

RadImageNet-VQA: A Large-Scale CT and MRI Dataset for Radiologic Visual Question Answering

Léo Butsanets¹

Charles Corbière¹

Julien Khlaut^{1,2,3}

Pierre Manceron¹

Corentin Dancette¹

LEO.BUTSANETS@RAIDIUM.EU

CHARLES.CORBIERE@RAIDIUM.EU

JULIEN.KHLAUT@RAIDIUM.EU

PIERRE.MANCERON@RAIDIUM.EU

CORENTIN.DANCETTE@RAIDIUM.EU

¹*Radium, Paris, France.*

²*Université de Paris Cité, Hôpital Européen Georges Pompidou, AP-HP, Paris, France.*

³*Department of Vascular and Oncological Interventional Radiology, INSERM, Paris, France.*

Editors: Under Review for MIDL 2026

Abstract

In this work, we introduce *RadImageNet-VQA*, a large-scale dataset designed to advance radiologic visual question answering (VQA) on CT and MRI exams. While existing medical VQA datasets are limited in scale, dominated by X-ray imaging or biomedical illustrations, and prone to text-based shortcuts, RadImageNet-VQA is built from expert-curated annotations and provides 750K images paired with 7.5M QA samples. It covers three key tasks—abnormality detection, anatomy recognition, and pathology identification—spanning 8 anatomical regions and 97 pathology categories, and supports open-ended, closed-ended, and multiple-choice questions. Extensive experiments show that state-of-the-art vision-language models still struggle with fine-grained pathology identification, especially in open-ended settings and even after fine-tuning. Text-only analysis further reveals that model accuracies collapse to near-random without image inputs, confirming that RadImageNet-VQA is free from linguistic shortcuts. The full dataset and benchmark are publicly available at <https://huggingface.co/datasets/raidium/RadImageNet-VQA>.

Keywords: Medical Visual Question Answering, Vision-Language Models, Radiology.

1. Introduction

Interpreting radiology exams requires integrating complex visual patterns with clinical reasoning, often with medical context (Yildirim et al., 2024; Dancette et al., 2025; Khlaut et al., 2025). Vision-language models (VLMs) (Li et al., 2023b; Bai et al., 2025; Chen et al., 2024; Liu et al., 2023) have recently achieved strong performance on general multimodal benchmarks. Medical variants, such as MedGemma (Sjellergren et al., 2025) and Lingshu (Xu et al., 2025b), now incorporate curated medical data to better align with clinical knowledge. These models are particularly attractive for radiology, where generating reports and supporting clinical reasoning require combining visual and textual information. Yet, evaluating free-text radiology reports remains difficult: commonly used similarity metrics often fail to capture clinical correctness, factual consistency, or semantic alignment with ground-truth findings (Irvin et al., 2019; Delbrouck et al., 2024; Xu et al., 2025a). In contrast, radiologic visual question answering (VQA) offers a more structured alternative to probe

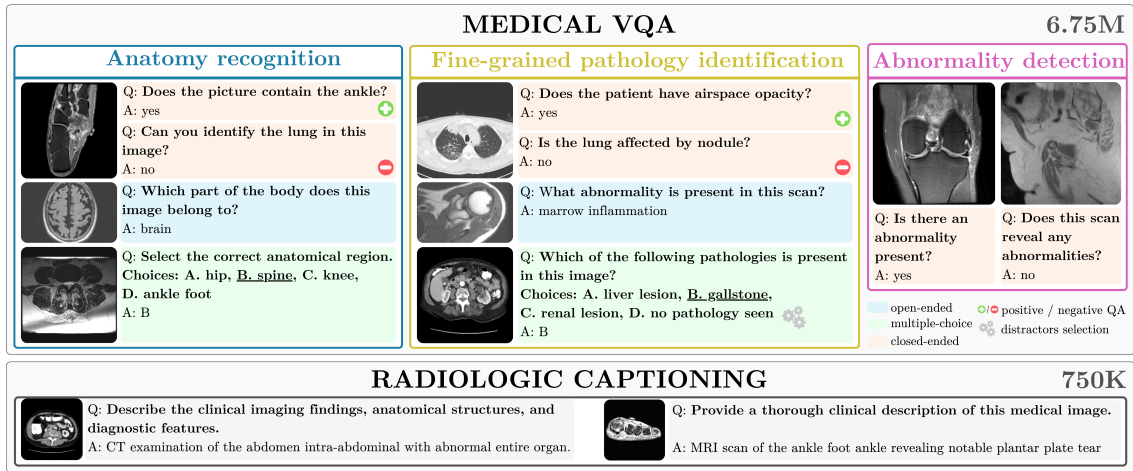


Figure 1: **Overview of the RadImageNet-VQA dataset**, which provides radiology-focused supervision across three VQA tasks (anatomy recognition, fine-grained pathology identification, and abnormality detection) using diverse open-ended, closed-ended, and multiple-choice formats. It also includes radiologic captioning pairs for image-text alignment.

model reasoning. Its question-answer format mirrors how radiologists interrogate imaging findings, providing an interpretable way to evaluate a model’s image-grounded reasoning.

Existing medical VQA datasets (Lau et al., 2018; Liu et al., 2021; Pal et al., 2025; Zhang et al., 2023) present several limitations that hinder proper evaluation of these models. Early efforts rely on limited expert annotation and thus offer small collections with narrow anatomical and pathological coverage (Lau et al., 2018; Liu et al., 2021). Larger recent datasets draw heavily from X-ray images or biomedical figures scraped from publications (Pal et al., 2025; Zhang et al., 2023), providing little representation of other common imaging modalities, in particular CT or MRI. Many also contain textual shortcuts that allow models to answer correctly without interpreting the image, as shown in Section 4.2. Beyond evaluation, progress in medical VLMs also requires access to large, well-curated training data aligned with diagnostic tasks. This emphasizes the need for large-scale VQA resources that spans multiple radiologic tasks while minimizing possible shortcut cues to offer a more rigorous foundation for assessing and advancing medical VLMs.

In this paper, we introduce *RadImageNet-VQA*, a large-scale dataset designed for training and benchmarking radiologic VQA on CT and MRI exams (see Figure 1). Built from the CT/MRI subset of RadImageNet (Mei et al., 2022) and its expert-curated anatomical and pathological annotations, RadImageNet-VQA provides 750K images with 7.5M generated samples, including 750K medical captions for visual-text alignment and 6.75M question-answer pairs that span three radiology tasks: fine-grained pathology identification, anatomy recognition, and abnormality detection. The dataset includes open-ended, closed-ended, and multiple-choice questions across 8 anatomical regions and 97 pathologies, generated with prompt-based templates and constructed to probe visual-grounded understanding while minimizing text-only shortcut answering. For evaluation, we construct a stratified benchmark of 1,000 images with 9,000 question-answer pairs covering all tasks

and question types. Through extensive zero-shot evaluation and fine-tuning experiments with RadImageNet-VQA, we show that while anatomy recognition is nearly solved for current VLMs, pathology identification remains a major bottleneck, especially for open-ended responses. Text-only analyses further confirm that RadImageNet-VQA greatly reduces shortcut cues present in existing datasets. Fine-tuning yields substantial gains across models, yet we observe that initializing models with medically pretrained vision encoders such as MedSigLIP (Sellersgren et al., 2025) provides no advantage over general-purpose encoders.

