

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.

A Refer-and-Ground Multimodal Large Language Model for Biomedicine

Xiaoshuang Huang^{1,2*}, Haifeng Huang¹, Lingdong Shen³, Yehui Yang^{1†},

Fangxin Shang¹, Junwei Liu¹, and Jia Liu¹(\boxtimes)

¹ Healthcare Group, Baidu Inc, Beijing 100085, China

² China Agricultural University, Beijing 100083, China

³ MAIS, Institute of Automation, Chinese Academy of Sciences (CASIA), Beijing

100086, China

huangxiaoshuang@cau.edu.cn, liujia9001cc@163.com

Abstract. With the rapid development of multimodal large language models (MLLMs), especially their capabilities in visual chat through refer and ground functionalities, their significance is increasingly recognized. However, the biomedical field currently exhibits a substantial gap in this area, primarily due to the absence of a dedicated refer and ground dataset for biomedical images. To address this challenge, we devised the Med-GRIT-270k dataset. It comprises 270k question-and-answer pairs and spans eight distinct medical imaging modalities. Most importantly, it is the first dedicated to the biomedical domain and integrating refer and ground conversations. The key idea is to sample large-scale biomedical image-mask pairs from medical segmentation datasets and generate instruction datasets from text using chatGPT. Additionally, we introduce a Refer-and-GrounD Multimodal Large Language Model for Biomedicine (BiRD) by using this dataset and multi-task instruction learning. Extensive experiments have corroborated the efficacy of the Med-GRIT-270k dataset and the multi-modal, fine-grained interactive capabilities of the BiRD model. This holds significant reference value for the exploration and development of intelligent biomedical assistants. The repository is at https://github.com/ShawnHuang497/BiRD

 ${\bf Keywords:} \ {\rm Referring \ and \ grounding} \cdot {\rm Instruction \ dataset} \cdot {\rm Biomedicine}.$

1 Introduction

Multimodal large language models (MLLMs) have become a popular area of research, with numerous applications in the field of visual languages, such as, Visual Question Answering (VQA), open vocabulary detection, and so on. Nonetheless, the unique challenges presented by the realm of biomedicine, which starkly

^{*} Work performed during an internship at Baidu Inc.

 $^{^{(\}boxtimes)}$ Corresponding author.

[†] Project Leader.

2 Xiaoshuang Huang et al.



Fig. 1: BiRD empowers multimodal large language models in biomedicine with sophisticated referring and grounding capabilities. For more equitable comparison, we append spatial information to each LLaVa-med test, such as "The image size is [w, h], and the origin of the coordinate system is located in the upper left corner of the image.", where w and h denote width and height, respectively.

contrasts with the natural world, often render conventional visual assistants inept. They may either refrain from responding to biomedical queries or, worse, provide inaccurate responses or entirely fabricated information [11].

Despite existing research within the realm of biomedical MLLMs, current studies have predominantly focused on image description and VQA, leaving a notable gap in capabilities concerning referring and grounding (shown in Fig. 1). The act of referring demands a model's accurate semantic comprehension of specified regions, while grounding necessitates the localization of regions based on semantic descriptions provided [27]. These fine-grained multimodal capabilities are essential for both the interaction process between intelligent biomedical assistants and patients and for biomedical education. This capability not only makes the information exchange process more intuitive but also significantly enhances the accuracy and efficiency of information exchange. A key factor hindering the development of this capability in the field of biomedicine is the lack of multimodal fine-grained interactive datasets.

To address these challenges, we develop the Bio**Med**ical **G**round-and-**R**efer Instruction-Tuning (Med-GRIT-270k) dataset by leveraging the medical segmentation dataset (SA-Med2D-20M [26]). Then a biomedical refer-and-ground multimodal large language model was explored with the Med-GRIT-270k and A Refer-and-Ground Multimodal Large Language Model for Biomedicine

multi-task instruction learning method. The paper principally contributes the following:

- Med-GRIT-270k Dataset. Large-scale biomedical image-mask pairs are transformed into multi-modal conversations by leveraging chatGPT [19] in a novel process. It is the first dataset in biomedicine to integrate referring, grounding, and conversations.
- The first **Bi**omedical **R**efer-and-groun**D** Multimodal Large Language Model (BiRD). It is fine-tuned by multi-task instruction learning for the biomedical domain with self-generated data. This validates the effectiveness of multitask instruction tuning and highlights best practices for adapting the MLLMs to the specialized domain.
- To advance biomedical multi-modal learning research, we will release the Med-GRIT-270k dataset and a comprehensive codebase for community use.

