

Medical Large Vision Language Models with Multi-Image Visual Ability

Xikai Yang¹, Juzheng Miao¹, Yuchen Yuan¹, Jiaze Wang¹, Qi Dou^{1,2}, Jinpeng Li^{1✉}, and Pheng-Ann Heng^{1,2}

¹ Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, China

² Institute of Medical Intelligence and XR, The Chinese University of Hong Kong, Hong Kong, China
jpli21@cse.cuhk.edu.hk

Abstract. Medical large vision-language models (LVLMs) have demonstrated promising performance across various single-image question answering (QA) benchmarks, yet their capability in processing multi-image clinical scenarios remains underexplored. Unlike single image based tasks, medical tasks involving multiple images often demand sophisticated visual understanding capabilities, such as temporal reasoning and cross-modal analysis, which are poorly supported by current medical LVLMs. To bridge this critical gap, we present the Med-MIM instruction dataset, comprising 83.2K medical multi-image QA pairs that span four types of multi-image visual abilities (temporal understanding, reasoning, comparison, co-reference). Using this dataset, we fine-tune Mantis and LLaVA-Med, resulting in two specialized medical VLMs: MIM-LLaVA-Med and Med-Mantis, both optimized for multi-image analysis. Additionally, we develop the Med-MIM benchmark to comprehensively evaluate the medical multi-image understanding capabilities of LVLMs. We assess eight popular LVLMs, including our two models, on the Med-MIM benchmark. Experimental results show that both Med-Mantis and MIM-LLaVA-Med achieve superior performance on the held-in and held-out subsets of the Med-MIM benchmark, demonstrating that the Med-MIM instruction dataset effectively enhances LVLMs’ multi-image understanding capabilities in the medical domain. The Med-MIM instruction dataset, benchmark, and fine-tuned models are available at [Med-MIM](#).

Keywords: Medical LVLMs · Multi-Image · Instruction Tuning.

1 Introduction

Large Vision-Language Models (LVLMs), with proper visual instruction tuning, such as LLaVA [14], have demonstrated remarkable potential in various image perception and QA tasks. To further adapt LVLMs to the medical domain, recent developments have introduced medical-specific LVLMs [24,26,12,2], such as LLaVA-Med [12] with impressive diagnostic capabilities. However, these models are typically designed for single-image based tasks, lacking instruction tuning

tailored for multi-image question answering. As a result, their performance is often limited when addressing medical multi-image tasks, which are more complex yet crucial in the field of medical image analysis. Unlike natural images, medical multi-image is often closely tied to clinical needs. During the diagnostic process, doctors typically analyze multi-visit, multi-view, and multi-modality data in a systematic manner to produce a comprehensive and accurate assessment [10,19]. For instance, given screening data from two consecutive visits of a patient, the goal might be to determine whether the patient’s condition is worsening [20]. In retrospective studies, such as monitoring Alzheimer’s disease, the question becomes whether we can predict a patient’s future cognitive status using his longitudinal neuroimaging data [9]. More commonly, when working with multiple imaging modalities, such as MRI from various sequences, the challenge is to integrate information from diverse sources to provide a reliable diagnosis [1,4].

While some pioneering works have started to explore the performance of LVLMs in natural multi-image settings [17,8,15], current medical LVLMs still overlook the development of multi-image visual capabilities. Due to the challenges associated with collecting and processing medical multi-image data, only a limited number of works attempted to develop multimodal large models trained on such data. Yang *et al.* [25] pretrained a medical vision-language framework to leverage spatial and temporal relationships across multiple X-ray views and historical records. However, their work primarily focused on multimodal alignment and pretraining, without providing the capability for open-ended visual question answering. Additionally, their dataset was restricted to chest X-ray images, limiting its applicability to other medical imaging scenarios. Other studies [22,21] have explored video-language models, achieving strong performance in analyzing surgical scenes and histopathology videos. Nevertheless, these methods are tailored to specific video domains and fail to address broader multi-image tasks, such as longitudinal temporal forecasting and multi-view understanding, which remain unexplored. Furthermore, existing medical VQA benchmarks, like VQA-RAD [11], do not adequately assess multi-image comprehension, leaving a critical gap in evaluating models’ clinical multi-image visual understanding capabilities.