Our main contributions are as follows:

- We release RadImageNet-VQA, a radiologic dataset that includes both a curated benchmark and a large training corpus. It spans 8 anatomical regions, 97 pathologies, and three task families (abnormality, anatomy, pathology) with multiple question formats, providing 750K images paired with 7.5M generated samples.
- We conduct an extensive evaluation of state-of-the-art VLMs showing that existing models perform well on anatomy and basic abnormality detection but still struggle with fine-grained pathology identification. Text-only ablations confirm that RadImageNet-VQA minimizes shortcut cues.
- Fine-tuning on RadImageNet-VQA produces substantial gains across all model families, while medical pretrained vision encoders do not improve downstream performance.

2. Related Work

2.1. Vision-language models in radiology

Recent VLMs, such as LLaVA (Liu et al., 2023), QwenVL (Bai et al., 2025), InternVL (Chen et al., 2024), typically combine a vision encoder, often initialized from CLIP (Radford et al., 2021) or SigLIP (Zhai et al., 2023) backbones, with a large language model (LLM) via a lightweight adapter module. In the medical domain, progress has largely been driven by data-centric adaptation, where general-purpose VLMs are aligned to medical images through image-text pretraining and multimodal instruction tuning (Li et al., 2023a; Sellersgren et al., 2025; Xu et al., 2025b). While several 3D radiology VLMs have recently been proposed (Blankemeier et al., 2024; Wu et al., 2023; Xin et al., 2025; Ates et al., 2025), their development is constrained by the scarcity of large volumetric datasets, making broad and diverse 2D datasets still essential for training and evaluation. Beyond training, reliable evaluation also remains challenging. Radiology report generation (Bannur et al., 2024) is appealing but difficult to assess automatically, as standard metrics correlate poorly with clinical correctness. On the other hand, visual-question answering (VQA) provides a more controlled and interpretable proxy framework for probing clinical reasoning grounded in images. This has motivated increasing interest in radiologic VQA benchmarks, particularly for CT and MRI, which remain underrepresented in existing resources.

2.2. Medical VQA datasets

Early radiologic VQA, such as VQA-RAD (Lau et al., 2018) and SLAKE (Liu et al., 2021) introduced structured QA annotations but are limited by scale and by narrow anatomical and pathological coverage. More recent efforts, such as the medical subset of MMMU (Yue et al., 2024) extend medical breadth, yet their radiology content is sparse and relies largely

Dataset	Modality	Question Types	# Images	# QAs Train	# QAs Test
VQA-RAD (Lau et al., 2018)	X-ray, CT	open, closed	0.3K	2.8K	0.45K
SLAKE (Liu et al., 2021)	X-ray, CT, MRI	open, closed	0.6K	4.9K	1.1K
MMMU-Med (Yue et al., 2024)	diverse	multiple-choice	1.9K	–	1.8K
ReXVQA (Pal et al., 2025)	X-ray	multiple-choice	160K	573K	40.8K
M3D-VQA (Bai et al., 2024)	3D CT	multiple-choice	96K	662K	2K
3D-RAD (Gai et al., 2025)	3D CT	open, closed	16K	136K	34K
RadImageNet-VQA (Ours)	CT, MRI	open, closed, multiple-choice	750K	7.5M	9K

Table 1: **Comparison of radiologic VQA datasets** compared by imaging modality, supported question types, and number of images, and the number of question-answer pairs for training and evaluation. RadImageNet-VQA provides the largest CT/MRI coverage and the most extensive QA corpus.

on images from textbooks or publications rather than CT or MRI. Programmatic datasets such as ReXVQA (Pal et al., 2025) improve scalability but are restricted to X-ray images and therefore cover only a small portion of radiology practice. Table 1 summarizes these datasets and highlights the narrow CT/MRI coverage of existing benchmarks. Additionally, a persistent issue across these benchmarks is the presence of linguistic shortcuts, label phrasing, or distributional biases that allow models to guess without relying on images. In parallel, several 3D medical VQA benchmarks, such as DeepTumorVQA (Chen et al., 2025), M3D-VQA (Bai et al., 2024), and 3D-RAD (Gai et al., 2025), extend to volumetric tasks. Although clinically meaningful, these datasets cover limited anatomy and pathologies, require extensive expert curation, and remain small in scale. Current 3D VLMs are trained on small volumetric corpora and lack the broad visual priors available to large 2D models. Strengthening 2D CT/MRI supervision therefore remains essential for building radiology VLMs, even as 3D approaches continue to mature.

3. RadImageNet-VQA Dataset

In this section, we describe the construction process of RadImageNet-VQA dataset, as illustrated in Figure 2, and detail the statistics and metrics of the benchmark subset.

3.1. Dataset construction pipeline

Data is sourced from RadImageNet (Mei et al., 2022), a large expert-annotated medical imaging dataset in which each image is associated with a modality (CT, MRI, US), a body part (e.g., abdomen, hip, brain) and a pathology label. From this resource, we use the CT and MRI subsets to form the basis for generating clinically meaningful captions and VQA samples across anatomy, abnormality, and fine-grained pathology tasks.

Radiologic captioning generation. VLMs are typically trained with an alignment stage in which the model learns to associate visual content with textual semantics using large collections of image-caption pairs. As presented in Figure 2, to support a clinically meaningful alignment dataset in training, we leverage the rich metadata in RadImageNet, including acquisition modality, anatomical region, and pathology category, and convert it into




Figure 2: **RadImageNet-VQA construction pipeline.** Expert-annotated CT/MRI images are converted into radiology-aware captions and VQA samples using task- and format-specific templates. The pipeline generates open-ended, closed-ended, and multiple-choice questions across anatomy recognition, abnormality detection, and pathology identification, with distractors designed to prevent shortcuts.

structured radiologic captions. Each image is paired with a synthesized textual description generated from these metadata fields sampling randomly from a diverse set of radiology-aware templates. For example, modality and anatomy tags are verbalized in formulations like "A [modality] scan of the [anatomy] showing [pathology]".

VQA sample generation. To support both instruction tuning and benchmark construction, each annotated image is converted into structured VQA samples using task-specific question-answer templates. As illustrated in Figure 2, we generate closed-ended, open-ended, and multiple-choice questions across three complementary tasks: *anatomy recognition*, *abnormality detection* and *pathology identification*. Anatomy recognition identifies the imaged region, abnormality detection determines whether an image contains any abnormal finding, and pathology identification—by far the most challenging—requires distinguishing between specific diseases within a given anatomical context.

For each task, we design a set of template-based question formulations for each question type (Appendix A.2). For every task and question type, we define between 2 and 7 linguistic variations, ensuring diversity while reducing opportunities for models to exploit textual shortcuts. For closed-ended anatomy and pathology questions, we generate both positive and negative variants, where the expected answer is respectively “yes” (the organ or pathology is present) or “no” (it is absent). These paired formulations allow us to probe VLMs more precisely, particularly with respect to biases such as hallucinating organs or pathologies even when they are not present.

Multiple-choice questions require careful construction as distractors, *i.e.* incorrect answer options, must be chosen in a way that prevents trivial elimination. For anatomy recognition, distractors are sampled from other existing anatomical regions in the dataset.




Figure 3: **Composition of the RadImageNet-VQA benchmark**, containing 1,000 CT/MRI images and 9,000 QA pairs: (a) most frequent pathology labels, (b) anatomy repartition, (c) question-type distribution.

For pathology identification, distractors are restricted to clinically plausible diseases from the same region, ensuring that the model cannot solve the question by matching anatomy alone. We also include the option “*no pathology seen*” in every pathology multiple-choice. As shown in Appendix B.1, increasing the frequency of this option in multiple-choices options causes a clear and consistent drop in accuracy for medical-oriented models (e.g., MedGemma and Lingshu), while general-purpose models remain comparatively stable. This suggests that medical VLMs often assume that pathology is present and treat “no pathology” as unlikely instead of confirming it visually. Including this option therefore reduce this bias and encourages models to distinguish true pathology from its absence, resulting in a more challenging and faithful evaluation.