2 Related Work

Biomedical Multi-modal Large Language Models. Amidst the rapid development of Large Language Models (LLMs) and the success of instruction-tuned LLMs within the general domain [25,17,28,30,5,27], researchers in the biomedical field have been fervently exploring the expansion of these models' capabilities. Recent studies have increasingly concentrated on the domain of MLLMs, with notable endeavors within the biomedical sector including BioMedGPT [18], RadFM [24], LLaVa-Med [12], and so on [29,16,8,23,7,21]. These methodologies have significantly propelled the development of MLLMs in the biomedical realm. For instance, LLaVa-Med [12], utilizing pre-trained LLMs for visual instruction tuning, has established a unique, end-to-end multi-modal biomedical chatbot capable of processing image inputs. RadFM [24] is a MLLM supporting 2D/3D radiographic imaging input for the medical domain. However, due to various challenges, biomedical MLLMs capable of supporting fine-grained interactions have yet to emerge.

MLLMs for Referring and Grounding. In natural images, the largescale public datasets have greatly supported the exploration into the sophisticated understanding abilities of multimodal large language models (MLLMs), such as Gpt4ROI [30], Ferret [27], QWen-VL [3], and so on. Although some work [14,9] has already begun to investigate grounding in biomedicine, it can only be applied to small models, as the amount of data is limited and there are only a few modalities. The paramount factor underlying the success of these initiatives is their access to pertinent, large-scale datasets. For instance, QWen-VL uses around 80M data for referring and grounding. However, the multi-modal fine-grained interactive dataset in biomedical is virtually nonexistent. 4 Xiaoshuang Huang et al.



Fig. 2: An instance of our generated instruction-following data. Top: the meta information is created according to rules in medical segmentation datasets, and the image caption was generated from chatGPT. Bottom: the instruction following data generated by chatGPT.

3 Med-GRIT-270k: Biomedical Ground-and-Refer Instruction-tuning Dataset

We've created the first biomedical refer-and-ground instruction-tuning dataset to address the lack of such resources. It was generated through the collaborative efforts of humans and Artificial Intelligence (AI), derived from large-scale biomedical image segmentation datasets. The generation process can be divided into three steps: (i) Manually generating instance-level meta information for each image based on its mask. (ii) Employing an AI assistant to generate global information for the images. (iii) Utilizing the AI assistant to craft fine-grained conversations based on the meta information and global image information obtained in the previous steps.

Generating Instance-level Meta Information. We first sampled biomedical image-mask pairs from the SA-Med2D-20M [26]. Ultimately, approximately 60K images were sampled from this dataset, considering the diversity of modality and redundancy. For instance, the original dataset includes a plethora of 2D slices from 3D data, leading to excessive data similarity. Subsequently, we calculated the coordinates of each instance based on the instance-level masks. Specifically, spatial locations are delineated via the textual representation in the format $[X_{topleft}, Y_{topleft}, X_{bottomright}, Y_{bottomright}]$, and normalize the coordinates to fall within the range [0,1]. Finally, we enrich the images with additional details to compile the meta information, which includes modality, scanned region, orientation, and object coordinates.

Generating Image Captions. We utilize meticulously designed prompts along with the meta information provided to ChatGPT [19], thereby acquiring the global information for each image.

Biomedical Instruction-Tuning Data. Spatial understanding is manifested through various task formats. This primarily encompasses two distinct types and their corresponding task names: (i) Region-in and Text-out: Referring Object Classification (ROC), Referring Captioning (RC), (ii) Text-in and Region-out: Visual Grounding (VG), and (iii) Text-in and Text-out: Medical Image Analysis (MIA). To reduce ambiguity and enhance the model's capability for fine-grained visual comprehension, some essential strategies are adopted. The special tokens (<ref> and </ref>) are introduced, marking the content referred to by the bounding box. This aptly associates bounding boxes with their corresponding descriptive words or sentences. Subsequently, we instructed ChatGPT to design a question and answer for each task.