To enhance multi-image visual abilities in medical LVLMs, we present the **Medical Multi-IMage** (Med-MIM) instruction dataset, along with the Med-MIM benchmark, designed for comprehensively evaluating LVLMs across various types of multi-image visual tasks. To the best of our knowledge, both the Med-MIM instruction dataset and Med-MIM benchmark represent the first focused effort on medical multi-image analysis. Our key contributions are summarized as follows: (1) We construct the Med-MIM instruction dataset, which contains 83.2K medical multi-image instruction samples. The constructed dataset is organized into four subsets designed to enhance VLMs’ co-reference, comparison, reasoning, and temporal understanding abilities. (2) We establish the Med-MIM benchmark, which includes a held-in evaluation subset derived from the Med-MIM instruction dataset. Additionally, we construct two held-out multi-image datasets to evaluate the models’ generalized zero-shot multi-image visual capabilities. (3) Using the Med-MIM instruction dataset, we develop two special-

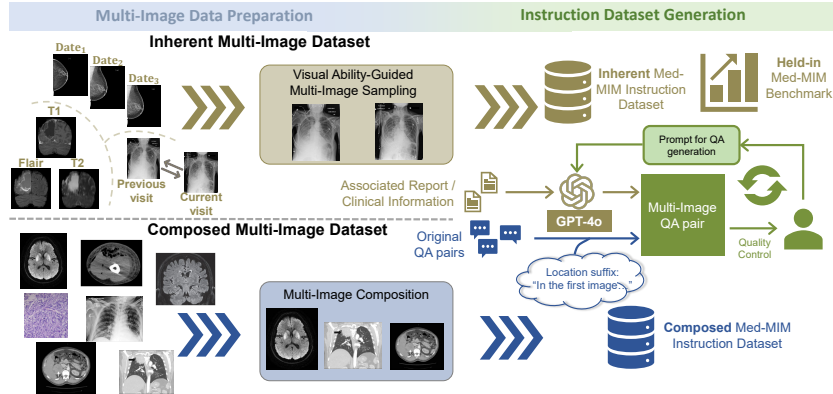


Fig. 1. Illustration of the pipeline to generate Med-MIM dataset and benchmark.

ized medical VLMs: MIM-LLaVA-Med and Med-Mantis, tailored for multi-image analysis. We compare these models against several state-of-the-art VLMs, and experimental results demonstrate that our models exhibit superior multi-image visual capabilities on both the held-in and held-out Med-MIM benchmarks.

2 Method

We create the Med-MIM Instruction dataset and benchmark following two steps: multi-image data preparation and instruction dataset generation (as illustrated in Fig. 1). During the data preparation phase, two types of multi-image samples (inherent and composed) are collected. Subsequently, GPT-4o [6], combined with iterative refinement, is employed to generate the inherent Med-MIM instruction dataset and the held-in benchmark. Meanwhile, the composed Med-MIM instruction dataset is created by appending location-specific suffixes to the original QA pairs, ensuring alignment with multi-image visual understanding.

Dataset Collection. Fig. 2 presents an overview of our constructed Med-MIM instruction dataset composed of 83.2K medical multi-image QA instruction samples from five main domains. Unlike previous medical instruction dataset like PMC-VQA [26] or LLaVA-Med VQA [12] that primarily based on published literature, we construct the Med-MIM instruction dataset from two sources: inherent multi-image datasets and composed multi-image dataset, where inherent multi-image datasets contain naturally occurring multi-image scenes such as longitudinal inspections, multi-modality or multi-view grounded diagnoses, etc. To curate high-quality multi-image dataset, we group multiple images based on tasks related to various visual abilities, utilizing three large-scale inherent multi-image datasets: MS-CXR-T [3], EMBED [7], and LUMIERE [23]. All three datasets include multi-visit images arranged in chronological order, while EMBED and LUMIERE further incorporate multi-view and multi-modality data. Additionally, to further enhance the richness of the dataset and include a broader

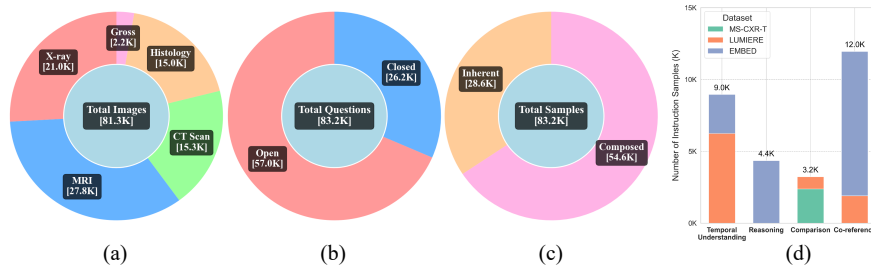


Fig. 2. Comprehensive analysis of the Med-MIM dataset: (a) Image modality composition; (b) Distribution of the question types; (c) Number of samples in the inherent and composed multi-image datasets; (d) **Inherent** multi-image dataset distribution categorized by four core visual abilities.