Training set creation. We apply the captioning and VQA generation pipeline to the full RadImageNet training split to construct a large-scale corpus for multimodal model training. This yields structured radiologic captions for medical visual-text alignment and extensive VQA data for instruction tuning. In total, the training corpus comprises approximately 750K images with 7.5M samples (750K image-captions and 6.75M QAs pairs). A complete list of anatomical regions and pathology categories is provided in Appendix A.1.

3.2. Benchmark

For evaluation, we sample 1,000 CT/MRI images from the RadImageNet test split while preserving the distribution of anatomical regions, pathology categories, and abnormal cases. Applying our pipeline yields 9,000 closed-ended, open-ended, and multiple-choice QA pairs. Figure 3 shows the distribution of pathology labels, question formats, anatomical regions, and normal/abnormal cases. Samples are evenly distributed across tasks and question type, and the dataset includes both normal (28.1%) and abnormal (71.9%) studies across multiple anatomical regions. The most common pathology categories include a wide range of soft-tissue, musculoskeletal, and neuro-abdominal findings represented in CT and MRI.

Evaluation metrics. Model predictions are scored as correct or incorrect, and mean accuracy is reported. Closed-ended (yes/no) questions are judged by checking for the expected token. Open-ended responses are evaluated using *LLM-as-a-judge* (Zheng et al., 2023)

Model	Anatomy				Abn	Pathology				Avg
	Open	Closed+	Closed−	MC	Closed	Open	Closed+	Closed−	MC	
General-purpose models										
LLaVA-OneVision-Qwen2-7B	48.4	82.7	81.3	88.7	49.8	16.0	55.3	61.3	33.6	57.5
Qwen2.5-VL-3B-Instruct	37.7	83.7	77.1	77.9	70.5	10.0	78.1	21.4	34.8	54.6
Qwen2.5-VL-7B-Instruct	37.5	84.9	79.1	80.5	69.5	9.8	69.2	47.4	30.1	56.4
InternVL3.5-8B	50.9	<u>98.1</u>	75.9	93.3	58.9	9.9	<u>85.9</u>	27.8	41.8	60.3
InternVL3.5-14B	56.6	98.2	74.4	89.9	74.4	11.7	86.7	33.7	47.1	63.6
GPT-5	44.3	72.4	81.8	89.3	27.5	15.8	54.9	68.3	41.2	54.9
Gemini 2.5 Pro	65.7	76.5	81.9	88.8	17.8	<u>21.1</u>	50.2	30.1	44.4	52.9
Medical-oriented models										
LLaVA-Med-v1.5-mistral-7b	44.3	89.9	55.3	58.1	22.4	10.2	41.8	66.6	26.4	48.2
HuatuoGPT-Vision-7B	45.4	82.5	<u>89.0</u>	88.3	60.6	13.6	65.5	69.2	<u>44.6</u>	48.9
medgemma-4b-it	<u>62.9</u>	76.4	82.5	84.8	55.4	30.6	54.2	77.4	36.8	51.5
Lingshu-7B	49.6	90.7	85.1	88.9	47.9	15.7	57.0	<u>78.8</u>	29.6	<u>60.4</u>
Lingshu-32B	45.2	75.5	92.1	<u>89.3</u>	54.5	14.4	46.4	88.8	31.7	59.8

Table 2: **Zero-shot accuracies (%) of VLMs on RadImageNet-VQA benchmark.** Results are reported across anatomy recognition, abnormality detection (*Abn*), and pathology identification using four question formats: *Open* (free-form), *Closed+* (always ‘yes’ as true answer), *Closed−* (always ‘no’), and *MC* (multiple-choice).

framework against ground-truth (Appendix A.3). For multiple-choice questions, we extract the predicted option letter with rule-based parser and compare it to the true answer.

4. Experiments

We evaluate state-of-the-art VLMs on RadImageNet-VQA under zero-shot and fine-tuned settings, and analyze text shortcuts, data-sampling strategies, and vision encoder choices.

4.1. Zero-Shot Evaluation

We benchmark a wide range of general-purpose and medical-oriented VLMs, including open-source models (e.g., LLaVA-OneVision, Qwen2.5-VL, InternVL) and proprietary systems (e.g., GPT-5, Gemini 2.5 Pro) on the full RadImageNet-VQA benchmark. Models generate free-text answers for open-ended questions and letter choices for multiple-choice questions. Accuracy is computed via exact match for closed and MC questions and LLM-as-a-judge scoring for open-ended responses. Results are displayed in Table 2.

First, pathology identification emerges as the primary bottleneck. While models perform well on anatomy recognition (e.g., InternVL3.5-8B (Chen et al., 2024) achieves 93.3% on anatomy multiple-choice), they struggle severely with fine-grained disease identification. Open-ended pathology questions prove particularly challenging, with most models scoring below 20% accuracy, and even the best performer on this task, MedGemma-4b (Sellergren et al., 2025), attains only 30.6%. These results underscore that discriminating between specific pathologies based on visual features remains a major unsolved problem.

Then, we observe a nuanced performance gap between general-purpose and medical-oriented models. General-purpose models achieve the highest average accuracy with In-

ternVL3.5-14B (63.6%) and lead in several categories, including abnormality detection and closed-ended pathology questions. Medical-oriented models outperform general ones only in some specific areas, such as MedGemma’s superior performance on open-ended pathology questions. This suggests that scale and broad visual-language pretraining still provide a stronger foundation than existing medical specialization strategies.

Finally, the results reveal that even the most capable general-purpose models lack essential radiologic skills. GPT-5 (27.5%) and Gemini 2.5 Pro (17.8%) perform poorly on abnormality detection, below random guessing. We intuit that aggressive safety alignment may suppress the model’s willingness to identify abnormalities, causing them to default to conservative “no abnormality” responses rather than rely on visual evidence.

4.2. Text shortcuts analysis

To assess whether RadImageNet-VQA can be solved from linguistic priors alone, we evaluate several VLMs in a text-only setting, where images are removed and models receive only the question. We compare performance on VQA-RAD (Lau et al., 2018), SLAKE (Liu et al., 2021), MMMU-Med-val (Yue et al., 2024), and RadImageNet-VQA benchmarks to quantify the extent to which different datasets allow text-based shortcutting. (see Figure 4)

On VQA-RAD and SLAKE, models reach 11.5–33% text-only accuracy, indicating that many answers can be partially inferred from linguistic regularities. By contrast, open-ended accuracy on RadImageNet-VQA drops to near-random levels (2–10%) across all models, confirming that shortcut cues are largely removed. The text-only evaluation also reveals differences in model behavior. Some models, e.g. Lingshu-7B and InternVL variants, maintain non-zero accuracy on RadImageNet-VQA, but qualitative inspection shows that these predictions reflect dataset priors rather than uncertainty awareness. Models rarely acknowledge missing visual information and instead commit to plausible guesses based on high-frequency anatomical regions or semantic variants of pathology terms, rather than image-grounded reasoning. Some examples are presented in Appendix B.4.

The same pattern holds for multiple-choice questions. On MMMU-Med-val, text-only accuracy remains noticeably above random, reflecting reliance on textual priors or weak distractor design. In contrast, RadImageNet-VQA collapses to the expected 25% baseline, indicating that models cannot exploit option distributions or question phrasing. We can further observe in Appendix 10 that residual correct answers are largely driven by conservative




Figure 4: **Text-only analysis of multiple VLMs’ accuracy** for open-ended and MC questions on RadImageNet-VQA, VQA-RAD, SLAKE, and MMU-Med-val.

defaults—such as selecting “no pathology seen” when no image is available—which yield correct predictions on normal cases but fail on abnormal ones. General-purpose models show a flatter spread across options, consistent with near-random guessing.