Finally, We mapped the coordinates within the range [0, 1000] and reformatted them as $(X_{topleft}, Y_{topleft})$, $(X_{bottomright}, Y_{bottomright})$. To differentiate between detection strings and regular text strings, two special tokens (<box> and </box>) are appended at the start and end of the bounding box string, respectively. Fig. 2 shows an example of our instruction-following data.

4 Multi-task Instruction Learning

We aim to imbue MLLMs with grounding and referring capacities via multi-task learning, simultaneously ensuring the retention of the MLLM's essential conversational proficiency. This section will henceforth elucidate from two perspectives: the architecture of the model and multi-task instruction training.

4.1 Model Architecture

We utilize Qwen-VL [3], a comprehensive multimodal conversational model, as the foundational general-domain language model. Specifically, the visual encoder employs the Vision Transformer (ViT) [6] architecture, initialized with pre-trained weights from OpenAI's CLIP ViT-BigG [10]. The vision-language adapter utilizes cross-attention with a trainable query. The large language model incorporates the pre-trained Qwen-7B [2].



Fig. 3: **Overreview.** Left: the training set (Top) and test set (Bottom) distribution of conversation turns in Med-GRIT-270k we collected. Right: the architecture of the Biomedical refer-and-ground multimodal large language model (**BiRD**), which is based on Qwen-VL [3]. We have developed it from the 240k data and evaluated it on 30k data.

4.2 Multi-task Instruction Training

Considering that the base model already possesses the capability to refer or ground within natural images, we employ only one stage to finetune it based on the pre-trained base model on the Med-GRIT-240k dataset. As illustrated in Fig. 3), We solely fine-tune the cross-attention and LLM parameters, while the visual encoder remains frozen. The input images are processed through the ViT-BigG [10] and vision-language adapter, yielding fixed-length sequences of visual features. We then append the markers ($\langle img \rangle$ and $\langle /img \rangle$) to the start and end of the image feature sequence, respectively, to denote the beginning and end of visual content. We fine-tuned the model using a dataset comprising 60k images and a total of 240k dialogue turns. The global training batch size is 128. The learning rate is 2e-5 and the scheduler is cosine. The multi-task instruction training just took 30 hours on $4 \times A100(40G)$ GPUs.

5 Experiments

In this section, we execute a thorough evaluation across diverse multimodal tasks to holistically gauge our models' proficiency in visual comprehension.

Table 1: Comparison with LLaVa-Med [12] and study on the multimodal dataset scales.

Model	Test dataset	VG (Recall@ $0.5\uparrow$)	ROC (Recall \uparrow)	RC (SPICE \uparrow)	MIA (mBMR↑)	$ $ Average \uparrow
LLaVa-Med [12]	Med-GRIT-Test30k	0	2.75	8.18	11.20	5.53
BiRD-Med-GRIT-20k	Med-GRIT-Test30k	38.59	47.94	29.02	27.22	35.69
BiRD-Med-GRIT-40k	Med-GRIT-Test30k	46.30	51.84	50.32	30.14	44.65
BiRD-Med-GRIT-80k	Med-GRIT-Test30k	52.87	52.02	52.84	44.83	50.64
BiRD-Med-GRIT-270k	Med-GRIT-Test30k	53.92	65.33	55.23	52.17	56.66
LLaVa-Med [12]	LLaVa-Med-qa0.2k	-	-	-	20.04	-
BiRD-Med-GRIT-270k	LLaVa-Med-qa0.2k	-	-	-	10.55	-

Where is the kidney

located in this image

<ref>Kidney</ref> is

Table 2: The capabilities of various biomedical MLLMs. Note that the "Modality" denotes image modality.

BiRD	~	1	~	~	8	hallucination in BiRD.
Med-palm m [22]	~	×	X	X	6	Fig. 4: The example of object
RadFM [24]	1	X	X	X	6	present in the image.
LLaVa-Med [12]	1	X	X	X	5	The kidney is not
	MIA	ROC	RC	\mathbf{VG}	Modality	<pre><box>(281,437),(399,52 5)</box>.</pre>
						located at the coordinates

Evaluation dataset. We randomly selected approximately 12% of the images and dialogues from the constructed Med-GRIT-270k dataset to serve as the test set. Given that a single 3D dataset contains multiple data slices, we extracted cases in their entirety to prevent leakage of test set data into the training set. This ensures that different slices from the same 3D dataset do not concurrently appear in both the training and test sets, thereby guaranteeing the reliability of the test results.

