. Radford, A., Narasimhan, K., Salimans, T., et al.: Improving language understanding by generative pre-training (2018)
25. Radford, A., Wu, J., Child, R., et al.: Language models are unsupervised multitask learners (2019)
26. Ramesh, A., Pavlov, M., Goh, G., et al.: Zero-shot text-to-image generation. In: International conference on machine learning. pp. 8821–8831. Pmlr (2021)
27. Redmon, J., Farhadi, A.: Yolov3: An incremental improvement. arXiv preprint arXiv:1804.02767 (2018)
28. Vaswani, A., Shazeer, N., Parmar, N., et al.: Attention is all you need. *Advances in neural information processing systems* **30** (2017)

29. Wang, F., Zhou, Y., Wang, S., et al.: Multi-granularity cross-modal alignment for generalized medical visual representation learning. *Advances in Neural Information Processing Systems* **35**, 33536–33549 (2022)
30. Wang, Q., Dai, S., Xu, B., et al.: Building chinese biomedical language models via multi-level text discrimination. *arXiv preprint arXiv:2110.07244* (2021)
31. Wang, Z., Wu, Z., Agarwal, D., et al.: Medclip: Contrastive learning from unpaired medical images and text. *Conference on Empirical Methods in Natural Language Processing* (2022)
32. Wu, X., Yang, S., Qiu, Z., et al.: Deltanet: Conditional medical report generation for covid-19 diagnosis. *International Conference On Computational Linguistics* (2022)
33. Xie, Y., Gu, L., Harada, T., et al.: Medim: Boost medical image representation via radiology report-guided masking. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 13–23. Springer (2023)
34. Yang, Y., Yu, J., Zhang, J., et al.: Joint embedding of deep visual and semantic features for medical image report generation. *IEEE Transactions on Multimedia* (2021)
35. Zhao, Z., Chen, H., Zhang, J., et al.: Uer: An open-source toolkit for pre-training models. *Conference on Empirical Methods in Natural Language Processing-International Joint Conference on Natural Language Processing* p. 241 (2019)
36. Zhao, Z., Li, Y., Hou, C., et al.: Tencentpretrain: A scalable and flexible toolkit for pre-training models of different modalities. *Annual Meeting of the Association for Computational Linguistics* p. 217 (2023)
37. Zhou, H.Y., Lian, C., Wang, L., et al.: Advancing radiograph representation learning with masked record modeling. *International Conference on Learning Representations* (2023)