Overall, these findings show that RadImageNet-VQA substantially suppresses linguistic shortcuts in both open-ended and multiple-choice formats. When deprived of images, models fail to recover anatomical or pathological information and instead rely on dataset frequency or generic medical priors. This confirms that RadImageNet-VQA requires image-grounded interpretation rather than text-driven guessing.

4.3. Fine-tuning on RadImageNet-VQA

4.3.1. EXPERIMENTAL SETUP

We fine-tune two widely used open-source VLMs, LLaVA-OneVision (Li et al., 2024a) and Qwen2.5-VL-7B-Instruct (Bai et al., 2025), and a medical-oriented model, Lingshu-7B (Xu et al., 2025b). All models employ of a a SigLIP (Zhai et al., 2023) vision encoder that produces image embeddings. LLaVA-OneVision pairs it with a Qwen2-7B (Yang et al., 2024) language model while Qwen2.5-VL-7B is based on the more recent Qwen2.5-7B language model Lingshu-7B share the same architecture as Qwen2.5-VL-7B but has been pretrained and fine-tuned on data including medical data. This allows us to apply the same training recipes across architectures while comparing general-purpose versus medically-aligned initialization.

We follow the standard two-phase training paradigm used in recent multimodal tuning consisting of (i) an *alignment stage*, where the LLM is frozen and only the vision encoder and projection layers are updated to adapt to radiologic features; and (ii) *instruction tuning*, where the full model is fine-tuned on supervised VQA data. As training data, we assemble a multimodal corpus combining: (1) RadImageNet-VQA train split, (2) CT/MRI datasets converted into 2D VQA pairs—KiTS22 (Cardoso et al., 2017) and AbdomenAtlas (Li et al., 2024b)—and (3) existing radiologic VQA datasets—VQA-RAD (Lau et al., 2018), SLAKE (Liu et al., 2021)—, and the radiology subset of LLaVA-Med (Li et al., 2023a). Additional details and workflow are provided in Appendix C.

4.3.2. EVALUATION OF FINE-TUNED MODELS

Fine-tuning on the composed radiologic corpus yields consistent and substantial performance gains across all models. As shown in Figure 5, every point shifts above the diagonal, indicating reliable gains regardless of architecture, with average increases of +19.5% to +22.5% (Appendix 7). Task-level results reveal that anatomy recognition is nearly saturated after fine-tuning (98.5–99.4% on multiple-choice),




Figure 5: Comparison of base and fine-tuned accuracies on RadImageNet-VQA.




Figure 6: **Impact of vision encoder and sampling with LLaVA-OneVision.** Alternating sampling with standard SigLIP yields the strongest performance.

abnormality detection shows the largest relative gains (+16.4–39.8%, reaching 85.9 – 87.7%), while pathology identification remains the primary bottleneck, clustering at lower accuracy even post-training. These trends hold across model families, demonstrating both the generality of our training strategy and the limits of class-level supervision for fine-grained disease identification. Finally, Lingshu-7B and Qwen2.5-VL-7B converge to nearly identical performance, indicating that radiologic supervision, and not prior medical pretraining, is the main driver of downstream capability.

4.3.3. ABLATIONS ON VISION ENCODER AND DATA SAMPLING

We compare four variants of LLaVA-OneVision variants using the standard SigLIP encoder or MedSigLIP (Sellersgren et al., 2025), and with two different data-sampling strategy during fine-tuning. *Mixed sampling* blends sample from all datasets within each batch, whereas *alternating sampling* draws each batch from a single dataset and the source alternates across batches to balance exposure over time. Performance is measured across VQA-RAD, SLAKE, and RadImageNet-VQA and reported in Figure 6. Models initialized with standard SigLIP outperform those using the medically pre-trained MedSigLIP, which may indicate that pretraining with multiple sources of medical images, including textbooks and X-ray, does not translate into systematic gains for VQA with CT and MRI. Alternating sampling data strategy yields slightly stronger and more stable results than mixed sampling, particularly early in training, likely because it reduces inter-dataset interference under small-batch constraints.

5. Conclusion

In this work, we introduce RadImageNet-VQA, a large-scale CT/MRI dataset built through a rigorous QA-generation pipeline for radiologic VQA. It includes both a training set and a curated benchmark, covering multiple tasks and anatomy regions. Our experiments show that fine-grained pathology identification remains challenging for current models. Fine-tuning existing VLMs on RadImageNet-VQA yields substantial performance gains, while greatly reducing text-based shortcuts, which demonstrate its potential to be a valuable resource for developing and assessing stronger medical VLMs for radiology.

References

- Mistral AI. Mistral large 2. <https://docs.mistral.ai/models/mistral-large-2-1-24-11>, 2024. Model card, accessed 2025-02-11.
- Gorkem Can Ates, Yu Xin, Kuang Gong, and Wei Shao. Dcformer: Efficient 3d vision-language modeling with decomposed convolutions. *arXiv preprint arXiv:2502.05091*, 2025.
- Fan Bai, Yuxin Du, Tiejun Huang, Max Q-H Meng, and Bo Zhao. M3d: Advancing 3d medical image analysis with multi-modal large language models. *arXiv preprint arXiv:2404.00578*, 2024.
- Shuai Bai, Keqin Chen, Xuejing Liu, Jialin Wang, Wenbin Ge, Sibao Song, Kai Dang, Peng Wang, Shijie Wang, Jun Tang, et al. Qwen2. 5-vl technical report. *arXiv preprint arXiv:2502.13923*, 2025.
- Shruthi Bannur, Kenza Bouzid, Daniel C. Castro, Anton Schwaighofer, Anja Thieme, Sam Bond-Taylor, Maximilian Ilse, Fernando Pérez-García, Valentina Salvatelli, Harshita Sharma, Felix Meissen, Mercy Ranjit, Shaury Srivastav, Julia Gong, Noel C. F. Codella, Fabian Falck, Ozan Oktay, Matthew P. Lungren, Maria Teodora Wetscherek, Javier Alvarez-Valle, and Stephanie L. Hyland. Maira-2: Grounded radiology report generation, 2024.
- Louis Blankemeier, Joseph Paul Cohen, Ashwin Kumar, Dave Van Veen, Syed Jamal Safdar Gardezi, Magdalini Paschali, Zhihong Chen, Jean-Benoit Delbrouck, Eduardo Reis, Cesar Truys, et al. Merlin: A vision language foundation model for 3d computed tomography. *Research Square*, pages rs–3, 2024.
- M Jorge Cardoso, Tal Arbel, Gustavo Carneiro, Tanveer Syeda-Mahmood, João Manuel RS Tavares, Mehdi Moradi, Andrew Bradley, Hayit Greenspan, João Paulo Papa, Anant Madabhushi, et al. *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings*, volume 10553. Springer, 2017.
- Yixiong Chen, Wenjie Xiao, Pedro RAS Bassi, Xinze Zhou, Sezgin Er, Ibrahim Ethem Hamamci, Zongwei Zhou, and Alan Yuille. Are vision language models ready for clinical diagnosis? a 3d medical benchmark for tumor-centric visual question answering. *arXiv preprint arXiv:2505.18915*, 2025.
- Zhe Chen, Jiannan Wu, Wenhai Wang, Weijie Su, Guo Chen, Sen Xing, Muyan Zhong, Qinglong Zhang, Xizhou Zhu, Lewei Lu, 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*, pages 24185–24198, 2024.
- Corentin Dancette, Julien Khlaut, Antoine Saporta, Helene Philippe, Elodie Ferreres, Baptiste Callard, Théo Danielou, Léo Alberge, Léo Machado, Daniel Tordjman, Julie Dupuis, Korentin Le Floch, Jean Du Terrail, Mariam Moshiri, Laurent Dercle, Tom Boeken, Jules