Evaluation metrics. The evaluation metrics for the four tasks are Recall@0.5, Recall, Spice [1], and mBMR, respectively. Recall@0.5 denotes a prediction as correct only when the intersection over union (IoU) between the predicted bounding box and the ground truth exceeds 0.5. The mBMR utilized for assessing the MIA task is the mean value of BLEU@4 [20], METEOR [4], and ROUGE-L [15], offering a more comprehensive evaluation of the prediction quality than a solitary metric.

Comparison. As shown in Table 2, we are the pioneers in developing a medical MLLM with referring and grounding capabilities, and existing MLLMs (such as Qwen-VL [3], GPT-4 [19], MiniGPT-v2 [5], etc.) have not seen medical referring and grounding data. So we will not compare them on evaluation metrics, as it would be profoundly unfair.

As illustrated in Table 1, we present the quantitative test outcomes for LLaVa-Med [12] and the impact of the data scale on these results. Between rows 3 and 6, we observe the performance of the BiRD-Med-GRIT model across varying data scales. With the expansion of training data, all metrics exhibit significant enhancements, with the average rising from 35.69 to 56.66. This underscores

8 Xiaoshuang Huang et al.

Table 3: The performance of the BiRD model across various tasks and modalities on the Med-GRIT-270k test dataset.

	Metric	CT	\mathbf{MR}	X-ray	PET	Endoscopy	Dermoscopy	Fundus	Ultrasound	Average
VG	Recall@0.5↑	44.47	29.26	41.73	56.46	53.60	75.63	84.15	46.04	53.92
ROC	Recall↑	34.76	61.79	53.74	-	60.40	96.61	-	84.65	65.33
\mathbf{RC}	Spice↑	41.88	51.69	37.39	47.95	54.07	77.44	48.73	82.65	55.23
MIA	$mBMR\uparrow$	47.01	49.35	37.17	57.15	39.91	72.13	48.87	65.78	52.17
Average	-	43.03	48.02	42.51	53.85	51.99	80.45	60.58	69.78	-

the efficacy of augmenting dataset size in bolstering the model's proficiency on multimodal datasets. Notably, at the 240k dataset level, the model achieved the highest scores across all metrics, showcasing optimal overall performance.

From the first and sixth rows of Table 1, it is evident that the LLaVa-Med [12] model demonstrates subpar performance on the Med-GRIT-Test30k dataset, particularly in terms of no efficacy in region-level visual content localization (with the Recall@0.5 of 0). Simultaneously, we evaluated our model on the LLaVa-Med qa-0.2k test set as well. As indicated in the last two rows of Table 1, due to not being trained on the LLaVa-Med [12] dataset, our performance metrics on its test set were marginally lower than its own. However, on similar MIA tasks within our test set, LLaVa-Med [12](with an mBMR of 11.20), significantly underperformed in comparison to our model (with an mBMR of 52.17).

Main Results. As shown in Table 3, we display the performance of the BiRD model across four distinct tasks in eight different medical imaging modalities. The **ROC** task tests the MLLM's understanding of text related to specific image areas and their visual details. The PET and Fundus, which focus on only one category, are not trained or evaluated. We find the recall of ROC mainly depends on the variety and distinctiveness of objects and features across image modalities. The **RC** task tests the model's ability to recognize image regions and describe them in words. The model does well with Ultrasound and Dermoscopy images but struggles with the more diverse CT images, where performance lags. The **VG** task tests how well the model matches text descriptions to image areas. MR modality performed the worst, likely because it mostly features tumor tissues, with far fewer anatomical structures. This issue is also seen in ultrasound images. The **MIA** task checks the model's understanding of medical images. The 4th row in Table 3 shows the model has some level of analysis and understanding across almost all modalities.

Across the four evaluated tasks, it is apparent that the Dermoscopy modality consistently exhibits the highest performance metrics. This can be attributed to the distinct visual features, a reduced number of object categories, and the substantial proportion of the image occupied by the object regions, collectively simplifying the task for this particular modality.