variety of medical images, we construct a composed multi-image dataset by manually grouping multiple images from the LLaVA-Med VQA [12] dataset. We cap the length of multi-image sequences at a maximum of three, allowing for grouping more multi-image QA pairs while ensuring feasibility with VLM fine-tuning under limited computational resources.

Four Multi-image Visual Abilities. Following the MANTIS [8], we can divide all collected multi-image samples into four categories corresponding to four different visual abilities (Fig.2 (d)). However, for medical multi-images, more specific tasks can be linked to various abilities as follows.

(1) Temporal Understanding: We focus on multi-visit data and construct QA pairs following the temporal forecasting task. For example, “Given the first two consecutive visits collected at {date1} and {date2}, will the subject become diseased in the next visit?” We filter 9K temporal understanding related samples from the LUMIERE and EMBED datasets.

(2) Reasoning: We restrict reasoning-type multi-image tasks to multi-view diagnosis in the EMBED dataset. A total of 4.4K samples are collected. The LVLM agent will be asked to diagnose patient from both two views (Cranio-Caudal & Medio-Lateral Oblique), such as “Considering the features observed in the two provided mammographic views, what is the BIRADS category for the findings?”

(3) Comparison: By comparing multiple images, we aim to find the similarities or differences. In our instruction dataset, we obtain comparison-type multi-image samples from both the MS-CXR-T and LUMIERE datasets. Note that although MS-CXR-T contains temporal relations, the actual task is to determine progression status based on previous and current visits, which requires the model to find differences between two visits. For instance, “Could you explain the differences between the X-ray images from the first visit and the second visit?”

(4) Co-reference: We construct a co-reference-related subset from both the LUMIERE and EMBED datasets by asking the model questions such as, “Which image, the first or the second, represents the view that captures the breast from above?” By incorporating location into the question, we aim to enable the VLM




Fig. 3. Example of generated multi-image QA pair in the MS-CXR-T dataset.

to understand both image location and content simultaneously. Additionally, for the composed multi-image dataset, we categorize it into the co-reference subset as well. To encode proper location information, the original single-image QA pairs are updated by adding prefixes such as “In the first image” or “In the second image.” (as illustrated in Fig.1).

Instruction Dataset Generation. After collecting the multi-image sets, we integrate the textual data for each image and generate the corresponding question-answer pairs. Unlike the composed Med-MIM instruction dataset, which can be created by simply appending location-specific suffixes (e.g., “in the first image”) to the original QA pairs, the inherent Med-MIM dataset requires a more sophisticated approach. For this, we prompt language-only GPT-4o to generate multi-image QA pairs by leveraging the medical reports or clinical information associated with each image. We also incorporate human reviewers as a quality control measure to iteratively refine the generated results, ensuring both accuracy and clarity. Fig.3 shows one multi-image QA example generated by GPT-4o based on the medical reports from two visits in the MS-CXR-T dataset.

Multi-image Visual Abilities Evaluation via Med-MIM Benchmark. Our constructed Med-MIM Benchmark comprises two parts. (a) Held-in part: To comprehensively evaluate four multi-image visual abilities, we construct the held-in Med-MIM Benchmark derived from the Med-MIM instruction dataset, which includes 2,968 closed-type examples (903, 454, 208, and 1,403 for temporal, reasoning, comparison, co-reference abilities, respectively) and 256 open-type examples (30, 30, 136, and 60, respectively). (b) Held-out part: Due to the limited availability of benchmarks for evaluating multi-image visual abilities in the medical domain, we constructed the held-out Med-MIM evaluation benchmark by curating multi-image cases from two complementary sources: VQA-RAD [11], which provides structured QA pairs but grounded in single medical images; ODIR [13], containing paired fundus images but lacking annotated QA pairs. To adapt the standard VQA-RAD benchmark for multi-image scenarios, we first group images based on their modalities, anatomical locations, and semantic coherence with the original QA pairs. We then construct the MIM-RAD by synthesizing composed multi-image instructions from the VQA-RAD dataset, similar to the composed Med-MIM instruction dataset. Additionally, we generate the MIM-ODIR by applying the same procedure used for inherent Med-MIM

dataset generation to the ODIR dataset. Specifically, MIM-ODIR comprises 300 closed-type cases and 300 open-type cases, while MIM-RAD contains 293 closed-type cases and 73 open-type cases.