- Gregory, Maxime Ronot, François Legou, Pascal Roux, Marc Sapoval, Pierre Manceron, and Paul Hérent. Curia: A multi-modal foundation model for radiology, 2025.
- Jean-Benoit Delbrouck, Pierre Chambon, Zhihong Chen, Maya Varma, Andrew Johnston, Louis Blankemeier, Dave Van Veen, Tan Bui, Steven Truong, and Curtis Langlotz. Radgraph-xl: A large-scale expert-annotated dataset for entity and relation extraction from radiology reports. In *Findings of the Association for Computational Linguistics: ACL 2024*, pages 12902–12915, 2024.
- Xiaotang Gai, Jiaxiang Liu, Yichen Li, Zijie Meng, Jian Wu, and Zuozhu Liu. 3d-rad: A comprehensive 3d radiology med-vqa dataset with multi-temporal analysis and diverse diagnostic tasks. *arXiv preprint arXiv:2506.11147*, 2025.
- Jeremy Irvin, Pranav Rajpurkar, Michael Ko, Yifan Yu, Silviana Ciurea-Ilcus, Chris Chute, Henrik Marklund, Behzad Haghighi, Robyn Ball, Katie Shpanskaya, et al. Chexpert: A large chest radiograph dataset with uncertainty labels and expert comparison. In *Proceedings of the AAAI conference on artificial intelligence*, volume 33, pages 590–597, 2019.
- Julien Khlaut, Elodie Ferreres, Daniel Tordjman, Hélène Philippe, Tom Boeken, Pierre Manceron, and Corentin Dancette. Radsam: Segmenting 3d radiological images with a 2d promptable model. *arXiv preprint arXiv:2504.20837*, 2025.
- Jason J Lau, Soumya Gayen, Asma Ben Abacha, and Dina Demner-Fushman. A dataset of clinically generated visual questions and answers about radiology images. *Scientific data*, 5(1):1–10, 2018.
- Bo Li, Yuanhan Zhang, Dong Guo, Renrui Zhang, Feng Li, Hao Zhang, Kaichen Zhang, Yanwei Li, Ziwei Liu, and Chunyuan Li. Llava-onevision: Easy visual task transfer. *arXiv preprint arXiv:2408.03326*, 2024a.
- Chunyuan Li, Cliff Wong, Sheng Zhang, Naoto Usuyama, Haotian Liu, Jianwei Yang, Tristan Naumann, Hoifung Poon, and Jianfeng Gao. Llava-med: Training a large language-and-vision assistant for biomedicine in one day. *Advances in Neural Information Processing Systems*, 36:28541–28564, 2023a.
- Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: bootstrapping language-image pre-training with frozen image encoders and large language models. In *Proceedings of the 40th International Conference on Machine Learning, ICML’23*. JMLR.org, 2023b.
- Wenxuan Li, Chongyu Qu, Xiaoxi Chen, Pedro RAS Bassi, Yijia Shi, Yuxiang Lai, Qian Yu, Huimin Xue, Yixiong Chen, Xiaorui Lin, et al. Abdomenatlas: A large-scale, detailed-annotated, & multi-center dataset for efficient transfer learning and open algorithmic benchmarking. *Medical Image Analysis*, 97:103285, 2024b.
- Bo Liu, Li-Ming Zhan, Li Xu, Lin Ma, Yan Yang, and Xiao-Ming Wu. Slake: A semantically-labeled knowledge-enhanced dataset for medical visual question answering. In *2021 IEEE 18th international symposium on biomedical imaging (ISBI)*, pages 1650–1654. IEEE, 2021.

- Haotian Liu, Chunyuan Li, Qingyang Wu, and Yong Jae Lee. Visual instruction tuning. *Advances in neural information processing systems*, 36:34892–34916, 2023.
- Xueyan Mei, Zelong Liu, Philip M Robson, Brett Marinelli, Mingqian Huang, Amish Doshi, Adam Jacobi, Chendi Cao, Katherine E Link, Thomas Yang, et al. Radimagenet: an open radiologic deep learning research dataset for effective transfer learning. *Radiology: Artificial Intelligence*, 4(5):e210315, 2022.
- Ankit Pal, Jung-Oh Lee, Xiaoman Zhang, Malaikannan Sankarasubbu, Seunghyeon Roh, Won Jung Kim, Meesun Lee, and Pranav Rajpurkar. Rexvqa: A large-scale visual question answering benchmark for generalist chest x-ray understanding. *arXiv preprint arXiv:2506.04353*, 2025.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International conference on machine learning*, pages 8748–8763. PmLR, 2021.
- Andrew Sellaergren, Sahar Kazemzadeh, Tiam Jaroensri, Atilla Kiraly, Madeleine Traverse, Timo Kohlberger, Shawn Xu, Fayaz Jamil, Cían Hughes, Charles Lau, et al. Medgemma technical report. *arXiv preprint arXiv:2507.05201*, 2025.
- Chaoyi Wu, Xiaoman Zhang, Ya Zhang, Yanfeng Wang, and Weidi Xie. Towards generalist foundation model for radiology by leveraging web-scale 2d&3d medical data. *arXiv preprint arXiv:2308.02463*, 2023.
- Yu Xin, Gorkem Can Ates, Kuang Gong, and Wei Shao. Med3dvlm: An efficient vision-language model for 3d medical image analysis. *arXiv preprint arXiv:2503.20047*, 2025.
- Justin Xu, Xi Zhang, Javid Abderezaei, Julie Bauml, Roger Boodoo, Fatemeh Haghighi, Ali Ganjizadeh, Eric Brattain, Dave Van Veen, Zaiqiao Meng, et al. Radeval: A framework for radiology text evaluation. In *Proceedings of the 2025 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 546–557, 2025a.
- Weiwen Xu, Hou Pong Chan, Long Li, Mahani Aljunied, Ruifeng Yuan, Jianyu Wang, Chenghao Xiao, Guizhen Chen, Chaoqun Liu, Zhaodonghui Li, et al. Lingshu: A generalist foundation model for unified multimodal medical understanding and reasoning. *arXiv preprint arXiv:2506.07044*, 2025b.
- An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, Guanting Dong, Haoran Wei, Huan Lin, Jialong Tang, Jialin Wang, Jian Yang, Jianhong Tu, Jianwei Zhang, Jianxin Ma, Jianxin Yang, Jin Xu, Jingren Zhou, Jinze Bai, Jinzheng He, Junyang Lin, Kai Dang, Keming Lu, Keqin Chen, Kexin Yang, Mei Li, Mingfeng Xue, Na Ni, Pei Zhang, Peng Wang, Ru Peng, Rui Men, Ruize Gao, Runji Lin, Shijie Wang, Shuai Bai, Sinan Tan, Tianhang Zhu, Tianhao Li, Tianyu Liu, Wenbin Ge, Xiaodong Deng, Xiaohuan Zhou, Xingzhang Ren, Xinyu Zhang, Xipin Wei, Xuancheng Ren, Xuejing Liu, Yang Fan, Yang Yao, Yichang Zhang, Yu Wan, Yunfei Chu, Yuqiong Liu, Zeyu Cui, Zhenru Zhang, Zhifang Guo, and Zhihao Fan. Qwen2 technical report, 2024.