Object Hallucination. As Fig. 4 shows, we have also observed instances of object hallucination in BiRD. This phenomenon is common and has also been observed in other MLLMs [13]. We believe this is attributed to the fact that

the model's visual encoder is frozen, and its initialized parameters have scarcely encountered medical imaging, resulting in a lack of comprehensive understanding of specific domains or topics in feature extraction. In a word, this phenomenon should receive increased attention in future research endeavors.

6 Conclusion

In this paper, to develop a single MLLM assistant capable of handling multiple vision-language tasks, we propose a Med-GRIT-270k dataset. By leveraging the dataset, we introduce the BiRD model, a **Bi**omedical **R**efer-and-Groun**D** Multi-modal Large Language Model. We verified BiRD on a diverse 30k question-and-answer test set, encompassing multimodal and multitask scenarios. The BiRD showcases a highly promising direction for developing intelligent biomedical assistants. To our knowledge, Med-GRIT-270k and BiRD are respectively the first refer-and-ground dataset and fine-grained interactive MLLM in the realm of biomedicine. We will release both the dataset and model to foster the development of intelligent biomedical assistants.

Limitations. Although this work developed a novel multimodal dataset in biomedicine, during the data construction process, most of the raw datasets only annotated certain organs or diseases for a sample. This makes it difficult to construct highly correlated negative samples. This issue will be a focus in the subsequent data construction work.

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

References

- Anderson, P., Fernando, B., Johnson, M., Gould, S.: Spice: Semantic propositional image caption evaluation. In: Computer Vision–ECCV 2016: 14th European Conference, Amsterdam, The Netherlands, October 11-14, 2016, Proceedings, Part V 14. pp. 382–398. Springer (2016)
- Bai, J., Bai, S., Chu, Y., Cui, Z., Dang, K., Deng, X., Fan, Y., Ge, W., Han, Y., Huang, F., et al.: Qwen technical report. arXiv preprint arXiv:2309.16609 (2023)
- Bai, J., Bai, S., Yang, S., Wang, S., Tan, S., Wang, P., Lin, J., Zhou, C., Zhou, J.: Qwen-vl: A versatile vision-language model for understanding, localization, text reading, and beyond (2023)
- 4. 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)
- Chen, J., Zhu, D., Shen, X., Li, X., Liu, Z., Zhang, P., Krishnamoorthi, R., Chandra, V., Xiong, Y., Elhoseiny, M.: Minigpt-v2: large language model as a unified interface for vision-language multi-task learning. arXiv preprint arXiv:2310.09478 (2023)