Instruction Tuning with Med-MIM Instruction Dataset. We conduct instruction tuning from two popular VLMs, LLaVA-Med [12] and MANTIS [8]. LLaVA-Med is a successful open-source VLM in the medical domain, while MANTIS has been trained on over 700K multi-natural-image instruction data. We fine-tune both VLMs for 3 epochs, running for about 72 hours on 4 NVIDIA A40 GPUs and yielded two specialized models: MIM-LLaVA-Med and Med-Mantis models, both of which excel at understanding medical multi-images. To accommodate the multi-image and question input, we adopt an interleaved image-text formation and restructure image embeddings as: “(image {id}: <Image> image embeddings </Image>)”. We employ the negative log-likelihood (NLL) loss function to measure the discrepancy between the predicted next token and the actual next token in the sequence. The specific formula is defined in Eq. (1), where K represents the total token length of the standard response text. \mathcal{I} consists of interleaved text token sequences (T_i) and image token sequences (I_i). $p(r^k|\mathcal{I}, q^{1:S}, r^{1:k-1})$ denotes the probability of generating the k th token given the interleaved token sequence \mathcal{I} , the question token sequence $q^{1:S}$, and the previous tokens in the response sequence $r^{1:k-1}$.

$$\mathcal{L}_{NLL} = - \sum_{k=1}^K \log p(r^k|\mathcal{I}, q^{1:S}, r^{1:k-1}), \mathcal{I} = \text{Concat}(T_1, I_1, \dots, T_N, I_N). \quad (1)$$

3 Experimental Results

Baselines and Metrics. We select six different SOTA VLMs for comparison including Med-Flamingo-9B [18], Deepseek-VL-7B [16], LLaVA-Med-7B [12], InternVL2-8B [5], Mantis-8B [8], and GPT-4o [6]. For close-ended (multiple choice questions) evaluation, accuracy is used as the evaluation metric. For open-ended (free-answer questioning) evaluation, the final score is calculated as the average of BERT-recall, BLEU score, and ROUGE-L recall.

Main Results on Med-MIM Benchmark. We evaluate the performance of various LVLMs in both held-in and held-out benchmarks. Table 1 presents the corresponding experimental results. The best result for each task is highlighted in green, while the second-best result is indicated with an underline. We present the performance improvement achieved after fine-tuning on the Med-MIM instruction dataset. For the held-in benchmark, after undergoing supervised fine-tuning on the Med-MIM instruction dataset, both MIM-LLaVA-Med and Med-Mantis show significant improvements across all visual ability subsets. Furthermore, compared to other LVLMs, our models secure the top-1 and top-2 performance across all tasks. We also evaluate the zero-shot performance of models on two held-out benchmarks. Experimental results demonstrate that GPT-4o, as one of the most powerful closed-source LVLM, consistently delivers superior performance across both held-out benchmarks. However, among all open-source

Table 1. Comparison of model performance on the held-in Med-MIM benchmark (temporal understanding, reasoning, comparison, co-reference) and the held-out Med-MIM benchmark (MIM-RAD, MIM-ODIR). Results are expressed as percentages (%). **C** and **O** refer to close and open scores, respectively. The \uparrow symbol indicates the performance improvement of our trained models compared to their original versions.