- Nur Yildirim, Hannah Richardson, Maria Teodora Wetscherek, Junaid Bajwa, Joseph Jacob, Mark Ames Pinnock, Stephen Harris, Daniel Coelho De Castro, Shruthi Bannur, Stephanie Hyland, Pratik Ghosh, Mercy Ranjit, Kenza Bouzid, Anton Schwaighofer, Fernando Pérez-García, Harshita Sharma, Ozan Oktay, Matthew Lungren, Javier Alvarez-Valle, Aditya Nori, and Anja Thieme. Multimodal healthcare ai: Identifying and designing clinically relevant vision-language applications for radiology. In *Proceedings of the 2024 CHI Conference on Human Factors in Computing Systems*, 2024.
- Xiang Yue, Yuansheng Ni, Kai Zhang, Tianyu Zheng, Ruoqi Liu, Ge Zhang, Samuel Stevens, Dongfu Jiang, Weiming Ren, Yuxuan Sun, et al. Mmmu: A massive multi-discipline multimodal understanding and reasoning benchmark for expert agi. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9556–9567, 2024.
- Xiaohua Zhai, Basil Mustafa, Alexander Kolesnikov, and Lucas Beyer. Sigmoid loss for language image pre-training. In *Proceedings of the IEEE/CVF international conference on computer vision*, pages 11975–11986, 2023.
- Xiaoman Zhang, Chaoyi Wu, Ziheng Zhao, Weixiong Lin, Ya Zhang, Yanfeng Wang, and Weidi Xie. Pmc-vqa: Visual instruction tuning for medical visual question answering. *arXiv preprint arXiv:2305.10415*, 2023.
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. Judging llm-as-a-judge with mt-bench and chatbot arena. In *Proceedings of the 37th International Conference on Neural Information Processing Systems*, 2023.

Appendix A. Dataset details

A.1. RadImageNet-VQA anatomy regions and pathology taxonomy

The full taxonomy of CT and MRI anatomical regions paired with fine-grained pathology labels used for caption and VQA generation is reported in Table 3.

Anatomy	Pathologies
abdomen/pelvis	Abnormal entire organ, adrenal pathology, arterial pathology, ascites, biliary dilation, bladder pathology, bowel abnormality, bowel inflammation, bowel mass, degenerative changes, dilated urinary tract, enlarged organ, fat-containing tumor, gallbladder pathology, gallstone, intraperitoneal mass, liver disease, liver lesion, marrow abnormality, osseous neoplasm, ovarian pathology, pancreatic lesion, post-operative state, prostate lesion, renal lesion, soft tissue collection, soft tissue mass, splenic lesion, urolithiasis, uterine pathology
ankle foot	Achilles pathology, anterior talofibular ligament pathology, bone inflammation, calcaneofibular ligament pathology, chondral abnormality, coalition (bone fusion), deltoid pathology, extensor pathology, fat-containing tumor, flexor pathology, hematoma, intra-articular abnormality, Lisfranc joint injury, osseous disruption, osseous neoplasm, peroneal tendon pathology, plantar fascia pathology, plantar plate tear, post-operative state, soft tissue edema, soft tissue fluid, soft tissue mass, spring ligament injury, syndesmosis pathology
brain	Acute infarct, arteriovenous anomaly, chronic infarct, edema, extra-axial lesion, focal FLAIR hyperintensity, intra-articular abnormality, pituitary lesion, white matter changes
hip	Abductor pathology, capsular pathology, chondral pathology, hamstring pathology, hematoma, labral pathology, marrow inflammation, osseous disruption, osseous lesion, post-operative state, soft tissue edema, soft tissue fluid, soft tissue mass
knee	Anterior cruciate ligament pathology, bone inflammation, chondral abnormality, fibular collateral ligament pathology, fracture, hematoma, intra-articular abnormality, medial collateral ligament pathology, meniscal abnormality, muscle strain, patella pathology, post-operative ACL reconstruction, posterior cruciate ligament pathology, quadriceps pathology, soft tissue edema, soft tissue fluid collection, soft tissue mass
chest	Airspace opacity, bronchiectasis, interstitial lung disease, nodule, parenchyma destruction
shoulder	Acromioclavicular joint osteoarthritis, biceps pathology, calcific tendinosis, glenohumeral joint osteoarthritis, infraspinatus pathology, labral pathology, marrow inflammation, osseous lesion, post-operative state, soft tissue edema, soft tissue fluid, subscapularis pathology, supraspinatus pathology
spine	Cystic lesions, disc pathology, dural or epidural abnormality, facet arthropathy, foraminal pathology, osseous abnormality, scoliosis, spinal cord pathology

Table 3: Anatomical regions and corresponding pathologies in RadImageNet-VQA

A.2. Template set for VQA and captioning generation

Figure 7 illustrates the template families used to generate radiologic captions and VQA samples for all tasks and question types.

ANATOMY	ABNORMALITY
<ul style="list-style-type: none"> Does the picture contain the {anatomy}? Does this image show the {anatomy}? Is the {anatomy} part of this scan? Is the {anatomy} visible in the image? Can you see the {anatomy}? Is the {anatomy} present in this image? Can you identify the {anatomy} in this image? 	<ul style="list-style-type: none"> Is this image abnormal? Is there an abnormality present? Does this scan reveal any abnormalities? Does the {anatomy} look abnormal?
<ul style="list-style-type: none"> Which body part is visible in this image? Which part of the body does this image belong to? What organ is highlighted in this scan? What anatomical region is shown? What specific anatomical structure is highlighted? Name the organ featured in this image. 	<h3>PATHOLOGY</h3> <ul style="list-style-type: none"> Is there evidence of {pathology} in this image? Does the patient have {pathology}? Is {pathology} present? Are there signs of {pathology}? Is the {anatomy} affected by {pathology}?
<ul style="list-style-type: none"> Which of the following anatomical regions is depicted in this scan? What is the correct anatomical region for this scan? Select the correct anatomical region. 	<ul style="list-style-type: none"> What abnormality is present in this scan? What condition is affecting the {anatomy}?
<h3>CAPTIONING</h3> <div> <ul style="list-style-type: none"> Provide a thorough clinical description of this medical image. What findings and structures are visible in this scan? Describe the medical imaging features and anatomical details. Give a detailed diagnostic description of this image. What does this scan show in terms of anatomy and pathology? Describe the radiological findings and anatomical structures. Provide a comprehensive clinical description of this scan. What imaging characteristics and findings are present? Describe the diagnostic features and anatomical structures. Give a detailed medical imaging description of this scan. What anatomical and pathological findings are demonstrated? Describe the clinical imaging findings in this medical scan. Provide a complete radiological description of this image. What structures, abnormalities, and findings are visible? Describe the medical imaging characteristics and clinical findings. What does this image reveal about the anatomy and pathology? Describe the imaging findings, anatomical structures, and clinical features. Provide a detailed clinical and radiological description of this scan. What anatomical structures, pathological findings, and imaging features are shown? Describe the diagnostic imaging findings and anatomical characteristics. Give a comprehensive medical imaging description of this clinical scan. What anatomical features, pathological changes, and imaging characteristics are visible? Describe the clinical imaging findings, anatomical structures, and diagnostic features. Give a detailed medical description of the imaging findings and anatomical details. What does this scan reveal about the anatomical structures and pathological findings? Describe the diagnostic imaging characteristics, anatomical features, and clinical findings. Provide a comprehensive radiological description of the medical imaging findings. What anatomical structures, pathological abnormalities, and imaging features are demonstrated? Describe the diagnostic imaging characteristics of this scan. </div> <div> <ul style="list-style-type: none"> Describe this medical image. Provide a detailed description of this scan. What do you see in this medical image? Give a comprehensive description of this medical scan. What is shown in this medical image? Describe the contents of this scan. Provide a clinical description of this image. What can you observe in this medical image? Give a detailed account of what this image shows. Describe the imaging findings in this medical image. Provide a thorough description of this medical scan. What structures and findings are visible in this image? Give an overview of what this scan reveals. Describe the radiological findings in this image. Provide a complete description of this medical image. What anatomical and pathological features are shown? Describe the imaging characteristics of this scan. Give a detailed medical description of this image. What can be observed in this diagnostic image? Provide a comprehensive overview of this image. What does this scan show anatomically and pathologically? Describe the medical imaging features present. What anatomical structures and abnormalities are visible? Describe the radiological characteristics of this image. Give a thorough medical description of this scan. What anatomical and pathological details are shown? Describe the imaging findings and anatomical structures. Give a comprehensive medical description of this scan. What structures and findings can be identified? Describe the anatomical and pathological features. Provide a detailed radiological description of this image. Describe the clinical and anatomical features shown. Give a complete description of the medical imaging findings. Describe the clinical findings, anatomical structures, and diagnostic imaging characteristics. Give a thorough medical imaging description of the anatomical and pathological features. What anatomical structures and abnormalities are demonstrated? </div>	