- 10 Xiaoshuang Huang et al.
- Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al.: An image is worth 16x16 words: Transformers for image recognition at scale. arXiv preprint arXiv:2010.11929 (2020)
- Eslami, S., Meinel, C., De Melo, G.: Pubmedclip: How much does clip benefit visual question answering in the medical domain? In: Findings of the Association for Computational Linguistics: EACL 2023. pp. 1151–1163 (2023)
- Han, T., Adams, L.C., Nebelung, S., Kather, J.N., Bressem, K.K., Truhn, D.: Multimodal large language models are generalist medical image interpreters. medRxiv pp. 2023–12 (2023)
- Huang, X., Li, H., Cao, M., Chen, L., You, C., An, D.: Cross-modal conditioned reconstruction for language-guided medical image segmentation. arXiv preprint arXiv:2404.02845 (2024)
- Ilharco, G., Wortsman, M., Wightman, R., Gordon, C., Carlini, N., Taori, R., Dave, A., Shankar, V., Namkoong, H., Miller, J., et al.: Openclip (2021). URL: https://doi.org/10.5281/zenodo 7439141
- Lee, P., Bubeck, S., Petro, J.: Benefits, limits, and risks of gpt-4 as an ai chatbot for medicine. New England Journal of Medicine p. 1233-1239 (Mar 2023). https:// doi.org/10.1056/nejmsr2214184, http://dx.doi.org/10.1056/nejmsr2214184
- Li, C., Wong, C., Zhang, S., Usuyama, N., Liu, H., Yang, J., Naumann, T., Poon, H., Gao, J.: Llava-med: Training a large language-and-vision assistant for biomedicine in one day. Advances in Neural Information Processing Systems 36 (2024)
- Li, Y., Du, Y., Zhou, K., Wang, J., Zhao, W.X., Wen, J.R.: Evaluating object hallucination in large vision-language models. arXiv preprint arXiv:2305.10355 (2023)
- Li, Z., Li, Y., Li, Q., Wang, P., Guo, D., Lu, L., Jin, D., Zhang, Y., Hong, Q.: Lvit: language meets vision transformer in medical image segmentation. IEEE transactions on medical imaging (2023)
- Lin, C.Y.: Rouge: A package for automatic evaluation of summaries. In: Text summarization branches out. pp. 74–81 (2004)
- Liu, F., Zhu, T., Wu, X., Yang, B., You, C., Wang, C., Lu, L., Liu, Z., Zheng, Y., Sun, X., et al.: A medical multimodal large language model for future pandemics. NPJ Digital Medicine 6(1), 226 (2023)
- 17. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. Advances in neural information processing systems **36** (2024)
- Luo, Y., Zhang, J., Fan, S., Yang, K., Wu, Y., Qiao, M., Nie, Z.: Biomedgpt: Open multimodal generative pre-trained transformer for biomedicine. arXiv preprint arXiv:2308.09442 (2023)
- 19. OpenAI, O.: Gpt-4 technical report (Mar 2023)
- Papineni, K., Roukos, S., Ward, T., Zhu, W.J.: Bleu: a method for automatic evaluation of machine translation. In: Proceedings of the 40th annual meeting of the Association for Computational Linguistics. pp. 311–318 (2002)
- Shen, L., Shang, F., Yang, Y., Huang, X., Xiang, S.: Segicl: A universal in-context learning framework for enhanced segmentation in medical imaging. arXiv preprint arXiv:2403.16578 (2024)
- 22. Tu, T., Azizi, S., Driess, D., Schaekermann, M., Amin, M., Chang, P.C., Carroll, A., Lau, C., Tanno, R., Ktena, I., et al.: Towards generalist biomedical ai. NEJM AI 1(3), AIoa2300138 (2024)
- 23. Wang, Z., Wu, Z., Agarwal, D., Sun, J.: Medclip: Contrastive learning from unpaired medical images and text. arXiv preprint arXiv:2210.10163 (2022)

A Refer-and-Ground Multimodal Large Language Model for Biomedicine

- Wu, C., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: Towards generalist foundation model for radiology. arXiv preprint arXiv:2308.02463 (2023)
- Wu, S., Fei, H., Qu, L., Ji, W., Chua, T.S.: Next-gpt: Any-to-any multimodal llm. arXiv preprint arXiv:2309.05519 (2023)
- Ye, J., Cheng, J., Chen, J., Deng, Z., Li, T., Wang, H., Su, Y., Huang, Z., Chen, J., Jiang, L., et al.: Sa-med2d-20m dataset: Segment anything in 2d medical imaging with 20 million masks. arXiv preprint arXiv:2311.11969 (2023)
- 27. You, H., Zhang, H., Gan, Z., Du, X., Zhang, B., Wang, Z., Cao, L., Chang, S.F., Yang, Y.: Ferret: Refer and ground anything anywhere at any granularity. arXiv preprint arXiv:2310.07704 (2023)
- Zhan, J., Dai, J., Ye, J., Zhou, Y., Zhang, D., Liu, Z., Zhang, X., Yuan, R., Zhang, G., Li, L., et al.: Anygpt: Unified multimodal llm with discrete sequence modeling. arXiv preprint arXiv:2402.12226 (2024)
- Zhang, S., Xu, Y., Usuyama, N., Bagga, J., Tinn, R., Preston, S., Rao, R., Wei, M., Valluri, N., Wong, C., et al.: Large-scale domain-specific pretraining for biomedical vision-language processing. arXiv preprint arXiv:2303.00915 2(3), 6 (2023)
- Zhang, S., Sun, P., Chen, S., Xiao, M., Shao, W., Zhang, W., Chen, K., Luo, P.: Gpt4roi: Instruction tuning large language model on region-of-interest. arXiv preprint arXiv:2307.03601 (2023)