Models	Held-in								Held-out			
	Temporal		Reasoning		Compare		Co-refer		MIM-RAD		MIM-ODIR	
	C	O	C	O	C	O	C	O	C	O	C	O
GPT-4o	47.95	37.28	16.96	30.18	29.81	37.20	42.30	33.81	69.00	39.52	45.67	34.60
Med-Flamingo-9B	39.65	24.11	15.64	22.01	32.21	27.03	28.72	27.20	4.00	28.69	16.00	30.83
deepseek-VL-7B	29.90	36.07	14.54	34.47	32.69	34.34	29.01	31.08	17.00	31.58	21.67	44.49
InternVL2-8B	46.07	35.85	16.08	32.12	34.13	36.09	41.05	33.92	56.33	35.62	31.67	41.01
LLaVA-Med-7B	38.87	37.41	14.32	34.90	33.17	36.65	41.13	33.83	48.67	34.86	29.67	53.89
Mantis-8B	68.89	37.32	15.20	24.05	38.94	31.47	35.71	22.72	59.67	28.08	23.00	21.20
MIM-LLaVA-Med	74.31	44.50	22.91	41.74	50.00	40.36	46.19	43.09	53.33	35.08	32.33	55.99
LLaVA-Med \uparrow	35.44	7.09	8.59	6.84	16.83	3.71	5.06	9.26	4.66	0.22	2.66	2.10
Med-Mantis	74.86	46.87	29.74	42.05	52.88	40.33	80.40	49.20	65.33	35.35	30.00	41.32
Mantis \uparrow	5.97	9.55	14.54	18.00	13.94	8.86	44.69	26.48	5.66	7.27	7.00	20.12

LVLMS, our Med-Mantis achieves the highest close-set score on the MIM-RAD benchmark, while MIM-LLaVA-Med attains the best close and open scores on the MIM-ODIR benchmark. Both fine-tuned models exhibit significant performance improvements across all metrics compared to their original versions. This highlights the effectiveness of incorporating Med-MIM instruction dataset, which enables LVLMS to develop advanced medical multi-image inference capabilities.

Ablation on Four Visual Abilities. To assess the impact of different multi-image visual ability subsets, we perform an ablation study using the Med-Mantis model. We systematically remove subsets corresponding to specific visual abilities from the Med-MIM instruction dataset and fine-tune the model on the remaining data. Fig.4 (a) presents the test results for five versions of the Med-Mantis model evaluated on the held-in Med-MIM benchmark. We find that excluding a specific visual ability subset from training leads to a noticeable drop in performance on the corresponding testing set, underscoring the importance of each visual ability subset in the Med-MIM instruction dataset.

Influence of the Size to the Med-MIM Instruction Dataset. Next, we examine the effect of varying Med-MIM instruction dataset sizes on the performance of the Med-Mantis model. Fig.4 (c) illustrates how the model’s performance changes when utilizing 0%, 25%, 50%, 75%, and 100% of the Med-MIM instruction dataset. The results reveal a consistent improvement in performance as the size of the instruction dataset increases, with notable gains in tasks requiring temporal understanding, comparison, and co-reference visual abilities.

Ablation on the Composed Med-MIM Instruction Dataset. We investigate the effectiveness of the composed instruction dataset in Fig.4 (b). We compare the performance of MIM-LLaVA-Med and Med-Mantis models with and without the composed instruction dataset on the overall held-in benchmark.

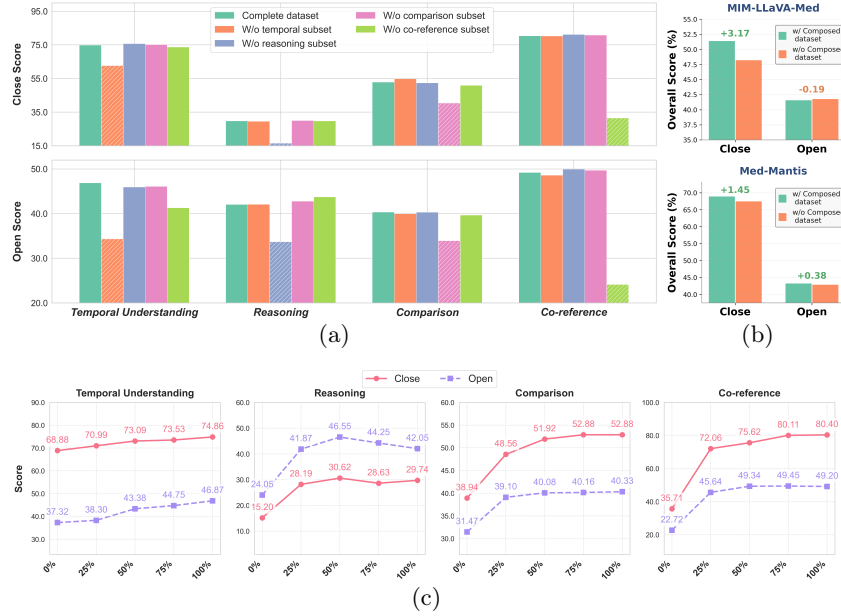


Fig. 4. Ablation results on (a) four visual ability subsets on the Med-Mantis, (b) the performance of Med-Mantis and MIM-LLaVA-Med with and without the composed Med-MIM instruction dataset, and (c) the impact of different sizes of the Med-MIM instruction dataset on the Med-Mantis.