Figure 7: Templates of questions for VQA samples and radiologic captions.

A.3. LLM-as-a-judge evaluation

We adopt the MedEvalKit framework¹ released with Lingshu-7B (Xu et al., 2025b) and use Mistral-Large 2.1 (AI, 2024) as the judge model.

```
Your task is to determine whether the user's answer is correct based on the provided questions
and standard answers (for example, if the user expresses a similar meaning to the standard
answer, or another interpretation of the standard answer, it is considered correct.)
The question is: {question}
The standard answer: {answer} The user's answer: {response}
Please strictly follow the following format for output (0 represents correct, 1 represents
incorrect): <think> your concise think step </think> <judge>0/1</judge>
For example:
<think> The standard answer is right, and the user's answer is right frontal lobe; they express
the same meaning, so it is correct. </think> <judge>0</judge>
```

Figure 8: Judgment prompt used in the LLM-as-a-judge module.

Appendix B. Additional Zero-Shot Analysis

B.1. Effect of “no pathology seen” distractors in pathology MC questions

To assess how the construction of multiple-choice distractors influences model behavior on RadImageNet-VQA, we vary the proportion of pathology questions whose answer set includes the option “no pathology seen”. We evaluate four representative VLMs—two medically oriented models (Lingshu-7B, MedGemma-4B-it) and two general-purpose models (InternVL3-8B, Qwen2.5-VL-7B)—under three configurations in which this option appears in 0%, 30%, or 100% of pathology multiple-choice questions.

As reported in Table 4, MedGemma and Lingshu show a progressive decrease in accuracy as the “no pathology” choice becomes more frequent, suggesting that they tend to assume the presence of pathology even when presented with an explicit normal alternative. In contrast, InternVL and Qwen remain comparatively stable or slightly improve, indicating that general models are more inclined to select this choice when uncertain or to rely on deduction from the visible options. Overall, these results show that distractor design can substantially alter model responses and that explicitly modelling the possibility of normal findings remains important for robust radiologic VQA.

Model	0%	30%	100%
General-purpose models			
InternVL3-8B	45.0	44.7	46.5
Qwen2.5-VL-7B	40.5	40.3	42.6
Medical specialized models			
Lingshu-7B	59.2	54.9	49.1
medgemma-4b-it	60.9	57.6	50.2

Table 4: Accuracy on pathology MC questions when varying the frequency of the “no pathology seen” option.

B.2. Per-pathology performance analysis

We observe in Table 2 that pathology identification is the most challenging task in the zero-shot setting; to further understand which specific conditions drive this difficulty, Figure 9 reports pathology-level performance by listing the ten best- and worst-solved categories. We

1. <https://github.com/alibaba-damo-academy/MedEvalKit>




Figure 9: Per-pathology multiple-choice accuracy across four VLMs on RadImageNet-VQA. *Left*: 10 best-solved pathologies; *right*: 10 worst-solved (≥ 10 samples per model).

find that common and visually salient abnormalities (e.g., organ enlargement, soft-tissue collections) are generally solved reliably, whereas fine-grained or anatomically localized findings—particularly musculoskeletal pathologies—remain difficult for all models. Although medically oriented models perform competitively on frequent categories, they do not consistently outperform general-purpose models on the harder ones, and InternVL often achieves the best results overall, suggesting that detailed disease characterization rather than abnormality detection constitutes the main bottleneck.

B.3. Extended results on text-only evaluation

Figure 4 show that text-only performance on RadImageNet-VQA drops close to the random baseline, suggesting that most multiple-choice questions cannot be solved without visual evidence. To better understand the residual accuracy that remains, Figure 10 breaks results down by abnormality status. When the image is removed, medical-oriented models such as MedGemma-4B almost always choose the “no pathology seen” option, which yields very high accuracy for normal cases but causes performance on abnormal cases to collapse. In contrast, general-purpose models (GPT-5 and




Figure 10: MC text-only accuracy on RadImageNet-VQA.

InternVL3.5-14B) do not rely on this conservative default and instead spread predictions across all choices, resulting in lower but more balanced scores. These results show that the small amount of text-only accuracy observed on RadImageNet-VQA primarily reflects conservative guessing strategies rather than the ability to infer pathology from the question text alone, further confirming that pathology identification requires image grounding.

B.4. Qualitative examples

The examples in Figure 11 illustrate how VQA-RAD and SLAKE contain strong textual cues that enable models to answer correctly even without visual input. In each case, the model produces essentially the same response whether it receives the image and question or the question alone, showing that the question text often provides enough information to determine the answer. For instance, in “Where does the left renal vein connect to?”, the mention of “renal vein” directly implies the inferior vena cava, a standard anatomical relationship independent of the specific CT slice. Likewise, questions such as “Which organ is part of the urinary system?” or “Which organs are part of the digestive system?” contain the key semantic identifiers—“urinary system”, “digestive system”—that uniquely determine the correct organ without requiring visual confirmation. Even pathology-oriented prompts (e.g., “What disease is shown on the right lung?”) can be shortcut through frequent-label priors, leading the model to answer “pneumonia” regardless of the underlying image.