The results indicate that the composed instruction dataset significantly improves closed-type QA performance for both models. Furthermore, as Med-Mantis is primarily trained on natural images, fine-tuning with additional composed medical multi-image dataset further boosts its open-type QA performance.




Fig. 5. Case studies for LVLMs in medical multi-image understanding.

Case Study. We present two examples in Fig.5 from the MS-CXR-T subset (left) and the EMBED subset (right) of the held-in Med-MIM benchmark. The results indicate that Med-Mantis and MIM-LLaVA-Med can produce more accurate response compared with original Mantis and LLaVA-Med models.

4 Conclusion

In this work, we present the Med-MIM instruction dataset, designed to enhance medical multi-image visual abilities of LVLMs. Leveraging this dataset, we fine-tune two models, Med-Mantis and MIM-LLaVA-Med, both of which demonstrate superior performance in answering medical questions involving multiple images. To further assess the multi-image understanding of LVLMs, we construct the Med-MIM benchmark. Experimental results comparing our models to six popular LVLMs reveal that fine-tuning with the Med-MIM instruction dataset significantly enhances multi-image visual abilities in the medical domain.

Acknowledgments. The work described in this paper was supported in part by the Research Grants Council of the Hong Kong Special Administrative Region, China, under Project T45-401/22-N; and by the Hong Kong Innovation and Technology Fund, under Project GHP/060/21GD.

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

References

1. Alleman, K., Knecht, E., Huang, J., Zhang, L., Lam, S., DeCuyper, M.: Multi-modal deep learning-based prognostication in glioma patients: a systematic review. *Cancers* **15**(2), 545 (2023)
2. Bai, F., Du, Y., Huang, T., Meng, M.Q.H., Zhao, B.: M3d: Advancing 3d medical image analysis with multi-modal large language models. *arXiv preprint arXiv:2404.00578* (2024)
3. Bannur, S., Hyland, S., Liu, Q., Pérez-García, F., Ilse, M., de Castro, D.C., Boecking, B., Sharma, H., Bouzid, K., Schwaighofer, A., et al.: Ms-cxr-t: Learning to exploit temporal structure for biomedical vision-language processing (2023)
4. Calhoun, V.D., Sui, J.: Multimodal fusion of brain imaging data: a key to finding the missing link (s) in complex mental illness. *Biological psychiatry: cognitive neuroscience and neuroimaging* **1**(3), 230–244 (2016)
5. Chen, Z., Wu, J., Wang, W., Su, W., Chen, G., Xing, S., Zhong, M., Zhang, Q., Zhu, X., Lu, L., et al.: Internvl: Scaling up vision foundation models and aligning for generic visual-linguistic tasks. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 24185–24198 (2024)
6. Hurst, A., Lerer, A., Goucher, A.P., Perelman, A., Ramesh, A., Clark, A., Ostrow, A., Welihinda, A., Hayes, A., Radford, A., et al.: Gpt-4o system card. *arXiv preprint arXiv:2410.21276* (2024)
7. Jeong, J.J., Vey, B.L., Bhimireddy, A., Kim, T., Santos, T., Correa, R., Dutt, R., Mosunjac, M., Oprea-Ilie, G., Smith, G., et al.: The emory breast imaging dataset (embed): A racially diverse, granular dataset of 3.4 million screening and diagnostic mammographic images. *Radiology: Artificial Intelligence* **5**(1), e220047 (2023)
8. Jiang, D., He, X., Zeng, H., Wei, C., Ku, M.W., Liu, Q., Chen, W.: Mantis: Interleaved multi-image instruction tuning. *Transactions on Machine Learning Research* **2024** (2024), <https://openreview.net/forum?id=skLtdUVaJa>

9. Johnson, D.K., Storandt, M., Morris, J.C., Galvin, J.E.: Longitudinal study of the transition from healthy aging to alzheimer disease. *Archives of neurology* **66**(10), 1254–1259 (2009)
10. Korot, E., Guan, Z., Ferraz, D., Wagner, S.K., Zhang, G., Liu, X., Faes, L., Pontikos, N., Finlayson, S.G., Khalid, H., et al.: Code-free deep learning for multi-modality medical image classification. *Nature Machine Intelligence* **3**(4), 288–298 (2021)
11. Lau, J.J., Gayen, S., Ben Abacha, A., Demner-Fushman, D.: A dataset of clinically generated visual questions and answers about radiology images. *Scientific data* **5**(1), 1–10 (2018)
12. 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)
13. Li, N., Li, T., Hu, C., Wang, K., Kang, H.: A benchmark of ocular disease intelligent recognition: One shot for multi-disease detection. In: *Benchmarking, Measuring, and Optimizing: Third BenchCouncil International Symposium, Bench 2020, Virtual Event, November 15–16, 2020, Revised Selected Papers 3*. pp. 177–193. Springer (2021)
14. Liu, H., Li, C., Wu, Q., Lee, Y.J.: Visual instruction tuning. *Advances in neural information processing systems* **36** (2024)
15. Liu, Z., Chu, T., Zang, Y., Wei, X., Dong, X., Zhang, P., Liang, Z., Xiong, Y., Qiao, Y., Lin, D., et al.: Mmdu: A multi-turn multi-image dialog understanding benchmark and instruction-tuning dataset for lvlms, 2024c. URL <https://arxiv.org/abs/2406.11833>
16. Lu, H., Liu, W., Zhang, B., Wang, B., Dong, K., Liu, B., Sun, J., Ren, T., Li, Z., Yang, H., et al.: Deepseek-vl: towards real-world vision-language understanding. *arXiv preprint arXiv:2403.05525* (2024)
17. Meng, F., Wang, J., Li, C., Lu, Q., Tian, H., Liao, J., Zhu, X., Dai, J., Qiao, Y., Luo, P., et al.: Mmiu: Multimodal multi-image understanding for evaluating large vision-language models. *arXiv preprint arXiv:2408.02718* (2024)
18. Moor, M., Huang, Q., Wu, S., Yasunaga, M., Dalmia, Y., Leskovec, J., Zakka, C., Reis, E.P., Rajpurkar, P.: Med-flamingo: a multimodal medical few-shot learner. In: *Machine Learning for Health (ML4H)*. pp. 353–367. PMLR (2023)
19. Poulakis, K., Pereira, J.B., Muehlboeck, J.S., Wahlund, L.O., Smedby, Ö., Volpe, G., Masters, C.L., Ames, D., Niimi, Y., Iwatsubo, T., et al.: Multi-cohort and longitudinal bayesian clustering study of stage and subtype in alzheimer’s disease. *Nature communications* **13**(1), 4566 (2022)
20. Rousan, L.A., Elobeid, E., Karrar, M., Khader, Y.: Chest x-ray findings and temporal lung changes in patients with covid-19 pneumonia. *BMC Pulmonary Medicine* **20**, 1–9 (2020)
21. Seyfioglu, M.S., Ikezogwo, W.O., Ghezloo, F., Krishna, R., Shapiro, L.: Quilt-llava: Visual instruction tuning by extracting localized narratives from open-source histopathology videos. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 13183–13192 (2024)
22. Srivastav, V., Navab, N., Padoy, N., et al.: Procedure-aware surgical video-language pretraining with hierarchical knowledge augmentation. In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*
23. Suter, Y., Knecht, U., Valenzuela, W., Notter, M., Hower, E., Schucht, P., Wiest, R., Reyes, M.: The lumiere dataset: Longitudinal glioblastoma mri with expert rano evaluation. *Scientific data* **9**(1), 768 (2022)

24. Wu, C., Zhang, X., Zhang, Y., Wang, Y., Xie, W.: Towards generalist foundation model for radiology. arXiv preprint arXiv:2308.02463 (2023)
25. Yang, J., Su, B., Zhao, W.X., Wen, J.R.: Unlocking the power of spatial and temporal information in medical multimodal pre-training. In: Proceedings of the 41st International Conference on Machine Learning. ICML'24, JMLR.org (2024)
26. Zhang, X., Wu, C., Zhao, Z., Lin, W., Zhang, Y., Wang, Y., Xie, W.: Pmc-vqa: Visual instruction tuning for medical visual question answering. arXiv preprint arXiv:2305.10415 (2023)