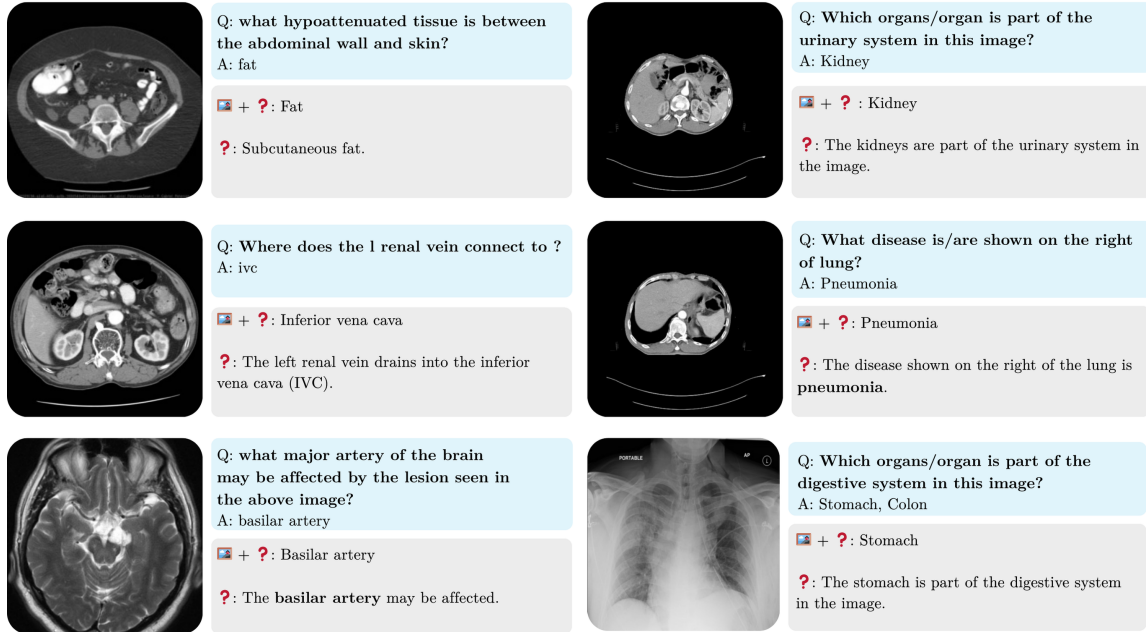


Figure 11: Examples of samples from VQA-RAD and SLAKE where current models do not need the image. The responses were obtain with MedGemma-4B.




Figure 12: Examples comparing Lingshu-7B’s answers when given the image and question versus the question alone on RadImageNet-VQA.

Building on the text-only results in Figure 4, where RadImageNet-VQA yields much lower shortcutability than VQA-RAD or SLAKE but where Lingshu-7B still attains 9.4% accuracy on open-ended questions, we further inspect its predictions with and without images. Figure 11 shows qualitative examples on RadImageNet-VQA for both pathology and anatomy questions. When the prompt explicitly mentions an organ or region (e.g., lung, knee, hip, brain), the text-only model often maps this cue to a high-frequency pathology for that location—such as pneumonia for the lung or osteoarthritis for the knee—which sometimes coincides with the image-based answer (e.g., pneumonia, cerebral/brain edema) but often diverges from it (e.g., osteoarthritis vs. meniscal tear, ascites vs. no abnormality, SCFE vs. avascular necrosis). This behavior indicates that, in the absence of visual evidence, Lingshu falls back on learned pathology priors tied to anatomical templates rather than genuinely reasoning about the specific case. For anatomy questions, multimodal predictions are typically correct and well localized (abdomen, chest, shoulder, ankle), whereas text-only outputs tend to overuse very frequent regions such as abdomen, chest, or liver, reflecting the underlying dataset distribution more than the individual question. Overall, these examples suggest that RadImageNet-VQA still leaves a narrow band of text-based shortcuts when anatomical regions are explicitly named, but that many failures in the multimodal setting stem from the intrinsic difficulty of fine-grained pathology recognition and from the model’s reliance on organ- and frequency-based priors under limited cues, rather than from purely linguistic shortcutting alone.

Appendix C. Extended Fine-Tuning Details

C.1. Training data composition

To train radiologic VLMs, we assemble a multimodal corpus composed of: (1) RadImageNet-VQA train split. (2) additional CT/MRI converted into 2D VQA pairs KiTS22 (Cardoso et al., 2017), and AbdomenAtlas (Li et al., 2024b), where 3D volumes are sliced into 2D images and paired with captions or VQA questions derived from organ and lesion annotations; and (3) existing radiologic VQA datasets such as VQA-RAD (Lau et al., 2018), SLAKE (Liu et al., 2021), and radiology subset of LLaVA-Med (Li et al., 2023a).

Dataset	# Images	Alignment	Instruction
RadImageNet (Mei et al., 2022)	750K	750K	6.75M
KiTS22 (Cardoso et al., 2017)	4.6K	4.6K	32.5K
AbdomenAtlas (Li et al., 2024b)	43.7K	43.7K	211K
LLaVA-Med (Li et al., 2023a)	38.4K	13.8K	72.3K
VQA-RAD (Lau et al., 2018)	0.3K	–	1.8K
SLAKE (Liu et al., 2021)	0.4K	–	4.9K
Total	837.5K	812.1K	7.07M

Table 5: Summary of datasets used for alignment and instruction tuning, including the number of images and approximate QA pair counts.

Together, these three branches feed into a unified training corpus comprising medical instruction data and medical alignment data, as depicted in Figure 13.




Figure 13: The overall data curation pipeline of radiologic multimodal data

C.2. Implementation Details

We fine-tune all models described in Section 4.3 with 4x **NVIDIA H100 GPUs**. Training follows a two-stage procedure: (1) a *visual alignment* phase in which only the vision tower and projector are updated while the language model remains frozen, and (2) an *instruction tuning* phase where all components are trainable. Evaluation uses a unified decoding configuration and a fixed LLM-as-a-judge setup for all models. The full set of hyperparameters and evaluation settings is summarized in Table 6.

C.3. Full fine-tuning results

Table 7 reports the exact fine-tuned values and deltas for each task and question type on RadImageNet-VQA. This shows that all models benefit substantially from radiologic instruction tuning, with very large gains on abnormality and open-ended pathology questions and near-ceiling performance on anatomy.

Category	Configuration
Training stages	Visual alignment (1 epoch) Instruction tuning (2 epochs)
Batching and scheduling	Global batch size: 4 (per-GPU batch = 1, gradient accumulation = 4) Learning rate: 1e-5 (LLM), 2e-6 (vision tower) Warm-up: 3% of total steps LR schedule: cosine decay
Optimization	Optimizer: AdamW ($\beta_1 = 0.9$, $\beta_2 = 0.999$) Weight decay: 0.0
Evaluation parameters	Max new tokens: 8192 Decoding: temperature = 0, top-p = 0.0001 Repetition penalty: 1.0 Judge backend: <code>mistral-large-latest</code> Random seed: 42.

Table 6: Training and evaluation configuration for RadImageNet-VQA experiments.

Task	LLaVA-One-Vision	Lingshu-7B	QwenVL-2.5-7B
Anatomy			
Open	99.4 $\uparrow(+51.0)$	99.4 $\uparrow(+49.8)$	99.3 $\uparrow(+61.8)$
Closed+	87.4 $\uparrow(+4.7)$	95.2 $\uparrow(+4.5)$	93.4 $\uparrow(+8.5)$
Closed-	97.8 $\uparrow(+16.5)$	97.6 $\uparrow(+12.5)$	97.3 $\uparrow(+18.2)$
MC	99.4 $\uparrow(+10.7)$	98.9 $\uparrow(+10.0)$	98.5 $\uparrow(+18.0)$
Abnormality			
Closed	87.6 $\uparrow(+37.8)$	87.7 $\uparrow(+39.8)$	85.9 $\uparrow(+16.4)$
Pathology			
Open	42.6 $\uparrow(+26.6)$	41.2 $\uparrow(+25.5)$	39.2 $\uparrow(+29.4)$
Closed+	67.2 $\uparrow(+11.9)$	84.3 $\uparrow(+27.3)$	74.7 $\uparrow(+5.5)$
Closed-	77.2 $\uparrow(+15.9)$	65.3 $\downarrow(-13.5)$	68.1 $\uparrow(+20.7)$
MC	60.1 $\uparrow(+26.5)$	49.9 $\uparrow(+20.3)$	53.4 $\uparrow(+23.3)$
Avg	79.9 $\uparrow(+22.4)$	79.9 $\uparrow(+19.5)$	78.9 $\uparrow(+22.5)$

Table 7: Fine-tuned performance on RadImageNet-VQA with delta values relative to zero-shot baselines. Improvements are shown in green and regressions in red.